

# Programming Exercise 1: Linear Regression

Carson Slater  
*Deep Learning - Spring 2026*

February 5, 2026

## 1 Introduction

The primary objective of this exercise is to implement linear regression to predict profits for a food truck based on city population (univariate) and house prices based on size and number of bedrooms (multivariate).

## 2 Univariate Linear Regression

### 2.1 Warm-up Exercise

The first task involves implementing a simple function that returns a 5x5 identity matrix.

```
1 function A = warmUpExercise()  
2 %WARMUPEXERCISE Returns the 5x5 identity matrix  
3 A = eye(5);  
4 end
```

Listing 1: warmUpExercise.m

### 2.2 Plotting the Data and Linear Fit

Visualizing the dataset along with the regression line allows for a qualitative assessment of the model's accuracy. Figure 1 shows the scatter plot of city population and profit, overlaid with the final linear fit.

### 2.3 Cost Function and Gradient Descent

The cost function  $J(\theta)$  is computed as the mean squared error. Gradient descent is then used to minimize  $J(\theta)$ .

```
1 function J = computeCost(X, y, theta)  
2 %COMPUTECOST Compute cost for linear regression  
3 m = length(y); % number of training examples  
4 predictions = X * theta;  
5 sqErrors = (predictions - y) .^ 2;  
6 J = 1 / (2 * m) * sum(sqErrors);  
7 end
```

Listing 2: computeCost.m

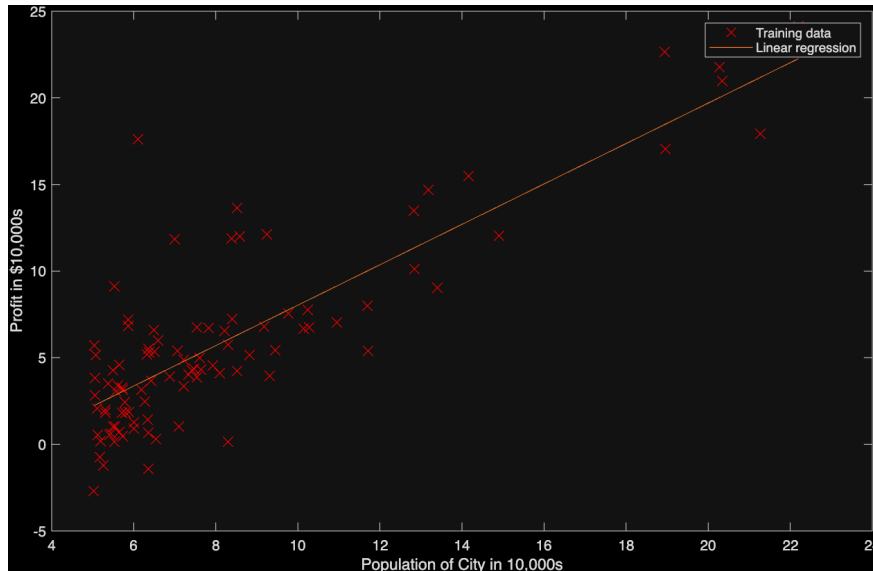


Figure 1: Scatter plot of training data with the linear regression fit.

```

1 function [theta, J_history] = gradientDescent(X, y, theta, alpha,
        num_iters)
2 %GRADIENDESCENT Performs gradient descent to learn theta
3 m = length(y);
4 J_history = zeros(num_iters, 1);
5
6 for iter = 1:num_iters
7     errors = X * theta - y;
8     gradient = (1 / m) * (X' * errors);
9     theta = theta - alpha * gradient;
10    J_history(iter) = computeCost(X, y, theta);
11 end
12 end

```

Listing 3: gradientDescent.m

Running the algorithm with  $\theta$  initialized to zeros and a learning rate  $\alpha = 0.01$  yields:

- Initial Cost: 32.0727
- Optimized  $\theta$ :  $[-3.6303, 1.1664]$

## 2.4 Visualizing the Cost Function

Figures 3 and 4 show the surface and contour plots of the cost function  $J(\theta)$  over a range of parameter values. The surface plot provides a 3D view of the error landscape, while the contour plot highlights the global minimum.

## 3 Multivariate Linear Regression

In this section, we predict house prices using two features: house size and number of bedrooms.




