


Section 04: Solutions

1. Gradient Descent

Like we've seen in lecture, gradient descent is an important algorithm commonly used to train machine learning models, particularly useful for when there is no closed form solution for the minimum of a loss function. Here, we'll go through short introduction to the algorithm.

Consider some function $f(w)$, which has some w_* for which $w_* = \arg \min_w f(w)$:




Let w_0 be some initial guess for the minimum of $f(w)$. Gradient descent will allow us to improve this solution.

- (a) For some w that is very close to w_0 , give the Taylor series approximation for $f(w)$ starting at $f(w_0)$.

Solution:

For w very close to w_0 , we see that $f(w) \approx f(w_0) + (w - w_0) \left(\frac{df(w)}{dw} \Big|_{w=w_0} \right)$.



- (b) Now, let us choose some $\eta > 0$ that is *very small*. With this very small η , let's assume that $w_1 = w_0 - \eta \left(\frac{df(w)}{dw} \Big|_{w=w_0} \right)$. Using your approximation from part (a), give an expression for $f(w_1)$.

Solution:

$$\begin{aligned}
 w_1 &= w_0 - \eta \left(\frac{df(w)}{dw} \Big|_{w=w_0} \right) \\
 f(w_1) &\approx f(w_0) + (w_1 - w_0) \left(\frac{df(w)}{dw} \Big|_{w=w_0} \right) \\
 &= f(w_0) + \left(w_0 - \eta \left(\frac{df(w)}{dw} \Big|_{w=w_0} \right) - w_0 \right) \left(\frac{df(w)}{dw} \Big|_{w=w_0} \right) \\
 &= f(w_0) - \eta \left(\frac{df(w)}{dw} \Big|_{w=w_0} \right)^2
 \end{aligned}$$

- (c) Given your expression for $f(w_1)$ from part (b), explain why, if η is small enough and if the function approximation is a good enough approximation, we are guaranteed to move in the “right” direction closer to the minimum w_* .

Solution:

Note that in part (b), the derivative is squared and will always be a nonnegative value. Therefore, $f(w_1) < f(w_0)$.

- (d) Building from your answer in part (c), write a general form for the gradient descent algorithm.


Solution:

Gradient descent is written as:

$$\text{For } k = 0, 1, 2, 3, \dots, w_{k+1} = w_k - \eta \left(\frac{df(w)}{dw} \Big|_{w=w_k} \right).$$

Note that as $k \rightarrow \infty$, $\left(\frac{df(w)}{dw} \Big|_{w=w_k} \right) \rightarrow 0$.

We visualize as:



2. Generalized Least Squares Regression

In class, we've seen linear regression and ridge regression. Here, we consider a problem that generalizes both of these. As a reminder, in linear regression, we seek a model that captures a linear relationship between input data and output data. The general case we consider imposes additional structure on the model.

Consider an experiment in which you have n data points $x_i \in \mathbb{R}^d$ and corresponding n observations y_i . We wish to come up with a model $\omega \in \mathbb{R}^d$ that satisfies the following properties: first, the error $\sum_{i=1}^n (x_i^\top \omega - y_i)^2$ should be small; second, we don't want small changes in training data resulting in large changes in solution; third, we want to put different weights in controlling the magnitude of different coordinates of ω . We therefore define

$$\hat{\omega}_{\text{general}} = \arg \min_{\omega} \sum_{i=1}^n (y_i - x_i^\top \omega)^2 + \lambda \sum_{i=1}^d D_{ii} \omega_i^2.$$

Here, D is a diagonal matrix, with positive entries on the diagonal. Observe that when D is the identity matrix, we recover ridge regression, and when $\lambda = 0$, we recover least squares regression. Different weights on D_{ii} cause the magnitudes of ω_i to be controlled differently.

2.1. Closed form in the general case

Deduce the closed form solution for $\hat{\omega}_{\text{general}}$. You should be comfortable with proofs in the "coordinate" form as well as the "matrix" form.

