Chapter 5

Thus far, we have focused on methods that are based on the notion of gradient: At every iteration, we compute first-order (=gradient) information about the objective, and we use this information to perform an educated step towards a local or a global minimum of the objective, as in gradient descent. We have shown—through theoretical analysis—what we can achieve by using gradients concerning convergence rates and what the best we can hope for (=lower bounds) is.

But what are some ways to accelerate this first set of algorithms regarding analytical complexity? We will present some approaches that deviate from simple gradient-based methods and provably and empirically outperform the methods studied thus far, including Newton's method and quasi-Newton variants. To complete the picture on the theory side, we will continue working in the convex world and compare the obtained bounds to understand what we gain and lose for each of these choices.

Newton's method | quasi-Newton variants | Natural gradient | derivative-free optimization

We first remind ourselves what the limits of gradient descent are. The following summarizes lower bounds we can expect, *by only using gradients in convex optimization*, for some types of objective functions we have previously discussed.

- For the class of L-smooth convex objective functions, one can prove the existence of functions f such that gradient descent satisfies:

$$f(x_T) - f(x^\star) \geq \frac{3L\|x_0 - x^\star\|_2^2}{32(T+1)^2} = O\left(\tfrac{1}{T^2}\right).$$

  Under this assumption, and only using gradients, we cannot achieve a better convergence rate than $O(1/T^2)$.
- For the class of convex objectives functions with both Lipschitz continuous gradients and strong convexity, one can prove the existence of functions f such that gradient descent satisfies:

$$\|x_T - x^\star\|_2^2 \geq \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^{2T} \|x_0 - x^*\|_2^2.$$

  where $\kappa = L/\mu > 1$. Here we observe that, while we have achieved the same convergence rate concerning the exponent—i.e., in both cases, we have c^T, for $c < 1$—in the lower bound case, we "deal with" the term $\sqrt{\kappa}$ instead of κ. This can be alternatively seen as $O(\sqrt{\kappa}\log\frac{1}{\varepsilon})$ iteration complexity, compared to $O(\kappa\log\frac{1}{\varepsilon})$ iteration complexity that we already proved for gradient descent.

But how can we achieve such lower bounds? Can we achieve something better? This chapter follows a path on different algorithmic approaches, some of which deviate from using first-order methods to show what we gain (and lose) in practice.

The (notorious) Newton's method. Newton's method has been and still is one of the most celebrated algorithms in the numerical scientific community. We mention the research community here to highlight that there are applications where we care about accurately resolving a problem. E.g., there might be cases where the estimation accuracy at the error level of 10^{-15} is essential for the problem and the algorithm. Getting a solution just 10^{-4}-close might be unacceptable.

Let us first derive the Newton's iteration. Newton's method, as another descent method, updates the current estimate x as:

$$x \leftarrow x + \Delta x,$$

where Δx abstractly defines a direction/update that moves x to a "better place", with respect to the objective we try to minimize. In this chapter, we focus on the unconstrained case:

$$\min_{x\in\mathbb{R}^p} \quad f(x),$$

where f is assumed to be twice differentiable, with gradient $\nabla f(\cdot)$ and Hessian $\nabla^2 f(\cdot)$. By taking the second-order Taylor expansion of f around $x + \Delta x$, we have:

$$\begin{aligned} f(x+\Delta x) &\approx f(x) + \langle \nabla f(x), (x+\Delta x) - x\rangle \\ &\quad + \tfrac{1}{2}\left\langle \nabla^2 f(x)\left((x+\Delta x)-x\right), \left((x+\Delta x)-x\right)\right\rangle \\ &\approx f(x) + \langle \nabla f(x), \Delta x\rangle + \tfrac{1}{2}\left\langle \nabla^2 f(x)\Delta x, \Delta x\right\rangle \end{aligned}$$

(*Similar reasoning is used for gradient descent to connect Δx with negative gradient, $-\nabla f(\cdot)$, as the best descent direction for first-order methods.*)

Using this characterization, we can locally find the best Δx by finding the root of the quadratic approximation. To see this, if we set $\Delta x \equiv y$, the above expression becomes:

$$f(x) + \langle \nabla f(x), y\rangle + \tfrac{1}{2}\left\langle \nabla^2 f(x) y, y\right\rangle,$$

which is a quadratic function with respect to y. Given the above, one could find Δx that makes the gradient of $f(x+\Delta x)$ be zero:

$$\begin{aligned} \nabla_{\Delta x} f(x+\Delta x) = 0 \overset{\text{approx.}}{\Rightarrow} \quad & \nabla f(x) + \nabla^2 f(x)\Delta x = 0 \\ \Rightarrow \quad & \Delta x = -\left(\nabla^2 f(x)\right)^{-1}\nabla f(x). \end{aligned}$$

Substituting Δx in $x + \Delta x$, we obtain the Newton's iteration:

Definition 26. (Newton's method) *For $x_t \in \mathbb{R}^p$ and $\eta_t \in \mathbb{R}$, we update our estimate x_t on each iteration as follows:*

$$x_{t+1} = x_t - \eta_t H_t^{-1}\nabla f(x_t), \quad H_t := \nabla^2 f(x_t).$$

Remark 1. *As we will show, $\eta_t = 1$ in theory. However, some cases require $\eta_t < 1$ at least at the beginning of the algorithm or when we initialize the algorithm badly. When $\eta_t < 1$, we call the method* damped Newton's method, *and its study is currently outside the scope of this course.*

Before we present some theory for Newton's method, we need to get the full picture of what we are proposing: With Newton's method, *we do gradient descent type-of motions.* But before applying the gradient $\nabla f(x_t)$, we "translate" it through the matrix $\nabla^2 f(x_t)^{-1}$; i.e., we use the transformed gradient $\widetilde{\nabla} f(x_t) := \nabla^2 f(x_t)^{-1}\nabla f(x_t)$.

Remark 2. *If $H_t^{-1} = \mathrm{diag}\{h_1, \ldots, h_p\}$, where not all $h_i = 1$, then using our interpretation above, we can see that in this condition we are doing gradient descent with coordinate-specific step sizes (i.e., for coordinate i we use step size $\eta_t h_i$). Notice by default, since η_t is scalar, we take equally-scaled steps in each direction for each descent step, which may be suboptimal depending on our domain and the distribution of x; customizing our step sizes coordinate-wise can perform better in practice in a variety of domains. Coordinate-specific step sizes are an active area of research at the center of machine learning research and the adaptive methods used in neural network training. This topic might be covered in later chapters of the course.*

Guarantees of Newton's method. Let us first study the behavior of Newton's method in general, even non-convex, scenarios.

Theorem 4. *Let $\min_x f(x)$ be the problem of interest, with f being twice differentiable. Assume f has Lipschitz continuous Hessians:*

$$\left\|\nabla^2 f(x) - \nabla^2 f(y)\right\|_2 \leq M \cdot \|x - y\|_2,$$

and f satisfies at the optimum point $x^\star$: $\nabla^2 f(x^\star) \succeq \mu I$. Assuming that we start from a point x_0 that is close enough to $x^\star$:

$$\|x_0 - x^\star\|_2 < \tfrac{2\mu}{3M},$$

Newton's method as in:

$$x_{t+1} = x_t - \eta_t H_t^{-1} \nabla f(x_t), \quad H_t := \nabla^2 f(x_t),$$

converges according to:

$$\|x_{t+1} - x^\star\|_2 \leq \frac{M\|x_t - x^\star\|_2^2}{2\left(\mu - M\|x_t - x^\star\|_2\right)}.$$

Proof: For the first step of the analysis, we need to prove a lemma that derives from the application of Taylor's theorem / mean value theorem / fundamental theorem of calculus (Part II) / Newton-Leibniz axiom, which we restate below for convenience.