Figure 2: Surface plot of the cost function  $J(\theta_0, \theta_1)$ .

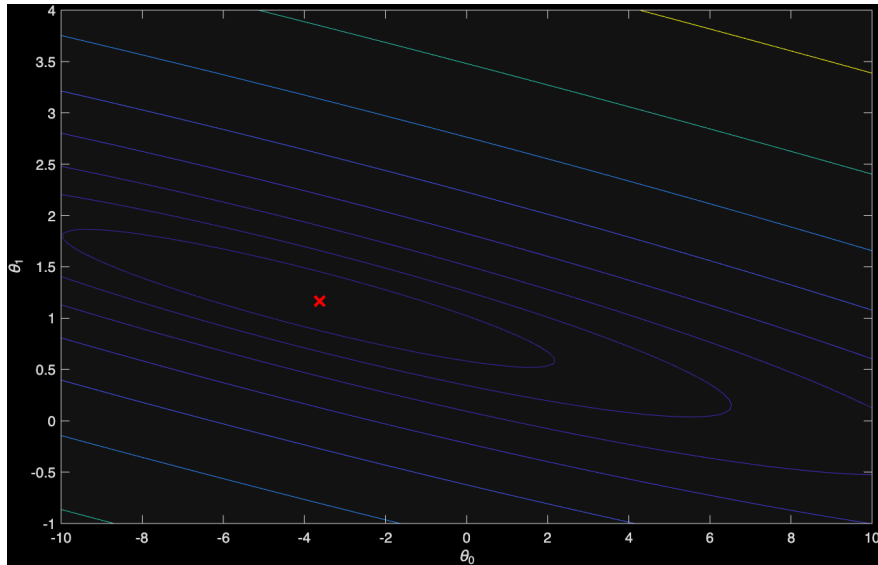


Figure 3: Contour plot of  $J(\theta_0, \theta_1)$ . The 'x' marks the global minimum found by gradient descent.

### 3.1 Feature Normalization

Since the features differ significantly in scale (house size in sq-ft vs. number of bedrooms), normalization is applied to ensure gradient descent converges efficiently.

```
1 function [X_norm, mu, sigma] = featureNormalize(X)
2     mu = mean(X, 1);
3     sigma = std(X, 0, 1);
4     X_norm = (X - mu) ./ sigma;
5 end
```

Listing 4: featureNormalize.m

### 3.2 Multi-variable Gradient Descent

The multivariate cost function and gradient descent are implemented below. These functions handle an arbitrary number of features using vectorized operations.

```
1 function J = computeCostMulti(X, y, theta)
2 %COMPUTECOSTMULTI Compute cost for linear regression with multiple
   variables
3     m = length(y);
4     predictions = X * theta;
5     sqErrors = (predictions - y) .^ 2;
6     J = 1 / (2 * m) * sum(sqErrors);
7 end
```

Listing 5: computeCostMulti.m

```
1 function [theta, J_history] = gradientDescentMulti(X, y, theta, alpha,
   num_iters)
2 %GRADIENDESCENTMULTI Performs gradient descent to learn theta
3     m = length(y);
4     J_history = zeros(num_iters, 1);
5
6     for iter = 1:num_iters
7         errors = X * theta - y;
8         gradient = (1 / m) * (X' * errors);
9         theta = theta - alpha * gradient;
10        J_history(iter) = computeCostMulti(X, y, theta);
11    end
12 end
```

Listing 6: gradientDescentMulti.m

### 3.3 Gradient Descent Convergence

The convergence of gradient descent for the multivariate case is tracked via the history of the cost function  $J(\theta)$ , ensuring that the objective value decreases steadily with each iteration.

### 3.4 Normal Equations

The normal equation provides a closed-form solution to minimizing  $J(\theta)$ , which does not require feature scaling.




Figure 4: Convergence of the cost function  $J$  over 400 iterations for the multivariate case.

```

1 function [theta] = normalEqn(X, y)
2   theta = pinv(X' * X) * X' * y;
3 end

```

Listing 7: normalEqn.m

Using both methods, we predict the price for a 1650 sq-ft house with 3 bedrooms to be approximately **\$293,081.46**.

