

THE LDL^T AND CHOLESKY DECOMPOSITIONS

The LDL^T decomposition is a variant of the LU decomposition that is valid for positive-definite symmetric matrices; the Cholesky decomposition is a tweak of the LDL^T decomposition.

Theorem. Let S be a positive-definite symmetric matrix. Then S has unique decompositions

$$S = LDL^T \quad \text{and} \quad S = L_1L_1^T$$

where:

- L is lower-unitriangular,
- D is diagonal with positive diagonal entries, and
- L_1 is lower-triangular with positive diagonal entries.

See [later in this note](#) for an efficient way to compute an LDL^T decomposition (by hand or by computer) and an example.

Remark. Any matrix admitting either decomposition is symmetric positive-definite by Problem 12(a) on Homework 12.

Remark. Since L_1^T has full column rank, taking $A = L_1^T$ shows that any positive-definite symmetric matrix S has the form $A^T A$.

Remark. Suppose that S has an LDL^T decomposition with

$$D = \begin{pmatrix} d_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & d_n \end{pmatrix}.$$

Then we define

$$\sqrt{D} = \begin{pmatrix} \sqrt{d_1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sqrt{d_n} \end{pmatrix},$$

so that $(\sqrt{D})^2 = D$, and we set $L_1 = L\sqrt{D}$. Then

$$L_1L_1^T = L(\sqrt{D})(\sqrt{D})^T L^T = LDL^T = S,$$

so $L_1L_1^T$ is the Cholesky decomposition of S .

Conversely, given a Cholesky decomposition $S = L_1L_1^T$, we can write $L_1 = LD'$, where D' is the diagonal matrix with the same diagonal entries as L_1 ; then $L = L_1D'^{-1}$ is the lower-unitriangular matrix obtained from L_1 by dividing each column by its diagonal entry. Setting $D = D'^2$, we have

$$S = (LD')(LD')^T = LD'^2L^T = LDL^T,$$

which is the LDL^T decomposition of S .

Since the LDL^T decomposition and the Cholesky decompositions are interchangeable, we will focus on the former.

Remark. The matrix $U = DL^T$ is upper-triangular with positive diagonal entries. In particular, it is in row echelon form, so $S = LU$ is the LU decomposition of S . This gives another way to interpret the Theorem: it says that every positive-definite symmetric matrix S has an LU decomposition (no row swaps are needed); moreover, U has positive diagonal entries, and if D is the diagonal matrix with the same diagonal entries as U , then $L^T = D^{-1}U$ (dividing each row of U by its pivot gives L^T).

This shows that one can easily compute an LDL^T decomposition from an LU decomposition: use the same L , and let D be the diagonal matrix with the same diagonal entries as U . However, we will see that one can compute LDL^T twice as fast as LU , by hand or by computer: see the end of this note.

Proof that the LDL^T decomposition exists and is unique. The idea is to do row and column operations on S to preserve symmetry in elimination. Suppose that E is an elementary matrix for a row replacement, say

$$R_2 \leftarrow 2R_1: \quad E = \begin{pmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Then EA performs the row operation $R_2 \leftarrow 2R_1$ on a 3×3 matrix A . On the other hand, AE^T performs the corresponding column operation $C_2 \leftarrow 2C_1$: indeed, taking transposes, we have $(AE^T)^T = EA^T$, which performs $R_2 \leftarrow 2R_1$ on A^T .

Starting with a positive-definite symmetric matrix S , first we note that the $(1, 1)$ -entry of S is positive: this is Problem 13(a) on Homework 12. Hence we can do row replacements (and no row swaps) to eliminate the last $n - 1$ entries in the first column. Multiplying the corresponding elementary matrices together gives us a lower-unitriangular matrix L_1 such that the last $n - 1$ entries of the first column of L_1S are zero. Multiplying L_1S by L_1^T on the right performs the same sequence of column operations; since S was symmetric, this has the effect of clearing the last $n - 1$ entries of the first row of S . In diagrams:

$$\begin{aligned} S &= \begin{pmatrix} a & b & c \\ b & * & * \\ c & * & * \end{pmatrix} & L_1 &= \begin{pmatrix} 1 & 0 & 0 \\ -b/a & 1 & 0 \\ -c/a & 0 & 1 \end{pmatrix} \\ L_1S &= \begin{pmatrix} a & b & c \\ 0 & * & * \\ 0 & * & * \end{pmatrix} & L_1SL_1^T &= \begin{pmatrix} a & 0 & 0 \\ 0 & * & * \\ 0 & * & * \end{pmatrix} \end{aligned}$$

