

Section 05: Subgradients, Convexity and Gradient Descent

1. K-fold Cross-Validation (Demonstrative code)

```
1 # Given dataset of 1000-by-50 feature
2 # matrix X, and 1000-by-1 labels vector
3 import numpy as np
4
5 X = np.random.random((1000,50))
6 y = np.random.random((1000,))
7
8 def fit(Xin, Yin, lbda):
9     mu = np.mean(Xin, axis=0)
10    Xin = Xin - mu
11    w = np.linalg.solve(np.dot(Xin.T, Xin) + lbda, np.dot(Xin.T, Yin))
12    b = np.mean(Yin) - np.dot(w, mu)
13    return w, b
14
15
16 def predict(w, b, Xin):
17    return np.dot(Xin, w) + b
18
19
20 # Note: X, y are all the train data and label for the entire experiments
21 N_TRAIN = 1000
22 idx = np.random.permutation(N_TRAIN)
23 K_FOLD = 5
24
25 NON_TEST = idx[0: 9 * N_TRAIN // 10]
26 N_PER_FOLD = len(NON_TEST) // K_FOLD
27 TEST = idx[9 * N_TRAIN // 10::]
28
29 # candidates regularization coefficient to choose from
30 lbdas = [0.1, 0.2, 0.3]
31 err = np.zeros(len(lbdas))
32
33
34 for lbda_idx, lbda in enumerate(lbdas):
35     for i in range(K_FOLD):
36         # CRUCIAL: we use slicing to calculate the index train set and val set should use!
37         # Using the ith fold as the validation set
38         VAL = NON_TEST[i * N_PER_FOLD:(i+1) * N_PER_FOLD]
39         # Using the rest as the train set
40         TRAIN = np.concatenate((NON_TEST[:i * N_PER_FOLD], NON_TEST[(i + 1) * N_PER_FOLD:]))
41
42         ytrain = y[TRAIN]
43         Xtrain = X[TRAIN]
44         yval = y[VAL]
45         Xval = X[VAL]
46
47         w, b = fit(Xtrain, ytrain, lbda)
48         yval_hat = predict(w, b, Xval)
49         # accumulate error from this fold of validation set
50         err[lbda_idx] += np.mean((yval_hat - yval)**2)
51
52     # calculate the error for the k-fold validation
53     err[lbda_idx] /= K_FOLD
```

```
54
55     lbda_best = lbdas[np.argmin(err)]
56
57     Xtot = np.concatenate((Xtrain, Xval), axis=0)
58     ytot = np.concatenate((ytrain, yval), axis=0)
59
60     w, b = fit(Xtot, ytot, lbda_best)
61
62     ytest = y[TEST]
63     Xtest = X[TEST]
64
65     ytot_hat = predict(w, b, Xtot)
66     train_error = np.mean((ytot_hat - ytot) ** 2)
67     ytest_hat = predict(w, b, Xtest)
68     test_error = np.mean((ytest_hat - ytest) ** 2)
69
70     print('Best choice of lambda = ', lbda_best)
71     print('Train error = ', train_error)
72     print('Test error = ', test_error)
```

