

Gradient descent is a very different algorithm for solving linear equations, and more. It is the main algorithm responsible for training machine learning models.

Gradient descent is an optimization algorithm. The goal of optimization is to find the best possible solution among many. How do we phrase solving linear equations as an optimization problem? The idea is to set up some “loss function” f that penalizes guesses which are far from being correct. The minimum of the loss function is attained only by the correct solution(s).

Let’s start with a silly but enlightening example. The equation is $2x = 3$. We want to cook up a function f that is uniquely minimized at the solution point $x = 3/2$. What would be the simplest function with this property? Linear functions won’t work as they favor one of the infinities. The next simplest type is quadratic and we are in luck: The function $(x - 3/2)^2$ has the solution as its unique minimizer. (Another option is the piecewise linear function $|x - 3/2|$. We’ll come back to it later.)

We want to set up f to solve the equation $2x = 3$. But in order to write up f we had to figure out the solution $3/2$ first. This isn’t as outlandish as it looks. We didn’t really have to solve any equation. If we “unsolve” it we get $(2(x - 3/2))^2 = (2x - 3)^2$ which is left minus right hand side squared. This function is minimized precisely at $x = 3/2$.

To set up a loss function for a problem with multiple constraints like solving a system of equations, it is natural to add up the losses of the individual constraints. The resulting function will be minimized exactly when all constraints are satisfied.

To summarize, one loss function for a linear system is the sum of left minus right hand sides squared. This is called the *sum of squares loss*. For last lecture’s example,

$$\begin{aligned} x - 3y - 3z &= -5 \\ x - 2y - 2z &= 0 \\ x + y - 5z &= 3 \end{aligned} \tag{1}$$


the sum of squares loss is

$$f(x, y, z) = (x - 3y - 3z + 5)^2 + (x - 2y - 2z)^2 + (x + y - 5z - 3)^2. \tag{2}$$


It is useful to think of (x, y, z) as the position of a particle. Gradient descent starts with this particle somewhere, say at $(0, 0, 0)$, and keeps moving it in a direction in which the sum of squares loss decreases. When it is convinced that no decrease is possible it declares that the minimum has been found. The name of the algorithm comes from direction of motion: the *gradient* of f , with a minus sign in front.

1 Gradient descent

Let’s start with $f(x) = (2x - 3)^2$. The gradient of a univariate function is the derivative $df/dx = 4(2x - 3)$. Here is a plot of $-0.05 \cdot df(x)/dx$ at various points along the curve $(x, f(x))$:



All arrows point towards the minimum. This suggests we should move x by $-0.05 \cdot df(x)/dx$ units in each step. If the particle is initially at $x = 0$ it will trace this path on the graph $(x, f(x))$, eventually reaching close to the minimum at $x = 1.5$.




This is the gradient descent algorithm. The gradient of an n -variate function $f(\mathbf{x}) = f(x, y, \dots, z)$ is the vector of partial derivatives $\nabla f = (\partial f/\partial x, \partial f/\partial y, \dots, \partial f/\partial z)$.

Algorithm GD (Simple Gradient Descent)

Input: A function f with multiple inputs.

- 1 Choose a starting point $\mathbf{x} = (x, y, \dots, z)$.
- 2 Choose a rate $\rho > 0$.
- 3 Until $\mathbf{x}$ is sufficiently small,
- 4 Move $\mathbf{x}$ to $\mathbf{x} - \rho \cdot \nabla f(\mathbf{x})$.

Gradient Descent is different from Gaussian Elimination in many ways. The first glaring difference is that it is not “set in stone”: We need to choose the initial $\mathbf{x}$ in line 1 and the rate ρ in line 2. The rate controls the speed of motion. Set it too small and the particle will take forever to budge. Set it too high and it will go in the wrong direction! Had I chosen a rate of 0.255 instead of 0.05 this is the path it would have taken:



Even if the parameters are chosen well, Gradient Descent never quite arrives at its destination. In line 3 we have to decide when to declare victory. There are guidelines that might govern these choices, but it is difficult to get a feel about Gradient Descent without some trial and error. Let's see what happens when we apply it to system (1).