Since S is symmetric, the matrix $S_1 = L_1SL_1^T$ is symmetric, and since S is positive-definite, the matrix S_1 is also positive-definite by Problem 12(a) on Homework 12. In particular, the $(2, 2)$ -entry of S_1 is nonzero, so we can eliminate the second column and the second row in the same way. We end up with another positive-definite symmetric matrix $S_2 = L_2S_1L_2^T = (L_2L_1)S(L_2L_1)^T$ where the only nonzero entries in the first two

rows/columns are the diagonal ones. Continuing in this way, we eventually get a diagonal matrix $D = S_{n-1} = (L_{n-1} \cdots L_1)S(L_{n-1} \cdots L_1)^T$ with positive diagonal entries. Setting $L = (L_{n-1} \cdots L_1)^{-1}$ gives $S = LDL^T$.

As for uniqueness,¹ suppose that $S = LDL^T = L'D'L'^T$. Multiplying on the left by L'^{-1} gives $L'^{-1}LDL^T = D'L'^T$, and multiplying on the right by $(DL^T)^{-1}$ gives

$$L'^{-1}L = (D'L'^T)(DL^T)^{-1}.$$

The left side is lower-unitriangular and the right side is upper-triangular. The only matrix that is both lower-unitriangular and upper-triangular is the identity matrix. It follows that $L'^{-1}L = I_n$, so $L' = L$. Then we have $(D'L'^T)(DL^T)^{-1} = I_n$, so $D'L^T = DL^T$ (using $L' = L$), and hence $D' = D$ (because L is invertible). \square

Remark (For experts). In abstract linear algebra, the expression $\langle x, y \rangle = x^T S y$ is called an *inner product*. This is a generalization of the usual dot product: $\langle x, y \rangle = x \cdot y$ when $S = I_n$. When S is positive-definite, one can run the Gram–Schmidt algorithm to turn the usual basis $\{e_1, \dots, e_n\}$ of \mathbf{R}^n into a basis $\{v_1, \dots, v_n\}$ which is orthogonal with respect to $\langle \cdot, \cdot \rangle$. The corresponding change-of-basis matrix is a lower-unitriangular matrix L' , and the matrix for $\langle \cdot, \cdot \rangle$ with respect to the orthogonal basis is a diagonal matrix D . This means $L'DL'^T = S$, so taking $L = L'^{-1}$, we have $S = LDL^T$.

Upshot: the LDL^T decomposition is exactly Gram–Schmidt as applied to the inner product $\langle x, y \rangle = x^T S y$.

Computational Complexity. The algorithm in the above proof appears to be the same as LU : the matrix $L = (L_{n-1} \cdots L_1)^{-1}$ is exactly what one would compute in an LU decomposition of an arbitrary matrix. However, one can save compute cycles by taking advantage of the symmetry of S .

In an ordinary LU decomposition, when clearing the first column, each row replacement involves $n - 1$ multiplications (scale the first row) and $n - 1$ additions (add to the i th row), for $2(n - 1)$ floating point operations (flops). Hence it takes $2(n - 1)^2$ flops to clear the first column. Clearing the second column requires $2(n - 2)^2$ flops, and so on, for a total of

$$2((n - 1)^2 + (n - 2)^2 + \cdots + 1) = 2 \frac{n(n - 1)(2n - 1)}{6} \approx \frac{2}{3}n^3$$

flops. However, when clearing the first row and column of a symmetric positive-definite matrix S , one only needs to compute the entries of $L_1 S L_1^T$ on or above the diagonal; the others are determined by symmetry. The first row replacement (the one that clears the $(2, 1)$ -entry) still needs $n - 1$ multiplications and $n - 1$ additions, but the second only needs $n - 2$ multiplications and $n - 2$ additions (because we don't need to compute the $(3, 2)$ -entry), and so on, for a total of

$$2((n - 1) + (n - 2) + \cdots + 1) = 2 \frac{n(n - 1)}{2} = n^2 - n$$

¹What follows is essentially the same proof that the LU decomposition is unique for an invertible matrix.

