

Topics in ML Lecture 2

Linear Least-Squares

Lénaïc Chizat*

September 15, 2025

In this class, we introduce and analyze Linear Least Squares Regression, a tool that can be traced back to Legendre (1805) and Gauss (1809) and which remains widely used in machine learning. Also, thanks to its simplicity and “theoretical tractability” (it can be analyzed with linear algebra), it is often used as a toy model to investigate phenomena that are observed in deep learning.

1 Introduction

- We recall the goal of supervised machine learning: given some observations $(x_i, y_i) \in \mathcal{X} \times \mathcal{Y}$, $i = 1, \dots, n$, of inputs/outputs, features/variables (training data), given a new $x \in \mathcal{X}$, predict $y \in \mathcal{Y}$ (testing data) with a *regression* function f such that $y \approx f(x)$.
- In today’s class, we consider empirical risk minimization with the square loss. This means that we choose a parameterized family of prediction functions $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ for $\theta \in \Theta$ and minimize the empirical risk

$$\frac{1}{n} \sum_{i=1}^n (y_i - f_\theta(x_i))^2.$$

- When $\mathcal{X} \subset \mathbb{R}^d$, the set of affine functions is a natural default choice. To simplify notations, we assume that the first component of the covariates x_i is 1 so that it is sufficient to consider linear functions. We thus consider minimizing

$$\hat{\mathcal{R}}(\theta) := \frac{1}{n} \sum_{i=1}^n (y_i - x_i^\top \theta)^2.$$

- This expression can be rewritten in matrix notations. Let $y = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$ be the vector of outputs and $X \in \mathbb{R}^{n \times d}$ the matrix of inputs, which rows are x_i^\top . It is called the *design* matrix. Its first column is a vector of 1 with our previous convention. In these notations, the empirical risk is

$$\hat{\mathcal{R}}(\theta) = \frac{1}{n} \|y - X\theta\|_2^2. \tag{1}$$

- In the expressions above, it is possible to replace the input x_i by any (potentially non-linear) function $\Phi(x_i)$. For instance, regression with degree 2 polynomials with $d = 2$ is obtained with $\Phi(x_i) = (1, x_{i,1}, x_{i,2}, x_{i,1}^2, x_{i,2}^2, x_{i,1}x_{i,2})$.

*EPFL lenaïc.chizat@epfl.ch

2 Ordinary least squares estimator

We make the assumption that X is injective (i.e. the rank of X is d). In particular, the problem is over-determined and $d \leq n$.

Definition 2.1. When X is injective, the minimizer of Eq. (1) is called the Ordinary least Squares (OLS) estimator.

2.1 Closed form solution

Proposition 2.2. If X is injective, then the OLS estimator exists and is unique. It is given by

$$\hat{\theta} = (X^\top X)^{-1} X^\top y.$$

With the covariance matrix $\Sigma := \frac{1}{n} X^\top X$, we have $\hat{\theta} = \frac{1}{n} \Sigma^{-1} X^\top y$.

Proof. Since $\hat{\mathcal{R}}$ is coercive and continuous, it admits at least a minimizer. Moreover, it is differentiable, so a minimizer $\hat{\theta}$ must satisfy $\nabla \hat{\mathcal{R}}(\hat{\theta}) = 0$. For all $\theta \in \mathbb{R}^d$, we have

$$\hat{\mathcal{R}}(\theta) = \frac{1}{n} \left(\|y\|_2^2 - 2\theta^\top X^\top y + \theta^\top X^\top X \theta \right) \quad \text{and} \quad \nabla \hat{\mathcal{R}}(\theta) = \frac{2}{n} \left(X^\top X \theta - X^\top y \right).$$

The condition $\nabla \hat{\mathcal{R}}(\hat{\theta}) = 0$ gives the so-called *normal equations*

$$X^\top X \hat{\theta} = X^\top y.$$

Under the assumption that X is injective, the matrix $X^\top X$ is invertible (Exercise: show that $X^\top X$ is positive definite). So the normal equations have a unique solution $\hat{\theta} = (X^\top X)^{-1} X^\top y$. This shows the uniqueness of the minimizer of $\hat{\mathcal{R}}$ as well as its closed form expression. \square

Another way to show uniqueness of the minimizer is by showing that $\hat{\mathcal{R}}$ is strongly convex since $\nabla^2 \hat{\mathcal{R}}(\theta) = 2\Sigma$ for all $\theta \in \mathbb{R}^d$ (convexity will be studied in the following lectures).

