

Lecture 12: Accelerating SGD with variance reduction and averaging.

CS4787/5777 — Principles of Large-Scale ML Systems

Continuing from last time, we'll talk about Adam (from last week's notes).

Recap: Recall that from the Lecture 4 notes, we had that (for strongly convex SGD)

$$\mathbf{E} [f(w_T) - f^*] \leq \exp(-\mu\alpha T) \cdot (f(w_0) - f^*) + \frac{\alpha\sigma^2\kappa}{2}.$$

Last week, we looked at the effect of the κ term on this and other algorithms; today, we'll look at the effect of the α and σ^2 parts of the noise ball, and discuss some ways to address the effect of these to accelerate SGD.

Polyak averaging. Intuition: SGD is converging to a “noise ball” where the iterates are randomly jumping around some space surrounding the optimum. We can think about these iterates as random samples that approximate the optimum.

What can we do when we have a bunch of random samples that approximate something to improve the precision of our estimate?

Technique: run regular SGD and just average the iterates. That is,

$$\begin{aligned} w_{t+1} &= w_t - \alpha_t \nabla f_{i_t}(w_t) \\ \bar{w}_{t+1} &= \frac{t}{t+1} \cdot \bar{w}_t + \frac{1}{t+1} \cdot w_{t+1}; \end{aligned}$$

eventually we output the average \bar{w}_T at the end of execution. This is equivalent to writing

$$\bar{w}_T = \frac{1}{T} \sum_{t=1}^T w_t.$$

The main idea here is that we're averaging out a bunch of iterations w_t that are all in the noise ball. When you average out a bunch of noisy things, often you are able to reduce the noise.

To gain intuition about Polyak averaging, let's look at a simple one-dimensional quadratic...

$$f(w) = \frac{1}{2}w^2$$

with example gradients

$$\nabla f_i(w) = w + u_i$$

where u_i considered as a random variable over all $i \in \{1, \dots, n\}$ has mean 0 and variance σ^2 , i.e.

$$\frac{1}{n} \sum_{i=1}^n u_i = 0 \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n u_i^2 = \sigma^2.$$

Suppose that we run SGD on this with a constant learning rate α . Our update step will be

$$w_{t+1} = w_t - \alpha w_t - \alpha u_{i_t} = (1 - \alpha)w_t - \alpha u_{i_t}.$$

Applying this recursively, we get

$$w_T = (1 - \alpha)^T w_0 - \underbrace{\alpha \sum_{t=0}^{T-1} (1 - \alpha)^{T-1-t} u_{i_t}}_{\text{noise term}}.$$

To tell how basic SGD compares with averaged SGD, we want to compare the expected loss f evaluated at both these points. For regular SGD, it is

$$\begin{aligned} \mathbf{E} [f(w_T)] &= \frac{1}{2} \mathbf{E} [w_T^2] \\ &= \frac{1}{2} \mathbf{E} \left[\left((1 - \alpha)^T w_0 - \alpha \sum_{t=0}^{T-1} (1 - \alpha)^{T-1-t} u_{i_t} \right)^2 \right] \\ &= \frac{1}{2} \mathbf{E} \left[\left((1 - \alpha)^T w_0 - \alpha \sum_{t=0}^{T-1} (1 - \alpha)^{T-1-t} u_{i_t} \right) \left((1 - \alpha)^T w_0 - \alpha \sum_{s=0}^{T-1} (1 - \alpha)^{T-1-s} u_{i_s} \right) \right] \\ &= \frac{1}{2} \left((1 - \alpha)^{2T} w_0^2 - 2\alpha (1 - \alpha)^T w_0 \sum_{t=0}^{T-1} (1 - \alpha)^{T-1-t} \cdot \mathbf{E} [u_{i_t}] \right. \\ &\quad \left. + \alpha^2 \sum_{t=0}^{T-1} \sum_{s=0}^{T-1} (1 - \alpha)^{T-1-t} \cdot (1 - \alpha)^{T-1-s} \cdot \mathbf{E} [u_{i_t} u_{i_s}] \right). \end{aligned}$$

Now, since the u_{i_t} are zero mean and independent, it follows that

$$\mathbf{E} [\mathbf{E} [u_{i_t}]] = 0.$$

For the same reason, it will also be the case that if $s \neq t$,

$$\mathbf{E} [u_{i_t} u_{i_s}] = 0.$$

On the other hand, if $s = t$, then

$$\mathbf{E} [u_{i_t} u_{i_s}] = \mathbf{E} [u_{i_t}^2] = \sigma^2.$$

As a result, all the cross terms in the above sum will cancel out to zero, and we can write the expected loss as

$$\mathbf{E} [f(w_T)] = \frac{1}{2} \left((1 - \alpha)^{2T} w_0^2 + \alpha^2 \sum_{t=0}^{T-1} (1 - \alpha)^{2(T-1-t)} \sigma^2 \right).$$

