Non-convex optimization.

When objective is non-convex, generally the best we can hope for is to find a stationary point: $\|\nabla f\| = 0$, or an ε-appx stationary point, $\|\nabla f\| \le \varepsilon$.

Stationary points need not be local min: $y = x^3$

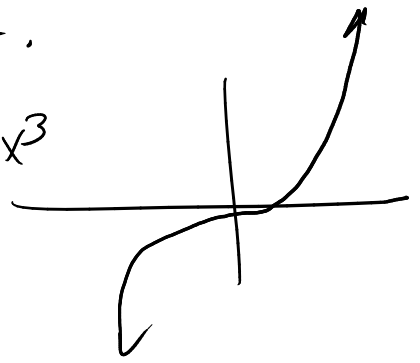

Recall Lemma 7 from last class:

Lemma 7: Let f be β-smooth (not nec. cvx). If $\alpha_t = \alpha < \frac{1}{\beta}$, then

$\forall T \ge 1$, G.D. satisfies $\min_{t<T} \|\nabla f(w_t)\|^2 \le \frac{2}{\alpha T}\left(f(w_0) - f(w_T)\right)$

So, GD on smooth objectives efficiently finds stationary points.

There are many problems where all stationary points are global optima.

Let's consider a few: Let $\mathcal{X}^* := \{x^* : f(x^*) = \min_x f(x)\}$.

Let $\Pi(x) :=$ projection of x onto $\mathcal{X}^*$.

① (Strong) convexity: for some $\mu > 0$ ($\mu > 0$: strong),

$$\forall x, y, \quad f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\mu}{2}\|x - y\|^2$$

② weak strong convexity: for some $\mu > 0$, if $f^* = \min_x f(x)$,

$$\forall x, \quad f^* \geq f(x) + \langle \nabla f(x), \Pi(x) - x \rangle + \frac{\mu}{2}\|x - \Pi(x)\|^2.$$

③ Restricted secant inequality:

$$\forall x, \quad \langle \nabla f(x), x - \Pi(x) \rangle \geq \mu \|x - \Pi(x)\|^2.$$

④ Error bound inequality:

$$\forall x, \quad \|\nabla f(x)\| \geq \mu \|x - \Pi(x)\|.$$

⑤ Polyak-Łojasiewicz: (PL)

$$\forall x, \quad \frac{1}{2}\|\nabla f(x)\|^2 \geq \mu \cdot (f(x) - f^*)$$

Clear that PL $\Rightarrow$ all local min are global. However:

Thm (Karimi et al '16): If f is β-smooth, then:

$$① \Rightarrow ② \Rightarrow ③ \Rightarrow (④ \Leftrightarrow ⑤).$$

So, PL inequality is quite crucial.

Thm. Suppose f is β-smooth, μ-PL, and $\exists x^*$ s.t. $f(x^*) = \min_x f(x)$.

Then: GD w/ $\alpha \le \frac{1}{\beta} \wedge \frac{1}{\mu}$ $f(x_k) - f^* \le (1-\alpha\mu)^k \cdot (f(x_0) - f^*)$.

Pf. By defn, $x_{k+1} = x_k - \alpha \nabla f(x_k)$.

By β-smoothness,

$$f(y) \le f(x) + \langle \nabla f(x), y - x \rangle + \frac{\beta}{2}\|x-y\|^2 \quad \forall x, y.$$

$$\Rightarrow f(x_{k+1}) - f(x_k) \le \langle \nabla f(x_k), x_{k+1} - x_k \rangle + \frac{\beta}{2}\|x_{k+1} - x_k\|^2$$

$$= -\alpha \|\nabla f(x_k)\|^2 + \frac{\beta\alpha^2}{2}\|\nabla f(x_k)\|^2$$

$$\le -\frac{\alpha}{2}\|\nabla f(x_k)\|^2 \quad \text{as } \beta\alpha \le \frac{1}{\beta}.$$

$$\le -\alpha\mu(f(x_k) - f^*) \quad \text{by PL}.$$

$$\Rightarrow f(x_{k+1}) - f^* \le (1-\alpha\mu) f(x_k) + \alpha\mu f^* - f^*$$

$$= (1-\alpha\mu) f(x_k) - (1-\alpha\mu) f^* = (1-\alpha\mu)(f(x_k) - f^*).$$

$\Rightarrow$ unrolling, we see $f(x_k) - f^* \le (1-\alpha\mu)^k (f(x_0) - f^*)$. ∎