Theorem 5. *Let f be a real-valued (continuous) function on $[\alpha, \beta]$ with anti-derivative F (i.e., $F'(x) = f(x)$). Then:*

$$\int_\alpha^\beta f(x)\, dx = F(\beta) - F(\alpha)$$

Lemma 7. *By the Fundamental Theorem of Calculus (i.e., above), we have*

$$\nabla f(x) - \nabla f(y) = \int_0^1 \nabla^2 f(y + \tau(x - y))(x - y)\, d\tau$$

Proof: Define $g'(\tau) = \nabla^2 f(y + \tau(x - y)) \cdot (x - y) = (\nabla f(y + \tau(x - y)))'$. Then:

$$\begin{aligned}
&\int_0^1 \nabla^2 f(y + \tau(x - y))(x - y)\, d\tau \\
&\quad = \int_0^1 g'(\tau)\, d\tau \\
&\quad = g(1) - g(0) \\
&\quad = \nabla f(y + 1 \cdot (x - y)) - \nabla f(y + 0 \cdot (x - y)) \\
&\quad = \nabla f(x) - \nabla f(y).
\end{aligned}$$

using $g' = f$ and $g = F$. ∎

Now, using the lemma we just proved to massage the main recursion of Newton's method, we have:

$$\begin{aligned}
&x_{t+1} - x^\star \\
&= x_t - \left(\nabla^2 f(x_t)\right)^{-1} \nabla f(x_t) - x^\star \\
&= x_t - \left(\nabla^2 f(x_t)\right)^{-1} \left(\nabla f(x_t) - \nabla f(x^\star)\right) - x^\star \\
&= x_t - \left(\nabla^2 f(x_t)\right)^{-1} \left(\int_0^1 \nabla^2 f\left(x^\star + \tau(x_t - x^\star)\right)(x_t - x^\star)\, d\tau\right) - x^\star \\
&= (x_t - x^\star) - \left(\nabla^2 f(x_t)\right)^{-1} \left(\int_0^1 \nabla^2 f\left(x^\star + \tau(x_t - x^\star)\right)(x_t - x^\star)\, d\tau\right) \\
&= \left(\nabla^2 f(x_t)\right)^{-1} \cdot G_t (x_t - x^\star)
\end{aligned}$$

where

$$G_t = \int_0^1 \left(\nabla^2 f(x_t) - \nabla^2 f\left(x^\star + \tau(x_t - x^\star)\right)\right) d\tau.$$

We proceed by bounding the terms on the right-hand side. (*Remember that $\|\cdot\|_2$ for matrices corresponds to the spectral norm, not the Frobenius norm.*)

$$\begin{aligned}
\|G_t\|_2 &= \left\| \int_0^1 \left(\nabla^2 f(x_t) - \nabla^2 f\left(x^\star + \tau(x_t - x^\star)\right)\right) d\tau \right\|_2 \\
&\leq \int_0^1 \left\|\nabla^2 f(x_t) - \nabla^2 f\left(x^\star + \tau(x_t - x^\star)\right)\right\|_2 d\tau \\
&\leq \int_0^1 M \cdot \|x_t - x^\star + \tau(x_t - x^\star)\|_2 d\tau \\
&= \tfrac{M\|x_t - x^\star\|_2}{2}
\end{aligned}$$

Moreover, we know that, by the Hessian Lipschtiz continuity:

$$\left\|\nabla^2 f(x) - \nabla^2 f(y)\right\|_2 \leq M \cdot \|x - y\|_2,$$

we have:

$$\nabla^2 f(x) - M\|x - y\|_2 \cdot I \preceq \nabla^2 f(y) \preceq \nabla^2 f(x) + M\|x - y\|_2 \cdot I,$$

$\forall x, y$, and thus holds for $x = x_t$ and $y = x^\star$:

$$\nabla^2 f(x_t) \succeq \nabla^2 f(x^\star) - M\|x_t - x^\star\|_2 \cdot I \succeq (\mu - M\|x_t - x^\star\|_2) \cdot I.$$

Assume that $\|x_t - x^\star\|_2 \leq \frac{\mu}{M}$ (*to be justified a posteriori*), we have:

$$\left\|\nabla^2 f(x_t)^{-1}\right\|_2 \leq \left(\mu - M\|x_t - x^\star\|_2\right)^{-1}.$$

Combining all the above, we get:

$$\begin{aligned}
\|x_{t+1} - x^\star\|_2 &\leq \left\| \left(\nabla^2 f(x_t)\right)^{-1} \cdot G_t \cdot (x_t - x^\star) \right\|_2 \\
&\leq \left\|\nabla^2 f(x_t)^{-1}\right\|_2 \cdot \|G_t\|_2 \cdot \|x_t - x^\star\|_2 \\
&\leq \left(\mu - M\|x_t - x^\star\|_2\right)^{-1} \cdot \tfrac{M\|x_t - x^\star\|_2}{2} \cdot \|x_t - x^\star\|_2 \\
&= \frac{M\|x_t - x^\star\|_2^2}{2\left(\mu - M\|x_t - x^\star\|_2\right)}.
\end{aligned}$$

Let us discuss the initialization assumption: $\|x_0 - x^\star\|_2 \leq \frac{2\mu}{3M}$. Using induction, we have the following two steps.

`Basis step:` We have:

$$\begin{aligned}
\|x_1 - x^\star\|_2 &\leq \frac{M\|x_0 - x^\star\|_2^2}{2\left(\mu - M\|x_0 - x^\star\|_2\right)} \\
&= \frac{M \cdot \frac{4\mu^2}{9M^2}}{2\left(\mu - M \cdot \frac{2\mu}{3M}\right)} \\
&= \frac{\frac{4\mu^2}{9M}}{2\left(\frac{3M\mu - 2M\mu}{3M}\right)} = \frac{\frac{4\mu^2}{9M}}{\frac{2\mu}{3}} = \tfrac{2\mu}{3M}.
\end{aligned}$$

`Induction step:` Assume that for some t, it holds $\|x_t - x^\star\|_2 \leq \frac{2\mu}{3M}$. This also justifies the assumption that $\|x_t - x^\star\|_2 \leq \frac{2\mu}{3M} \leq \frac{\mu}{M}$ which is used in the proof above as an assumption. Then:

$$\begin{aligned}
\|x_{t+1} - x^\star\|_2 &\leq \frac{M\|x_t - x^\star\|_2^2}{2\left(\mu - M\|x_t - x^\star\|_2\right)} \\
&= \cdots = \tfrac{2\mu}{3M}.
\end{aligned}$$

This completes the proof: i.e., assuming a good enough initialization, $\|x_0 - x^\star\|_2 \leq \frac{2\mu}{3M}$, all the assumptions in the proof are justified, leading the recursion in the theorem. ■

Before we proceed, let's first understand what this recursion means. By assumption of initialization, the recursion becomes:

$$\begin{aligned}\|x_{t+1} - x^\star\|_2 &\leq \frac{M\|x_t - x^\star\|_2^2}{2\left(\mu - M\|x_t - x^\star\|_2\right)} \\ &\leq \frac{M\|x_t - x^\star\|_2^2}{2\left(\mu - M\frac{2\mu}{3M}\right)} \\ &= \tfrac{3M}{2\mu} \cdot \|x_t - x^\star\|_2^2 \equiv c \cdot \|x_t - x^\star\|_2^2.\end{aligned}$$

Under the assumption that we start from a good initialization point where $\|x_t - x^\star\|_2 \leq 1$—i.e., $\frac{2\mu}{3M} \leq 1$—this translates that the new distance is *quadratically* decreased, rather than linearly. That is, if we want $\|x_T - x^\star\|_2 \leq \varepsilon$, then this can be achieved in $O\left(\log\log\frac{1}{\varepsilon}\right)$ iterations. See also the convergence rate figure.

What if we assume convexity of f? It turns out that, using convexity, we do not gain anything in terms of convergence rate. However, assuming convexity, we can achieve *global* convergence: irrespective of the initialization, there is an analysis that proves that Newton's method converges to the global minimum. *There is a caveat, though:* The quadratic convergence rate holds only locally! I.e., we are guaranteed a quadratic convergence rate after we perform some steps at a slower rate. Only after we get inside a region close enough to the global minimum is the quadratic rate activated!

Some comments on Newton's method

- Newton's method exploits the local curvature of the function. This is depicted in the following figures, borrowed from Boyd's and Vandenberghe's book. In the first case, the gradient descent method is myopic, and the gradient suggests a direction almost perpendicular to the direction we should move.

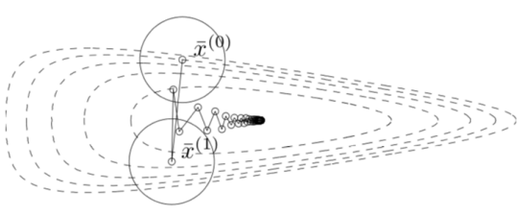

Fig. 33. Gradient descent behavior in function valleys.

On the other hand, Newton's method "warps" the function landscape, where the gradient direction moves more towards the optimum.

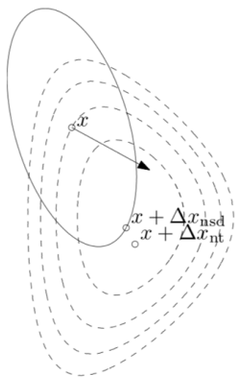

Fig. 34. Newton's method first changes the function landscape and then performs gradient descent on this space.