Next, if we let $k = T - 1 - t$, then summing from $t = 0$ to $T - 1$ is equivalent to summing from $k = 0$ to $T - 1$, so

$$\begin{aligned} \mathbf{E} [f(w_T)] &= \frac{1}{2} \left((1 - \alpha)^{2T} w_0^2 + \alpha^2 \sum_{k=0}^{T-1} (1 - \alpha)^{2k} \sigma^2 \right) \\ &= \frac{1}{2} \left((1 - \alpha)^{2T} w_0^2 + \alpha^2 \sigma^2 \frac{1 - (1 - \alpha)^{2T}}{1 - (1 - \alpha)^2} \right) \\ &= \frac{1}{2} \left((1 - \alpha)^{2T} w_0^2 + \alpha \sigma^2 \frac{1 - (1 - \alpha)^{2T}}{2 - \alpha} \right) \\ &= (1 - \alpha)^{2T} \left(f(w_0) - \frac{\alpha \sigma^2}{2(2 - \alpha)} \right) + \frac{\alpha \sigma^2}{2(2 - \alpha)} \end{aligned}$$

where in the derivation above, we used the fact that for any x ,

$$\sum_{n=0}^{N-1} x^n = \frac{1 - x^N}{1 - x}.$$

Note that this expression is exact! We didn't use any approximations or inequalities here, but if we want a somewhat simpler expression we can upper bound this with

$$\mathbf{E}[f(w_T)] = (1 - \alpha)^{2T} \cdot f(w_0) + \frac{\alpha\sigma^2}{2(2 - \alpha)}$$

Now let's do the same analysis for the Polyak-averaged updates. If we use averaging, the output of the algorithm becomes

$$\begin{aligned} \bar{w}_T &= \frac{1}{T} \sum_{k=0}^{T-1} \left((1 - \alpha)^k w_0 - \alpha \sum_{t=0}^{k-1} (1 - \alpha)^{k-1-t} u_{i_t} \right) \\ &= \frac{1}{T} \left(\sum_{k=0}^{T-1} (1 - \alpha)^k w_0 - \alpha \sum_{k=0}^{T-1} \sum_{t=0}^{k-1} (1 - \alpha)^{k-1-t} u_{i_t} \right). \end{aligned}$$

The pair of sums in the second term here results in the value $(1 - \alpha)^{k-1-t} u_{i_t}$ being summed up for all pairs (k, t) that satisfy

$$0 \leq k < T \quad \text{and} \quad 0 \leq t < k.$$

This condition is equivalent to $0 \leq t < k < T$, and so we could also write it as

$$0 \leq t < T - 1 \quad \text{and} \quad t < k < T.$$

That is, these inequalities correspond to the same ordered pairs of (k, t) . As a result, we can re-order our summations above into

$$\bar{w}_T = \frac{1}{T} \left(\sum_{k=0}^{T-1} (1 - \alpha)^k w_0 - \alpha \sum_{t=0}^{T-2} \sum_{k=t+1}^{T-1} (1 - \alpha)^{k-1-t} u_{i_t} \right),$$

since the resulting expression sums up the exact same terms. Next, if we substitute $l = k - 1 - t$ in the inner sum, then

$$\begin{aligned} \bar{w}_T &= \frac{1}{T} \left(\sum_{k=0}^{T-1} (1 - \alpha)^k w_0 - \alpha \sum_{t=0}^{T-2} \sum_{l=0}^{T-t-2} (1 - \alpha)^l u_{i_t} \right) \\ &= \frac{1}{T} \left(\frac{1 - (1 - \alpha)^T}{\alpha} \cdot w_0 - \sum_{t=0}^{T-2} (1 - (1 - \alpha)^{T-t-1}) \cdot u_{i_t} \right). \end{aligned}$$