So, with PL on smooth objectives, can get "linear convergence" to get $f(w_t) - f^* \leq \varepsilon$, only need $O(\log \frac{1}{\varepsilon})$ iterations.

Exercise Show that there exists a β-smooth convex function on $\mathbb{R}$ for which PL is not satisfied.

So establishing PL is clearly very powerful. Unsurprisingly, it often does not hold, and if it does, can be very hard to show. Moreover, for many problems of interest, we do not have β-smoothness. Eg, neural networks with ReLU activations, since $\varphi(t) = \max(0, t)$ has $\varphi'(t) = \mathbb{1}(t \geq 0)$ (subderivative).

Clearly φ' is not Lipschitz, so φ is not β-smooth.

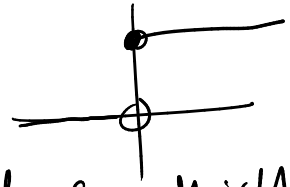

We'll next describe some alternative non-convex approaches which
(1) allow for non-smooth objectives, &/or
(2) allow for situations where "all local min are good" (but not nec. global opt.)

Surrogate objectives.

In classification, we are most interested in the 0-1 loss:

$$\mathbb{1}(y \neq f(x)) = \begin{cases} 1: & y \neq f(x), \\ 0: & y = f(x) \end{cases} = \mathbb{1}(y \cdot f(x) < 0)$$

for $y, f(x) \in \{-1, 1\}$

[In general: $\mathbb{1}(A) = \begin{cases} 1 \text{ if } A \text{ occurs,} \\ 0 \text{ otherwise.} \end{cases}$]

Unfortunately, $\mathbb{1}(\cdot)$ is discontinuous, piecewise constant, so its derivative is zero a.e $\rightarrow$ gradient descent is useless.

So instead, we minimize convex surrogates for the 0-1 loss:

Call $\ell(t)$ a convex surrogate for $g(t)$ if:

① · ℓ is convex

② · $g(t) \leq \ell(t)$ $\forall t$.

Thus, minimizing ℓ implies minimizing g.

We will typically also require ℓ is differentiable/smooth so we can apply G.D. to ℓ.

Common convex surrogates to 0-1 loss $\mathbb{1}(t < 0) = \begin{cases} 1: t < 0, \\ 0: t \geq 0 : \end{cases}$

- $\exp(-t)$ "exponential loss"
- $\max(0, 1-t)$ "hinge loss" ~~✗~~
- $\log(1 + \exp(-t))$ "logistic / cross entropy"

(note ② is satisfied, but $\mathbb{1}(t < 0) \leq 2 \cdot \log(1 + \exp(-t))$)

More generally, there are a large family of surrogates to 0-1 loss.

Let $\ell(t)$ be convex, decreasing, and twice differentiable. Then:

Assume $\ell'(0) \neq 0$.

① $\ell''(t) \geq 0$ (so $-\ell'$ is decreasing)

② $-\ell'(t) \geq 0$.

