

Section 06: Stochastic Gradient Descent, SVMs

1. Stochastic Gradient Descent

Consider minimizing an average of functions:

$$\min_w \frac{1}{n} \sum_{i=1}^n \ell_i(w),$$

where w is a d -dimensional vector (or the feature dimension is d). The minimization of the negative of a log-likelihood function can serve as an example. Recall that the (full) gradient descent step is given by

$$w^{(t+1)} = w^{(t)} - \eta \cdot \frac{1}{n} \sum_{i=1}^n \nabla \ell_i(w^{(t)}).$$

The computational cost of a single step here is $\mathcal{O}(dn)$. To reduce cost, one idea is to just use a subset of all samples to approximate the full gradient. Specifically, consider revising the gradient descent step as follows:

$$w^{(t+1)} = w^{(t)} - \eta \cdot \nabla \ell_{I_t}(w^{(t)}),$$

where I_t is chosen randomly within $\{1, 2, \dots, n\}$ with equal probabilities. This is called **stochastic gradient descent (SGD)**, and the computational cost of a single step now reduces to $\mathcal{O}(d)$.

(a) The following two results provide intuitions or foundations for why SGD works.

- $\mathbb{E}_{I_t}(\nabla \ell_{I_t}(w^{(t)})) = \frac{1}{n} \sum_{i=1}^n \nabla \ell_i(w^{(t)})$, which is the full gradient. Hence the estimate of gradient is unbiased.
- Let $\ell(w) = \frac{1}{n} \sum_i \ell_i(w)$ and $w^* = \arg \min_w \ell(w)$. Assume $\|w^{(1)} - w^*\|_2^2 \leq R$ and $\sup_w \max_i \|\nabla \ell_i(w)\|_2^2 \leq G$. Then

$$\mathbb{E}[\ell(\bar{w}) - \ell(w^*)] \leq \sqrt{\frac{RG}{T}},$$

where $\bar{w} := \frac{1}{T} \sum_{t=1}^T w^{(t)}$. Therefore, the expected error over T iterations is $\mathcal{O}(\frac{1}{\sqrt{T}})$. (The proof of this result is provided in the solution part for reference.)

(b) What disadvantages can SGD have? How can we balance between the noise in updates and computational cost?

(c) Gradient descent requires the full gradient when updating while (standard) SGD utilizes the gradient of one sample when updating. **Mini-batching** is somewhere between the two extremes. That is, we choose a random subset $I_t \subseteq \{1, \dots, n\}$ with size $|I_t| = b \ll n$ in the stochastic gradient descent step:

$$w^{(t+1)} = w^{(t)} - \eta \cdot \frac{1}{b} \sum_{i_t \in I_t} \nabla \ell_{i_t}(w^{(t)}).$$

With mini-batching, we have the following results:

- $\mathbb{E}_{I_t} \left(\frac{1}{b} \sum_{i_t \in I_t} \nabla \ell_{i_t}(w^{(t)}) \right) = \frac{1}{n} \sum_{i=1}^n \nabla \ell_i(w^{(t)})$: we still have an unbiased estimate of the full gradient.
- Compared to standard SGD, variance of the gradient estimate is reduced approximately by $\frac{1}{b}$.
- Computational cost for each step now becomes $\mathcal{O}(db)$.

Remark: By matrix computations (computing b gradients at a time) and parallelization, we can denoise the estimated gradients without increasing much computational cost (for batch size b that is not large).

(d) How should we choose the batch size? Are there other extensions or variants of the basic stochastic gradient descent algorithm?

2. SVMs

Consider the dataset consisting of 7 data points, 4 with positive labels $\{0, 1, 2, 3\}$, and 3 with negative labels $\{-3, -2, -1\}$. Suppose we want to learn a linear SVM with slack variables for this dataset. Recall we can formalize this as a constrained optimization problem:

$$\begin{aligned} & \min_{w, b, \xi} \|w\|^2 + C \sum_i \xi_i \\ & \text{subject to} \\ & y_i(w \cdot x_i + b) \geq 1 - \xi_i \quad \forall i \\ & \xi_i \geq 0 \quad \forall i \end{aligned}$$

where C is a regularization parameter that balances the size of the margin (smaller $\|w\|^2$) vs. the violation of the margin (smaller $\sum_i \xi_i$).




Figure 1: Dataset

Also recall that a support vector is a datapoint that lies on a margin.

- (a) If $C = 0$, which means we only care about the size of the margin, how many support vectors do we have?

- (b) If $C \rightarrow \infty$, which means we only care about the violation of the margin, how many support vectors do we have?