- Each iteration of Newton's method is *more expensive computationally* than simple gradient descent. Thus, there is a trade-off: while we need a much smaller number of iterations to get to optimum (after good initialization), we pay much more per iteration. (*Think of the case where computing the Hessian does not fit in the computer's main memory*). Remember that if $\nabla f(x) \in \mathbb{R}^p$, then $\nabla^2 f(x) \in \mathbb{R}^{p\times p}$; if $\nabla f(X) \in \mathbb{R}^{p\times p}$, then $\nabla^2 f(X) \in \mathbb{R}^{p^2\times p^2}$. Setting $p = 10^6$, we get an idea of how things could scale in practice.
- Theory so far assumes a good initialization point to achieve quadratic convergence rate *–this is an active research area even recently– hopefully, more notes will be added to this bullet in the future.*
- Newton's method is rarely used in machine learning applications because we often need to care about exact solutions. Newton's method is critical in cases where accuracy is key, such as numerical analysis and scientific computing *–this is an active research area even recently– hopefully, more notes will be added to this bullet in the future.*
- Comparing to what we can achieve with gradient descent, Newton's method "breaks" the lower bound

$$\|x_T - x^\star\|_2^2 \geq \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^{2T}\|x_0 - x^*\|_2^2 \equiv c^T \cdot \|x_0 - x^*\|_2^2,$$

since for Newton's method, we have:

$$\|x_T - x^\star\|_2^2 \leq c^T \cdot \|x_0 - x^*\|_2^4.$$

Spanning the space between gradient descent and Newton's method. Newton's method proposes a different way of performing gradient descent: instead of just taking the gradient per iteration, we compute the Hessian to weigh the gradient. This raises the question: *Does only the true Hessian work as a weighting factor for the gradient? Can we generate some approximate Hessian H_t and use it in*

$$x_{t+1} = x_t - \eta_t H_t^{-1}\nabla f(x_t)$$

to similar success?

This leads to the class of general preconditioning matrices and preconditioning methods. These methods are often called

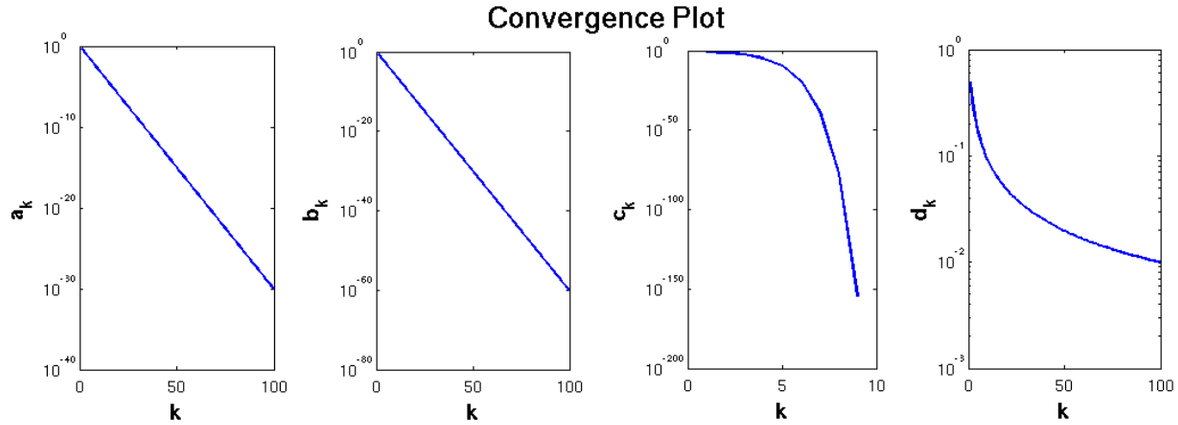

Fig. 35. Borrowed from Wikipedia. Illustration of different convergence rates. Note that the y-axis is in logarithmic scale for all the plots, while the x-axis has a linear scale. The y-axis denotes a metric that dictates the optimum point; for example, $\|x_k - x^\star\|_2$ (*We use k as an iteration subscript here*). The x-axis represents the iteration count k. The first two plots represent *linear* convergence rates: it is called linear as a convention to match the linear curve in the *logarithmic* y-axis scale. While the second plot depicts a preferable behavior, the two plots are equivalent in the big-Oh notation. For an error level ε, linear convergence rate implies $O\left(\log \frac{1}{\varepsilon}\right)$. The third plot depicts a *quadratic* convergence rate. For an error level ε, linear convergence rate implies $O\left(\log\log \frac{1}{\varepsilon}\right)$. Finally, the fourth plot represents the *sublinear* convergence rate, much slower than the linear rate. Some typical rates are: $O\left(1/\varepsilon^2\right)$, $O\left(1/\varepsilon\right)$, $O\left(1/\sqrt{\varepsilon}\right)$.

quasi-Newton methods, as we do not use the exact Hessian information per iteration.

Definition 27. (Quasi-Newton method) *For $x_t \in \mathbb{R}^p$, we update our estimate x_t on each iteration as follows:*

$$x_{t+1} = x_t - \eta_t B_t \nabla f(x_t), \quad B_t \in \mathbb{R}^{p\times p}.$$

where $B_t \approx H_t^{-1}$ is some approximation to the inverse of the true Hessian.

There are numerous ways to perform this step—i.e., there are various ways to generate B_t per iteration—but we will focus on two of them for now:

- The (L)BFGS approximation;
- The SR1 approximation.

Both of them handle the unconstrained case:

$$\min_{x\in\mathbb{R}^p} f(x).$$

The Broyden-Fletcher-Goldfarb-Shanno approximation, a.k.a. BFGS. The BFGS approximation is based on the following reasoning:

- We know by Taylor's theorem that we can approximate the objective $f(\cdot)$ around x_t as:

  $$g_t(\Delta x) := f(x_t) + \langle \nabla f(x_t),\ \Delta x\rangle + \tfrac{1}{2}\langle H_t \Delta x,\ \Delta x\rangle,$$

  where H_t represents the actual Hessian matrix. Note here that $g_t(\Delta x)$ represents a local quadratic approximation of f, and Δx is a vector that defines the direction we want to take: $x_{t+1} = x_t + \Delta x$. Thus, iteratively, we will generate the sequence $\ldots, g_{t-1}(\cdot), g_t(\cdot), g_{t+1}(\cdot), \ldots$, where at each iteration we compute a new Δx.
- Instead of using the exact Hessian in $H_t := \nabla^2 f(x_t)$, we look for an approximation of the Hessian. Remember that we use $g(\cdot)$ to compute the new Δx. We need some conditions that this function should satisfy:

  1. *When we take the gradient of $g_{t+1}(\cdot)$ at the zero point—meaning that we do not move at all—we should get back the gradient of the original function.* This condition makes sure that the quadratic approximation of f around its original point x_{t+1} gives back the original gradient of the function, $\nabla f(x_{t+1})$:

     $$\nabla g_{t+1}(0) = \nabla f(x_{t+1})$$

  2. *When we take the gradient of the new function approximation $g_{t+1}(\cdot)$, evaluated at the point after reversing the direction $-\Delta x$, then we should obtain back the gradient of f at the previous iteration.* I.e.,

     $$\nabla g_{t+1}(-\Delta x) = \nabla f(x_t)$$

  3. Inspired by making the local approximation quadratic, we also require per iteration to have:

     $$H_{t+1} \succ 0$$

- Let us use the above information to generate some useful equations. First, observe that by taking gradient of $\nabla g_{t+1}(-\Delta x)$ and using the above equation, we get:

  $$\nabla g_{t+1}(-\Delta x) = \nabla f(x_t) \ \Rightarrow\ H_{t+1}\Delta x = \nabla f(x_{t+1}) - \nabla f(x_t)$$

  This is known as the *secant equation*. Further, by the assumption that $H_{t+1} \succ 0$, the above becomes:

  $$\langle \Delta x,\ \nabla f(x_{t+1}) - \nabla f(x_t)\rangle > 0.$$

- The above lead to a recipe: per iteration, we are looking for a matrix $H_{t+1} \succ 0$ such that the secant equation is satisfied. But, *how many such H_{t+1} exist? Quite a lot!* To restrict the search space, the BFGS method solves the following optimization problem per iteration:

  $$\begin{aligned} \min_{H\succ 0} \quad & \|H - H_t\|_F^2 \\ \text{subject to} \quad & H = H^\top \\ & H\Delta x = \nabla f(x_{t+1}) - \nabla f(x_t) \end{aligned}$$

  Solving this problem, we obtain H_{t+1}; to use H_{t+1}, we further need to invert it and use it as:

  $$x_{t+2} = x_{t+1} - \eta_{t+1} H_{t+1}^{-1} \nabla f(x_{t+1}).$$

In other words, we have found a way to compute a matrix H_{t+1}, but we still need to invert it, just like Newton's method! If so, why don't we compute $\nabla^2 f(x_{t+1})$, which we know is optimal?