2. Extra: Lasso and CV

```
1 LR = 0.01
2 NUM_ITERATIONS = 500
3
4 def predict(w, b, Xin):
5     return np.dot(Xin, w) + b
6
7 def fit(Xin, Yin, l1_penalty) :
8     # no_of_training_examples, no_of_features
9     m, n = Xin.shape
10
11     # weight initialization
12     w = np.zeros(n)
13     b = 0
14
15     # gradient descent learning
16     for i in range(NUM_ITERATIONS) :
17         w, b = update_weights(w, b, Xin, Yin, l1_penalty)
18
19     return w, b
20
21 def update_weights(w, b, Xin, Yin, l1_penalty) :
22     m, n = Xin.shape
23     Y_pred = predict(w, b, Xin)
24
25     # calculate gradients
26     dW = np.zeros(n)
27     for j in range(n) :
28         if w[j] > 0 :
29             dW[j] = ( - ( 2 * ( Xin[:, j] ).dot(Yin - Y_pred))
30                     + l1_penalty ) / m
31         else :
32             dW[j] = ( - ( 2 * ( Xin[:, j] ).dot(Yin - Y_pred))
33                     - l1_penalty ) / m
34
35     db = - 2 * np.sum(Yin - Y_pred) / m
36
37     # update weights
38     w = w - LR * dW
39     b = b - LR * db
40
41     return w, b
42
43 def rmse_lasso(w, b, Xin, Yin):
44     Y_pred = predict(w, b, Xin)
45     return rmse(Yin, Y_pred)
46
47 def rmse(a, b):
48     return np.sqrt(np.mean(np.square(a - b)))
49
50 l1_penalties = [] #some values here
51 err = np.zeros(len(l1_penalties))
52
53 k_fold = 10
54 num_samples = len(X) / k_fold
55
56 for idx, l1_penalty in enumerate(l1_penalties):
57     for k in range(k_fold): #10-fold CV
58         x_train_fold = X[0: k*num_samples] + X[(k+1)*num_samples: ]
59         y_train_fold = Y[0: k*num_samples] + Y[(k+1)*num_samples: ]
```

```

60     x_val_fold = X[k*num_samples : (k+1)*num_samples]
61     y_val_fold = Y[k*num_samples : (k+1)*num_samples]
62
63
64     w, b = fit(x_train_fold, y_train_fold)
65
66     #accumulate error from this fold of validation set
67     err[idx] += rmse_lasso(w, b, x_val_fold, y_val_fold)
68
69     #calculate error for kth fold
70     err[idx]/=k_fold
71
72 l1_penalty_best = l1_penalties[np.argmin(err)]
73
74 print('Best choice of l1_penalty = ', l1_penalty_best)

```

3. Subgradients

We start with the definition of subgradients before discussing the motivation and its usefulness.

Definition 1 (subgradients). A vector $g \in \mathbb{R}^d$ is a subgradient of a convex function $f : D \rightarrow \mathbb{R}$ at $x \in D \subseteq \mathbb{R}^d$ if

$$f(y) \geq f(x) + g^T(y - x) \quad \text{for all } y \in D.$$

One interpretation of subgradient g is that the affine function (of y) $f(x) + g^T(y - x)$ is a global underestimator of f . Note that if a convex function f is differentiable at x (i.e., $\nabla f(x)$ exists), then $f(y) \geq f(x) + \nabla f(x)^T(y - x)$ is true for all $y \in D$, meaning that $\nabla f(x)$ is a subgradient of f at x . But a subgradient can exist even when f is not differentiable at x .

- (a) Why are subgradients useful in optimization? If $g = 0$ is a subgradient of a function f at x^* , what does it imply?
- (b) What are the subgradients of $f(x) = \max(x, x^2)$ at 0, with $x \in \mathbb{R}$? (Hint: draw a picture and note that subgradients at a point might not be unique)
- (c) Some important results about subgradients are
 - If f is convex, then a subgradient of f at $x \in \text{int}(D)$ (interior of the domain of f) always exists.
 - If f is convex, then f is differentiable at x if and only if $\nabla f(x)$ is the only subgradient of f at x .
 - A point x^* is a global minimizer of a function f (not necessarily convex) if and only if $g = 0$ is a subgradient of f at x^* .

4. Convexity

Convexity is defined for both sets and functions. For today we'll focus on discussing the convexity of functions.

Definition 2 (convex functions). A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is **convex** on a set A if for all $x, y \in A$ and $\lambda \in [0, 1]$:

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$$

When this definition holds with the inequality being reversed, then f is said to be concave. From the definition, it is clear that a function f is convex if and only if $-f$ is concave.

(a) Why do we care whether a function is convex or not?

