

Coursework

Instructions

- Due date: 10 March, Monday, at 12pm
- The submission is through Gradescope <https://www.gradescope.com/courses/946198>.
- It's best to typeset your answers, but it is fine to submit hand-written answers.
- For Q2, you do not need to submit the source code of the entire file. Copy-paste the snippets of your implementation, and submit 1 PDF document together with the answer in Q1.

Questions

1. In this question, we are going to work out the convergence rate of running gradient descent on the mean-squared error.

Suppose we have a data set $\{(x_1, y_1), \dots, (x_n, y_n)\}$, where $x_i \in \mathbb{R}^d$ and $y_i \in \mathbb{R}$ for $i = 1, \dots, n$. Recall that in linear regression, the goal is to find the minimum of

$$L(w) = \|Xw - y\|_2^2, \quad (1)$$

where

$$X = \begin{bmatrix} - & x_1 & - \\ - & x_2 & - \\ & \vdots & \\ - & x_n & - \end{bmatrix} \quad y = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}. \quad (2)$$

We assume the last dimension of x_i is a constant 1, so we do not need to worry about the bias term.

- (a) Show that the gradient and the Hessian of L are

$$\nabla L(w) = 2X^\top Xw - 2X^\top y \quad (3)$$

$$H = 2X^\top X. \quad (4)$$

[2 marks]

- (b) Recall that the optimal solution is the w^* where $\nabla L(w^*) = 2X^\top Xw^* - 2X^\top y = 0$. Show that

$$\frac{1}{2}Hw^* = X^\top y. \quad (5)$$

[3 marks]

(c) To optimize L with gradient descent, we iteratively perform

$$w_t = w_{t-1} - \eta \nabla L(w_{t-1}). \quad (6)$$

Use the results in (a) and (b) and show that

$$w_t = w_{t-1} - \eta H(w_{t-1} - w^*). \quad (7)$$

[5 marks]

(d) If we subtract both sides of (7) with w^* , show that

$$w_t - w^* = (I - \eta H)(w_{t-1} - w^*). \quad (8)$$

Now unroll the result in (d) and show that

$$w_t - w^* = (I - \eta H)^t (w_0 - w^*). \quad (9)$$

[5 marks]

(e) We will need to take a small detour here. Recall that in Optimization 4, we've talked about how solving

$$\max_w \frac{w^\top A w}{w^\top w} \quad (10)$$

is equivalent to solving

$$\max_w w^\top A w \quad \text{s.t.} \quad \|w\|_2^2 = 1. \quad (11)$$

The solution is to form the Lagrangian

$$F(w) = w^\top A w - \lambda(w^\top w - 1) \quad (12)$$

and find the optimal solution by setting $\nabla F(w) = 0$. The optimal solution satisfies $Aw = \lambda w$, implying that w is an eigenvector of A and λ is its corresponding eigenvalue. Use this result and show that, for any matrix A whose largest eigenvalue is λ_{\max} ,

$$w^\top A w \leq \lambda_{\max} \|w\|_2^2, \quad (13)$$

for all w .

[5 marks]

(f) Apply the results in (d) to (e) and show

$$\|w_t - w^*\|_2^2 = (w_0 - w^*)^\top (I - \eta H)^{2t} (w_0 - w^*) \leq \nu^{2t} \|w_0 - w^*\|_2^2, \quad (14)$$

where ν is the largest eigenvalue of $I - \eta H$.

[5 marks]

(g) Use the results in (a) and show that

$$L(w_1) = L(w_2) + \nabla L(w_2)^\top (w_1 - w_2) + \frac{1}{2}(w_1 - w_2)^\top H(w_1 - w_2) \quad (15)$$

for all $w_1 \in \mathbb{R}^d$ and $w_2 \in \mathbb{R}^d$.¹

[5 marks]

(h) If we choose w_1 to be an arbitrary $w \in \mathbb{R}^d$ and w_2 to be an optimal solution w^* where $\nabla L(w^*) = 0$, show that

$$L(w) - L(w^*) = \frac{1}{2}(w - w^*)^\top H(w - w^*). \quad (16)$$

[5 marks]

(i) Use (h) to get

$$L(w_t) - L(w^*) = \frac{1}{2}(w_t - w^*)^\top H(w_t - w^*). \quad (17)$$

Apply (e) again, plug in (f), and show that

$$L(w_t) - L(w^*) \leq \frac{\lambda_{\max}}{2} \nu^{2t} \|w_0 - w^*\|_2^2, \quad (18)$$

where λ_{\max} is the largest eigenvalue of H .