2.2 Numerical resolution

While the closed form $\hat{\theta} = (X^\top X)^{-1} X^\top y$ is convenient for analysis, inverting $X^\top X$ is sometimes unstable and has a large computational cost when d is large. Other methods are usually preferred, such as those QR factorization or (standard/stochastic/conjugate) gradient descent.

QR factorization. The QR decomposition factorizes the matrix X as $X = QR$ where $Q \in \mathbb{R}^{n \times d}$ has orthonormal columns and $R \in \mathbb{R}^{d \times d}$ is upper triangular and full rank. Computing a QR decomposition is faster (time complexity of $O(nd^2)$ flops) and more stable than inverting a matrix. One has

$$\begin{aligned} (X^\top X)\hat{\theta} = X^\top y &\Leftrightarrow R^\top Q^\top QR\hat{\theta} = R^\top Q^\top y \\ &\Leftrightarrow R^\top R\hat{\theta} = R^\top Q^\top y \\ &\Leftrightarrow R\hat{\theta} = Q^\top y. \end{aligned}$$

It only remains to solve a triangular linear system which is easy.

Gradient descent. We can completely bypass the need of matrix inversion or factorization using gradient descent. It consists in minimizing $\hat{\mathcal{R}}$ by taking an initial point $\theta_0 \in \mathbb{R}^d$ and iteratively going towards the minimizer by following the opposite of the gradient

$$\theta_{k+1} = \theta_k - \eta \nabla \hat{\mathcal{R}}(\theta_k) \quad \text{for } k \geq 0,$$

where $\eta > 0$ is the step-size. When these iterates converge, it is towards the OLS estimator since a fixed-point θ satisfies $\nabla \hat{\mathcal{R}}(\theta) = 0$. We'll discuss such methods at depth in the second half of the class.

3 Statistical analysis of OLS

We now prove guarantees on the performance of the OLS estimator. As before, we assume that X is injective.

3.1 Linear model and risk decomposition

Any kind of guarantee requires assumptions about how the data is generated. We assume that:

- there exists a vector $\theta^* \in \mathbb{R}^d$ such that the relationship between input and output is for $i \in [n]$

$$Y_i = x_i^\top \theta^* + Z_i.$$

- Z_i are independent of expectation $\mathbf{E}Z_i = 0$ and variance $\mathbf{E}Z_i^2 = \sigma^2$.

As before, we assume that $x_{i,1} = 1$ for all $i \in [n]$. We write Y and Z with capital letters to remind ourselves that (from now on) they are random variables. The vector Z accounts for variabilities in the output which are due to unobserved factors or to noise. From here, there are two settings of analysis for least squares:

- *Random design.* In this setting, both the input and the output are random. This is the classical setting of supervised machine learning, where the goal is *generalization*.
- *Fixed design.* In this setting, we assume that the input data (x_1, \dots, x_n) is *not* random and we are interested in obtaining a small prediction error on those input points. Our goal is to minimize

$$\mathcal{R}_X(\theta) = \mathbf{E}_Y \left[\frac{1}{n} \sum_{i=1}^n (Y_i - x_i^\top \theta)^2 \right].$$

This assumption allows a complete analysis with basic linear algebra. It is justified in some settings, e.g. when the input is a fixed grid, but is otherwise just a simplifying assumption. It can be understood as learning the vector $X\theta^* \in \mathbb{R}^n$ of best predictions, instead of a function.

In today's class, we consider the fixed design setting.

