## THE $L D L^{T}$ AND CHOLESKY DECOMPOSITIONS

The $L D L^{T}$ decomposition is a variant of the $L U$ decomposition that is valid for positivedefinite symmetric matrices; the Cholesky decomposition is a tweak of the $L D L^{T}$ decomposition.

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

$$
S=L D L^{T} \quad \text { and } \quad S=L_{1} L_{1}^{T}
$$

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 $L D L^{T}$ decomposition (by hand or by computer) and an example.

Remark. Any matrix admitting either decomposition is symmetric positive-definite by 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 $L D L^{T}$ decomposition with

$$
D=\left(\begin{array}{rrr}
d_{1} & \cdots & 0 \\
\vdots & \ddots & \vdots \\
0 & \cdots & d_{n}
\end{array}\right)
$$

Then we define

$$
\sqrt{D}=\left(\begin{array}{rrr}
\sqrt{d_{1}} & \cdots & 0 \\
\vdots & \ddots & \vdots \\
0 & \cdots & \sqrt{d_{n}}
\end{array}\right)
$$

so that $(\sqrt{D})^{2}=D$, and we set $L_{1}=L \sqrt{D}$. Then

$$
L_{1} L_{1}^{T}=L(\sqrt{D})(\sqrt{D})^{T} L^{T}=L D L^{T}=S,
$$

so $L_{1} L_{1}^{T}$ is the Cholesky decomposition of $S$.
Conversely, given a Cholesky decomposition $S=L_{1} L_{1}^{T}$, we can write $L_{1}=L D^{\prime}$, where $D^{\prime}$ is the diagonal matrix with the same diagonal entries as $L_{1}$; then $L=L_{1} D^{\prime-1}$ is the lower-unitriangular matrix obtained from $L_{1}$ by dividing each column by its diagonal entry. Setting $D=D^{\prime 2}$, we have

$$
S=\left(L D^{\prime}\right)\left(L D^{\prime}\right)^{T}=L D^{\prime 2} L^{T}=L D L^{T}
$$

which is the $L D L^{T}$ decomposition of $S$.

Since the $L D L^{T}$ decomposition and the Cholesky decompositions are interchangeable, we will focus on the former.

Remark. The matrix $U=D L^{T}$ is upper-triangular with positive diagonal entries. In particular, it is in row echelon form, so $S=L U$ is the $L U$ decomposition of $S$. This gives another way to interpret the Theorem: it says that every positive-definite symmetric matrix $S$ has an $L U$ 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 $L D L^{T}$ decomposition from an $L U$ 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 $L D L^{T}$ twice as fast as $L U$, by hand or by computer: see the end of this note.
Proof that the $L D L^{T}$ 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

$$
R_{2}=2 R_{1}: \quad E=\left(\begin{array}{rrr}
1 & 0 & 0 \\
-2 & 1 & 0 \\
0 & 0 & 1
\end{array}\right)
$$

Then $E A$ performs the row operation $R_{2}-=2 R_{1}$ on a $3 \times 3$ matrix $A$. On the other hand, $A E^{T}$ performs the corresponding column operation $C_{2}-=2 C_{1}$ : indeed, taking transposes, we have $\left(A E^{T}\right)^{T}=E A^{T}$, which performs $R_{2}-=2 R_{1}$ on $A^{T}$.

Starting with a positive-definite symmetric matrix $S$, first we note that the ( 1,1 )-entry of $S$ is positive: this was a problem 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 $L_{1} S$ are zero. Multiplying $L_{1} S$ 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:

$$
\begin{array}{cc}
S=\left(\begin{array}{lll}
a & b & c \\
b & * & * \\
c & * & *
\end{array}\right) & L_{1}=\left(\begin{array}{rrr}
1 & 0 & 0 \\
-b / a & 1 & 0 \\
-c / a & 0 & 1
\end{array}\right) \\
L_{1} S=\left(\begin{array}{lll}
a & b & c \\
0 & * & * \\
0 & * & *
\end{array}\right) & L_{1} S L_{1}^{T}=\left(\begin{array}{lll}
a & 0 & 0 \\
0 & * & * \\
0 & * & *
\end{array}\right)
\end{array}
$$

Since $S$ is symmetric, the matrix $S_{1}=L_{1} S L_{1}^{T}$ is symmetric, and since $S$ is positive-definite, the matrix $S_{1}$ is also positive-definite by a problem 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 $S_{2}=L_{2} S_{1} L_{2}^{T}=$ $\left(L_{2} L_{1}\right) S\left(L_{2} L_{1}\right)^{T}$ 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 $D=S_{n-1}=$ $\left(L_{n-1} \cdots L_{1}\right) S\left(L_{n-1} \cdots L_{1}\right)^{T}$ with positive diagonal entries. Setting $L=\left(L_{n-1} \cdots L_{1}\right)^{-1}$ gives $S=L D L^{T}$ 。

As for uniqueness, ${ }^{1}$ suppose that $S=L D L^{T}=L^{\prime} D^{\prime} L^{\prime T}$. Multiplying on the left by $L^{\prime-1}$ gives $L^{\prime-1} L D L^{T}=D^{\prime} L^{\prime T}$, and multiplying on the right by $\left(D L^{T}\right)^{-1}$ gives

$$
L^{\prime-1} L=\left(D^{\prime} L^{\prime T}\right)\left(D L^{T}\right)^{-1} .
$$

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 $L^{\prime-1} L=I_{n}$, so $L^{\prime}=L$. Then we have $\left(D^{\prime} L^{\prime T}\right)\left(D L^{T}\right)^{-1}=I_{n}$, so $D^{\prime} L^{T}=D L^{T}$ (using $L^{\prime}=L$ ), and hence $D^{\prime}=D$ (because $L$ is invertible).