Solution:

We first give the proof using "matrix" notation. The objective function can be expressed as

$$\begin{aligned} f(\omega) &= \|X\omega - y\|_2^2 + \lambda \omega^\top D \omega \\ &= (X\omega - y)^\top (X\omega - y) + \lambda \omega^\top D \omega \\ &= (X\omega)^\top X\omega - (X\omega)^\top y - y^\top X\omega + y^\top y + \lambda \omega^\top D \omega \\ &= \omega^\top X^\top X \omega - 2\omega^\top X^\top y + y^\top y + \lambda \omega^\top D \omega \\ &= \omega^\top (X^\top X + \lambda D) \omega - 2\omega^\top X^\top y + y^\top y \end{aligned}$$

The gradient of f is

$$\begin{aligned} \nabla f(\omega) &= \nabla_{\omega} (\omega^\top (X^\top X + \lambda D) \omega - 2\omega^\top X^\top y + y^\top y) \\ &= \nabla_{\omega} (\omega^\top (X^\top X + \lambda D) \omega) - 2\nabla_{\omega} (\omega^\top X^\top y) + \nabla_{\omega} (y^\top y) \\ &= 2(X^\top X + \lambda D) \omega - 2X^\top y \end{aligned}$$

Here note that $X^\top X + \lambda D$ is a symmetric matrix, which explains the factor 2 in the gradient term. Setting the gradient $\nabla f(\omega)$ to zero, we can conclude that

$$(X^\top X + \lambda D) \hat{\omega}_{\text{general}} = X^\top y$$

If $X^\top X + \lambda D$ is full rank then we can get a unique solution:

$$\hat{\omega}_{\text{general}} = (X^\top X + \lambda D)^{-1} X^\top y$$

Since D is already given to be a diagonal matrix with strictly positive entries on the diagonal, any strictly positive λ will make the matrix $X^\top X + \lambda D$ invertible.

Solution:

We now give a solution in the "coordinate" form. The objective, when written in coordinate form, is $f(\omega) = \sum_{i=1}^n (y_i - x_i^\top \omega)^2 + \lambda \sum_{i=1}^d D_{ii} \omega_i^2$. As in the previous proof, we first simplify it as follows and then set it zero:

$$\begin{aligned}
 \nabla_{\omega} \left[\sum_{i=1}^n (y_i - x_i^\top \omega)^2 + \lambda \sum_{i=1}^d D_{ii} \omega_i^2 \right] &= \nabla_{\omega} \sum_{i=1}^n (y_i - x_i^\top \omega)^2 + \nabla_{\omega} \lambda \sum_{i=1}^d D_{ii} \omega_i^2 \\
 &= \sum_{i=1}^n \nabla_{\omega} (y_i - x_i^\top \omega)^2 + 2\lambda D \omega \\
 &= - \sum_{i=1}^n 2x_i (y_i - x_i^\top \omega) + 2\lambda D \omega \\
 &= - \sum_{i=1}^n 2x_i y_i + \sum_{i=1}^n 2x_i x_i^\top \omega + 2\lambda D \omega \\
 &= -2 \sum_{i=1}^n x_i y_i + 2 \left(\sum_{i=1}^n x_i x_i^\top + \lambda D \right) \omega \\
 &= 0 \quad (\text{set it to be } 0)
 \end{aligned}$$

$$\hat{\omega}_{\text{general}} = \left(\sum_{i=1}^n x_i x_i^\top + \lambda D \right)^{-1} \left(\sum_{i=1}^n x_i y_i \right)$$

Note that, as expected, this exactly matches the answer we got from the previous approach (because x_i 's are all the rows of X , and therefore $\sum_i x_i y_i = X^\top y$, and $\sum_i x_i x_i^\top = X^\top X$).

2.2. Special cases: linear regression and ridge regression

- (a) In the simple least squares case ($\lambda = 0$ above), what happens to the resulting $\hat{\omega}$ if we double all the values of y_i ?

Solution:

As can be seen from the formula $\hat{\omega} = (X^\top X)^{-1} X^\top y$, doubling y doubles ω as well. This makes sense intuitively as well because if the observations are scaled up, the model should also be.

- (b) In the simple least squares case ($\lambda = 0$ above), what happens to the resulting $\hat{\omega}$ if we double the data matrix $X \in \mathbb{R}^{n \times d}$?

Solution:

As can be seen from the formula $\hat{\omega} = (X^\top X)^{-1} X^\top y$, doubling X halves ω . This also makes sense intuitively because the error we are trying to minimize is $\|X\omega - y\|_2^2$, and if the X has doubled, while y has remained unchanged, then ω must compensate for it by reducing by a factor of 2.

- (c) Suppose $D = I$ (that is, it is the identity matrix). That is, this is the *ridge* regression setting. Explain why $\lambda > 0$ ensures that the solution exists and the matrix can be inverted.

Solution:

The solution is $\hat{\omega} = (X^\top X + \lambda I)^{-1} X^\top y$. We already saw in a previous part that $X^\top X$ is always positive semidefinite, that is, its eigenvalues are at least zero. Adding λI , where $\lambda > 0$, ensures that $X^\top X + \lambda I$ is in fact positive *definite*. This helps us have a good condition number.