Proposition 3.1 (Risk decomposition). *Under the linear model and fixed design assumptions, for any $\theta \in \mathbb{R}^d$, it holds*

$$\mathcal{R}_X(\theta) - \mathcal{R}^* = \|\theta - \theta^*\|_\Sigma^2$$

where $\Sigma := \frac{1}{n} X^\top X$ is the input covariance and $\|\theta\|_\Sigma^2 := \theta^\top \Sigma \theta$. If now θ is a random variable (such as an estimator of θ^*), then

$$\mathbf{E}[\mathcal{R}_X(\theta)] - \mathcal{R}^* = \underbrace{\|\mathbf{E}[\theta] - \theta^*\|_\Sigma^2}_{\text{Bias}} + \underbrace{\mathbf{E}\left[\|\theta - \mathbf{E}[\theta]\|_\Sigma^2\right]}_{\text{Variance}}.$$

Proof. The Bayes predictor f^* for the square loss satisfies $f^*(x_i) = \mathbf{E}[Y_i] = x_i^\top \theta^*$ (see Exercices series 1) and the Bayes risk is $\mathcal{R}^* = \frac{1}{n} \sum_{i=1}^n \mathbf{E}[(Y_i - f^*(x_i))^2] = \frac{1}{n} \sum_{i=1}^n \mathbf{E}[Z_i^2] = \sigma^2$. Now for any fixed $\theta \in \mathbb{R}^d$, we have

$$\begin{aligned} \mathcal{R}_X(\theta) &= \frac{1}{n} \sum_{i=1}^n \mathbf{E}_{Y_i} [(Y_i - x_i^\top \theta)^2] = \frac{1}{n} \sum_{i=1}^n \mathbf{E}_{Y_i} [(Y_i - x_i^\top \theta^* + x_i^\top \theta^* - x_i^\top \theta)^2] \\ &= \frac{1}{n} \sum_{i=1}^n \left(\mathbf{E}_{Y_i} [(Y_i - x_i^\top \theta^*)^2] + \mathbf{E}_{Y_i} [(Y_i - x_i^\top \theta^*)(x_i^\top \theta^* - x_i^\top \theta)] + \mathbf{E}_{Y_i} [(x_i^\top \theta^* - x_i^\top \theta)^2] \right) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{E}_{Y_i} [Z_i^2] + 0 + \frac{1}{n} \sum_{i=1}^n (x_i^\top (\theta - \theta^*))^2 \\ &= \sigma^2 + \frac{1}{n} \|X(\theta - \theta^*)\|_2^2 = \mathcal{R}^* + \|\theta - \theta^*\|_\Sigma^2. \end{aligned}$$

This shows the first claim. Now if θ is random, we perform the usual bias/variance decomposition:

$$\begin{aligned} \mathbf{E}[\mathcal{R}_X(\theta)] - \mathcal{R}^* &= \mathbf{E} \left[\|\theta - \mathbf{E}[\theta] + \mathbf{E}[\theta] - \theta^*\|_\Sigma^2 \right] \\ &= \mathbf{E} \left[\|\theta - \mathbf{E}[\theta]\|_\Sigma^2 \right] + \mathbf{E} \left[(\theta - \mathbf{E}[\theta])^\top \Sigma (\mathbf{E}[\theta] - \theta^*) \right] + \mathbf{E} \left[\|\mathbf{E}[\theta] - \theta^*\|_\Sigma^2 \right] \\ &= \mathbf{E} \left[\|\theta - \mathbf{E}[\theta]\|_\Sigma^2 \right] + \|\mathbf{E}[\theta] - \theta^*\|_\Sigma^2. \end{aligned}$$

□

- The quantity $\|\theta\|_\Sigma$ is called the Mahalanobis norm (it is a “true” norm whenever Σ is positive definite). It is the norm on the parameter space induced by the input data.
- The same risk decomposition holds in the random design setting, with replacing Σ with $\mathbf{E}[xx^\top]$ the *population covariance*.

3.2 Statistical properties of the OLS estimator

We now analyze the properties of the OLS estimator.

