

1. Gradient Descent Mechanics

Gradient descent is the primary algorithm to search optimal parameters for our models. Typically, we want to solve optimization problems stated as

$$\min_{\theta \in \Theta} \mathcal{L}(f_{\theta}, \mathcal{D})$$

where $\mathcal{L}$ are differentiable functions. In this example, we look at a simple supervised learning problem where given a dataset $\mathcal{D} = \{(x_i, y_i)\}_i^N$, we want to find the optimal parameters θ that minimizes some loss. We consider different models for learning the mapping from input to output, and examine the behavior of gradient descent for each model.

- (a) The simplest parametric model entails learning a single-parameter constant function, where we set $\hat{y}_i = \theta$. We wish to find

$$\hat{\theta}_{const} = \operatorname{argmin}_{\theta \in \mathbb{R}} \mathcal{L}(f_{\theta}, \mathcal{D}) = \operatorname{argmin}_{\theta \in \mathbb{R}} \frac{1}{N} \sum_{i=1}^N (y_i - \theta)^2$$

- i. What is the gradient of $\mathcal{L}$ w.r.t. θ ?
- ii. What is the optimal value of θ ?
- iii. Write the gradient descent update.
- iv. *Stochastic Gradient Descent (SGD)* is an alternative optimization algorithm, where instead of using all N samples, we use single sample per optimization step to update the model. What is the contribution of each data-point to the full gradient update?

- (b) Instead of constant functions, we now consider a single-parameter **linear** model $\hat{y}(x_i) = \theta x_i$, where we search for θ such that

$$\hat{\theta} = \min_{\theta \in \mathbb{R}} \frac{1}{N} \sum_{i=1}^N (y_i - \theta x_i)^2$$

- i. What is the gradient of $\mathcal{L}$ w.r.t. θ ?
- ii. What is the optimal value of θ ?
- iii. Write the gradient descent update.
- iv. Do all points get the same *vote* in the update? Why or why not? (*Hint*: repeat the derivation of the last bullet point in part(a).)

2. Two Matrix Norms from the SVD

Let $A \in \mathbb{R}^{m \times n}$ have rank r and singular value decomposition

$$A = U \Sigma V^\top,$$

where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal ($U^\top U = I_m$, $V^\top V = I_n$) and $\Sigma \in \mathbb{R}^{m \times n}$ is zero off the main diagonal, with diagonal entries the singular values

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0, \quad \sigma_i = 0 \text{ for } i > r.$$

Write $u_1, \dots, u_m$ and $v_1, \dots, v_n$ for the columns of U and V .

We will use two ways of measuring the “size” of A : the **Frobenius norm**, which treats A as a long vector of its entries,

$$\|A\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n A_{ij}^2,$$

and the **spectral norm**, which measures how much A can stretch a vector,

$$\|A\|_2 = \max_{x \neq 0} \frac{\|Ax\|_2}{\|x\|_2} = \max_{\|x\|_2=1} \|Ax\|_2.$$

The goal of this problem is to show that the SVD reads off both at a glance:

$$(1) \quad \|A\|_F^2 = \sum_{i=1}^r \sigma_i^2, \quad (2) \quad \|A\|_2^2 = \sigma_1^2.$$

- (a) The Frobenius norm is defined entrywise, which on its face has nothing to do with the SVD. The trace gives us the bridge, and everything else in this problem follows from it.

Show that $\|A\|_F^2 = \text{Tr}(A^\top A)$, and use this to show that for any orthogonal $Q \in \mathbb{R}^{m \times m}$ and orthogonal $P \in \mathbb{R}^{n \times n}$,

$$\|QA\|_F = \|A\|_F = \|AP\|_F.$$

Show also that $\|Qx\|_2 = \|x\|_2$ for every $x \in \mathbb{R}^m$.

- (b) **Using part (a), show that $\|A\|_F^2 = \sum_{i=1}^r \sigma_i^2$.**

- (c) **Using part (a), show that $\|A\|_2^2 = \sigma_1^2$, and identify a unit vector x that attains the maximum.**

(d) **Combine the two results to show that $\|A\|_2 \leq \|A\|_F \leq \sqrt{r} \|A\|_2$. When is each inequality tight?**

3. ReLU SGD Visualization

Work through the notebook to explore how a simple network with ReLU non-linearities adapts to model a function with SGD updates. Training the networks takes 5-10 minutes depending on whether you run locally or on a server, so you should start the training process (ie. run through the *train all layers* cell) then return to the theory part of the discussion while training occurs. The link to run the notebook on Google Colab is [q_relu_vis.ipynb](#)

As you walk through the notebook, pay attention to how the slopes and the elbows of the ReLU functions change during training and how they impact the shape of the final learned function.

Contributors:

- Ashwin Pananjady.
- Josh Sanz.
- Yichao Zhou.
- Anant Sahai.
- Gireeja Ranade.
- Matteo Guarrera.
- Mert Cemri.
- Sizhe Chen.
- Kumar Krishna Agrawal.
- Kevin Li.