Then: $\{t < 0\} = \{-\ell'(t) \geq -\ell'(0)\}$ by ①

If we can show $-\ell'(t) \to 0$, then by

$$\frac{1}{n}\sum_{i=1}^{n} \mathbb{1}(y_i f(x_i)) = \frac{1}{n}\sum_{1}^{n} \mathbb{1}(y_i \cdot f(x_i) < 0)$$

$$= \hat{P}(y \cdot f(x) < 0)$$

$$= \hat{P}(-\ell'(y \cdot f(x)) \geq -\ell'(0))$$

$$\leq \frac{\hat{E}(-\ell'(y \cdot f(x)))}{-\ell'(0)} \quad \text{by Markov}$$

$\ell'(0) \neq 0$.

$$= \frac{1}{-\ell'(0)} \cdot \frac{1}{n}\sum_{1}^{n} -\ell'(y_i f(x_i))$$

Thus, for any convex, decreasing, differentiable (twice) ℓ s.t. $\ell'(0) \neq 0$,
$-\ell'(t)$ is surrogate for 0-1 loss.

① $\ell(t) = \exp(-t) \Rightarrow -\ell'(t) = \ell(t)$.

② $\ell(t) = \log(1+\exp(-t)) \Rightarrow -\ell'(t) = \frac{\exp(-t)}{1+\exp(-t)} = \frac{1}{1+\exp(t)}$.

Here $-\ell'$ is not convex.

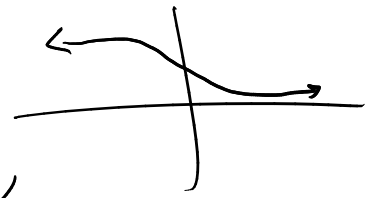

However, if we do G.D. on ℓ (which is convex),
i.e. G.D. on $\hat{L}(W) := \frac{1}{n}\sum_1^n \ell(y_i f(x_i; W))$,
and we somehow found that the loss
$\hat{G}(W) := \frac{1}{n}\sum_1^n -\ell'(y_i f(x_i; W))$ were small, then we
would know that the 0-1 loss is small as well.
(In the homework, we'll see why this is useful.)

Def. $f:\mathbb{R}^p$ is called (g,h,ν)-proxy convex if, $\exists g, h:\mathbb{R}^p \to \mathbb{R}$ s.t.
$\exists \nu$ s.t. $\forall w$, $h(\nu) \geq g(w) + \langle \nabla f(w), \nu - w \rangle$.

Thm: If f is (g,h,ν)-proxy-convex and $\|\nabla f(w)\| \leq L_1$ for all w,
then for any $w(0)$, $\eta > 0$, $T \geq 1$

$$\min_{t<T} g(w_t) \leq h(\nu) + \frac{2L_1^2}{2} + \frac{\|w(0) - \nu\|^2}{2\eta T}.$$

Remark. For small step sizes η and large T, G.D. on f finds $g(w_t)$ which is close to $h(\nu)$.

useful if we can show there are families of related objective functions $\{g,h\}$ for which are "meaningful". i.e., $\nu = \operatorname{argmin} f$,

- $h(w) = \sqrt{f(w)}$ and $g(w) = \frac{1}{2} f(w)$ would imply

$$\min_{t<T} \frac{1}{2} f(w_t) \leq \min_w \sqrt{f(w)} + \varepsilon.$$

NOT a global min, but a fcn of it. If $\min_w f(w)$ is very small, then good.

In remainder, we will prove that GD converges to a global min for 2-layer leaky ReLU nets on linearly sep. data.

Two proof techniques:

① Proxy convexity] lecture

② perceptron argument.] homework

2-layer leaky nets: $w_j \in \mathbb{R}^d$, $W \in \mathbb{R}^{m \times d}$ (rows w_j),

$a_j \in \{\pm 1/\sqrt{m}\}$, $m = \#$ neurons. $\varphi(t) = \max(t, \alpha t)$, $\alpha \in [0,1]$

$$f(x; W) = \sum_{j=1}^{m} a_j \varphi(\langle w_j, x \rangle).$$

only training first layer.

$\ell(t) = \log(1 + \exp(-t))$; $\|x_i\| \le 1$

$\exists v^* \in \mathbb{R}^d$ s.t. $\|v^*\| = 1$, $\forall i$, $\langle y_i x_i, v^* \rangle \ge \gamma$, $\gamma \in (0,1]$.