3. Convexity

Convexity is defined for both sets and functions. For today we'll focus on discussing the convexity of functions.

Definition 1 (convex functions). A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is **convex** on a set A if for all $x, y \in A$ and $\lambda \in [0, 1]$:

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$$

When this definition holds with the inequality being reversed, then f is said to be concave. From the definition, it is clear that a function f is convex if and only if $-f$ is concave.

- (a) Why do we care whether a function is convex or not?

Solution:

Many numerical methods or algorithms were developed for finding local minima while in machine learning we are typically interested in finding the global minimum. Convex functions are useful because local minima are always the global minimum. Here is a short proof of this result: let x^* is a local minimizer for a convex function f and suppose $f(x_0) < f(x^*)$. Now note that there exists a $\lambda \in (0, 1)$ such that $y = \lambda x^* + (1 - \lambda)x_0$ and $f(y) \geq f(x^*)$ (i.e., y is close to x^* enough). We now have a contradiction:

$$f(y) \leq \lambda f(x^*) + (1 - \lambda)f(x_0) < f(x^*) \leq f(y).$$

In words, a line segment between any arbitrary point x_0 and a local minimizer x^* should be entirely above the function by definition of convexity, ensuring that $f(x_0) < f(x^*)$ cannot happen.

- (b) Which of the following functions are convex? (Hint: draw a picture!)

- (i) $|x|$
- (ii) $\cos(x)$
- (iii) $x^T x$

Solution:

$|x|$ and $x^T x$ are both convex. $\cos(x)$ is not convex since we can draw a line at two points (from say $\frac{\pi}{2}$ to $2\pi + \frac{\pi}{2}$) that is not entirely above the function.

Proof that $|x|$ is convex:

$$\begin{aligned} f(\lambda x + (1 - \lambda)y) &= |\lambda x + (1 - \lambda)y| \\ &\leq \lambda|x| + (1 - \lambda)|y| \end{aligned}$$

Proof that $x^T x$ is convex:

We begin by examining the definition: whenever $\lambda \in [0, 1]$, we have

$$\begin{aligned} (\lambda x + (1 - \lambda)y)^T (\lambda x + (1 - \lambda)y) &= \lambda^2 x^T x + (1 - \lambda)^2 y^T y + 2\lambda(1 - \lambda)x^T y \\ &= \lambda x^T x + (1 - \lambda)y^T y - \lambda(1 - \lambda)(x^T x - 2x^T y + y^T y) \\ &= \lambda x^T x + (1 - \lambda)y^T y - \lambda(1 - \lambda)(x - y)^T (x - y) \\ &\leq \lambda x^T x + (1 - \lambda)y^T y, \end{aligned}$$

where the inequality holds because $(x - y)^T (x - y) = \|x - y\|_2^2 \geq 0$. So our function is convex.

- (c) Can a function be both convex and concave on the same set? If so, give an example. If not, describe why not.

Solution:

Linear functions (i.e. functions such that $f(\lambda x + (1 - \lambda)y) = \lambda f(x) + (1 - \lambda)f(y)$) are both convex and concave.

4. Practical Methods for Checking Convexity

Using the definition to check whether a function is convex or not can be a tedious task in many situations. Some basic methods that can help us achieve the task in an efficient way are introduced below:

- for differentiable function, examine $f(y) \geq f(x) + \nabla f(x)^T(y - x)$ for any x, y in the domain of f .
- for twice differentiable functions, examine $\nabla^2 f(x) \succeq 0$ (i.e., the Hessian matrix is positive semidefinite).
- nonnegative weighted sum
- composition with affine function
- pointwise maximum and supremum

Note: there are even more such methods, which are covered in a convex optimization course or textbook.

- (a) If f is differentiable, then f is convex if and only if $f(y) \geq f(x) + \nabla f(x)^T(y - x)$ for any x, y in the domain of f . A geometric interpretation of this characterization is that any tangent plane of a convex function f must lie entirely below f . One interesting application of this characterization is one of the most important inequalities in probability and statistics: the Jensen's inequality, which states that $\mathbb{E}f(X) \geq f(\mathbb{E}(X))$ when f is convex. Prove Jensen's inequality using the other inequality mentioned here.

Solution:

Let $\mu = \mathbb{E}(X)$, then since f is convex, we have

$$f(X) \geq f(\mu) + \nabla f(\mu)^T(X - \mu)$$

with probability 1. This means that taking expectation on both sides preserves the inequality: $\mathbb{E}f(X) \geq f(\mu) = f(\mathbb{E}X)$.

- (b) If f is twice differentiable with convex domain, then f is convex if and only if

$$\nabla^2 f(x) \succeq 0,$$

for any x in the domain of f . Use this method to show that the objective function in linear regression is convex.