- BFGS method goes a bit further to handle this case: Instead of computing in the H domain and then performing inversion, we define $B := H^{-1}$, and we substitute that in the above expression:

$$\begin{aligned} \min_{B \succ 0} \quad & \|B - B_t\|_F^2 \\ \text{subject to} \quad & B = B^\top \\ & \Delta x = B\left(\nabla f(x_{t+1}) - \nabla f(x_t)\right) \end{aligned}$$

I.e., *we approximate the inverse directly so that* $x_{t+1} = x_t - \eta_t B_t \nabla f(x_t)$*!*

- But, how easy is it to solve the above problem? It turns out (*remember that various matrices satisfy what we need*) that the above problem has a closed-form solution:

$$B_{t+1} = \left(I - \frac{s_t y_t^\top}{s_t^\top y_t}\right) B_t \left(I - \frac{y_t s_t^\top}{s_t^\top y_t}\right) + \frac{s_t s_t^\top}{s_t^\top y_t}$$

where

$$\begin{aligned} s_t &:= \Delta x \\ y_t &:= \nabla f(x_{t+1}) - \nabla f(x_t) \end{aligned}$$

Remark 3. *How do we initialize? In other words, how do we set* B_0*? Standard configurations assume* $B_0 = I$*.*

- What is the computational complexity of the above operations? First, observe that we have all the ingredients computed as if we were performing gradient descent: at the $t+1$ iteration, we have $\nabla f(x_t)$ and $\nabla f(x_{t+1})$. Because we do matrix-matrix multiplication, this algorithm is still $O(n^3)$, the same asymptotic complexity as inverting the Hessian! However, we perform only inner and outer product operations, often much faster than computing the actual Hessian and inverting it (e.g., via SVD); the big-O notation hides this speedup in the constants it elides.
- The BFGS method achieves a convergence rate of

$$\|x_{t+1} - x^\star\|_2 \le c_t \|x_t - x^\star\|_2 \qquad \text{where } c_t \to 0$$

We call this a *super-linear* method. It is faster than sub-linear as the convergence constant shrinks with t but still slower than the quadratic convergence given by actual second-order methods (e.g., Newton's method). Still, that's okay, given that BFGS is a method that uses only first-order information to approximate the second-order information!

The symmetric, rank-1 approximation, a.k.a. SR1/ The BFGS method described above does not assume anything about the function f other than being differentiable. Thus, BFGS can be used both for convex and non-convex optimization, where we force the secant equation + positive definiteness to find the new preconditioner at each iteration. This means that we approximate the function f per iteration with a second-order function that always looks upwards! In other words, we locally approximate f with a "bowl" per iteration, even if f initially might look locally as a saddle. While this never happens in the convex case (and thus BFGS sounds like a great choice when we minimize a convex function), there might be cases where we minimize a non-convex function, and it would be great to have different approaches to handle these cases.

This is where SR1 approximation could be handy. As its name indicates, the SR1 approximation approximates a preconditioner matrix through successive rank-1 updates. In particular, if H_t is the current approximation of second-order information, SR1 is based on the following approximation:

$$H_{t+1} = H_t + \sigma v v^\top,$$

for some vector v with appropriate dimensions, and $\sigma \in \{\pm 1\}$. *Key property is that such updates do not guarantee that the new approximation is positive definite, which could be a nice feature when we approximate non-convex functions.*

Assuming the secant equation is satisfied, the combination of the two equations leads to the following update:

$$B_{t+1} = B_t + \frac{(s_t - B_t y_t)(s_t - B_t y_t)^\top}{(s_t - B_t y_t)^\top y_t}.$$

Natural gradient: entering methodology in modern ML. Here, we will discuss the notion of *natural gradient*, relate it to the idea of Hessian in optimization, and set the scene for adaptive methods in training neural networks. To do so, we will need the following notions.

Let $\theta \in \mathbb{R}^p$ denote a set of variables that are unknown to us and we want to estimate. Here, we will follow more of a probabilistic approach where given these parameters θ, we observe $x \in \mathbb{R}^d$, according to the distribution $p(x|\theta)$. To give a concrete example, assume that θ models the space of human faces: then, given fixed $\theta := \theta_0$, the probability of observing face #1 over face #2 could be:

$$p(x_1|\theta = \theta_0) > p(x_2|\theta = \theta_0),$$

while, for a different θ realization, $\theta := \theta_1$, it might be:

$$p(x_1|\theta = \theta_1) < p(x_2|\theta = \theta_1).$$

Now, assume that we have a data set $\{x_i\}_{i=1}^n$. One way to learn θ is through maximum log-likelihood: we define the log-likelihood as $\log p(x|\theta)$, and we are interested in:

$$\widehat{\theta} \in \arg\max_{\theta \in \mathbb{R}^p} \left\{ \mathcal{L}(\theta) := \mathbb{E}_{p(x|\theta)} \left[\log p(x|\theta) \right] \right\}$$

Let us compute the gradient and the Hessian of this new function. For the gradient, we first compute:

$$\nabla \log p(x|\theta) = \tfrac{1}{p(x|\theta)} \cdot \nabla p(x|\theta)$$

and thus,

$$\nabla \mathcal{L}(\theta) = \mathbb{E}_{p(x|\theta)} \left[\tfrac{1}{p(x|\theta)} \cdot \nabla p(x|\theta) \right].$$

For the Hessian, as the Jacobian of the gradient, we have:

$$\begin{aligned} H_{\log p(x|\theta)} &= \nabla \left(\tfrac{1}{p(x|\theta)} \cdot \nabla p(x|\theta) \right) \\ &= \tfrac{H_{p(x|\theta)} \cdot p(x|\theta) - \nabla p(x|\theta) \cdot \nabla p(x|\theta)^\top}{p(x|\theta) \cdot p(x|\theta)} \\ &= \frac{H_{p(x|\theta)}}{p(x|\theta)} - \left(\frac{\nabla p(x|\theta)}{p(x|\theta)} \right) \cdot \left(\frac{\nabla p(x|\theta)}{p(x|\theta)} \right)^\top \end{aligned}$$

where $H_{p(x|\theta)}$ is the Hessian with respect to $p(x|\theta)$. Computing the expectation with respect to $p(x|\theta)$, we have:

$$\begin{aligned} &\mathbb{E}_{p(x|\theta)} \left[H_{\log p(x|\theta)} \right] \\ &= \mathbb{E}_{p(x|\theta)} \left[\frac{H_{p(x|\theta)}}{p(x|\theta)} \right] - \mathbb{E}_{p(x|\theta)} \left[\left(\frac{\nabla p(x|\theta)}{p(x|\theta)} \right) \cdot \left(\frac{\nabla p(x|\theta)}{p(x|\theta)} \right)^\top \right] \\ &= \int \frac{H_{p(x|\theta)}}{p(x|\theta)} p(x|\theta) dx - \mathbb{E}_{p(x|\theta)} \left[\left(\frac{\nabla p(x|\theta)}{p(x|\theta)} \right) \cdot \left(\frac{\nabla p(x|\theta)}{p(x|\theta)} \right)^\top \right] \end{aligned}$$

Before we proceed, observe that:

$$\begin{aligned}\int \frac{H_{p(x|\theta)}}{p(x|\theta)} p(x|\theta)dx &\equiv \int \nabla\left(\tfrac{1}{p(x|\theta)} \cdot \nabla p(x|\theta)\right) \cdot p(x|\theta)dx \\ &= \int \nabla\left(\tfrac{1}{p(x|\theta)} \cdot \nabla p(x|\theta) \cdot p(x|\theta)\right) dx \\ &= \int \nabla^2 \left(p(x|\theta)\right) dx \\ &\overset{\text{(mild assumptions)}}{=} \nabla^2 \left(\int p(x|\theta)dx\right) = \nabla^2(1) = 0\end{aligned}$$

Thus,

$$\mathbb{E}_{p(x|\theta)}\left[H_{\log p(x|\theta)}\right] = -\mathbb{E}_{p(x|\theta)}\left[\left(\frac{\nabla p(x|\theta)}{p(x|\theta)}\right) \cdot \left(\frac{\nabla p(x|\theta)}{p(x|\theta)}\right)^\top\right]$$

where the quantity on the right hand side is the *Fisher information*, usually denoted with F. We can see the Fisher information as a measure of curvature for the log-likelihood function.