Proposition 3.2 (Estimation properties of OLS). *The OLS estimator $\hat{\theta}$ has the following properties:*

1. it is unbiased $\mathbf{E}\hat{\theta} = \theta^*$,
2. its variance is $\text{var}(\hat{\theta}) = \mathbf{E}[(\hat{\theta} - \theta^*)(\hat{\theta} - \theta^*)^\top] = \frac{\sigma^2}{n} \Sigma^{-1}$ (Σ^{-1} is often called the precision matrix).

Proof. 1. Since $\mathbf{E}Y = X\theta^*$, we have directly $\mathbf{E}\hat{\theta} = (X^\top X)^{-1} X^\top X\theta^* = \theta^*$.

2. It follows that $\hat{\theta} - \theta^* = (X^\top X)^{-1} X^\top Z$. Thus, using that $\mathbf{E}ZZ^\top = \sigma^2 I_d$, it holds

$$\text{var} \hat{\theta} = \mathbf{E} \left[(X^\top X)^{-1} X^\top Z Z^\top X (X^\top X)^{-1} \right] = \sigma^2 (X^\top X)^{-1} (X^\top X) (X^\top X)^{-1} = \sigma^2 (X^\top X)^{-1} = \frac{\sigma^2}{n} \Sigma^{-1}.$$

□

Proposition 3.3 (Risk of OLS). *The excess risk of the OLS estimator is*

$$\mathbf{E} \left[\mathcal{R}_X(\hat{\theta}) \right] - \mathcal{R}^* = \frac{\sigma^2 d}{n}.$$

Proof. Using the risk decomposition of Proposition 3.1 and the fact that $\mathbf{E}[\hat{\theta}] = \theta^*$, we have

$$\mathbf{E} \left[\mathcal{R}_X(\hat{\theta}) \right] - \mathcal{R}^* = \mathbf{E} \|\hat{\theta} - \theta^*\|_{\Sigma}^2.$$

Using the identity $\hat{\theta} - \theta^* = (X^\top X)^{-1} X^\top Z$, we get

$$\begin{aligned} \mathbf{E}[\mathcal{R}_X(\hat{\theta})] - \mathcal{R}^* &= \mathbf{E} \|(X^\top X)^{-1} X^\top Z\|_{\Sigma}^2 \\ &= \frac{1}{n} \mathbf{E} \left[Z^\top X (X^\top X)^{-1} X^\top X (X^\top X)^{-1} X^\top Z \right] \\ &= \frac{1}{n} \mathbf{E} \left[\text{tr}(Z^\top X (X^\top X)^{-1} X^\top Z) \right] \\ &= \frac{1}{n} \mathbf{E} \left[\text{tr}(X^\top Z Z^\top X (X^\top X)^{-1}) \right] \\ &= \frac{\sigma^2}{n} \text{tr}(X^\top X (X^\top X)^{-1}) \\ &= \frac{\sigma^2 d}{n}. \end{aligned}$$

where we used that $\text{tr}(a) = a$ for $a \in \mathbb{R}$, the invariance of the trace under cyclic permutations, and $\mathbf{E}[ZZ^\top] = \sigma^2 I_d$. \square

4 Ridge least squares regression

- When d/n approaches 1, we are essentially memorizing the observations Y_i . Also when $d > n$, then $X^\top X$ is not invertible and the normal equations admit a linear subspace of solutions. These behaviors of OLS in high dimension (d large) are often undesirable (but there is nuance to this story, see the notion of “benign overfitting” in Bach’s book).
- Several solutions exist to fix these issues. The most common is to regularize the least squares objective, either by adding $\|\theta\|_1$ (*LASSO* regression, see later in the course) or $\|\theta\|_2^2$ (*ridge* regression, this lecture) to the empirical risk.