The first task is to calculate the gradient of the square loss. Applying the usual rules from calculus,

$$\begin{aligned}\partial f / \partial x &= 2 \cdot (x - 3y - 3z + 5) + 2 \cdot (x - 2y - 2z) + 2(x + y - 5z - 3) \\ \partial f / \partial y &= 2 \cdot (-3)(x - 3y - 3z + 5) + 2 \cdot (-2)(x - 2y - 2z) + 2 \cdot (x + y - 5z - 3) \\ \partial f / \partial z &= 2 \cdot (-3)(x - 3y - 3z + 5) + 2 \cdot (-2)(x - 2y - 2z) + 2 \cdot (-5)(x + y - 5z - 3).\end{aligned}$$

Each partial derivative is a linear combination of the equations, with the coefficient equal twice to the one in front of the corresponding variable. This can be written succinctly in matrix notation. If the system has the form $A\mathbf{x} = b$ and $f(\mathbf{x}) = \|A\mathbf{x} - b\|^2$ then

$$\nabla f(\mathbf{x}) = 2A^T(A\mathbf{x} - b), \quad (3)$$

where A^T is the matrix transpose of A . (Later in the course we will derive this formula as an instance of backpropagation.) Now we can start playing.

Let the olympic games begin!

We first need to initialize $\mathbf{x}$. It is reasonable to expect that the closer we start to the minimum the faster we will get there. From last week's problem formulation we know that the variables represent people's ages so they should not be negative or too large. Let's set the initial $\mathbf{x}$ to $(0, 0, 0)$.

How about the rate? Let's try to get a sense of scale. The rate is the multiple of the gradient by which the particle moves. The solution should be single or double-digit numbers. This suggests that the rate should be roughly inverse-proportional to the norm (the length) of the gradient. We can get a rough estimate by evaluating the gradient at the initial point. In our example, formula (3) gives $\nabla f(0, 0, 0) = 2A^Tb = (-4, 36, 0)^T$, which has norm about 36. This suggests $\rho \approx 1/36 \approx 0.03$.

A few trial runs reveal that gradient descent diverges on system (1) with $\rho = 0.03$, but seems to make progress when ρ is set to 0.02. Here are the first ten steps:

step	x	root(f(x))
1	: [0. 0. 0.]	5.831
2	: [-0.08 0.72 0.]	3.937
3	: [-0.035 1.024 -0.262]	3.178
4	: [-0.052 1.249 -0.205]	2.912
5	: [-0.008 1.327 -0.314]	2.822
6	: [-0.004 1.403 -0.265]	2.788
7	: [0.038 1.422 -0.311]	2.771

```

8   : [ 0.057  1.451 -0.278]  2.760
9   : [ 0.091  1.457 -0.297]  2.751
10  : [ 0.114  1.471 -0.275]  2.742

```

The particle moves rapidly in the first two steps but then progress slows down. It takes a long time for it to get close to the target solution $(10, 3, 2)$.


step	x	root(f(x))
200 :	[4.26 2.113 0.674]	1.592
400 :	[6.761 2.5 1.252]	0.898
600 :	[8.172 2.718 1.578]	0.507
800 :	[8.969 2.841 1.762]	0.286
1000 :	[9.418 2.91 1.866]	0.161
2000 :	[9.994 2.999 1.999]	0.002

Gradient descent looks very fickle, and takes thousands of steps to complete what an average sixth-grader can do in five minutes. Is it worth the trouble? Absolutely!

2 Gradient Descent fast and slow

Analyzing gradient descent is tricky. Fortunately we can learn quite a bit by looking at the one-variable equation $ax = b$. To visualize more easily let's assume that a and b are positive.

The most pressing question is why Gradient Descent works at all. The proof by picture from Section 1 should be convincing enough but let's try to quantify which rates ρ result in convergence. There is some precise value ρ^* at which the particle keeps bouncing between two points on the curve:



We can find ρ^* by solving an equation. The displacement between these two points is $2(b/a - x)$. Gradient Descent moves x by $-\rho \cdot df(x)/dx = -\rho \cdot 2a(ax - b)$, so ρ^* must be the value of ρ at which these two expressions are equal, namely

$$\rho^* = \frac{2(b/a - x)}{-2a(ax - b)} = \frac{1}{a^2}.$$

As long as $\rho < 1/a^2$, Gradient Descent will make progress.

