

Chapter 7

Matrix Factorization and Analysis

Matrix factorizations are an important part of the practice and analysis of signal processing. They are at the heart of many signal-processing algorithms. Their applications include solving linear equations (LU), decorrelating random variables (LDLT, Cholesky), orthogonalizing sets of vectors (QR), and finding low-rank matrix approximations (SVD). Their usefulness is often two-fold: they allow efficient computation of important quantities and they are (often) designed to minimize round-off error due to finite-precision calculation. An algorithm is called *numerically stable*, for a particular set of inputs, if the error in the final solution is proportional to the round-off error in the elementary field operations.

7.1 Triangular Systems

A square matrix $L \in F^{n \times n}$ is called **lower triangular** (or **upper triangular**) if all elements above (or below) the main diagonal are zero. Likewise, a triangular matrix (lower or upper) is a **unit triangular** if it has all ones on the main diagonal. A system of linear equations is called *triangular* if it can be represented by the matrix equation $A\underline{x} = \underline{b}$ where A is either upper or lower triangular.

7.1.1 Solution by Substitution

Let $L \in F^{n \times n}$ be a lower triangular matrix with entries $l_{ij} = [L]_{ij}$. The matrix equation $L\underline{y} = \underline{b}$ can be solved efficiently using **forward substitution**, which is

defined by the recursion

$$y_j = \frac{1}{l_{jj}} \left(b_j - \sum_{i=1}^{j-1} l_{ji} y_i \right), \quad j = 1, 2, \dots, n.$$

Example 7.1.1. Consider the system

$$\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 2 & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 9 \end{bmatrix}.$$

Applying the above recursion gives

$$\begin{aligned} y_1 &= \frac{1}{1} = 1 \\ y_2 &= \frac{1}{1}(2 - 1 \cdot 1) = 1 \\ y_3 &= \frac{1}{1}(9 - 1 \cdot 1 - 2 \cdot 1) = 6. \end{aligned}$$

Let $U \in F^{n \times n}$ be an upper triangular matrix with entries $u_{ij} = [U]_{ij}$. The matrix equation $U\mathbf{x} = \mathbf{y}$ can be solved efficiently using **backward substitution**, which is defined by the recursion

$$x_j = \frac{1}{u_{jj}} \left(y_j - \sum_{i=j+1}^n u_{ji} x_i \right), \quad j = n, n-1, \dots, 1.$$

Example 7.1.2. Consider the system

$$\begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 3 \\ 0 & 0 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 6 \end{bmatrix}.$$

Applying the above recursion gives

$$\begin{aligned} x_3 &= \frac{6}{2} = 3 \\ x_2 &= \frac{1}{1}(1 - 3 \cdot 3) = -8 \\ x_1 &= \frac{1}{1}(1 - 1 \cdot 6 - 1 \cdot (-8)) = 3. \end{aligned}$$

The computational complexity of each substitution is roughly $\frac{1}{2}n^2$ operations.

Problem 7.1.3. Show that set of upper triangular matrices is a subalgebra of the set of all matrices. Since it is clearly a subspace, only two properties must be verified:

1. that the product of two upper triangular matrices is upper triangular
2. that the inverse of an upper triangular matrix is upper triangular

7.1.2 The Determinant

The determinant $\det(A)$ of a square matrix $A \in F^{n \times n}$ is a scalar which captures a number of important properties of that matrix. For example, A is invertible iff $\det(A) \neq 0$ and the determinant satisfies $\det(AB) = \det(A)\det(B)$ for square matrices A, B . Mathematically, it is the unique function mapping matrices to scalars that is (i) linear in each column, (ii) negated by column transposition, and (iii) satisfies $\det(I) = 1$.

The determinant of a square matrix can be defined recursively using the fact that $\det([a]) = a$. Let $A \in F^{n \times n}$ be an arbitrary square matrix with entries $a_{ij} = [A]_{ij}$. The (i, j) -minor of A is the determinant of the $(n - 1) \times (n - 1)$ matrix formed by deleting the i -th row and j -th column of A .

Fact 7.1.4 (Laplace's Formula). *The determinant of A is given by*

$$\det(A) = \sum_{j=1}^n a_{ij}(-1)^{i+j} M_{ij} = \sum_{i=1}^n a_{ij}(-1)^{i+j} M_{ij},$$

where M_{ij} is the (i, j) -minor of A .