G.D. on empirical risk $\hat{L}(W) = \frac{1}{n} \sum_{1}^{n} \ell(y_i f(x_i; W))$.

Remark We will consider parameterizing the objective for $\hat{L}(W)$ by a matrix W, while we typically think of vectors in $\mathbb{R}^d$ for some d.

You can cast the problem to that form by considering $\text{vec}(W) \in \mathbb{R}^{d \cdot m}$: stack each column of W one after another.

Everything works out — can check via matrix calculus.

- $\langle W, V \rangle := \text{tr}(W^T V) \quad (= \text{tr}(V^T W) = \text{tr}(W V^T))$

  $W, V \in \mathbb{R}^{m \times d}$

  $= \langle \text{vec}(W), \text{vec}(V) \rangle \in \mathbb{R}^{dm}$

- $\|W\|_F^2 = \|\text{vec}(W)\|_2^2$

Lemma. If $g:\mathbb{R}^d \to \mathbb{R}$ is L-positively homogeneous, ($g(\alpha w) = \alpha^L g(w)$ for $\alpha \geq 0$)
then $\langle \nabla g(w), w \rangle = L\, g(w)$

pf: Exercise.

Remark: $\varphi(t) = \max(t, \alpha t)$ is 1-positively homogeneous.

Our goal will be to show $\hat{L}(w)$ is proxy convex for appropriate proxies g, h, i.e. want to show, for some g, h,

$$\langle \nabla \hat{L}(w), w - v \rangle \geq g(w) - h(v).$$

$$\nabla \hat{L}(w) = \frac{1}{n} \sum_{i=1}^{n} \nabla \ell(y_i f(x_i; w))$$

$$= \frac{1}{n} \sum_{i}^{n} \ell'(y_i f(x_i; w))\, y_i \nabla f(x_i; w).$$

Now: $f(x_i; w) = \sum_j a_j \varphi(\langle w_j, x \rangle)$ is 1-homog. since φ is. So,

$$\langle \nabla f(x_i; w), w \rangle = f(x_i, w).$$

Thus: $$\langle \nabla \hat{L}(w), W - V \rangle = \frac{1}{n} \sum_i^n \ell'(y_i f(x_i; w))\, y_i \langle \nabla f(x_i; w), W - V \rangle$$
$$= \frac{1}{n} \sum_i^n \ell'(y_i f(x_i; w)) \cdot \left(y_i f(x_i; w) - y_i \langle \nabla f(x_i; w), V \rangle \right)$$

So now, which V do we want to choose?

Calculus $\Rightarrow$ $\nabla f(x; w) = D_x^w a x^T$, $D_x^w = \mathrm{diag}\left(\varphi'(\langle w_j, x \rangle)\right)$.

$D_x^w a x^T \in \mathbb{R}^{m \times d}$

$$y_i \langle \nabla f(x_i; w), V \rangle = y_i \cdot \mathrm{tr}\left((D_{x_i}^w a x_i^T)^T V \right)$$
$$= y_i \, \mathrm{tr}\left(x_i a^T D_{x_i}^w V \right)$$
$$= y_i \, \mathrm{tr}\left(a^T D_{x_i}^w V x_i \right)$$
$$= y_i \, a^T D_{x_i}^w V x_i$$
$$= \sum_j a_j \varphi'(\langle w_j, x_i \rangle) \langle v_j, y_i x_i \rangle .$$

Know: $\langle v^*, y_i x_i \rangle \geq \gamma > 0$. So let $v_j := \rho \cdot a_j \cdot v^*$, $\rho > 0$.