To say something similar about general systems we need a single number that measures the “magnitude” of the matrix A . This is the *spectral norm* $\|A\|$. As long as $\rho < 1/\|A\|^2$, Gradient Descent is guaranteed to make progress. We'll define the spectral norm shortly.

The next question is how fast this convergence is. How much closer to the minimum does each step take us? If the particle moves from x to x' , then

$$x' = x - \rho \cdot 2a(ax - b).$$

Multiplying by a and subtracting b from both sides, we get

$$ax' - b = (ax - b) - a \cdot \rho \cdot 2a(ax - b) = (1 - 2\rho a^2)(ax - b).$$

The value $y = ax - b$ is the square root of the loss. In each step the root of the loss scales by a factor of $(1 - 2\rho a^2)$. When $|1 - 2\rho a^2|$ is less than one, or $\rho < 1/a^2$, the scaling causes the root-loss to shrink. This shrinking occurs at an *exponential* rate: After t steps, the root of the loss will be a $(1 - 2\rho a^2)^t$ -fraction of its initial value.

Exponential convergence rates are very desirable. They mean that every few steps we get closer to the solution by an extra digit of accuracy.

3 The spectral norm

For more equations and more unknowns, we can mimic the same analysis, but with vectors and matrices instead of numbers. The particle's position evolves by

$$\mathbf{x}' = \mathbf{x} - 2\rho A^T(A\mathbf{x} - b).$$

Multiplying by A and subtracting b from both sides gives

$$A\mathbf{x}' - b = (A\mathbf{x} - b) - A \cdot \rho \cdot 2A^T(A\mathbf{x} - b) = (I - 2\rho AA^T)(A\mathbf{x} - b).$$

The quantity $\mathbf{y} = A\mathbf{x} - b$ is now a vector, and $B = I - 2\rho AA^T$ is a matrix. We would again like to say that if ρ is small enough, this matrix B shrinks the vector $\mathbf{y}$. In general, the effect of applying a matrix on a vector $\mathbf{y}$ may depend not only on its magnitude, but also on its direction.

The spectral norm $\|B\|$ of a matrix measures the worst amount by which a matrix can grow a vector, regardless of its direction.

$$\|B\| = \max_{\mathbf{v} \neq 0} \frac{\|B\mathbf{v}\|}{\|\mathbf{v}\|}.$$

If the spectral norm is less than one, then every vector shrinks. In conclusion.

$$\frac{\sqrt{f(\mathbf{x}')}}{\sqrt{f(\mathbf{x})}} \leq \|I - 2\rho AA^T\|. \quad (4)$$

The progress is again exponential: After t steps, the square loss is at most $\|I - 2\rho AA^T\|^t$ of where it started.




FIGURE 1: Loss as a function of time at (a) regular (b) regular-log scale.

When A is the matrix of system (1) and $\rho = 0.02$, the matrix $B = I - 2\rho AA^T$ is

$$B = \begin{pmatrix} 0.24 & -0.52 & -0.52 \\ -0.52 & 0.64 & -0.36 \\ -0.52 & -0.36 & -0.08 \end{pmatrix}$$

and its spectral norm is $\|B\| \approx 0.997$, which is just under one. After t steps, the root-loss can be at most 0.997^t times the initial root loss.

The actual loss and this upper bound are plotted in Figure 1(a). There is some gap between the two, but the shape of the two curves is very similar. When we take the log of the loss, the dependence on the step is clearly linear as seen in Figure 1(b).