But how will we use this information in optimization? One can think of the immediate application of F as a replacement for Hessian in second-order methods. Nevertheless, it is not apparent why and under which settings.

One key difference so far in our narrative is the introduction of a probability distribution $p(x|\theta)$. In our course thus far, we have discussed about *deterministic optimization*: We are given an objective $f(x)$ that is usually differentiable, and we try to find the minimum/maximum. However, "under the rag", we have implied that any step we perform will be measured in the Euclidean space. To see this, assume that $\mathcal{L}(\theta)$ denotes the negative log-likelihood we want to minimize. Gradient descent is one way: we compute the direction d on the parameter space θ that minimizes the objective. Formally, we can find that the best direction is the negative gradient:

$$\lim_{\epsilon\to 0}\left(\tfrac{1}{\epsilon}\arg\min_{\|d\|\leq\epsilon} \mathcal{L}(\theta+d)\right) = -\frac{\nabla\mathcal{L}(\theta)}{\|\nabla\mathcal{L}(\theta)\|_2}.$$

i.e., the direction (*that is why we have normalization; we care about the direction, not how far we go on this direction*) of minimum drop on $\mathcal{L}(\cdot)$ is the negative gradient.

By definition of this steepest descent direction, though, we use the *Euclidean norm*. Thus, the optimization in gradient descent depends on the Euclidean geometry of the parameter space.

Though we have introduced the notion of likelihoods and expectations, stated differently, we have different objectives by minimizing the negative log-likelihood loss function, and it is natural to think of steps in the space of all possible likelihood, realizable by parameter θ.

The Kullback-Leibler divergence. First, we need to define the notion of the Kullback-Leibler (KL) divergence metric.

Definition 28. (Kullback-Leibler divergence) *Let $p_1(\cdot)$, $p_2(\cdot)$ be two distributions. Then, the KL divergence is given by:*

$$D_{KL}\left(p_1(\cdot)||p_2(\cdot)\right) = \mathbb{E}_{p_1(\cdot)}\left[\log\tfrac{p_1(\cdot)}{p_2(\cdot)}\right].$$

Remark 4. *Intuitively, the KL-divergence measures the "closeness" of two distributions. A little more rigorously, the KL-divergence is closely related to information theory; indeed, it is precisely the relative entropy between distributions $p(\cdot)$ and $q(\cdot)$. Under this lens, another interpretation of the metric is the information we gain when using $p(\cdot)$ instead of $q(\cdot)$.*

Remark 5. *The KL-divergence "metric" is not a valid distance metric in the measure-theoretic sense, as it is not symmetric nor satisfies the triangle inequality.*

To give an example, consider two Gaussians, with their means fixed in the two plots to follow (Figure 36; borrowed from `https://wiseodd.github.io/techblog/2018/03/14/natural-gradient`), but with different variances.

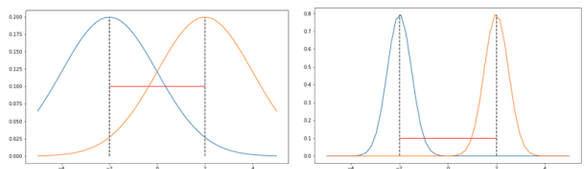

Fig. 36. Graphical illustration of distance metrics between distributions.

If we were to use the Euclidean distance between the means as a metric between the two distributions, it is clear that the metric would be the same between the two cases: since the means remain the same, their ℓ_2-norm distance is the same. Nevertheless, the distributions in these two plots are different; thus, we should find a metric that mirrors this in its definition.

Properties of the KL divergence: Observe that the KL divergence, by definition, satisfies:

$$\begin{aligned}D_{\mathrm{KL}}\left(p_1(\cdot)||p_2(\cdot)\right) &= \mathbb{E}_{p_1(\cdot)}\left[\log\tfrac{p_1(\cdot)}{p_2(\cdot)}\right] \\ &= \mathbb{E}_{p_1(\cdot)}\left[\log p_1(\cdot)\right] - \mathbb{E}_{p_1(\cdot)}\left[\log p_2(\cdot)\right].\end{aligned}$$

To help our discussion and make connections with our problem so far, we use $p_1(\cdot) = p(x|\theta)$ and $p_2(\cdot) = p(x|\theta')$. Then, the gradient of KL with respect to θ' satisfies:

$$\begin{aligned}&\nabla_{\theta'}\left(D_{\mathrm{KL}}\left(p(x|\theta)||p(x|\theta')\right)\right) \\ &= \nabla_{\theta'}\mathbb{E}_{p(x|\theta)}[\log p(x|\theta)] - \nabla_{\theta'}\mathbb{E}_{p(x|\theta)}[\log p(x|\theta')] \\ &= -\mathbb{E}_{p(x|\theta)}[\nabla_{\theta'}\log p(x|\theta')] \\ &= -\int p(x|\theta)\nabla_{\theta'}\log p(x|\theta')\,\mathrm{d}x.\end{aligned}$$

and the second derivative satisfies:

$$\nabla^2_{\theta'}\left(D_{\mathrm{KL}}\left(p(x|\theta)||p(x|\theta')\right)\right) = -\int p(x|\theta)\,\nabla^2_{\theta'}\log p(x|\theta')\,\mathrm{d}x$$

Then, the Hessian evaluated at $\theta' = \theta$ is:

$$\begin{aligned}\mathrm{H}_{D_{\mathrm{KL}}(p(x|\theta)||p(x|\theta'))} &= -\int p(x|\theta)\ \nabla^2_{\theta'}\log p(x|\theta')\big|_{\theta'=\theta}\,\mathrm{d}x \\ &= -\int p(x|\theta)\,\mathrm{H}_{\log p(x|\theta)}\,\mathrm{d}x \\ &= -\mathop{\mathbb{E}}_{p(x|\theta)}[\mathrm{H}_{\log p(x|\theta)}] \\ &= F,\end{aligned}$$

which is what we have shown above; i.e., the expected Hessian of the log function is the Fisher information matrix.

2nd-oder Taylor expansion of KL divergence and natural gradient: Let us now connect the dots. Following similar reasoning to classical optimization, given an objective function, we can locally approximate the objective with its second-order Taylor approximation (which involves both the gradient and the Hessian information) and then locally minimize that approximation; then, we iterate.

The second-order approximation of the KL divergence metric satisfies:

$$\begin{aligned} &D_{\mathrm{KL}}\left(p(x|\theta)||p(x|\theta+d)\right) \\ &\approx D_{\mathrm{KL}}\left(p(x|\theta)||p(x|\theta)\right) + \left\langle \nabla D_{\mathrm{KL}}\left(p(x|\theta)||p(x|\theta)\right), d\right\rangle + \frac{1}{2}\left\langle Fd, d\right\rangle \\ &= \frac{1}{2}\left\langle Fd, d\right\rangle. \end{aligned}$$

Similar to the Euclidean case, we seek an update vector d that minimizes the loss function $\mathcal{L}(\theta)$ in the *distribution space*. Analogously to steepest descent:

$$d^\star = \mathrm{argmin}_{D_{\mathrm{KL}}(p(x|\theta)||p(x|\theta+d))=c}\mathcal{L}(\theta+d),$$

where c is some constant. Compare this with the Euclidean case where:

$$d^\star = \mathrm{argmin}_{\|d\|_2\leq\epsilon}\mathcal{L}(\theta+d).$$

The purpose of fixing the KL-divergence to some constant is to ensure that we move along the space of distributions with constant speed, regardless of the curvature.

How do we solve this part? If we write the above minimization in Lagrangian form, we get:

$$\begin{aligned} d^\star &= \arg\min_d \left\{\mathcal{L}(\theta+d) + \lambda\cdot\left(D_{\mathrm{KL}}\left(p(x|\theta)||p(x|\theta+d)\right) - c\right)\right\} \\ &\approx \arg\min_d \left\{\mathcal{L}(\theta) + \left\langle\nabla\mathcal{L}(\theta),\ d\right\rangle + \frac{1}{2}\lambda\left\langle Fd,\ d\right\rangle - \lambda c\right\}. \end{aligned}$$

To solve this minimization, we set its derivative with respect to d to zero; this will lead to the solution:

$$d = -\tfrac{1}{\lambda}F^{-1}\nabla\mathcal{L}(\theta).$$

The above leads to the definition of the *natural gradient* descent method:

- **Repeat:**
  1. Compute the gradient $\nabla\mathcal{L}(\theta_t)$.
  2. Compute the Fisher information matrix F_t.
  3. Compute the natural gradient direction: $d_t = F_t^{-1}\nabla\mathcal{L}(\theta_t)$.
  4. For a step size η, compute $\theta_{t+1} = \theta_t - \eta d_t$.