Definition 4.1 (Ridge least-squares regression). *For a regularization parameter $\lambda > 0$, we define the ridge least squares estimator $\hat{\theta}_\lambda$ as the minimizer of*

$$\min_{\theta \in \mathbb{R}^d} \frac{1}{n} \|Y - X\theta\|_2^2 + \lambda \|\theta\|_2^2.$$

Proposition 4.2. *Let us define $\Sigma_\lambda = \frac{1}{n} X^\top X + \lambda I_d$, and we recall that $\Sigma = \frac{1}{n} X^\top X$. It holds*

$$\hat{\theta}_\lambda = \frac{1}{n} \Sigma_\lambda^{-1} X^\top Y.$$

Proof. Left as an exercise (similar to the proof of Proposition 2.2). \square

As for the OLS, we can analyze the statistical properties of this estimator under the linear model and fixed design assumptions. To simplify the derivations and without loss of generality, we assume that the axes are aligned with the eigenvectors of Σ and we have $\Sigma = \text{diag}(\lambda_1, \dots, \lambda_d)$. In this section, we do not assume that X is injective, so $\lambda_j \geq 0$ is potentially 0 and d may be larger than n . Note that we have $\Sigma_\lambda = \text{diag}(\lambda_1 + \lambda, \dots, \lambda_d + \lambda)$

Proposition 4.3. *Under the linear model assumption, the ridge least squares estimator $\hat{\theta}_\lambda = \frac{1}{n} \Sigma_\lambda^{-1} X^\top Y$ has the following excess risk*

$$\mathbf{E} \left[\mathcal{R}_X(\hat{\theta}_\lambda) \right] - \mathcal{R}^* = \sum_{j \geq 1} (\theta_j^*)^2 \frac{\lambda^2 \lambda_j}{(\lambda_j + \lambda)^2} + \frac{\sigma^2}{n} \sum_{j \geq 1} \frac{\lambda_j^2}{(\lambda_j + \lambda)^2}.$$

Observe how this converges to the OLS estimator (when it is defined) as $\lambda \rightarrow 0$.

Proof. We use the risk decomposition of Proposition 3.1 into a bias B and variance V terms. Since it holds $\mathbf{E}[\hat{\theta}_\lambda] = \frac{1}{n}\Sigma_\lambda^{-1}X^\top X\theta^* = \Sigma_\lambda^{-1}\Sigma\theta^*$, and $\theta^* = \Sigma_\lambda^{-1}(\Sigma + \lambda I_d)\theta^*$ it follows $\mathbf{E}[\hat{\theta}_\lambda] - \theta^* = -\lambda\Sigma_\lambda^{-1}\theta^*$ (i.e. the ridge penalization induces a “shrinking” effect). Hence

$$\begin{aligned} B &= \|\mathbf{E}[\hat{\theta}_\lambda] - \theta^*\|_\Sigma^2 \\ &= \lambda^2(\theta^*)^\top \Sigma_\lambda^{-1}\Sigma\Sigma_\lambda^{-1}\theta^* \\ &= (\theta^*)^\top (\Sigma_\lambda^{-1}\Sigma - I_d)\Sigma(\Sigma_\lambda^{-1}\Sigma - I_d)\theta^* \\ &= \sum_{j \geq 1} (\theta_j^*)^2 \left(\frac{\lambda_j}{\lambda_j + \lambda} - 1 \right)^2 \lambda_j = \sum_{j \geq 1} (\theta_j^*)^2 \frac{\lambda^2 \lambda_j}{(\lambda + \lambda_j)^2}. \end{aligned}$$

For the variance term, using the fact that $ZZ^\top = \sigma^2 I_d$ we have

$$\begin{aligned} V &= \mathbf{E} \left[\|\hat{\theta}_\lambda - \mathbf{E}[\hat{\theta}_\lambda]\|_\Sigma^2 \right] \\ &= \mathbf{E} \left[\left\| \frac{1}{n}\Sigma_\lambda^{-1}X^\top Z \right\|_\Sigma^2 \right] \\ &= \mathbf{E} \left[\frac{1}{n^2} \text{tr} \left(Z^\top X \Sigma_\lambda^{-1} \Sigma \Sigma_\lambda^{-1} X^\top Z \right) \right] \\ &= \mathbf{E} \left[\frac{1}{n^2} \text{tr} \left(X^\top Z Z^\top X \Sigma_\lambda^{-1} \Sigma \Sigma_\lambda^{-1} \right) \right] \\ &= \frac{\sigma^2}{n} \text{tr} \left(\Sigma \Sigma_\lambda^{-1} \Sigma \Sigma_\lambda^{-1} \right) \\ &= \frac{\sigma^2}{n} \sum_{j \geq 1} \frac{\lambda_j^2}{(\lambda_j + \lambda)^2}. \end{aligned}$$

The proposition follows by summing the bias and variance terms. \square

Choosing λ in theory. Based on this expression for the risk, we can tune the parameter λ to obtain a potentially better bound than with the OLS (which corresponds to $\lambda = 0$).