## 4 Main Execution Scripts

The following scripts drive the execution of the functions implemented above.

```

1 %% Machine Learning Online Class - Exercise 1: Linear Regression
2 % Initialization
3 clear ; close all; clc
4
5 %% Part 1: Basic Function
6 fprintf('Running warmUpExercise ... \n');
7 warmUpExercise()
8
9 %% Part 2: Plotting
10 data = load('ex1data1.txt');
11 X = data(:, 1); y = data(:, 2);
12 m = length(y);
13 plotData(X, y);
14
15 %% Part 3: Cost and Gradient descent
16 X = [ones(m, 1), data(:,1)];
17 theta = zeros(2, 1);
18 iterations = 1500;
19 alpha = 0.01;
20 J = computeCost(X, y, theta);

```

```

21 theta = gradientDescent(X, y, theta, alpha, iterations);
22
23 % Plot the linear fit
24 hold on;
25 plot(X(:,2), X*theta, '-')
26 legend('Training data', 'Linear regression')
27 hold off

```

Listing 8: ex1.m (Univariate Part)

```

1 %% Machine Learning Online Class
2 % Exercise 1: Linear regression with multiple variables
3
4 %% Initialization
5 clear ; close all; clc
6
7 %% Part 1: Feature Normalization
8 data = load('ex1data2.txt');
9 X = data(:, 1:2);
10 y = data(:, 3);
11 m = length(y);
12
13 [X mu sigma] = featureNormalize(X);
14 X = [ones(m, 1) X];
15
16 %% Part 2: Gradient Descent
17 alpha = 0.3;
18 num_iters = 400;
19 theta = zeros(3, 1);
20 [theta, J_history] = gradientDescentMulti(X, y, theta, alpha, num_iters);
21
22 %% Part 3: Normal Equations
23 data = csvread('ex1data2.txt');
24 X = data(:, 1:2);
25 y = data(:, 3);
26 X = [ones(m, 1) X];
27 theta = normalEqn(X, y);

```

Listing 9: ex1\_multi.m (Multivariate Part)

## 5 Execution Output

The following is the terminal output from running the main execution scripts `ex1.m` and `ex1_multi.m`, demonstrating successful completion of both parts of the exercise.

```

1 Running warmUpExercise ...
2 5x5 Identity Matrix:
3 Program paused. Press enter to continue.
4 Plotting Data ...
5 Program paused. Press enter to continue.
6
7 Testing the cost function ...
8 With theta = [0 ; 0]
9 Cost computed = 32.072734

```

```

10 Expected cost value (approx) 32.07
11
12 With theta = [-1 ; 2]
13 Cost computed = 54.242455
14 Expected cost value (approx) 54.24
15 Program paused. Press enter to continue.
16
17 Running Gradient Descent ...
18 Theta found by gradient descent:
19 -3.630291
20 1.166362
21 Expected theta values (approx)
22 -3.6303
23 1.1664
24
25 For population = 35,000, we predict a profit of 4519.767868
26 For population = 70,000, we predict a profit of 45342.450129
27 Program paused. Press enter to continue.
28 Visualizing J(theta_0, theta_1) ...

```

Listing 10: ex1.m Output

```

1 Loading data ...
2 First 10 examples from the dataset:
3 x = [2104 3], y = 399900
4 x = [1600 3], y = 329900
5 x = [2400 3], y = 369000
6 x = [1416 2], y = 232000
7 x = [3000 4], y = 539900
8 x = [1985 4], y = 299900
9 x = [1534 3], y = 314900
10 x = [1427 3], y = 198999
11 x = [1380 3], y = 212000
12 x = [1494 3], y = 242500
13 Program paused. Press enter to continue.
14 Normalizing Features ...
15 Running gradient descent ...
16 Theta computed from gradient descent:
17 340412.659574
18 110631.050279
19 -6649.474271
20
21 Predicted price of a 1650 sq-ft, 3 br house (using gradient descent):
22 $293081.464335
23 Program paused. Press enter to continue.
24 Solving with normal equations...
25 Theta computed from the normal equations:
26 89597.909543
27 139.210674
28 -8738.019112
29
30 Predicted price of a 1650 sq-ft, 3 br house (using normal equations):
31 $293081.464335

```

Listing 11: ex1\_multi.m Output

## 6 Conclusion

The implementation successfully demonstrates the application of linear regression. Gradient descent and the normal equations provide consistent results, confirming the correctness of the vectorized MATLAB implementation.