Theorem 7.1.5. *The determinant of a triangular matrix is the product of its diagonal elements.*

Proof. For upper (lower) triangular matrices, this can be shown by expanding the determinant along the first column (row) to compute each minor. \square

7.2 LU Decomposition

7.2.1 Introduction

LU decomposition is a generalization of Gaussian elimination which allows one to efficiently solve a system of linear equations $A\underline{x} = \underline{b}$ multiple times with different

right-hand sides. In its basic form, it is numerically stable only if the matrix is positive definite or diagonally dominant. A slight modification, known as *partial pivoting*, makes it stable for a very large class of matrices.

Any square matrix $A \in F^{n \times n}$ can be factored as $A = LU$, where L is a unit lower-triangular matrix and U is an upper-triangular matrix. The following example uses elementary row operations to cancel, in each column, all elements below the main diagonal. These elementary row operations are represented using left multiplication by a unit lower-triangular matrix.

$$\begin{aligned} & \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix} \\ & \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 3 \\ 0 & 2 & 8 \end{bmatrix} \\ & \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -2 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 3 \\ 0 & 0 & 2 \end{bmatrix} \end{aligned}$$

This allows one to write

$$\begin{aligned} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix} &= \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix}^{-1} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -2 & 1 \end{bmatrix}^{-1} \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 3 \\ 0 & 0 & 2 \end{bmatrix} \\ \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix} &= \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 2 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 3 \\ 0 & 0 & 2 \end{bmatrix} \\ \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix} &= \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 2 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 3 \\ 0 & 0 & 2 \end{bmatrix} \end{aligned}$$

LU decomposition can also be used to efficiently compute the determinant of A . Since $\det(A) = \det(LU) = \det(L) \det(U)$, the problem is reduced to computing the determinant of triangular matrices. Using Theorem 7.1.5, it is easy to see that $\det(L) = 1$ and $\det(U) = \prod_{i=1}^n u_{ii}$.

7.2.2 Formal Approach

To describe LU decomposition formally, we first need to describe the individual operations that are used to zero out matrix elements.

Definition 7.2.1. Let $A \in F^{n \times n}$ be an arbitrary matrix, $\alpha \in F$ be a scalar, and $i, j \in \{1, 2, \dots, n\}$. Then, adding α times the j -th row to the i -th row is an **elementary row-addition operation**. Moreover, $I + \alpha E_{ij}$, where $E_{ij} \triangleq \underline{e}_i \underline{e}_j^T$ and \underline{e}_k is the k -th standard basis vector, is the **elementary row-addition matrix** which effects this operation via left multiplication.

Example 7.2.2. For example, elementary row operations are used to cancel the $(2, 1)$ matrix entry in

$$(I - E_{2,1})A = \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 3 \\ 1 & 3 & 9 \end{bmatrix}.$$

Lemma 7.2.3. The following identities capture the important properties of elementary row-operation matrices:

- (i) $E_{ij}E_{kl} = \delta_{j,k}E_{il}$
- (ii) $(I + \alpha E_{ij})(I + \beta E_{kl}) = I + \alpha E_{ij} + \beta E_{kl}$ if $j \neq k$
- (iii) $(I + \alpha E_{ij})^{-1} = (I - \alpha E_{ij})$ if $i \neq j$.

Proof. This proof is left as an exercise. □

Now, consider the process for computing the LU decomposition of A . To initialize the process, we let $A^{(1)} = A$. In each round, we let

$$L_j^{-1} = \prod_{i=j+1}^n \left(I - \frac{a_{i,j}^{(j)}}{a_{j,j}^{(j)}} E_{i,j} \right)$$

be the product of elementary row operation matrices which cancel the subdiagonal elements of the j -th column. The process proceeds by defining $A^{(j+1)} = L_j^{-1}A^{(j)}$ so that $A^{(j+1)}$ has all zeros below the diagonal in the first j columns. After $n - 1$ rounds, the process terminates with

$$U = A^{(n)} = L_{n-1}^{-1}L_{n-2}^{-1} \cdots L_1^{-1}A$$

where $L = L_1L_2 \cdots L_{n-1}$ is unit lower triangular.