[5 marks]

(j) If we choose the step size to be $\eta = \frac{1}{K}$ for a constant K , then ν becomes the largest eigenvalue of $I - \frac{1}{K}H$, which is $1 - \lambda_{\min}/K$, where λ_{\min} is the smallest eigenvalue of H . We arrive at

$$L(w_t) - L(w^*) \leq \frac{\lambda_{\max}}{2} \left(1 - \frac{\lambda_{\min}}{K}\right)^{2t} \|w_0 - w^*\|_2^2. \quad (19)$$

When we choose $K > \lambda_{\min}$, show that

$$L(w_t) - L(w^*) \leq O(r^t), \quad (20)$$

for some $0 < r < 1$.

[5 marks]

(k) Given the result in (j), what type of convergence is running gradient descent on mean-squared error? Sublinear, linear, or quadratic?

[5 marks]

2. In this question, we are going to write a small neural network library.

Download <https://homepages.inf.ed.ac.uk/htang2/mlg2024/coursework/nn.py>. The file `nn.py` provides an initial scaffolding.

¹Since L is quadratic in w , the right hand side of (15) is exactly the second-order Taylor expansion of L . We did not talk about Taylor expansion in class, so please do not use it to solve this question.

A neural network is essentially a sequence of function compositions. To train a neural network with stochastic gradient descent, we need to be able to compute the gradient of arbitrary function compositions. The algorithm is called backpropagation, and the data structure is called computation graphs.


Take a two-layer ReLU network for example, where the network has the form

$$f(x) = \text{ReLU}(xW_1 + b_1)W_2 + b_2, \tag{21}$$

and $\text{ReLU}(x) = \max(0, x)$. To compute $f(x)$, we hope to be able to write the following snippet.

```
g = Graph()
output = g.add(g.matmul(g.relu(g.add(g.matmul(x, w1), b1)), w2), b2)
```

The result of running the snippet is the following graph.




The graph is called a computation graph, where every vertex is either a variable or an operation. When values are given to the variables, we can run the computation graph to get the result of the computation. We hope to write the following snippet that prints the value of the output.

```
g.run_eval(output)
print(g.value[output])
```

To represent a computation graph, the `Graph` class has a set of vertices with integer IDs. Suppose we have a `Graph` instance `g`. A vertex with ID `i` has a single parent vertex `g.parent[i]`, a list of children vertices `g.children[i]`, its associated value `g.value[i]`, its associated gradient `g.grad[i]`, and finally the name of the vertex `g.name[i]`.

- (a) The computation needs to be done in a particular order, we need a concept called topological ordering of vertices. Suppose there are n vertices in the graph. A topological ordering of vertices is a sequence v_1, \dots, v_n in which for every pair v_i and v_j with $i < j$, if there is an edge between the two, then v_i needs to be a child of v_j . In simple terms, if we line up the vertices in topological order, the edges can only point forward, not backward. For example, below is a topological order of the graph above.



Have a look at `add` and `_new_vertex`. Suppose we have n vertices. If every time an operation is called (e.g., `c = g.add(a, b)`), we create a new vertex (`c` in this case), connect the children (`a` and `b`), and assign a new ID to it. Use induction and show that the IDs of the vertices follow a topological order.

[10 marks]

- (b) Have a look at `run_eval`. Where is topological order used in `run_eval`? What happens if the IDs of the vertices do not follow a topological order?

[5 marks]

- (c) Have a look at `eval_add`. The forward computation is straightforward, so let's focus on implementing the backward computation `grad_add`.