Take-away messages:

- The natural gradient descent is a generalization of Newton's method, as there are cases where we can obtain Newton's iteration from the natural gradient descent.
- Remember that, after all, the Fisher information matrix is computed per iteration and inverted. It turns out that on expectation, it corresponds to the expected Hessian of the objective.
- How do we implement the natural gradient method in reality? Remember the definition of the Fisher information:

$$F = -\mathop{\mathbb{E}}_{p(x|\theta)}[H_{\log p(x|\theta)}].$$

In realistic scenarios, we do not have access to the distribution $p(x|\theta)$ but rather have data from that distribution. In that case, we refer to the *empirical* Fisher information matrix, defined as:

$$F = \tfrac{1}{n}\sum_{i=1}^{n}\nabla\log p(x_i|\theta)\cdot\nabla\log p(x_i|\theta)^\top$$

where $\{x_i\}_{i=1}^n$ denote the training set of examples. In that case, the main recursion of natural gradient descent becomes:

$$\theta_{t+1} = \theta_t - \eta\left(\tfrac{1}{n}\sum_{i=1}^{n}\nabla\log p(x_i|\theta_t)\cdot\nabla\log p(x_i|\theta_t)^\top\right)^{-1}\cdot\nabla\widehat{\mathcal{L}}(\theta_t),$$

where also the objective and its gradient are evaluated in their empirical form:

$$\nabla\widehat{\mathcal{L}}(\theta_t) := \tfrac{1}{n}\sum_{i=1}^{n}\nabla\log p(x_i|\theta_t).$$

- The main reason we studied natural gradient descent is to motivate our discussion later on regarding algorithms in training neural networks. The empirical Fisher information matrix appears in almost all modern algorithms in ML, usually further approximated to be easily computed (e.g., one way to get around computing the exact empirical Fisher information matrix is to constrain it to be a diagonal matrix, a technique that is heavily used in algorithms such as AdaGrad, AdaDelta, RMSprop, Adam, AMSGrad, Yogi, etc. In other words, it is the most used algorithm in neural network training).

———— ∞ ————

Zeroth order methods, a.k.a. derivative-free methods. Until now, we have only analyzed algorithms adhering to the "black-box" optimization model, where we have some local oracle O that we rely on for information at each update step. In this chapter, we looked at oracles providing us $H_t = \nabla^2 f(x_t)$, either by directly computing it (Newton's method) or approximating it (quasi-Newton methods).

We now briefly turn to a class of algorithms with barely any oracle: we are only allowed a zeroth-order oracle. In other words, given a query point x, we can only query $f(x)$ from the oracle; we cannot access gradients or other higher-order information. Optimization problems fitting this criterion arise incredibly commonly in practice: perhaps our objective function is not differentiable, or we have no way of characterizing it in an analytical form (and thus cannot compute higher-order information analytically), or it is simply too computationally expensive to compute gradients. In these cases, we still must find a way to guarantee convergence using our limited zeroth-order oracle. Some examples of zeroth-order optimization algorithms are the bisection method, genetic algorithms, simulated annealing, Metropolis methods, etc.

So, how do we even start with such little information to work with? As before, calculus comes to the rescue yet again. Recall the definition of the derivative:

$$f'(x) = \lim_{h\to 0}\frac{f(x+h)-f(x)}{h}$$

A natural approximation arises from the definition, for some $\epsilon \ll 1$:

$$f'(x) \approx \frac{f(x+\epsilon)-f(x)}{\epsilon}$$

So, we can approximate the derivative using just the function itself! This realization forms the core of the *finite differences method.*

Definition 29. (Finite differences method) *For $x_t \in \mathbb{R}^p$, we update each iteration of via the rule*

$$x_{t+1} = x_t - \eta_t\left(\frac{f(x_t+\mu_t u)-f(x_t)}{\mu_t}\right)\cdot u$$

Application: Adversarial examples in neural net training. Searching for generalizable models is a holy grail topic in modern-day machine learning research. Unfortunately, we still have a long way to go — despite many recent advances in neural net architecture and algorithms, they still fail to generalize flexibly to examples one would intuitively "expect" an ideal, generalized model to classify accurately. As a damning example of this claim, we briefly investigate the idea of adversarial examples.

The idea of adversarial examples is as follows: if you take a valid input but then do a small perturbation, the classically trained neural net models will suddenly go nuts and no longer give the correct answer.

More specifically, given a valid input x, we can create an adversarial example x_{adv} by the update rule

$$x_{\text{adv}} = x + \epsilon \cdot \text{sign}(\nabla f(x, y_{\text{true}}))$$

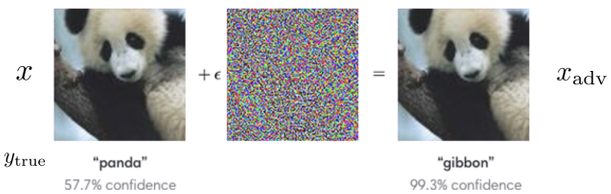

Fig. 37. An illustration of how an adversarial perturbation can lead to misclassification.

Intuitively, one can read this as the adversary is looking to move the input in directions away from the optimal minimum. The problem we are trying to solve now is to defend against a series of adversarial attacks. In other words, we wish to make our model robust against maliciously crafted input. One such defense mechanism is to "obfuscate" the gradient information. In other words, we prevent access to the black box (e.g., the back-propagation gradient computation).

However, the forward operations remain intact, so the function evaluations are usually computed. Then, given this, we can do the finite differences method as an attacker to approximate the gradient and still do our attack normally. This is the basis of the SPSA attack (Simultaneous Perturbation Stochastic Approximation).