Lemma 7.2.4. *From the structure of elementary row operation matrices, we see*

$$\prod_{j=1}^{n-1} \prod_{i=j+1}^n (I + \alpha_{ij} E_{i,j}) = I + \sum_{j=1}^{n-1} \sum_{i=j+1}^n \alpha_{ij} E_{i,j}.$$

Proof. First, we notice that

$$\prod_{i=j+1}^n (I + \alpha_{ij} E_{i,j}) = I + \sum_{i=j+1}^n \alpha_{ij} E_{i,j}$$

for $j = 1, 2, \dots, n-1$. Expanding the product shows that any term with two E matrices must contain a product $E_{i,j} E_{l,j}$ with $l > i > j$. By Lemma 7.2.3i, we see that this term must be zero because $j \neq l$.

Now, we can prove the main result via induction. First, we assume that

$$\prod_{j=1}^k \prod_{i=j+1}^n (I + \alpha_{ij} E_{i,j}) = I + \sum_{j=1}^k \sum_{i=j+1}^n \alpha_{ij} E_{i,j}.$$

Next, we find that if $k \leq n-2$, then

$$\begin{aligned} \prod_{j=1}^{k+1} \prod_{i=j+1}^n (I + \alpha_{ij} E_{i,j}) &= \left(\prod_{j=1}^k \prod_{i=j+1}^n (I + \alpha_{ij} E_{i,j}) \right) \left(\prod_{l=k+2}^n (I + \alpha_{l,k+1} E_{l,k+1}) \right) \\ &= \left(I + \sum_{j=1}^k \sum_{i=j+1}^n \alpha_{ij} E_{i,j} \right) \left(I + \sum_{l=k+2}^n \alpha_{l,k+1} E_{l,k+1} \right) \\ &= I + \sum_{j=1}^{k+1} \sum_{i=j+1}^n \alpha_{ij} E_{i,j} + \sum_{j=1}^k \sum_{i=j+1}^n \sum_{l=k+2}^n \alpha_{ij} \alpha_{l,k+1} E_{i,j} E_{l,k+1} \\ &= I + \sum_{j=1}^{k+1} \sum_{i=j+1}^n \alpha_{ij} E_{i,j} + \sum_{j=1}^k \sum_{i=j+1}^n \sum_{l=k+2}^n \alpha_{ij} \alpha_{l,k+1} E_{i,k+1} \delta_{j,l} \\ &= I + \sum_{j=1}^{k+1} \sum_{i=j+1}^n \alpha_{ij} E_{i,j}. \end{aligned}$$

Finally, we point out that the base case $k = 1$ is given by the initial observation. \square

Theorem 7.2.5. *This process generates one column of L per round because*

$$[L]_{ij} = \begin{cases} \frac{a_{i,j}^{(j)}}{a_{j,j}^{(j)}} & \text{if } 1 \leq i < j \\ 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Proof. First, we note that

$$\begin{aligned}
L &= L_1 L_2 \cdots L_{n-1} \\
&= \prod_{j=1}^{n-1} \left(\prod_{i=j+1}^n \left(I - \frac{a_{i,j}^{(j)}}{a_{j,j}^{(j)}} E_{i,j} \right) \right)^{-1} \\
&\stackrel{(a)}{=} \prod_{i=1}^{n-1} \prod_{i=j+1}^n \left(I - \frac{a_{i,j}^{(j)}}{a_{j,j}^{(j)}} E_{i,j} \right)^{-1} \\
&\stackrel{(b)}{=} \prod_{i=1}^{n-1} \prod_{i=j+1}^n \left(I + \frac{a_{i,j}^{(j)}}{a_{j,j}^{(j)}} E_{i,j} \right) \\
&= I + \sum_{i=1}^{n-1} \sum_{i=j+1}^n \frac{a_{i,j}^{(j)}}{a_{j,j}^{(j)}} E_{i,j},
\end{aligned}$$

where (a) follows from Lemma 7.2.3ii (i.e., all matrices in the inside product commute) and (b) follows from Lemma 7.2.3iii. Picking off the (i, j) entry of L (e.g., with $\underline{e}_i^T L \underline{e}_j$) gives the stated result. \square

Finally, we note that the LU decomposition can be computed in roughly $\frac{2}{3}n^3$ field operations.