Unlike for univariate equations, there are multivariate systems for which the spectral norm of $B = I - 2\rho AA^T$ can never drop below one, no matter which rate ρ we choose. One example is the system

$$\begin{aligned} x + y &= 1 \\ x + y &= 1 \end{aligned} \tag{5}$$

This system is overdetermined. The matrix

$$B = I - 2\rho AA^T = \begin{pmatrix} 1 - 4\rho & -4\rho \\ -4\rho & 1 - 4\rho \end{pmatrix}$$

has spectral norm at least one, even when ρ is small. Formula (4) does not guarantee any progress for such systems. Nevertheless, it turns out that gradient is still guaranteed to converge at a solution, albeit at a rate that is only *inverse linear* in the time step.

4 The condition number

The reason that gradient descent may struggle with systems like (5) is that there are many possible solutions. To understand why we will shift attention from the loss function to the solution. The loss function measures how close the right-hand side $A\mathbf{x}$ is to the target vector b . Another measure of progress is the distance of the state $\mathbf{x}$ to the sought solution $\mathbf{x}^*$ for which $A\mathbf{x}^*$ equals b .

For a system like (5), $\mathbf{x}^*$ is not even well-defined. There are many possible solutions. If we slightly perturb the instance, however, $\mathbf{x}^*$ becomes unique, as in these examples:

$$S: \begin{aligned} x + 0.9 \cdot y &= 1 \\ x + 1.1 \cdot y &= 1 \end{aligned} \quad \text{and} \quad T: \begin{aligned} 0.9 \cdot x + y &= 1 \\ 1.1 \cdot x + y &= 1 \end{aligned}$$

Systems S and T look very close. Their right-hand sides are identical, and all the coefficients are close. Their solutions, however, are dramatically different: S has $\mathbf{x}^* = (1, 0)$ as its unique solution, while in T the roles are reversed: $\mathbf{x}^* = (0, 1)$. If two systems of equations are close, it doesn't mean that so must be their solutions. Yet Gradient Descent (unlike Gaussian Elimination again) is a “smooth” algorithm: If we slightly perturb the instance, it should take many steps for the particle to notice. When does this happen?

Suppose our initial guess for a solution in S is $(1/2, 1/2)$. The actual solution is $\mathbf{x}^* = (1, 0)$. We can write a matrix-vector recurrence like in the one-variable case. The gradient descent update rule is

$$\mathbf{x}' = \mathbf{x} - \rho \nabla f(\mathbf{x}) = \mathbf{x} - 2\rho A^T(A\mathbf{x} - b).$$

We now manipulate this equation differently: We substitute $b = A\mathbf{x}^*$ to get

$$\mathbf{x}' = \mathbf{x} - \rho \nabla f(\mathbf{x}) = \mathbf{x} - 2\rho A^T A(\mathbf{x} - \mathbf{x}^*).$$

After subtracting $\mathbf{x}^*$ from both sides, we get

$$\mathbf{x}' - \mathbf{x}^* = (\mathbf{x} - \mathbf{x}^*) - 2\rho A^T A(\mathbf{x} - \mathbf{x}^*)$$

We can factor the right-hand side to obtain vector recurrence

$$\mathbf{x}' - \mathbf{x}^* = (I - 2\rho A^T A)(\mathbf{x} - \mathbf{x}^*).$$

The relevant matrix $C = I - 2\rho A^T A$ in this recurrence is different from $B = I - 2\rho A A^T$. In general, $A A^T$ and $A^T A$ are not the same matrix. When A is not square, even their sizes do not match.

By the same logic, we can conclude that the rate of progress towards $\mathbf{x}^*$ is governed by the spectral norm of C . In this example,

$$A = \begin{bmatrix} 1 & 0.9 \\ 1 & 1.1 \end{bmatrix} \longrightarrow C = \begin{bmatrix} 1 - 4\rho & -4\rho \\ -4\rho & 1 - 4.04\rho \end{bmatrix}.$$

The spectral norm of C is about $1 - 0.02\rho$. We will learn how to calculate it next time. Initially, $\mathbf{x} - \mathbf{x}^* = (1/2, 1/2) - (1, 0) = (-1/2, 1/2)$ which has norm about 0.707. After t steps, $\mathbf{x}$ and $\mathbf{x}^*$ are $(1 - 0.02\rho)^t \cdot 0.707$ -close. If we set $\rho = 0.2$, an acceptable rate (less than $1/\|A\|^2 \approx 0.25$), $\mathbf{x}$ and $\mathbf{x}^*$ would be guaranteed to be within 0.1 after 500 steps. Indeed $\mathbf{x}_{500} = (0.932, 0.068)$, which is 0.096-close to $(1, 0)$. The spectral norm of C looks like an excellent predictor of convergence speed.