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 $\left\{e_{1}, \ldots, e_{n}\right\}$ of $\mathbf{R}^{n}$ into a basis $\left\{v_{1}, \ldots, v_{n}\right\}$ which is orthogonal with respect to $\langle\cdot, \cdot\rangle$. The corresponding change-of-basis matrix is a lower-unitriangular matrix $L^{\prime}$, and the matrix for $\langle\cdot, \cdot\rangle$ with respect to the orthogonal basis is a diagonal matrix $D$. This means $L^{\prime} D L^{\prime T}=S$, so taking $L=L^{\prime-1}$, we have $S=L D L^{T}$.

Upshot: the $L D L^{T}$ 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 $L U$ : the matrix $L=\left(L_{n-1} \cdots L_{1}\right)^{-1}$ is exactly what one would compute in an $L U$ decomposition of an arbitrary matrix. However, one can save compute cycles by taking advantage of the symmetry of $S$.

In an ordinary $L U$ 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

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

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

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

[^0]flops to clear the first column. Clearing the second column requires $(n-1)^{2}-(n-1)$ flops, and so on, for a total of
\[

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

flops: half what was required for a full $L U$ decomposition!

An Algorithm. The above discussion tells us how to modify the $L U$ algorithm to compute the $L D L^{T}$ decomposition. We use the two-column method as for an $L U$ decomposition, but instead of keeping track of $L_{1} S, L_{2} L_{1} S, \ldots$ in the right column, we keep track of the symmetric matrices $S_{1}=L_{1} S L_{1}^{T}, S_{2}=L_{2} L_{1} S\left(L_{2} L_{1}\right)^{T}, \ldots$, 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 $S_{1}=L_{1} S L_{1}^{T}$ 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 $L U$ decomposition, by recording the column divided by the pivot at each step.

Example. Let us compute the $L D L^{T}$ decomposition of the positive-definite symmetric matrix

$$
S=\left(\begin{array}{rrrr}
2 & 4 & -2 & 2 \\
4 & 9 & -1 & 6 \\
-2 & -1 & 14 & 13 \\
2 & 6 & 13 & 35
\end{array}\right)
$$

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 $L U$ decomposition.

|  | $L$ | $S_{i}$ |
| :---: | :---: | :---: |
| start | $\left(\begin{array}{llll}1 & 0 & 0 & 0 \\ ? & 1 & 0 & 0 \\ ? & ? & 1 & 0 \\ ? & ? & ? & 1\end{array}\right)$ | $\left(\begin{array}{rrrr}2 & 4 & -2 & 2 \\ 4 & 9 & -1 & 6 \\ -2 & -1 & 14 & 13 \\ 2 & 6 & 13 & 35\end{array}\right)$ |
| clear first column (then row) | $\left(\begin{array}{rrrr}1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ -1 & ? & 1 & 0 \\ 1 & ? & ? & 1\end{array}\right)$ | $\left(\begin{array}{rrrr}2 & 0 & 0 & 0 \\ 0 & 1 & 3 & 2 \\ 0 & 3 & 12 & 15 \\ 0 & 2 & 15 & 33\end{array}\right)$ |
| clear second column (then row) | $\left(\begin{array}{rrrr}1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ -1 & 3 & 1 & 0 \\ 1 & 2 & ? & 1\end{array}\right)$ | $\left(\begin{array}{rrrr}2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 3 & 9 \\ 0 & 0 & 9 & 29\end{array}\right)$ |
| clear third column (then row) | $\left(\begin{array}{rrrr}1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ -1 & 3 & 1 & 0 \\ 1 & 2 & 3 & 1\end{array}\right)$ | $\left(\begin{array}{llll}2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 2\end{array}\right)$ |

Hence $S=L D L^{T}$ for

$$
L=\left(\begin{array}{rrrr}
1 & 0 & 0 & 0 \\
2 & 1 & 0 & 0 \\
-1 & 3 & 1 & 0 \\
1 & 2 & 3 & 1
\end{array}\right) \quad D=\left(\begin{array}{llll}
2 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 \\
0 & 0 & 3 & 0 \\
0 & 0 & 0 & 2
\end{array}\right)
$$

The Determinant Criterion We can also use the $L D L^{T}$ 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:

$$
S=\left(\begin{array}{llll} 
& & & * \\
& S_{1} & & * \\
& & & * \\
* & * & * & *
\end{array}\right)
$$