7.2.3 Partial Pivoting

Sometimes the *pivot element* $a_{j,j}^{(j)}$ can be very small or zero. In this case, the algorithm will either fail (e.g., divide by zero) or return a very unreliable result. The algorithm can be easily modified to avoid this problem by swapping rows of $A^{(j)}$ to increase the magnitude of the pivot element before each cancellation phase. This results in a decomposition of the form $PA = LU$, where P is a permutation matrix.

In this section, we will describe LU decomposition with partial pivoting using the notation from the previous section. The main difference is that, in each round, we will define $A^{(j+1)} = M_j^{-1} P_j A^{(j)}$ where P_j is a permutation matrix. In particular, left multiplication by P_j swaps row j with row p_j , where

$$p_j = \arg \max_{i=j,j+1,\dots,n} |a_{i,j}^{(j)}|.$$

The matrix M_j^{-1} is now chosen to cancel the subdiagonal elements in j -th column of $P_j A^{(j)}$. After $n - 1$ rounds, the resulting decomposition has the form

$$A^{(n)} = M_{n-1}^{-1} P_{n-1} M_{n-2}^{-1} P_{n-2} \cdots M_1^{-1} P_1 A = U.$$

To show this can also be written in the desired form, we need to understand some properties of the permutations. First, we point that swapping two rows is a transposition and therefore $P_j^2 = I$. Next, we will show that the permutations can be moved to the right.

Lemma 7.2.6. *Let $M = I + \sum_{j=1}^k \sum_{i=j+1}^n \alpha_{ij} E_{ij}$ and Q be a permutation matrix which swaps row $l \geq k + 1$ and row $m > l$. Then, $QM = \widetilde{M}Q$ where*

$$\widetilde{M} = I + \sum_{j=1}^k \sum_{i=j+1}^n \alpha_{ij} Q E_{ij}.$$

Therefore, we can write

$$A^{(n)} = \underbrace{\widetilde{M}_{n-1}^{-1} \widetilde{M}_{n-2}^{-1} \cdots \widetilde{M}_1^{-1}}_{L^{-1}} \underbrace{P_{n-1} \cdots P_2 P_1}_P A = U$$

and $PA = LU$.

Proof. The proof is left as an exercise. □

7.3 LDLT and Cholesky Decomposition

If the matrix $A \in \mathbb{C}^{n \times n}$ is Hermitian, then the LU decomposition allows the factorization $A = LDL^H$, where L is unit lower triangular and D is diagonal. Since this factorization is typically applied to real matrices, it is referred to as **LDLT decomposition**. If A is also positive definite, then the diagonal elements of D are positive and we can write $A = (LD^{1/2})(LD^{1/2})^H$. The form $A = \widetilde{L}\widetilde{L}^H$, where \widetilde{L} is lower triangular, is known as **Cholesky factorization**.

To see this, we will describe the LDLT decomposition using the notation from LU decomposition starting from $A^{(1)} = A$. In the j -th round, define L_j^{-1} to be the product of elementary row-operation matrices which cancel the subdiagonal elements of the j -th column $A^{(j)}$. Then, define $A^{(j+1)} = L_j^{-1} A^{(j)} L_j^{-H}$ and notice that $A^{(j+1)}$ is Hermitian because $A^{(j)}$ is Hermitian. Next, notice that $A^{(j+1)}$ has zeros below the diagonal in the first j columns and zeros to the right of diagonal in the first j rows. This follows from the fact that the first j rows of $A^{(j)}$ are not affected by applying L_j^{-1} on left. Therefore, applying L_j^{-H} on the right also cancels

the elements to the right of the diagonal in the j -th row. After $n - 1$ rounds, we find that $D = A^{(n)}$ is a diagonal matrix.

There are a number of redundancies in the computation described above. First off, the L matrix computed by LU decomposition is identical to the L matrix computed by LDLT decomposition. Therefore, one can save operations by defining $A^{(j+1)} = L_j^{-1}A^{(j)}$. Moreover, the elements to the right of the diagonal in $A^{(j)}$ do not affect the computation at all. So, one can roughly half the number of additions and multiplies by only updating the lower triangular part of $A^{(j)}$. The resulting computational complexity is roughly $\frac{1}{3}n^3$ field operations.