Given that `c = g.add(a, b)` which represents $c = a + b$, we can write down the chain rule

$$\frac{\partial f}{\partial a} = \frac{\partial f}{\partial c} \frac{\partial c}{\partial a} = \frac{\partial f}{\partial c} \cdot 1 = \frac{\partial f}{\partial c}. \quad (22)$$

Note that `a` is the first child of `c`. We can write the following line in `grad_add` for an arbitrary vertex `n` (in this case, `c`).

```
self.grad[self.children[n][0]] = self.grad[n]
```

Derive $\frac{\partial f}{\partial b}$ and show your final implementation of `g.grad_add` by completing the gradient from `n` to the second child.

[10 marks]

- (d) Have a look at `eval_matmul`. The forward computation is again straightforward, so let's focus on implementing the backward computation `grad_matmul`.

Given that `y = g.matmul(w, A)` which represents $y = wA$, we can write down the chain rule

$$\frac{\partial f}{\partial w_i} = \sum_j \frac{\partial f}{\partial y_j} \frac{\partial y_j}{\partial w_i} = \sum_j \frac{\partial f}{\partial y_j} a_{ij}. \quad (23)$$

In the matrix form, we have

$$\frac{\partial f}{\partial w} = \frac{\partial f}{\partial y} A^\top. \quad (24)$$

Note that `w` here is the first child of `y`, so we can write the following line to implement the chain rule for an arbitrary vertex `n` (in this case, `y`).

```
self.grad[self.children[n][0]] = self.grad[n] @ self.value[self.children[n][1]].T
```

Derive $\frac{\partial f}{\partial a_{ij}}$ and its matrix form. Show your final implementation of `grad_matmul` by completing the gradient from `n` to the second child.

[10 marks]

- (e) Have a look at `run_grad`. Where is topological order used in `run_grad`? What happens if the IDs of the vertices do not follow a topological order?

[5 marks]

- (f) Now that we have implemented backpropagation, we need a way to verify whether the implementation is correct. A simple approach to estimating gradients is called the finite difference method. The idea is that we can simulate derivative by making a tiny change in the function input. For example,

$$\frac{\partial f}{\partial w_i} \approx \frac{f(w + e_i \epsilon) - f(w)}{\epsilon}, \quad (25)$$

where e_i is a vector that has 1 on the i -th coordinate and 0 everywhere else, and ϵ is a small number, say 10^{-3} .

Below is an implementation of the finite difference method to check the gradient to the first child in `grad_matmul`.

```
eps = 1e-3

g1 = Graph()
w1 = g1.tensor(np.array([1, 2, 3]))
a1 = g1.tensor(np.array([[1, 2], [3, 4], [5, 6]]))
y1 = g1.matmul(w1, a1)
g1.run_eval(y1)

g2 = Graph()
w2 = g2.tensor(np.array([1 + eps, 2, 3]))
a2 = g2.tensor(np.array([[1, 2], [3, 4], [5, 6]]))
y2 = g2.matmul(w2, a2)
g2.run_eval(y2)

g3 = Graph()
w3 = g3.tensor(np.array([1, 2 + eps, 3]))
a3 = g3.tensor(np.array([[1, 2], [3, 4], [5, 6]]))
y3 = g3.matmul(w3, a3)
g3.run_eval(y3)

g4 = Graph()
w4 = g4.tensor(np.array([1, 2, 3 + eps]))
a4 = g4.tensor(np.array([[1, 2], [3, 4], [5, 6]]))
y4 = g4.matmul(w4, a4)
g4.run_eval(y4)

g1.run_grad(y1, np.array([1, 0]))

g5 = Graph()
w5 = g5.tensor(np.array([1, 2, 3]))
a5 = g5.tensor(np.array([[1, 2], [3, 4], [5, 6]]))
```

```
y5 = g5.matmul(w5, a5)

g5.run_grad(y5, np.array([0, 1]))

print((g2.value[y2] - g1.value[y1]) / eps)
print((g3.value[y3] - g1.value[y1]) / eps)
print((g4.value[y4] - g1.value[y1]) / eps)
print(g1.grad[w1])
print(g5.grad[w5])
```

Use the finite difference method to check your implementation of the gradient to the second child in `grad_matmul`. Show your implementation of the finite difference method and the respective output to confirm that the gradient is implemented correctly.

[10 marks]