If we compute the expected value of the loss function f evaluated at this point, then as before, all the cross terms will cancel, and we'll just be left with

$$\begin{aligned} \mathbf{E}[f(\bar{w}_T)] &= \frac{1}{2} \mathbf{E}[\bar{w}_T^2] = \frac{1}{2} \mathbf{E} \left[\left(\frac{1}{T} \left(\frac{1 - (1 - \alpha)^T}{\alpha} \cdot w_0 - \sum_{t=0}^{T-2} (1 - (1 - \alpha)^{T-t-1}) \cdot u_{i_t} \right) \right)^2 \right] \\ &= \frac{1}{2T^2} \left(\left(\frac{1 - (1 - \alpha)^T}{\alpha} \cdot w_0 \right)^2 + \sum_{t=0}^{T-2} (1 - (1 - \alpha)^{T-t-1})^2 \cdot \sigma^2 \right). \end{aligned}$$

Next, I'm going to simplify the above expression by noting that for sufficiently small α (i.e. $0 < \alpha \leq 1$) the term $1 - (1 - \alpha)^{T-t-1} \leq 1$, and so

$$\sum_{t=0}^{T-2} (1 - (1 - \alpha)^{T-t-1})^2 \leq T.$$

Using this, and the fact that $1 - (1 - \alpha)^T \leq 1$, we can get the simplified upper bound

$$\mathbf{E} [f(\bar{w}_T)] = \frac{1}{2T^2} \left(\left(\frac{1}{\alpha} \cdot w_0 \right)^2 + T \cdot \sigma^2 \right) = \frac{w_0^2}{2\alpha^2 T^2} + \frac{\sigma^2}{2T} = \frac{f(w_0)}{\alpha^2 T^2} + \frac{\sigma^2}{2T}.$$

Now, let's compare our rates. For SGD and averaged SGD, we have

$$\mathbf{E} [f(w_T)] \leq (1 - \alpha)^{2T} \cdot f(w_0) + \frac{\alpha \sigma^2}{2(2 - \alpha)} \quad \text{and} \quad \mathbf{E} [f(\bar{w}_T)] \leq \frac{f(w_0)}{\alpha^2 T^2} + \frac{\sigma^2}{2T}.$$

When we compare, we notice that the second term (the term that depends on σ^2 , i.e. the *variance term*) is actually going to zero as T increases. That is, even with a constant learning rate α , we can get loss values that are arbitrarily close to the actual optimal loss $f^* = f(0) = 0$. (Unfortunately, this is not guaranteed to happen for general convex losses, and only happens here because our loss is a quadratic function.) On the other hand, the term that depends on the initialization w_0 (the first term) is no longer decaying exponentially.

One issue with this is that we are averaging with equal weight iterates from the very start of training, when we have not reached the noise ball. In order to address this, we often run averaged SGD by first using a **warm-up period** during which we do not average, and then only starting to average after the warm-up period is over. Long warm-up periods (such as half of the total number of epochs ran) usually produce better results in practice than averaged SGD without any warmup.

For **Polyak averaging on non-convex problems**, we run into the same issue that we ran into with AdaGrad: as we move through a non-convex landscape, we may get very far away from the parameter values we were at during previous iterations. If this happens, averaging together with those parameter values doesn't really make sense, and can hurt the performance of our algorithm. So, instead, a standard approach is to use an *exponential moving average*, just like was done to transform AdaGrad into RMSProp. That is, we pick some decay factor $0 < \rho < 1$ and run

$$\begin{aligned} w_{t+1} &= w_t - \alpha_t \nabla f_{i_t}(w_t) \\ \bar{w}_{t+1} &= \rho \cdot \bar{w}_t + (1 - \rho) \cdot w_{t+1}. \end{aligned}$$

This running average approach often outperforms other averaging methods in non-convex settings.

Stochastic variance-reduced gradient. Idea: reduce the variance of the gradient estimators by using an infrequent full-gradient step.

Procedure SVRG

Parameters update frequency m and learning rate η
Initialize \tilde{w}_0
Iterate: for $s = 1, 2, \dots$

$$\begin{aligned} \tilde{w} &= \tilde{w}_{s-1} \\ \tilde{\mu} &= \frac{1}{n} \sum_{i=1}^n \nabla \psi_i(\tilde{w}) \\ w_0 &= \tilde{w} \end{aligned}$$

Iterate: for $t = 1, 2, \dots, m$

Randomly pick $i_t \in \{1, \dots, n\}$ and update weight

$$w_t = w_{t-1} - \eta (\nabla \psi_{i_t}(w_{t-1}) - \nabla \psi_{i_t}(\tilde{w}) + \tilde{\mu})$$

end

option I: set $\tilde{w}_s = w_m$
option II: set $\tilde{w}_s = w_t$ for randomly chosen $t \in \{0, \dots, m - 1\}$
end