Where did the factor of 0.02 come from? Here is a “pattern-matching” explanation. The spectral norm of $C = I - 2\rho A^T A$ is about $1 - 0.02\rho = 1 - 2\rho \cdot 0.01$. There is a 2ρ in both expressions, the 1 comes from the identity matrix, so 0.01 must be the contribution of A .

The matrix A is close to the all-ones matrix. If A were the all-ones matrix, its rows would be linearly dependent, there would be many solutions, and it would be hopeless to expect convergence to any particular $\mathbf{x}^*$. The only possible explanation is that the contribution of the all-ones matrix is zero. Any matrix with linearly dependent rows contributes zero. The *condition number* of a matrix measures how close it is to having linearly dependent rows. You will investigate it in Homework 1. Here is a summary of all this.

Theorem 1. *If A is a square matrix with linearly independent rows, $\rho < 1/\|A\|^2$, after t steps with initial guess $\mathbf{x}$, the state $\mathbf{x}_t$ of gradient descent satisfies*

$$\|\mathbf{x}_t - \mathbf{x}^*\| \leq |1 - 2\rho \cdot \|A\|^2 \kappa^{-1}|^t \cdot \|\mathbf{x} - \mathbf{x}^*\|,$$

where $\mathbf{x}^*$ is the unique solution to $A\mathbf{x}^* = b$, and κ is the condition number of $A^T A$.

Theorem 1 and equation (4) look similar, but measure different notions of progress. Equation (4) bounds the rate of decrease of the sum of squares loss $f(\mathbf{x})$, but it tells us little about how close to target the solution $\mathbf{x}$ is. For example, after running gradient descent on S for 100 steps (with the above parameters), $\sqrt{f(\mathbf{x}_{100})}$ is 0.048. This looks great! But the corresponding solution is $\mathbf{x}_{100} = (0.664, 0.334)$, which is still a long way from $\mathbf{x}^* = (1, 0)$.

On the other hand, if the system happens to be underdetermined—say we have 10 unknowns but only 7 equations—then Theorem 1 does not apply because $\mathbf{x}^*$ cannot be uniquely defined. Equation (4), however, is relevant. It tells us that eventually the system will be almost solved by gradient descent. It doesn’t say which of the many candidate solutions the algorithm picks. This is not a bug but an important feature of gradient descent. By adding extra terms to the “objective” f we can guide gradient descent to favor some solutions over others, for instance ones that tend to be short. We’ll explain why this is useful in Section 7.

5 Linear regression

Gaussian Elimination struggled mightily with data analysis. This is where Gradient Descent shines. Let’s try it to infer grade weights from student marks on the same data:

	h	m	p	f
Alice	87.78	96.67	80.00	87.33
Bob	100.83	100.00	95.00	98.25
Dave	84.44	76.67	80.00	80.33
Charlie	48.89	86.67	72.50	69.67
Eve	84.72	96.67	85.00	88.42

The initial guess is that all three components weigh equally: $\mathbf{x} = (1/3, 1/3, 1/3)$. As for the rate, the data is at the scale of hundreds, so we might expect ρ to be some small multiple of $1/100^2$. With $\rho = 0.05 \cdot (1/100)^2$ the result is

step	x	root(f(x))
1000:	[0.31 0.32 0.368]	0.333
2000:	[0.305 0.31 0.385]	0.161
3000:	[0.302 0.305 0.393]	0.078
4000:	[0.301 0.302 0.397]	0.038
5000:	[0.3 0.301 0.398]	0.019

This is a very close approximation to the actual grade weights (0.3, 0.3, 0.4)! Let's do another experiment. When the data was rounded to the closest integer, Gaussian Elimination produced a very poor answer. Here is how Gradient Descent performs on the rounded data:

step	x	root(f(x))
1000 :	[0.287 0.32 0.387]	0.608
2000 :	[0.279 0.305 0.413]	0.415
3000 :	[0.275 0.297 0.425]	0.355
4000 :	[0.273 0.293 0.431]	0.339
5000 :	[0.272 0.291 0.434]	0.335