1. J. Nocedal and S. Wright. Numerical optimization. Springer Science & Business Media, 2006.
2. Y. Nesterov. Introductory lectures on convex optimization: A basic course, volume 87. Springer Science & Business Media, 2013.
3. S. Boyd and L. Vandenberghe. Convex optimization. Cambridge university press, 2004.
4. D. Bertsekas. Convex optimization algorithms. Athena Scientific Belmont, 2015.
5. Sébastien Bubeck. Convex optimization: Algorithms and complexity. Foundations and Trends® in Machine Learning, 8(3-4):231–357, 2015.
6. S. Weisberg. Applied linear regression, volume 528. John Wiley & Sons, 2005.
7. T. Hastie, R. Tibshirani, and M. Wainwright. Statistical learning with sparsity: the lasso and generalizations. CRC press, 2015.
8. J. Friedman, T. Hastie, and R. Tibshirani. The elements of statistical learning, volume 1. Springer series in statistics New York, 2001.
9. M. Paris and J. Rehacek. Quantum state estimation, volume 649. Springer Science & Business Media, 2004.
10. M. Daskin. A maximum expected covering location model: formulation, properties and heuristic solution. Transportation science, 17(1):48–70, 1983.
11. I. Goodfellow, Y. Bengio, and A. Courville. Deep learning. MIT press, 2016.
12. L. Trefethen and D. Bau III. Numerical linear algebra, volume 50. Siam, 1997.
13. G. Strang. Introduction to linear algebra, volume 3. Wellesley-Cambridge Press Wellesley, MA, 1993.
14. G. Golub. Cmatrix computations. The Johns Hopkins, 1996.
15. Yann LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In Neural networks: Tricks of the trade, pages 9–50. Springer, 2002.
16. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In Proceedings of the IEEE international conference on computer vision, pages 1026–1034, 2015.
17. Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. Advances in neural information processing systems, 31, 2018.
18. A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.
19. K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.
20. S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In Advances in neural information processing systems, pages 91–99, 2015.
21. T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In Advances in neural information processing systems, pages 3111–3119, 2013.
22. Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473, 2014.
23. Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. In Proceedings of the 34th International Conference on Machine Learning-Volume 70, pages 1243–1252. JMLR. org, 2017.
24. Haşim Sak, Andrew Senior, and Françoise Beaufays. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In Fifteenth annual conference of the international speech communication association, 2014.
25. Tom Sercu, Christian Puhrsch, Brian Kingsbury, and Yann LeCun. Very deep multilingual convolutional neural networks for LVCSR. In 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 4955–4959. IEEE, 2016.
26. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. page arXiv:1706.03762, 2017.
27. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. page arXiv:1810.04805, 2018.
28. Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and VQA. In AAAI, pages 13041–13049, 2020.
29. Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. arXiv preprint arXiv:2005.14165, 2020.
30. Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using gpu model parallelism. arXiv preprint arXiv:1909.08053, 2019.
31. Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. arXiv preprint arXiv:1910.10683, 2019.
32. Gary Marcus, Ernest Davis, and Scott Aaronson. A very preliminary analysis of DALL-E 2. arXiv preprint arXiv:2204.13807, 2022.
33. John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with AlphaFold. Nature, 596(7873):583–589, 2021.
34. Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020.
35. Or Sharir, Barak Peleg, and Yoav Shoham. The cost of training nlp models: A concise overview. arXiv preprint arXiv:2004.08900, 2020.
36. H. Karimi, J. Nutini, and M. Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In Joint European Conference on Machine Learning and Knowledge Discovery in Databases, pages 795–811. Springer, 2016.
37. Philip Wolfe. Convergence conditions for ascent methods. SIAM review, 11(2):226–235, 1969.
38. Larry Armijo. Minimization of functions having Lipschitz continuous first partial derivatives. Pacific Journal of mathematics, 16(1):1–3, 1966.
39. Stephen Wright and Jorge Nocedal. Numerical optimization. Springer Science, 35(67-68):7, 1999.
40. B. Polyak. Introduction to optimization. Inc., Publications Division, New York, 1, 1987.
41. Stephen Boyd, Lin Xiao, and Almir Mutapcic. Subgradient methods. lecture notes of EE392o, Stanford University, Autumn Quarter, 2004:2004–2005, 2003.
42. Marguerite Frank, Philip Wolfe, et al. An algorithm for quadratic programming. Naval research logistics quarterly, 3(1-2):95–110, 1956.
43. M. Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In Proceedings of the 30th international conference on machine learning, number CONF, pages 427–435, 2013.
44. J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the ℓ_1-ball for learning in high dimensions. In Proceedings of the 25th international conference on Machine learning, pages 272–279, 2008.
45. Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. Computer, (8):30–37, 2009.
46. A. Mnih and R. Salakhutdinov. Probabilistic matrix factorization. In Advances in neural information processing systems, pages 1257–1264, 2008.
47. T. Booth and J. Gubernatis. Improved criticality convergence via a modified Monte Carlo power iteration method. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.
48. S. Zavriev and F. Kostyuk. Heavy-ball method in nonconvex optimization problems. Computational Mathematics and Modeling, 4(4):336–341, 1993.
49. E. Ghadimi, H. Feyzmahdavian, and M. Johansson. Global convergence of the heavy-ball method for convex optimization. In 2015 European control conference (ECC), pages 310–315. IEEE, 2015.
50. Y. Nesterov. A method of solving a convex programming problem with convergence rate $O(\frac{1}{k^2})$. In Soviet Mathematics Doklady, volume 27, pages 372–376, 1983.
51. B. O'Donoghue and E. Candes. Adaptive restart for accelerated gradient schemes. Foundations of computational mathematics, 15(3):715–732, 2015.
52. O. Devolder, F. Glineur, and Y. Nesterov. First-order methods of smooth convex optimization with inexact oracle. Mathematical Programming, 146(1-2):37–75, 2014.
53. L. Bottou, F. Curtis, and J. Nocedal. Optimization methods for large-scale machine learning. Siam Review, 60(2):223–311, 2018.
54. S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. SIAM review, 43(1):129–159, 2001.
55. R. Tibshirani. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society: Series B (Methodological), 58(1):267–288, 1996.
56. P. Hoff. Lasso, fractional norm and structured sparse estimation using a Hadamard product parametrization. Computational Statistics & Data Analysis, 115:186–198, 2017.
57. S. Becker, J. Bobin, and E. Candès. NESTA: A fast and accurate first-order method for sparse recovery. SIAM Journal on Imaging Sciences, 4(1):1–39, 2011.
58. T. Blumensath and M. Davies. Iterative hard thresholding for compressed sensing. Applied and computational harmonic analysis, 27(3):265–274, 2009.
59. D. Needell and J. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. Applied and computational harmonic analysis, 26(3):301–321, 2009.
60. S. Foucart. Hard thresholding pursuit: an algorithm for compressive sensing. SIAM Journal on Numerical Analysis, 49(6):2543–2563, 2011.
61. J. Tanner and K. Wei. Normalized iterative hard thresholding for matrix completion. SIAM Journal on Scientific Computing, 35(5):S104–S125, 2013.
62. K. Wei. Fast iterative hard thresholding for compressed sensing. IEEE Signal processing letters, 22(5):593–597, 2014.
63. Rajiv Khanna and Anastasios Kyrillidis. Iht dies hard: Provable accelerated iterative hard thresholding. In International Conference on Artificial Intelligence and Statistics, pages 188–198. PMLR, 2018.
64. Jeffrey D Blanchard and Jared Tanner. GPU accelerated greedy algorithms for compressed sensing. Mathematical Programming Computation, 5(3):267–304, 2013.
65. A. Kyrillidis, G. Puy, and V. Cevher. Hard thresholding with norm constraints. In 2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 3645–3648. Ieee, 2012.
66. A. Kyrillidis and V. Cevher. Recipes on hard thresholding methods. In Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2011 4th IEEE International Workshop on, pages 353–356. IEEE, 2011.