Solution:

Let $f(w) = (Y - Xw)^T(Y - Xw)$, then

$$\nabla^2 f(w) = 2(X^T X),$$

which is clearly a positive semidefinite matrix.

- (c) Let $\alpha \geq 0$ and $\beta \geq 0$, and if f and g are convex, then αf , $f + g$, $\alpha f + \beta g$ are all convex. One application: When a (possibly complicated) objective function can be expressed as a sum (e.g., the negative log-likelihood function), then showing the convexity of each individual term is typically easier.

Solution:

From definition, αf and $f + g$ are easily proved convex. To show that $\alpha f + \beta g$ is convex, first note that αf and βg are both convex, hence their sum is convex as well.

- (d) Suppose $f(\cdot)$ is convex, then $g(x) := f(Ax + b)$ is convex. Use this method to show that $\|Ax + b\|_1$ is convex (in x), where $\|z\|_1 = \sum_i |z_i|$.

Solution:

With this method, we only need to show the convexity of $\|x\|_1$. This is true from definition by observing that

$$\|\lambda x + (1 - \lambda)y\|_1 = \sum_i |\lambda x_i + (1 - \lambda)y_i| \leq \sum_i \lambda |x_i| + (1 - \lambda)|y_i| = \lambda \|x\|_1 + (1 - \lambda)\|y\|_1,$$

where the inequality holds because of triangular inequality.

- (e) Suppose you know that f_1 and f_2 are convex functions on a set A . The function $g(x) := \max\{f_1(x), f_2(x)\}$ is also convex on A . In general, if $f(x, y)$ is convex in x for each y , then $g(x) := \sup_y f(x, y)$ is convex. Use this method to show that the largest eigenvalue of a matrix X , $\lambda_{\text{Max}}(X)$, is convex in X (Using the definition of convexity would make this question quite difficult).

Solution:

Consider $f(v, X) := v^T X v$, then for each v , we have

$$f(v, \lambda X + (1 - \lambda)Y) = \lambda f(v, X) + (1 - \lambda)f(v, Y),$$

suggesting that $f(v, X)$ is convex in X for each v . Then $g(X) := \lambda_{\text{Max}}(X) = \sup_{\|v\|_2=1} f(v, X)$ is convex in X using this method.

- (f) Does the same result hold for $h(x) := \min\{f_1(x), f_2(x)\}$? If so, give a proof. If not, provide convex functions f_1, f_2 such that h is not convex.

Solution:

No, consider $f_1(x) = x^2$, $f_2(x) = (x - 1)^2$. Then $h(0) = h(1) = 0$, but $h(0.5) = 0.25$, so $h(0.5 * 0 + 0.5 * 1) = 0.25 > 0 = 0.5 * h(0) + 0.5 * h(1)$. So the minimum of two convex functions is not convex in general.

5. MAP as Regularization

Recall the regularization techniques that were presented in class this week and ponder their objectives:

(a) **Ridge-Regression:** $\hat{w}_{ridge} = \arg \min_w \sum_{i=1}^n (y_i - x_i^\top w)^2 + \lambda \|w\|_2^2$

(b) **LASSO:** $\hat{w}_{LASSO} = \arg \min_w \sum_{i=1}^n (y_i - x_i^\top w)^2 + \lambda \|w\|_1$

Reminder: don't ever regularize your bias term. This term doesn't add any complexity to the model (since it just shifts), so we'd like it to take on any value that best fits our training data.

The two types of regularization above can be derived from a statistical perspective in which we assume some prior belief about what the weights of our model should be and then observe data to further update the belief.

More specifically, let w denote our weights and Y, X our data (Y represents the labels and X the inputs). As before, $p(X, Y|w)$ represents the **likelihood function**. We specify our belief of what the weights should be through a **prior distribution** over $p(w)$. Using Bayes' Rule, we can write our updated belief of what the weights ought to be after observing the data as:

$$p(w|X, Y) = \frac{p(X, Y|w)p(w)}{p(X, Y)} = \frac{p(X, Y|w)p(w)}{\int_{w'} p(X, Y|w')p(w')dw'}$$

where we call $p(w|X, Y)$ the **posterior distribution** and $p(X, Y)$ the **evidence**.