This doesn't look bad at all. There is a reason behind it. Among all algorithms you could try, Gradient Descent is guaranteed to give the best possible approximation, the one that minimizes the sum of squares error. Theorem 1 generalizes to Theorem 2:

Theorem 2. *For any matrix A , $\rho < 1/\|A\|^2$, after t steps with initial guess $\mathbf{x}$, the state $\mathbf{x}_t$ of gradient descent satisfies*

$$\|\mathbf{x}_t - \mathbf{x}^*\| \leq (1 - 2\rho \cdot \kappa^{-1})^t \|\mathbf{x} - \mathbf{x}^*\|,$$

assuming $\mathbf{x}^$ is the unique input that minimizes $f(\mathbf{x}^*) = \|A\mathbf{x}^* - b\|^2$.*

A linear system that has more linearly independent equations than unknowns is called *overconstrained*. Taken literally such systems have no solutions. Yet solving them is a pressing need in data analysis and machine learning. The price of a house is likely to depend on the square footage, the median income in the neighborhood, and years since the last renovation. We do not expect it to be an exact linear combination of those three numbers. Nevertheless, a retail agent might still want to model

$$\text{price} = x \cdot \text{footage} + y \cdot \text{income} + z \cdot \text{years} + \text{noise}.$$

This data-dependent noise could be very large. We should expect it to be large for such a simplistic model of the property market. Nevertheless, Gradient Descent produces the model (x, y, z) that minimizes the average square noise based on available data.

If this model makes lousy predictions it is not because Gradient Descent underperformed. It is the model's fault. To improve predictions we should consider adding variables (turnover rate, distance to subway) or look into more complex, nonlinear models. A surprising discovery of modern machine learning is that "fine-tuned" variants of gradient descent are just as effective on complex nonlinear models.

6 Variants and applications

Yet another way in which Gradient Descent differs from Gaussian Elimination is that its performance is highly dependent on the input. Input parameters like the spectral norm and the condition number greatly affect its rate of convergence. These parameters can sometimes be as hard to estimate as solving the system

itself. Testing the algorithm with different initializations and rates can be helpful in figuring out what is best for a given input. A natural extension of *GD* picks the rate adaptively depending on the state:

Algorithm *AGD* (Adaptive Gradient Descent)

Input: A function f with multiple inputs.

- 1 Choose a starting point $\mathbf{x} = (x, y, \dots, z)$.
- 3 Until $\mathbf{x}$ is sufficiently small,
- 2 Choose a rate ρ depending on $\mathbf{x}$, $\nabla f(\mathbf{x})$, $f(\mathbf{x})$, and the previous value of ρ .
- 4 Move $\mathbf{x}$ to $\mathbf{x} - \rho \cdot \nabla f(\mathbf{x})$.

There are several reasonable rules for picking the rate. In *steepest descent* ρ is chosen to minimize the target $\|\mathbf{x} - \rho \cdot \nabla f(\mathbf{x})\|$ among all possible choices of ρ . The advantage is that progress towards the minimum in any given step is as good as it can get. The disadvantage is that finding the best possible ρ requires time-consuming calculations. Another variant called *conjugate gradient descent* reduces the dependence on the condition number from κ^{-1} to $\kappa^{-1/2}$, and is guaranteed to terminate in n steps (assuming κ is finite).

To further improve convergence, it is sometimes possible to apply a transformation on the rows of the original system to make them less dependent and thereby decrease the condition number. The transformation has to be simple enough so that one can easily recover the solution to the original system from a solution to the transformed one. This is called *preconditioning*.

Gradient descent interacts with its input only by evaluating the gradient $2A^T(A\mathbf{x} - b) = 2A^T A\mathbf{x} - 2A^T b$. It is effectively solving the system $S\mathbf{x} = b'$, with $S = A^T A$ and $b' = A^T b$. Matrices of the form $A^T A$ are called *symmetric positive semidefinite* (PSD). There are many applications in which the input itself is already of this type. Gradient descent tends to be a lot more efficient on PSD inputs.