flops to clear the first column. Clearing the second column requires $(n-1)^2 - (n-1)$ flops, and so on, for a total of

$$\begin{aligned} & (n^2 - n) + ((n-1)^2 - (n-1)) + \cdots + (1^2 - 1) \\ &= (n^2 + (n-1)^2 + \cdots + 1) - (n + (n-1) + \cdots + 1) \\ &= \frac{n(n+1)(2n+1)}{6} - \frac{n(n+1)}{2} \approx \frac{1}{3}n^3 \end{aligned}$$

flops: half what was required for a full LU decomposition!

An Algorithm. The above discussion tells us how to modify the LU algorithm to compute the LDL^T decomposition. We use the two-column method as for an LU decomposition, but instead of keeping track of L_1S, L_2L_1S, \dots in the right column, we keep track of the symmetric matrices $S_1 = L_1SL_1^T, S_2 = L_2L_1S(L_2L_1)^T, \dots$, for which we only have to compute the entries on or above the diagonal. Instead of ending up with the matrix U in the right column, we end up with D .

Very explicitly: to compute $S_1 = L_1SL_1^T$ from S , first do row operations to eliminate the entries below the first pivot, then do column operations to eliminate the entries to the right of the first pivot; since the entries below the first pivot are zero after doing the row operations, this only changes entries in the first row. We end up with a symmetric matrix, so we only need to compute the entries on and above the diagonal. Now clear the second row/column, and continue recursively. Computing L is done the same way as in the LU decomposition, by recording the column divided by the pivot at each step.

Example. Let us compute the LDL^T decomposition of the positive-definite symmetric matrix

$$S = \begin{pmatrix} 2 & 4 & -2 & 2 \\ 4 & 9 & -1 & 6 \\ -2 & -1 & 14 & 13 \\ 2 & 6 & 13 & 35 \end{pmatrix}.$$

The entries in blue came for free by symmetry and didn't need to be calculated; the entries in green come from dividing the column by the pivot, as in the usual LU decomposition.

	L	S_i
start	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ ? & 1 & 0 & 0 \\ ? & ? & 1 & 0 \\ ? & ? & ? & 1 \end{pmatrix}$	$\begin{pmatrix} 2 & 4 & -2 & 2 \\ 4 & 9 & -1 & 6 \\ -2 & -1 & 14 & 13 \\ 2 & 6 & 13 & 35 \end{pmatrix}$
clear first column (then row)	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ -1 & ? & 1 & 0 \\ 1 & ? & ? & 1 \end{pmatrix}$	$\begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 1 & 3 & 2 \\ 0 & 3 & 12 & 15 \\ 0 & 2 & 15 & 33 \end{pmatrix}$
clear second column (then row)	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ -1 & 3 & 1 & 0 \\ 1 & 2 & ? & 1 \end{pmatrix}$	$\begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 3 & 9 \\ 0 & 0 & 9 & 29 \end{pmatrix}$
clear third column (then row)	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ -1 & 3 & 1 & 0 \\ 1 & 2 & 3 & 1 \end{pmatrix}$	$\begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}$

Hence $S = LDL^T$ for

$$L = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ -1 & 3 & 1 & 0 \\ 1 & 2 & 3 & 1 \end{pmatrix} \quad D = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}.$$

The Determinant Criterion We can also use the LDL^T decomposition to prove the determinant criterion that we discussed in class.

Theorem. A symmetric matrix S is positive-definite if and only if all upper-left determinants are positive.