Since $\varphi'(t) \geq \alpha > 0$ $\forall t$,

$$y_i \langle \nabla f(x_i; w), V \rangle = \sum_{j=1}^m a_j \varphi'(\langle w_j, x_i \rangle) \cdot \langle v_j, y_i x_i \rangle = \sum_{j=1}^m a_j^2 \cdot \rho \cdot \varphi'(\langle w_j, x_i \rangle) \langle v^*, y_i x_i \rangle$$
$$\geq \frac{\rho}{m} \sum_{j=1}^m \alpha \gamma = \rho \alpha \gamma .$$

Continuing from [ ]: Since ℓ is decreasing, $-\ell'(z) < 0$. So,

$$-\ell'(y_i f(x_i; w)) \cdot y_i \langle \nabla f(x_i, w), v \rangle \geq -\ell'(y_i f(x_i; w)) \cdot \rho\alpha\gamma.$$

Thus,

$$\langle \nabla \hat{L}(w), w - v \rangle \geq \frac{1}{n} \sum_i^n \ell'(y_i f(x_i; w)) \cdot (y_i f(x_i; w) - \rho\alpha\gamma).$$

(since ℓ is convex, $\ell'(z_1) \cdot (z_1 - z_2) \geq \ell(z_1) - \ell(z_2)$. →

$$\geq \frac{1}{n} \sum_i^n [\ell(y_i f(x_i; w)) - \ell(\rho\alpha\gamma)].$$

$$= \hat{L}(w) - \ell(\rho\alpha\gamma).$$

Therefore, $\hat{L}(w)$ is $(\hat{L}, \ell(\rho\alpha\gamma), v)$-proxy convex with $v = v(\rho)$ having rows $v_j = \rho a_j v^*$.

Since $\|\nabla f(x_i; w)\|_F = \|\mathrm{Diag}(\varphi'(\langle w_j, x_i \rangle)_j) \cdot x_i^+\|_F$

$$\leq 1 \quad (\underline{\text{Exercise}}!)$$

Thus the theorem guarantees that for any $T \geq 1$, $w(0)$, $\eta > 0$, $\varepsilon > 0$,

$$\min_{t<T} \hat{L}(w_t) \leq \ell(\rho\alpha\gamma) + \frac{\eta}{2} + \frac{\|w(0) - V\|_F^2}{2\eta T}.$$

Since $\ell(t) \leq 2\exp(-t)$, taking $\rho = \alpha^{-1}\gamma^{-1}\log\left(\frac{6}{\varepsilon}\right)$ guarantees $\ell(\rho\alpha\gamma) \leq \frac{\varepsilon}{3}$.

For $\eta \leq \frac{\varepsilon}{2}$, this means

$$\min_{t<T} \hat{L}(w_t) \leq \frac{2\varepsilon}{3} + \frac{\|w(0) - V\|_F^2}{2\eta T}.$$

For $T \geq 2\eta^{-1}\|w(0) - V\|_F^2 \varepsilon^{-1}$, $\min_{t<T} \hat{L}(w_t) \leq \varepsilon$.

All together:

Theorem. Suppose $\|x_i\| \leq 1$, $y_i \langle v^*, x_i \rangle \geq \gamma$ for all i for some $\|v^*\| = 1$. Then GD on 2-layer leaky ReLU net satisfies the following. For any $\varepsilon > 0$, if $\eta \leq \varepsilon/2$ then for $V \in \mathbb{R}^{m \times d}$ with rows $v_j = a_j \alpha^{-1}\gamma^{-1}\log\left(\frac{6}{\varepsilon}\right)$, if $T \geq 2\eta^{-1}\|w(0) - V\|_F^2 \varepsilon^{-1}$, then $\min_{t<T} \hat{L}(w_t) \leq \varepsilon$.

Remark I'm not aware of proofs with any of:

① - ϕ is ReLU rather than leaky ReLU

② - bias terms included

③ - training both first and second layer weights

④ - training > 2 layers net.

difficulty: 4 > 1 > 3 > 2.

If you can come up with a proof in any of these settings, this would probably be sufficient for a NeurIPS paper! (except maybe ②) Requires:

① Random initialization

② Only assume linear separability of training data.

③ No "NTK". Ideally allow for constant number of neurons.

You can work on this instead of choosing paper for final project if you'd like.