67. X. Zhang, Y. Yu, L. Wang, and Q. Gu. Learning one-hidden-layer ReLU networks via gradient descent. In The 22nd International Conference on Artificial Intelligence and Statistics, pages 1524–1534, 2019.
68. Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. IEEE Transactions on information theory, 52(2):489–509, 2006.
69. Joachim Dahl, Lieven Vandenberghe, and Vwani Roychowdhury. Covariance selection for nonchordal graphs via chordal embedding. Optimization Methods & Software, 23(4):501–520, 2008.
70. Joseph B Altepeter, Daniel FV James, and Paul G Kwiat. 4 qubit quantum state tomography. In Quantum state estimation, pages 113–145. Springer, 2004.
71. Jens Eisert, Dominik Hangleiter, Nathan Walk, Ingo Roth, Damian Markham, Rhea Parekh, Ulysse Chabaud, and Elham Kashefi. Quantum certification and benchmarking. arXiv preprint arXiv:1910.06343, 2019.
72. Masoud Mohseni, AT Rezakhani, and DA Lidar. Quantum-process tomography: Resource analysis of different strategies. Physical Review A, 77(3):032322, 2008.
73. D. Gross, Y.-K. Liu, S. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. Physical review letters, 105(15):150401, 2010.
74. Y.-K. Liu. Universal low-rank matrix recovery from Pauli measurements. In Advances in Neural Information Processing Systems, pages 1638–1646, 2011.
75. K Vogel and H Risken. Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. Physical Review A, 40(5):2847, 1989.
76. Miroslav Ježek, Jaromír Fiurášek, and Zdeněk Hradil. Quantum inference of states and processes. Physical Review A, 68(1):012305, 2003.
77. Konrad Banaszek, Marcus Cramer, and David Gross. Focus on quantum tomography. New Journal of Physics, 15(12):125020, 2013.
78. A. Kalev, R. Kosut, and I. Deutsch. Quantum tomography protocols with positivity are compressed sensing protocols. Nature partner journals (npj) Quantum Information, 1:15018, 2015.
79. Giacomo Torlai, Guglielmo Mazzola, Juan Carrasquilla, Matthias Troyer, Roger Melko, and Giuseppe Carleo. Neural-network quantum state tomography. Nat. Phys., 14:447–450, May 2018.
80. Matthew JS Beach, Isaac De Vlugt, Anna Golubeva, Patrick Huembeli, Bohdan Kulchytskyy, Xiuzhe Luo, Roger G Melko, Ejaaz Merali, and Giacomo Torlai. Qucumber: wavefunction reconstruction with neural networks. SciPost Physics, 7(1):009, 2019.
81. Giacomo Torlai and Roger Melko. Machine-learning quantum states in the NISQ era. Annual Review of Condensed Matter Physics, 11, 2019.
82. M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu. Efficient quantum state tomography. Nat. Comm., 1:149, 2010.
83. BP Lanyon, C Maier, Milan Holzäpfel, Tillmann Baumgratz, C Hempel, P Jurcevic, Ish Dhand, AS Buyskikh, AJ Daley, Marcus Cramer, et al. Efficient tomography of a quantum many-body system. Nature Physics, 13(12):1158–1162, 2017.
84. D. Gonçalves, M. Gomes-Ruggiero, and C. Lavor. A projected gradient method for optimization over density matrices. Optimization Methods and Software, 31(2):328–341, 2016.
85. E. Bolduc, G. Knee, E. Gauger, and J. Leach. Projected gradient descent algorithms for quantum state tomography. npj Quantum Information, 3(1):44, 2017.
86. Jiangwei Shang, Zhengyun Zhang, and Hui Khoon Ng. Superfast maximum-likelihood reconstruction for quantum tomography. Phys. Rev. A, 95:062336, Jun 2017.
87. Zhilin Hu, Kezhi Li, Shuang Cong, and Yaru Tang. Reconstructing pure 14-qubit quantum states in three hours using compressive sensing. IFAC-PapersOnLine, 52(11):188 – 193, 2019. 5th IFAC Conference on Intelligent Control and Automation Sciences ICONS 2019.
88. Zhibo Hou, Han-Sen Zhong, Ye Tian, Daoyi Dong, Bo Qi, Li Li, Yuanlong Wang, Franco Nori, Guo-Yong Xiang, Chuan-Feng Li, et al. Full reconstruction of a 14-qubit state within four hours. New Journal of Physics, 18(8):083036, 2016.
89. C. Riofrío, D. Gross, S.T. Flammia, T. Monz, D. Nigg, R. Blatt, and J. Eisert. Experimental quantum compressed sensing for a seven-qubit system. Nature Communications, 8, 2017.
90. Martin Kliesch, Richard Kueng, Jens Eisert, and David Gross. Guaranteed recovery of quantum processes from few measurements. Quantum, 3:171, 2019.
91. S. Flammia, D. Gross, Y.-K. Liu, and J. Eisert. Quantum tomography via compressed sensing: Error bounds, sample complexity and efficient estimators. New Journal of Physics, 14(9):095022, 2012.
92. A. Kyrillidis, A. Kalev, D. Park, S. Bhojanapalli, C. Caramanis, and S. Sanghavi. Provable quantum state tomography via non-convex methods. npj Quantum Information, 4(36), 2018.
93. B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. SIAM review, 52(3):471–501, 2010.
94. N. Srebro, J. Rennie, and T. Jaakkola. Maximum-margin matrix factorization. In Advances in neural information processing systems, pages 1329–1336, 2004.
95. J. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In Proceedings of the 22nd international conference on Machine learning, pages 713–719. ACM, 2005.
96. D. DeCoste. Collaborative prediction using ensembles of maximum margin matrix factorizations. In Proceedings of the 23rd international conference on Machine learning, pages 249–256. ACM, 2006.
97. J. Bennett and S. Lanning. The Netflix prize. In Proceedings of KDD cup and workshop, volume 2007, page 35, 2007.
98. M. Jaggi and M. Sulovsk. A simple algorithm for nuclear norm regularized problems. In Proceedings of the 27th International Conference on Machine Learning (ICML-10), pages 471–478, 2010.
99. R. Keshavan. Efficient algorithms for collaborative filtering. PhD thesis, Stanford University, 2012.
100. R. Agrawal, A. Gupta, Y. Prabhu, and M. Varma. Multi-label learning with millions of labels: Recommending advertiser bid phrases for web pages. In Proceedings of the 22nd international conference on World Wide Web, pages 13–24. International World Wide Web Conferences Steering Committee, 2013.
101. K. Bhatia, H. Jain, P. Kar, M. Varma, and P. Jain. Sparse local embeddings for extreme multi-label classification. In Advances in Neural Information Processing Systems, pages 730–738, 2015.
102. G. Carneiro, A. Chan, P. Moreno, and N. Vasconcelos. Supervised learning of semantic classes for image annotation and retrieval. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 29(3):394–410, 2007.
103. A. Makadia, V. Pavlovic, and S. Kumar. A new baseline for image annotation. In Computer Vision–ECCV 2008, pages 316–329. Springer, 2008.
104. C. Wang, S. Yan, L. Zhang, and H.-J. Zhang. Multi-label sparse coding for automatic image annotation. In Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on, pages 1643–1650. IEEE, 2009.
105. J. Weston, S. Bengio, and N. Usunier. WSABIE: Scaling up to large vocabulary image annotation. In IJCAI, volume 11, pages 2764–2770, 2011.
106. Andrew I. Schein, Lawrence K. Saul, and Lyle H. Ungar. A generalized linear model for principal component analysis of binary data. In AISTATS, 2003.
107. K.-Y. Chiang, C.-J. Hsieh, N. Natarajan, I. Dhillon, and A. Tewari. Prediction and clustering in signed networks: A local to global perspective. The Journal of Machine Learning Research, 15(1):1177–1213, 2014.
108. C. Johnson. Logistic matrix factorization for implicit feedback data. Advances in Neural Information Processing Systems, 27, 2014.
109. Koen Verstrepen. Collaborative Filtering with Binary, Positive-only Data. PhD thesis, University of Antwerpen, 2015.
110. N. Gupta and S. Singh. Collectively embedding multi-relational data for predicting user preferences. arXiv preprint arXiv:1504.06165, 2015.
111. Y. Liu, M. Wu, C. Miao, P. Zhao, and X.-L. Li. Neighborhood regularized logistic matrix factorization for drug-target interaction prediction. PLoS Computational Biology, 12(2):e1004760, 2016.
112. S. Aaronson. The learnability of quantum states. In Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, volume 463, pages 3089–3114. The Royal Society, 2007.
113. E. Candes, Y. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. SIAM Review, 57(2):225–251, 2015.
114. I. Waldspurger, A. d'Aspremont, and S. Mallat. Phase recovery, MaxCut and complex semidefinite programming. Mathematical Programming, 149(1-2):47–81, 2015.
115. P. Biswas, T.-C. Liang, K.-C. Toh, Y. Ye, and T.-C. Wang. Semidefinite programming approaches for sensor network localization with noisy distance measurements. IEEE transactions on automation science and engineering, 3(4):360, 2006.
116. K. Weinberger, F. Sha, Q. Zhu, and L. Saul. Graph Laplacian regularization for large-scale semidefinite programming. In Advances in Neural Information Processing Systems, pages 1489–1496, 2007.
117. F. Lu, S. Keles, S. Wright, and G. Wahba. Framework for kernel regularization with application to protein clustering. Proceedings of the National Academy of Sciences of the United States of America, 102(35):12332–12337, 2005.
118. H. Andrews and C. Patterson III. Singular value decomposition (SVD) image coding. Communications, IEEE Transactions on, 24(4):425–432, 1976.
119. M. Fazel, H. Hindi, and S. Boyd. Rank minimization and applications in system theory. In American Control Conference, 2004. Proceedings of the 2004, volume 4, pages 3273–3278. IEEE, 2004.
120. E. Candès and B. Recht. Exact matrix completion via convex optimization. Foundations of Computational mathematics, 9(6):717–772, 2009.
121. P. Jain, R. Meka, and I. Dhillon. Guaranteed rank minimization via singular value projection. In Advances in Neural Information Processing Systems, pages 937–945, 2010.
122. S. Becker, V. Cevher, and A. Kyrillidis. Randomized low-memory singular value projection. In 10th International Conference on Sampling Theory and Applications (Sampta), 2013.
123. L. Balzano, R. Nowak, and B. Recht. Online identification and tracking of subspaces from highly incomplete information. In Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on, pages 704–711. IEEE, 2010.
124. K. Lee and Y. Bresler. ADMiRA: Atomic decomposition for minimum rank approximation. Information Theory, IEEE Transactions on, 56(9):4402–4416, 2010.
125. A. Kyrillidis and V. Cevher. Matrix recipes for hard thresholding methods. Journal of mathematical imaging and vision, 48(2):235–265, 2014.
126. Z. Lin, M. Chen, and Y. Ma. The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices. arXiv preprint arXiv:1009.5055, 2010.
127. S. Becker, E. Candès, and M. Grant. Templates for convex cone problems with applications to sparse signal recovery. Mathematical Programming Computation, 3(3):165–218, 2011.
128. J. Cai, E. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. SIAM Journal on Optimization, 20(4):1956–1982, 2010.
129. Y. Chen, S. Bhojanapalli, S. Sanghavi, and R. Ward. Coherent matrix completion. In Proceedings of The 31st International Conference on Machine Learning, pages 674–682, 2014.

130. A. Yurtsever, Q. Tran-Dinh, and V. Cevher. A universal primal-dual convex optimization framework. In Advances in Neural Information Processing Systems 28, pages 3132–3140. 2015.
131. F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. Psychological Review, 65(6):386–408, 1958.
132. Robin M. Schmidt, Frank Schneider, and Philipp Hennig. Descending through a crowded valley - benchmarking deep learning optimizers. CoRR, abs/2007.01547, 2020.
133. John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. J. Mach. Learn. Res., 12(null):2121–2159, jul 2011.
134. Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Marc' aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, Quoc Le, and Andrew Ng. Large scale distributed deep networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, Advances in Neural Information Processing Systems, volume 25. Curran Associates, Inc., 2012.
135. Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015.