Proof. First we show that if S is positive-definite then all upper-left determinants are positive. Let S_1 be an $r \times r$ upper-left submatrix:

$$S = \begin{pmatrix} & & & * \\ & S_1 & & * \\ & & & * \\ * & * & * & * \end{pmatrix}$$

Let $x = (x_1, \dots, x_r, 0, \dots, 0)$ be a nonzero vector in $\text{Span}\{e_1, \dots, e_r\}$. Then $x^T S x$ only depends on S_1 :

$$(x_1 \ x_2 \ x_3 \ 0) \begin{pmatrix} & * \\ S_1 & * \\ & * \\ * & * & * \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ 0 \end{pmatrix} = (x_1 \ x_2 \ x_3 \ 0) \begin{pmatrix} S_1 \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \\ * \end{pmatrix} = (x_1 \ x_2 \ x_3) S_1 \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}.$$

Since $x^T S x > 0$, this shows that S_1 is also positive-definite. Hence the eigenvalues of S_1 are positive, so $\det(S_1) > 0$.

We will prove the converse by *induction*: that is, we'll prove it for 1×1 matrices, then for 2×2 matrices using that the 1×1 case is true, then for 3×3 matrices using that the 2×2 case is true, etc. The 1×1 case is easy: it says that the matrix (a) is positive-definite if and only if $\det(a) = a > 0$. Suppose then that we know that an $(n-1) \times (n-1)$ matrix with positive upper-left determinants is positive-definite. Let S be an $n \times n$ matrix with positive upper-left determinants, and let S_1 be the upper-left $(n-1) \times (n-1)$ submatrix of S :

$$S = \begin{pmatrix} & & & a_1 \\ & S_1 & & \vdots \\ & & & a_{n-1} \\ a_1 \ \cdots \ a_{n-1} & & & a_n \end{pmatrix}$$

Then we already know that S_1 is positive-definite, so it has an LDL decomposition: say $S_1 = L_1' D_1 L_1'^T$. Taking $L_1 = L_1'^{-1}$, we have $L_1 S_1 L_1^T = D_1$. Let L be the matrix obtained from L_1 by adding the vector e_n to the right and e_n^T to the bottom, and likewise for D and D_1 :

$$L = \begin{pmatrix} & & 0 \\ & L_1 & \vdots \\ & & 0 \\ 0 \ \cdots \ 0 & & 1 \end{pmatrix} \quad D = \begin{pmatrix} & & 0 \\ & D_1 & \vdots \\ & & 0 \\ 0 \ \cdots \ 0 & & 1 \end{pmatrix}.$$

Multiplying out LSL^T and keeping track of which entries are multiplied by which gives

$$LSL^T = \begin{pmatrix} & & & b_1 \\ & L_1 S_1 L_1^T & & \vdots \\ & & & b_{n-1} \\ b_1 & \cdots & b_{n-1} & b_n \end{pmatrix} = \begin{pmatrix} d_1 & & & b_1 \\ & \ddots & & \vdots \\ & & d_{n-1} & b_{n-1} \\ b_1 & \cdots & b_{n-1} & b_n \end{pmatrix},$$

where $d_1, \dots, d_{n-1} > 0$ are the diagonal entries of D_1 , $(b_1, \dots, b_{n-1}) = L_1(a_1, \dots, a_{n-1})$, and $b_n = a_n$. Since the d_i are nonzero, we can do $n - 1$ row operations $R_n \leftarrow \frac{b_i}{d_i} R_i$ to clear the entries in the last row. Doing the same column operations $C_n \leftarrow \frac{b_i}{d_i} C_i$ clears the entries in the last column as well. Letting E be the product of the (lower-unitriangular) elementary matrices for these row operations, we get

$$E(LSL^T)E^T = \begin{pmatrix} d_1 & & & 0 \\ & \ddots & & 0 \\ & & d_{n-1} & 0 \\ 0 & 0 & 0 & d_n \end{pmatrix}.$$

Note that $\det(E) = \det(L) = \det(L^T) = \det(E^T) = 1$, so $\det(S) = d_1 \cdots d_{n-1} d_n$. Since $\det(S) > 0$ and $d_1, \dots, d_{n-1} > 0$, we have $d_n > 0$ as well. Setting $L_2 = (EL)^{-1}$ and $D_2 = ELSL^T E^T$, this gives $S = L_2 D_2 L_2^T$. Since S has an LDL^T decomposition, it is positive-definite, as desired. \square