Lecture Notes for September 23th

Dmitry Berenson

Projected Gradient Descent Regret Bounds

Distance between optimal weight vector w^* and the weight vector at time t (w_t).

$$D(w_t, w^*) = (w_t - w^*)^T (w_t - w^*) \tag{1}$$

then the distance after the update for the next time step is:

$$D(w_{t+1}, w^*) - D(w_t, w^*) = ||w_t - \alpha \nabla l_t - w^*||^2 - ||w_t - w^*||^2 \tag{2}$$

Now we substitute $z_t = w_t - w^*$ and get

$$D(w_{t+1}, w^*) - D(w_t, w^*) = z_t^2 - 2\alpha \nabla l_t^T z_t + \alpha^2 \nabla l_t^2 - z_t^2 \tag{3}$$

canceling and substituting in for z_t gives us:

$$D(w_{t+1}, w^*) - D(w_t, w^*) = -2\alpha \nabla l_t^T (w_t - w^*) + \alpha^2 \nabla l_t^2 \tag{4}$$

Now we can sum over time to get the distance up until time t.

$$\sum_t (D(w_{t+1}, w^*) - D(w_t, w^*)) = -2\alpha \sum_t \left(\nabla l_t^T (w_t - w^*)\right) + \sum_t \left(\alpha^2 \nabla l_t^2\right) \tag{5}$$

$$\sum_t (D(w_{t+1}, w^*) - D(w_t, w^*)) = D(w_n, w^*) - D(w_0, w^*) \tag{6}$$

The regret is bounded by:

$$R(w^*) \leq \sum_t \left(\nabla l_t^T (w_t - w^*)\right) = -\frac{1}{2\alpha} \left[D(w_n, w^*) - D(w_0, w^*) - \sum_t \alpha^2 \nabla l_t^2\right] \tag{7}$$

Next, we need to assume that the loss at any time step is bounded by a constant G:

$$\nabla l_t^2 \leq G \tag{8}$$

The bound on the reward then becomes:

$$R(w^*) \leq \frac{\alpha}{2} G^2 n + \frac{1}{2\alpha} D(w_0, w^*) - \frac{1}{2\alpha} D(w_n, w^*) \tag{9}$$

Since we want an upper bound and the last term can only decrease this bound, we can eliminate it.

$$R(w^*) \leq \frac{\alpha}{2} G^2 n + \frac{1}{2\alpha} D(w_0, w^*) \tag{10}$$

Since we are dealing with a bounded convex set of w vectors, let us bound $D(w_0, w^*)$ by the size of the set F^2. Now we need to pick α to minimize regret. To do this, we take the derivative of the above equation w.r.t α and set it equal to 0. Some algebra gives us

$$\alpha = \frac{F}{G\sqrt{n}} \tag{11}$$

and the regret is thus bounded by[1]

$$R(w^*) \leq FG\sqrt{n} \tag{12}$$

note that the regret grows sub-linearly with time which means that as $t \rightarrow \infty$:

$$\frac{R_t}{t} \longrightarrow 0 \tag{13}$$

Thus this algorithm is a "no-regret" algorithm.

Portfolio Optimization

We want to do nearly as well as the best Constantly Rebalancing Portfolio. The weight vector (percent invested in a given stock) is subject to the constraints:

$$w_i \geq 0 \tag{14}$$

$$\sum_i w_i = 1 \tag{15}$$

The loss is:

$$l_t(w^*) = \log \sum_i w_i^* r_i^t \tag{16}$$

We make the no-junk-bond assumption which is:

$$r_{\max} \geq r_i \geq r_{\min} > 0 \tag{17}$$

the regret is bounded by:

$$R(\text{Investment}) \leq \sqrt{2}\sqrt{n}\left(\frac{r_{\max}}{r_{\min}}\right)^2 \sqrt{T} \tag{18}$$

[1] Drew says the right bound is actually $R(w^*) \leq 2FG\sqrt{n}$