What **Maximum A Posteriori Estimation (MAP)** does is compute the weights which maximize the posterior distribution, $p(w|X, Y)$. This type of estimation differs from MLE (which maximizes the likelihood function $p(X, Y|w)$) by taking into account our prior belief of what the weights are, namely $p(w)$. More specifically, the MAP estimate is:

$$\begin{aligned}\hat{w}_{MAP} &= \arg \max_w p(w|X, Y) \\ &= \arg \max_w \frac{p(X, Y|w)p(w)}{p(X, Y)} \\ &= \arg \max_w p(X, Y|w)p(w) \\ &= \arg \max_w \log p(X, Y|w) + \log p(w)\end{aligned}$$

where we dispose of the denominator because it doesn't depend on w . Contrast this with the MLE which is:

$$\hat{w}_{MLE} = \arg \max_w p(X, Y|w)$$

Let us now study how we can obtain the Ridge and LASSO regression objectives from this perspective:

(a) Suppose the elements of w are independently distributed according to a Laplacian distribution:

$$p(w_i) = \frac{\lambda}{4\sigma^2} \exp(-|w_i| \frac{\lambda}{2\sigma^2}).$$

Show that under this prior on w , MAP estimation of the linear measurement model recovers the LASSO objective.

Solution:

We work in the argmax space, which allows us to drop and add constants or monotonically increasing functions as necessary.

$$\begin{aligned}
{}_w p(w | Y, X) &= {}_w \log p(w | Y, X) \\
&= {}_w \log p(X, Y | w) + \log p(w) \\
&\stackrel{*}{=} {}_w \log \left\{ (2\pi)^{-\frac{n}{2}} (\sigma^2)^{-\frac{n}{2}} \exp\left(-\frac{1}{2}(Y - Xw)^\top (\sigma^2 \mathbb{I})^{-1} (Y - Xw)\right) \right\} \\
&\quad + \sum_{i=1}^n \log \frac{\lambda}{4\sigma^2} \exp\left(-|w_i| \frac{\lambda}{2\sigma^2}\right) \\
&= {}_w -\frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} (Y - Xw)^\top (Y - Xw) \\
&\quad + n \log \frac{\lambda}{4\sigma^2} - \frac{\lambda}{2\sigma^2} \sum_{i=1}^n |w_i| \\
&\stackrel{**}{=} {}_w -\frac{1}{2\sigma^2} \left[(Y - Xw)^\top (Y - Xw) + \lambda \sum_{i=1}^n |w_i| \right] \\
&\stackrel{***}{=} {}_w -\|Y - Xw\|_2^2 + \lambda \|w\|_1 \\
&= \arg \min_w \|Y - Xw\|_2^2 + \lambda \|w\|_1,
\end{aligned}$$

where the first starred equality follows from applying the PDFs, and the second and third follow from dropping constant terms or multipliers. In other words, solving MAP with a Laplacian prior also solves the LASSO regression problem.

- (b) Derive an expression for the prior on w that corresponds to the ridge regression objective. What is the significance of this result?

Solution:

Our high-level approach to this problem is to expand the terms of the objective to a form that resembles the core of a PDF, then attach the additive/multiplicative constants necessary to recover the full form of the PDF. We do this first with the likelihood term (since we know its form), and then with the prior.

$$\begin{aligned}
\arg \min_w \|Y - Xw\|_2^2 + \lambda \|w\|_2^2 &= {}_w -(Y - Xw)^\top (Y - Xw) + -\lambda w^\top w \\
&= {}_w -\frac{1}{2\sigma^2} (Y - Xw)^\top (Y - Xw) - \frac{\lambda}{2\sigma^2} w^\top w \\
&= {}_w -\frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} (Y - Xw)^\top (Y - Xw) - \frac{\lambda}{2\sigma^2} w^\top w \\
&= {}_w \mathcal{N}(Y; Xw, \sigma^2 \mathbb{I}) - \frac{\lambda}{2\sigma^2} w^\top w \\
&= {}_w \mathcal{N}(Y; Xw, \sigma^2 \mathbb{I}) - \frac{n}{2} \log \left(2\pi \frac{\sigma^2}{\lambda} \right) - \frac{1}{2} w^\top \left(\frac{\sigma^2}{\lambda} \mathbb{I} \right)^{-1} w \\
&= {}_w \mathcal{N}(Y; Xw, \sigma^2 \mathbb{I}) + \mathcal{N} \left(w; 0, \frac{\sigma^2}{\lambda} \mathbb{I} \right).
\end{aligned}$$

In other words, our prior is that $w \sim \mathcal{N} \left(0, \frac{\sigma^2}{\lambda} \mathbb{I} \right)$.

This means that when we apply ℓ_2 regularization to our linear regression problem (i.e., ridge regression), we make the implicit assumption that our weight vector is drawn from a Gaussian prior. More generally, we can see that applying various forms of regularization correspond to different prior assumptions on w .