Unlike Gaussian Elimination, Gradient Descent is friendly to sparse matrices. Gradient calculation can be implemented in time proportional to the number of nonzero entries of the input matrix. This provides dramatic speedup in some applications.

Fairly recently, theorists developed ingenious algorithms for an important class of positive semidefinite systems called Laplacians. Laplacian systems come up in calculating the *hitting time* it takes a random walk to reach one graph vertex from another, for example how many steps it would take you to reach my web page if you surfed the internet at random. Hitting times measure how important a given connection is in a computer network. These fast Laplacian solvers tread close to the limit of algorithmic efficiency. Their execution time is only a few factors larger than the time it takes to read the matrix, at least in theory. Their implementation and analysis uses all of the above ideas (conjugate gradients, sparsity, preconditioning), and then some.

Finally, a simple twist on Gradient Descent is quite effective in machine learning applications. Owing to the humongous number of parameters in modern machine learning models, the matrices there can be very large. Evaluating the gradient can be a significant bottleneck. In this context rows of matrices represent data records like the grades of a student. Instead of evaluating the gradient exactly we can try to approximate it using a small subsample of the records. This variant is called *stochastic gradient descent*.

7 The unreasonable effectiveness of Gradient Descent

The spectacular success of modern machine learning probably comes from a confluence of several innovations: Data availability, extensive and cheap processing power, painstaking trial and error. Yet the one truly indispensable algorithm that powers it all is gradient descent. Object recognition, language translation, and chatbot interaction are all trained with gradient descent. How can an algorithm designed for solving $2x + 3 = 7$ be so powerful?

The reason is that a great many problems from all walks of engineering, science, and life can be turned into minimizing some objective function. It is not guaranteed that gradient descent will succeed on them but it can at least be attempted and played with. This paradigm has been especially effective in modern

machine learning.




FIGURE 2: Non-convex functions. (a) $f(x) = 1 - e^{-x^2}$: The average of $f(1)$ and $f(3)$ is smaller than the value $f(2)$ at the midpoint. Gradient descent can get stuck on the flat part. (b) $g(x) = (x^2 - 1)^2 - 0.1x$: Gradient descent could converge to the local minimum $x_- \approx -1$ or to the global one $x_+ \approx 1$ depending on the choice of initialization.

The gradient descent strategies we discussed generalize almost verbatim to any *convex* objective function. A function is convex if the averaging two evaluations is never smaller than evaluating the average. For example, x^2 and $|x|$ are convex, but $1 - e^{-x^2}$ and $(x^2 - 1)^2$ are not. A consequence of convexity is that not only is there a unique minimum but the gradient always points towards it.¹ Bound (??) and Theorem 2 apply to any convex f for a suitable generalization of “spectral norm” and “condition number”.

A fantastic feature of convex functions is that they can be added together. This is incredibly useful for model selection. Suppose you want to build a linear model but don’t have enough data. This is the problem of *overfitting*. There are many models consistent with the data so which one should you trust? A reasonable strategy is to select the one that least complex by some criterion. If this criterion is convex, gradient descent can handle it. A common complexity measure for solutions to linear systems is the sum of magnitudes, namely $g(\mathbf{x}) = |x| + |y| + \dots + |z|$. At the very least it penalizes solutions with unusually large numbers, like one that predicts the weight of the homeworks in the final grade is 157.

One surprising more recent realization is that gradient descent is often effective even for non-convex optimization. On non-convex instances gradient descent could converge to a local minimum that is not the true, global minimum, or it might not converge at all (see Figure 2). Not only is there no guarantee that it will succeed, but it is sometimes even hard to interpret if the outcome was a success or a failure. Yet highly non-convex complex machine learning models are successfully trained using gradient descent every day. It remains a great challenge for theorists to explain why gradient descent works so far beyond the humble problem that it was designed for—minimizing sums of squares.

¹For the math nerds: If the function is not differentiable a generalized *subgradient* can be defined.