

Programming Exercise 2: Logistic Regression

Carson Slater
Deep Learning - Spring 2026

February 5, 2026

1 Introduction

The primary objective of this exercise is to implement logistic regression and regularized logistic regression. Logistic regression is a classification algorithm used when the target variable is categorical. This report explores two cases: one linearly separable classroom dataset and a non-linearly separable microchip dataset that requires feature mapping and regularization.

2 Logistic Regression

In this part of the exercise, we build a logistic regression model to predict whether a student gets admitted into a university based on their results of two exams.

2.1 Sigmoid Function

The logistic regression hypothesis is defined as:

$$h_{\theta}(x) = g(\theta^T x)$$

where g is the sigmoid function:

$$g(z) = \frac{1}{1 + e^{-z}}$$

The implementation in `sigmoid.m` is as follows:

```
1 function g = sigmoid(z)
2 %SIGMOID Compute sigmoid function
3 g = 1 ./ (1 + exp(-z));
4 end
```

Listing 1: `sigmoid.m`

2.2 Cost Function and Gradient

The cost function for logistic regression is:

$$J(\theta) = \frac{1}{m} \sum_{i=1}^m [-y^{(i)} \log(h_{\theta}(x^{(i)})) - (1 - y^{(i)}) \log(1 - h_{\theta}(x^{(i)}))]$$

The gradient of the cost is a vector where the j^{th} element is defined as:

$$\frac{\partial J(\theta)}{\partial \theta_j} = \frac{1}{m} \sum_{i=1}^m (h_{\theta}(x^{(i)}) - y^{(i)}) x_j^{(i)}$$

The implementation in `costFunction.m` is as follows:

```
1 function [J, grad] = costFunction(theta, X, y)
2   m = length(y);
3   h = sigmoid(X * theta);
4   J = (1 / m) * sum(-y .* log(h) - (1 - y) .* log(1 - h));
5   grad = (1 / m) * (X' * (h - y));
6 end
```

Listing 2: `costFunction.m`

2.3 Learning Parameters using `fminunc`

Instead of gradient descent, we use MATLAB's `fminunc` to find the optimal parameters θ that minimize the cost function. This built-in function is more efficient and does not require choosing a learning rate.

2.4 Results and Visualization

After optimization, the model achieves a training accuracy of **89%**. Figure 1 shows the training data, and Figure 2 displays the learned decision boundary.




Figure 1: Scatter plot of university admission data.

3 Regularized Logistic Regression

In the second part, we implement regularized logistic regression to predict whether microchips from a fabrication plant pass quality assurance (QA) tests.




Figure 2: Decision boundary for logistic regression on the university admission dataset.

3.1 Feature Mapping

The data is not linearly separable, so we map the features into all polynomial terms of x_1 and x_2 up to the sixth power to create more features.

3.2 Regularized Cost Function and Gradient

The cost function for regularized logistic regression is:

$$J(\theta) = \frac{1}{m} \sum_{i=1}^m [-y^{(i)} \log(h_{\theta}(x^{(i)})) - (1 - y^{(i)}) \log(1 - h_{\theta}(x^{(i)}))] + \frac{\lambda}{2m} \sum_{j=2}^n \theta_j^2$$

Note that we do not regularize θ_0 . The partial derivative for the regularized gradient is:

$$\frac{\partial J(\theta)}{\partial \theta_0} = \frac{1}{m} \sum_{i=1}^m (h_{\theta}(x^{(i)}) - y^{(i)}) x_0^{(i)} \quad \text{for } j = 0$$

$$\frac{\partial J(\theta)}{\partial \theta_j} = \left(\frac{1}{m} \sum_{i=1}^m (h_{\theta}(x^{(i)}) - y^{(i)}) x_j^{(i)} \right) + \frac{\lambda}{m} \theta_j \quad \text{for } j \geq 1$$

The implementation in `costFunctionReg.m` is as follows:

```

1 function [J, grad] = costFunctionReg(theta, X, y, lambda)
2   m = length(y);
3   h = sigmoid(X * theta);
4   theta_reg = [0; theta(2:end)];
5   J = (1 / m) * sum(-y .* log(h) - (1 - y) .* log(1 - h)) + (lambda / (2 *
6     m)) * sum(theta_reg .^ 2);
7   grad = (1 / m) * (X' * (h - y)) + (lambda / m) * theta_reg;
8 end

```

Listing 3: `costFunctionReg.m`

3.3 Results and Regularization Effects

With $\lambda = 1$, the model found a decision boundary (Figure 3) that balances fitting the data and preventing overfitting, achieving a training accuracy of approximately **83.1%**.