(b) Which of the following functions are convex? (Hint: draw a picture)

(i) $|x|$ (ii) $\cos(x)$ (iii) $x^T x$

(c) Can a function be both convex and concave on the same set? If so, give an example. If not, describe why not.

5. Practical Methods for Checking Convexity

Using the definition to check whether a function is convex or not can be a tedious task in many situations. Some basic methods that can help us achieve the task in an efficient way are introduced below:

- for differentiable function, examine $f(y) \geq f(x) + \nabla f(x)^T(y - x)$ for any x, y in the domain of f .
- for twice differentiable functions, examine $\nabla^2 f(x) \succeq 0$ (i.e., the Hessian matrix is positive semidefinite).
- nonnegative weighted sum
- composition with affine function
- pointwise maximum and supremum

Note: there are even more such methods, which are covered in a convex optimization course or textbook.

- (a) If f is differentiable, then f is convex if and only if $f(y) \geq f(x) + \nabla f(x)^T(y - x)$ for any x, y in the domain of f . A geometric interpretation of this characterization is that any tangent plane of a convex function f must lie entirely below f . One interesting application of this characterization is one of the most important inequalities in probability and statistics: the Jensen's inequality, which states that $\mathbb{E}f(X) \geq f(\mathbb{E}(X))$ when f is convex.

- (b) If f is twice differentiable with convex domain, then f is convex if and only if

$$\nabla^2 f(x) \succeq 0,$$

for any x in the domain of f . Use this method to show that the objective function in linear regression is convex.

(c) Let $\alpha \geq 0$ and $\beta \geq 0$, and if f and g are convex, then αf , $f + g$, $\alpha f + \beta g$ are all convex. One application: When a (possibly complicated) objective function can be expressed as a sum (e.g., the negative log-likelihood function), then showing the convexity of each individual term is typically easier.

(d) Suppose $f(\cdot)$ is convex, then $g(x) := f(Ax + b)$ is convex. Use this method to show that $\|Ax + b\|_1$ is convex (in x), where $\|z\|_1 = \sum_i |z_i|$.

(e) Suppose you know that f_1 and f_2 are convex functions on a set A . The function $g(x) := \max\{f_1(x), f_2(x)\}$ is also convex on A . In general, if $f(x, y)$ is convex in x for each y , then $g(x) := \sup_y f(x, y)$ is convex. Use this method to show that the largest eigenvalue of a matrix X , $\lambda_{\text{Max}}(X)$, is convex in X (Using the definition of convexity would make this question quite difficult).

- (f) Does the same result hold for $h(x) := \min\{f_1(x), f_2(x)\}$? If so, give a proof. If not, provide convex functions f_1, f_2 such that h is not convex.

6. Gradient Descent

We will now examine gradient descent algorithm and study the effect of learning rate α on the convergence of the algorithm. Recall from lecture that Gradient Descent takes on the form of $x_{t+1} = x_t - \alpha \nabla f$

- (a) Assume that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex and differentiable, and additionally,

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\| \text{ for any } x, y$$

I.e., ∇f is Lipschitz continuous with constant $L > 0$

Show that:

Gradient descent with fixed step size $\eta \leq \frac{1}{L}$ satisfies

$$f(x^{(k)}) - f(x^*) \leq \frac{\|x^{(0)} - x^*\|^2}{2\eta k}$$

I.e., gradient descent has convergence rate $O(\frac{1}{k})$

Hints:

- (i) ∇f is Lipschitz continuous with constant $L > 0 \rightarrow f(y) \leq f(x) + \nabla f(x)(y - x) + \frac{L}{2}\|y - x\|^2$ for all x, y .
- (ii) f is convex $\rightarrow f(x) \leq f(x^*) + \nabla f(x^*)(x - x^*)$, where x^* is the local minima that the gradient descent algorithm is converging to.
- (iii) $2\eta \nabla f(x)(x - x^*) - \eta^2 \|\nabla f(x)\|^2 = \|x - x^*\|^2 - \|x - \eta \nabla f(x) - x^*\|^2$