Let $x=\left(x_{1}, \ldots, x_{r}, 0, \ldots, 0\right)$ be a nonzero vector in $\operatorname{Span}\left\{e_{1}, \ldots, e_{r}\right\}$. Then $x^{T} S x$ only depends on $S_{1}$ :

$$
\left(\begin{array}{llll}
x_{1} & x_{2} & x_{3} & 0
\end{array}\right)\left(\begin{array}{rrr}
S_{1} & & * \\
* & * & *
\end{array}\right)\left(\begin{array}{r}
x_{1} \\
x_{2} \\
x_{3} \\
0
\end{array}\right)=\left(\begin{array}{llll}
x_{1} & x_{2} & x_{3} & 0
\end{array}\right)\binom{x_{1}\left(\begin{array}{r}
x_{1} \\
x_{2} \\
x_{3}
\end{array}\right)}{*}=\left(\begin{array}{lll}
x_{1} & x_{2} & x_{3}
\end{array}\right) S_{1}\left(\begin{array}{l}
x_{1} \\
x_{2} \\
x_{3}
\end{array}\right) .
$$

Since $x^{T} S x>0$, this shows that $S_{1}$ is also positive-definite. Hence the eigenvalues of $S_{1}$ are positive, so $\operatorname{det}\left(S_{1}\right)>0$.

We will prove the converse by induction: that is, we'll prove it for $1 \times 1$ matrices, then for $2 \times 2$ matrices using that the $1 \times 1$ case is true, then for $3 \times 3$ matrices using that the $2 \times 2$ case is true, etc. The $1 \times 1$ case is easy: it says that the matrix (a) is positive-definite if and only if $\operatorname{det}(a)=a>0$. 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 $S$ :

$$
S=\left(\begin{array}{rrrr} 
& & & a_{1} \\
& S_{1} & & \vdots \\
& & & a_{n-1} \\
a_{1} & \cdots & a_{n-1} & a_{n}
\end{array}\right)
$$

Then we already know that $S_{1}$ is positive-definite, so it has an $L D L$ decomposition: say $S_{1}=L_{1}^{\prime} D_{1} L_{1}^{\prime T}$. Taking $L_{1}=L_{1}^{\prime-1}$, 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 $D_{1}$ :

$$
L=\left(\begin{array}{cccc} 
& & & 0 \\
& L_{1} & & \vdots \\
& & & 0 \\
0 & \cdots & 0 & 1
\end{array}\right) \quad D=\left(\begin{array}{cccc} 
& & & 0 \\
& D_{1} & & \vdots \\
& & & 0 \\
0 & \cdots & 0 & 1
\end{array}\right)
$$

Multiplying out $L S L^{T}$ and keeping track of which entries are multiplied by which gives

$$
L S L^{T}=\left(\begin{array}{rrrr} 
& & & b_{1} \\
& L_{1} S_{1} L_{1}^{T} & & \vdots \\
b_{1} & \cdots & b_{n-1} & b_{n-1} \\
b_{n}
\end{array}\right)=\left(\begin{array}{rrrr}
d_{1} & & & b_{1} \\
& \ddots & & \vdots \\
& & d_{n-1} & b_{n-1} \\
b_{1} & \cdots & b_{n-1} & b_{n}
\end{array}\right)
$$

where $d_{1}, \ldots, d_{n-1}>0$ are the diagonal entries of $D_{1},\left(b_{1}, \ldots, b_{n-1}\right)=L_{1}\left(a_{1}, \ldots, a_{n-1}\right)$, and $b_{n}=a_{n}$. Since the $d_{i}$ are nonzero, we can do $n-1$ row operations $R_{n}-=\frac{b_{i}}{d_{i}} R_{i}$ to clear the entries in the last row. Doing the same column operations $C_{n}-=\frac{b_{i}}{d_{i}} C_{i}$ 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

$$
E\left(L S L^{T}\right) E^{T}=\left(\begin{array}{rrrrr}
d_{1} & & & & 0 \\
& \ddots & & 0 \\
& & d_{n-1} & 0 \\
0 & 0 & 0 & d_{n}
\end{array}\right)
$$

Note that $\operatorname{det}(E)=\operatorname{det}(L)=\operatorname{det}\left(L^{T}\right)=\operatorname{det}\left(E^{T}\right)=1$, so $\operatorname{det}(S)=d_{1} \cdots d_{n-1} d_{n}$. Since $\operatorname{det}(S)>0$ and $d_{1}, \ldots, d_{n-1}>0$, we have $d_{n}>0$ as well. Setting $L_{2}=(E L)^{-1}$ and $D_{2}=E L S L^{T} E^{T}$, this gives $S=L_{2} D_{2} L_{2}^{T}$. Since $S$ has an $L D L^{T}$ decomposition, it is positivedefinite, as desired.


[^0]:    ${ }^{1}$ What follows is essentially the same proof that the $L U$ decomposition is unique for an invertible matrix.