7.3.1 Cholesky Decomposition

For a positive-definite matrix A , we can first apply the LDLT decomposition and then define $\tilde{L} = LD^{1/2}$. This gives the Cholesky decomposition $\tilde{L}\tilde{L}^H = LDL^H = A$.

The Cholesky decomposition is typically used to compute whitening filters for random variables. For example, one can apply it to the correlation matrix $R = E[\underline{X}\underline{X}^H]$ of a random vector \underline{X} . Then, one can define $\underline{Y} = \tilde{L}^{-1}\underline{X}$ and see that

$$E[\underline{Y}\underline{Y}^H] = E[\tilde{L}^{-1}\underline{X}\underline{X}^H\tilde{L}^{-H}] = \tilde{L}^{-1}R\tilde{L}^{-H} = I.$$

From this, one sees that \underline{Y} is a vector of uncorrelated (or white) random variables.

7.3.2 QR decomposition

A complex matrix $Q \in \mathbb{C}^{n \times n}$ is called **unitary** if $Q^H Q = Q Q^H = I$. If all elements of the matrix are real, then it is called **orthogonal** and $Q^T Q = Q Q^T = I$.

Theorem 7.3.1. Any matrix $A \in \mathbb{C}^{m \times n}$ can be factored as

$$A = QR,$$

where Q is an $m \times m$ unitary matrix, $Q Q^H = I$, and R is an $m \times n$ upper-triangular matrix.

Proof. To show this decomposition, we start by applying Gram-Schmidt Orthogonalization to the columns $\underline{a}_1, \dots, \underline{a}_n$ of A . This results in orthonormal vectors

$\{\underline{q}_1, \dots, \underline{q}_l\}$, where $l = \min(m, n)$, such that

$$\underline{a}_j = \sum_{i=1}^{\min(j,l)} r_{i,j} \underline{q}_i \quad \text{for } j = 1, 2, \dots, n.$$

This gives an $m \times l$ matrix $Q = [\underline{q}_1 \ \dots \ \underline{q}_l]$ and an $l \times n$ upper-triangular matrix R , with entries $[R]_{i,j} = r_{i,j}$, such that $A = QR$. If $m \leq n$, then $l = m$, Q is unitary, and the decomposition is complete. Otherwise, we must extend the orthonormal set $\{\underline{q}_1, \dots, \underline{q}_l\}$ to an orthonormal basis $\{\underline{q}_1, \dots, \underline{q}_m\}$ of \mathbb{C}^m . This gives an $m \times m$ unitary matrix $Q' = [\underline{q}_1 \ \dots \ \underline{q}_m]$. Adding $m - n$ rows of zeros to the previous R matrix gives an $m \times n$ matrix R' such that $A = Q'R'$. \square

7.4 Hermitian Matrices and Complex Numbers

Definition 7.4.1. A square matrix $Q \in \mathbb{R}^{n \times n}$ is **orthogonal** if $Q^T Q = Q Q^T = I$.

Definition 7.4.2. A square matrix $U \in \mathbb{C}^{n \times n}$ is **unitary** if $U^H U = U U^H = I$.

It is worth noting that, for unitary (resp. orthogonal) matrices, it suffices to check only that $U^H U = I$ (resp. $Q^T Q = I$) because U is invertible (e.g., it has linearly independent columns) and

$$U^H U = I \implies I = U U^{-1} = U (U^H U) U^{-1} = U U^H.$$

A useful analogy between matrices and complex numbers is as follows.

- *Hermitian matrices* satisfying $A^H = A$ are analogous to real numbers, whose complex conjugates are equal to themselves.
- *Unitary matrices* satisfying $U^H U = I$ are analogous to complex numbers on the unit circle, satisfying $\bar{z}z = 1$.
- *Orthogonal matrices* satisfying $Q^T Q = I$ are analogous to the real numbers $z = \pm 1$, such that $z^2 = 1$.

The transformation

$$z = \frac{1 + jr}{1 - jr}$$

maps real number r into the unit circle $|z| = 1$. Analogously, by *Cayley's formula*,

$$U = (I + jR)(I - jR)^{-1},$$

a Hermitian matrix R is mapped to a unitary matrix.