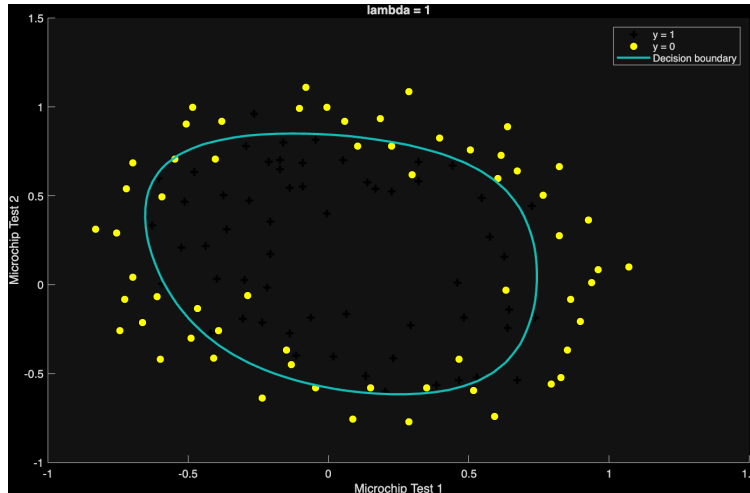


Figure 3: Decision boundary with $\lambda = 1$ for the regularized logistic regression case.

4 Execution Scripts

The following scripts provide the full context for the experiments conducted.

```

1 %% Initialization
2 clear ; close all; clc
3 data = load('ex2data1.txt');
4 X = data(:, [1, 2]); y = data(:, 3);
5 [m, n] = size(X);
6 X = [ones(m, 1) X];
7 initial_theta = zeros(n + 1, 1);
8 [cost, grad] = costFunction(initial_theta, X, y);
9 options = optimset('GradObj', 'on', 'MaxIter', 400);
10 [theta, cost] = fminunc(@(t)(costFunction(t, X, y)), initial_theta,
    options);
11 plotDecisionBoundary(theta, X, y);

```

Listing 4: ex2.m (Simplified)

```

1 %% Initialization
2 clear ; close all; clc
3 data = load('ex2data2.txt');
4 X = data(:, [1, 2]); y = data(:, 3);
5 X = mapFeature(X(:,1), X(:,2));
6 initial_theta = zeros(size(X, 2), 1);
7 lambda = 1;
8 [cost, grad] = costFunctionReg(initial_theta, X, y, lambda);
9 options = optimset('GradObj', 'on', 'MaxIter', 400);
10 [theta, J, exit_flag] = fminunc(@(t)(costFunctionReg(t, X, y, lambda)),
    initial_theta, options);
11 plotDecisionBoundary(theta, X, y);

```

Listing 5: ex2_reg.m (Simplified)

5 Console Output

The following is the console output from running the main scripts, verifying the correctness of the cost function and optimized parameters.

```
1 Plotting data with + indicating (y = 1) examples and o indicating (y = 0)
  examples.
2
3 Program paused. Press enter to continue.
4 Cost at initial theta (zeros): 0.693147
5 Expected cost (approx): 0.693
6 Gradient at initial theta (zeros):
7   -0.100000
8   -12.009217
9   -11.262842
10 Expected gradients (approx):
11   -0.1000
12   -12.0092
13   -11.2628
14
15 Cost at test theta: 0.218330
16 Expected cost (approx): 0.218
17 Gradient at test theta:
18   0.042903
19   2.566234
20   2.646797
21 Expected gradients (approx):
22   0.043
23   2.566
24   2.647
25
26 Program paused. Press enter to continue.
27
28 Local minimum found.
29
30 Optimization completed because the size of the gradient is less than
31 the value of the optimality tolerance.
32
33 Cost at theta found by fminunc: 0.203498
34 Expected cost (approx): 0.203
35 theta:
36   -25.161343
37   0.206232
38   0.201472
39 Expected theta (approx):
40   -25.161
41   0.206
42   0.201
43
44 Program paused. Press enter to continue.
45 For a student with scores 45 and 85, we predict an admission probability
  of 0.776291
46 Expected value: 0.775 +/- 0.002
47
```

```
48 Train Accuracy: 89.000000
49 Expected accuracy (approx): 89.0
```

Listing 6: ex2.m Console Output

```
1 Cost at initial theta (zeros): 0.693147
2 Expected cost (approx): 0.693
3 Gradient at initial theta (zeros) - first five values only:
4 0.008475
5 0.018788
6 0.000078
7 0.050345
8 0.011501
9 Expected gradients (approx) - first five values only:
10 0.0085
11 0.0188
12 0.0001
13 0.0503
14 0.0115
15
16 Program paused. Press enter to continue.
17
18 Cost at test theta (with lambda = 10): 3.164509
19 Expected cost (approx): 3.16
20 Gradient at test theta - first five values only:
21 0.346045
22 0.161352
23 0.194796
24 0.226863
25 0.092186
26 Expected gradients (approx) - first five values only:
27 0.3460
28 0.1614
29 0.1948
30 0.2269
31 0.0922
32
33 Program paused. Press enter to continue.
34
35 Local minimum found.
36
37 Optimization completed because the size of the gradient is less than
38 the value of the optimality tolerance.
39
40 Train Accuracy: 83.050847
41 Expected accuracy (with lambda = 1): 83.1 (approx)
```

Listing 7: ex2_reg.m Console Output

6 Conclusion

The implementation of logistic regression and regularization was successful. The models were able to capture the underlying patterns in both linearly separable and non-linearly separable datasets.

The optimized parameters yields high accuracy, and the regularized results demonstrate how the model complexity can be controlled to improve generalization.