Proposition 4.4. *With the choice $\lambda^* = \frac{\sigma\sqrt{\text{tr}(\Sigma)}}{\|\theta^*\|_2\sqrt{n}}$, we have*

$$\mathbf{E} \left[\mathcal{R}_X(\hat{\theta}_{\lambda^*}) \right] - \mathcal{R}^* \leq \frac{\sigma\sqrt{\text{tr}(\Sigma)}\|\theta^*\|_2}{\sqrt{n}}.$$

Proof. Let us upper bound the bias and variance terms using the inequality $(a + b)^2 \geq 2ab$ for $a, b \geq 0$. We have

$$\begin{aligned} B &\leq \sum_{j \geq 1} (\theta_j^*)^2 \frac{\lambda_j}{2\lambda_j/\lambda} = \frac{\lambda\|\theta^*\|_2^2}{2} \\ V &\leq \frac{\sigma^2}{n} \sum_{j \geq 1} \frac{\lambda_j^2}{2\lambda_j\lambda} = \frac{\sigma^2 \text{tr} \Sigma}{2\lambda n}. \end{aligned}$$

Plugging in λ^* (which was chosen to minimize the upper bound¹ on $B+V$) gives the result. \square

¹One can check that a bound of the form $a\lambda + b/\lambda$ for $a, b > 0$ is minimized for $\lambda = \sqrt{b/a}$ with optimal value $2\sqrt{ab}$.

- Observe that if we write $R = \max_i \|x_i\|_2$, then we have

$$\text{tr}(\Sigma) = \sum_{j \geq 1} \Sigma_{jj} = \frac{1}{n} \sum_{i=1}^n \sum_{j \geq 1} x_{i,j}^2 = \frac{1}{n} \sum_{i=1}^n \|x_i\|_2^2 \leq R^2.$$

Thus in the excess risk bound, the dimension d plays no role and it could even be infinite (given R and $\|\theta^*\|_2$ remain finite). This type of bounds are called *dimension free* bounds.

- Comparing this bound with that of the OLS estimator, we see that it converges slower to 0 as a function of n (from n^{-1} to $n^{-1/2}$) but it has a milder dependence on the noise (from σ^2 to σ).
- The value of λ^* involves quantities which we typically do not know in practice (σ and $\|\theta^*\|_2$).

Choosing λ in practice. The regularization λ is an example of a *hyper-parameter*. This term refers broadly to any quantity that influences the behavior of a machine learning algorithm and that is left to choose by the practitioner. While theory often offers guidelines and qualitative understanding on how to best choose the hyper-parameters, their precise numerical value depends on quantities which are often difficult to know or even guess. In practice, we typically resort to the two following approaches, which are attempts to estimate the test performance without using the test set:

- *Validation.* We divide the training data into two categories: a training set and a validation set. Given a choice of hyper-parameters, we run the algorithm on the training set and then measure the performance on the validation set. We repeat this process with different hyper-parameters and we select the hyper-parameters with the best validation performance.
- *Cross-validation.* A similar, but slightly more elaborate method is to partition the training data into k smaller data sets (typically $k = 5$). Each of these subsets plays successively the role of the validation set while the other form the training set. This gives, for each choice of hyper-parameters, k validation performances that can be averaged to choose the best hyper-parameter.

References

Alnur Ali, Edgar Dobriban, and Ryan Tibshirani. The implicit regularization of stochastic gradient flow for least squares. In *International conference on machine learning*, pages 233–244. PMLR, 2020.