Gaussian Elimination without/with Pivoting and Cholesky Decomposition

Gaussian Elimination WITHOUT pivoting

Notation: For a matrix $A \in \mathbb{R}^{n \times n}$ we define for $k \in \{1,\dots,n\}$ the **leading principal submatrix**

$$A_{(k)} := \begin{bmatrix} a_{11} & \cdots & a_{1k} \\ \vdots & & \vdots \\ a_{k1} & \cdots & a_{kk} \end{bmatrix}$$

We found out that Gaussian elimination without pivoting can fail even if the matrix A is nonsingular.

Example: For $A = \begin{bmatrix} 4 & -2 & 2 \\ -2 & 1 & 3 \\ 2 & -2 & 2 \end{bmatrix}$ we use $\ell_{21} = \frac{-2}{4}, \ell_{31} = \frac{2}{4}$ and obtain after step 1 of the elimination $U = \begin{bmatrix} 4 & -2 & 2 \\ 0 & 0 & 4 \\ 0 & -1 & 2 \end{bmatrix}$

Now we have $u_{22} = 0$ and the algorithm fails in column 2. Note that a linear system with the matrix $A_{(2)} = \begin{bmatrix} 4 & -2 \\ -2 & 1 \end{bmatrix}$ is

equivalent to a linear system with the matrix $U_{(2)} = \begin{bmatrix} 4 & -2 \\ 0 & 0 \end{bmatrix}$. This matrix is singular, hence the matrix $A_{(2)}$ is singular.

This explains why Gaussian elimination fails in column 2: The matrix $A_{(1)} = [4]$ was nonsingular, so elimination worked in column 1. But the matrix $A_{(2)}$ is singular, so the elimination fails in column 2.

For a matrix $A \in \mathbb{R}^{n \times n}$ we consider the submatrices $A_{(1)}, \dots, A_{(n)}$. If *all* of these matrices are nonsingular, then Gaussian elimination WITHOUT pivoting succeeds, and we obtain an upper triangular matrix U with nonzero elements on the diagonal. If one of these submatrices is singular: let $A_{(k)}$ be the first submatrix which is singular. Then Gaussian elimination WITHOUT pivoting works in columns $1,\dots,k-1$, but fails in column k.

Theorem: For a matrix $A \in \mathbb{R}^{n \times n}$ the following three statements are equivalent:

1. All the submatrices $A_{(1)}, \dots, A_{(n)}$ are nonsingular

2. Gaussian elimination WITHOUT pivoting succeeds and yields $u_{jj} \neq 0$ for $j = 1,\dots,n$

3. The matrix A has a decomposition $A = LU$ where L is lower triangular with 1's on the diagonal and U is upper triangular with nonzero diagonal elements.

Proof:

$(1.) \Longrightarrow (2.)$: Assume Gaussian elimination fails in column k, yielding a matrix U with $u_{kk} = 0$. Then a linear system with the matrix $A_{(k)}$ is equivalent to a linear system with the matrix

$$U_{(k)} = \begin{bmatrix} \circledast & * & \cdots & * \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \circledast & * \\ 0 & \cdots & 0 & 0 \end{bmatrix}$$

This matrix is singular, hence the matrix $A_{(k)}$ is singular.

$(2.) \Longrightarrow (3.)$: The row operations given by the multipliers ℓ_{ij} turn the matrix A into the matrix U. Hence reversing these row operations turns the matrix U back into the matrix A, yielding the equation $A = LU$.

$(3.) \Longrightarrow (1.)$ Let $j \in \{1,\dots,n\}$. $A = LU$ implies $A_{(j)} = L_{(j)}U_{(j)}$. Hence solving a linear system with the matrix $A_{(j)}$ is equivalent to solving a linear system with the matrix $U_{(j)}$. Since $U_{(j)}$ is nonsingular the matrix $A_{(j)}$ is nonsingular.

Gaussian Elimination WITH Pivoting

We now have for each column several **pivot candidates**: the diagonal element and all elements below it.

If one of the pivot candidates is nonzero we use a row interchange to move it to the diagonal position, and we can perform elimination in this column.

If all pivot candidates are zero the algorithm breaks down. If we can perform elimination for columns $1, \ldots, n-1$ and then obtain $u_{nn} = 0$ we say that the algorithm breaks down in column n.

The algorithm succeeds if we can perform elimination in columns $1, \ldots, n-1$ and we obtain $u_{nn} \neq 0$. In this case we obtain an upper triangular matrix U with nonzero diagonal elements.

If we start with a nonsingular matrix A, then Gaussian elimination with pivoting succeeds. Solving a linear system $Ax = b$ is then equivalent to a linear system $Ux = y$ which we can solve by back substitution.

Theorem: For a matrix $A \in \mathbb{R}^{n \times n}$ the following three statements are equivalent:

1. The matrix $A \in \mathbb{R}^{n \times n}$ is nonsingular

2. Gaussian elimination WITH pivoting succeeds and yields $u_{jj} \neq 0$ for $j = 1, \ldots, n$

3. There exists a decomposition
$$\begin{bmatrix} \text{row } p_1 \text{ of } A \\ \vdots \\ \text{row } p_n \text{ of } A \end{bmatrix} = LU$$
   where L is lower triangular with 1's on the diagonal, U is upper triangular with nonzero diagonal elements, and $p_1, \ldots, p_n$ is a permutation of the numbers $1, \ldots, n$.

Proof:

$(1.) \Longrightarrow (2.)$: Assume Gaussian elimination fails in column k, yielding a matrix U where *all* pivot candidates $u_{kk}, \ldots, u_{nk}$ are zero. Since elimination succeeded in columns $1, \ldots, k-1$ the diagonal elements $u_{11}, \ldots, u_{k-1,k-1}$ are all nonzero:

$$U = \begin{bmatrix}
\circledast & * & \cdots & \cdots & \cdots & \cdots & * \\
0 & \ddots & \ddots & & & & \vdots \\
\vdots & \ddots & \circledast & * & & & \vdots \\
\vdots & & 0 & {\color{red}0} & * & \cdots & * \\
\vdots & & \vdots & {\color{red}\vdots} & \vdots & & \vdots \\
\vdots & & \vdots & {\color{red}\vdots} & \vdots & & \vdots \\
0 & \cdots & 0 & {\color{red}0} & * & \cdots & *
\end{bmatrix}$$

We want to show that the linear system $Ax = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$ has a solution $x \neq \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$: The linear system $Ax = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$ is equivalent

to the linear system $Ux = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$. Consider a vector x with $x_k = 1$, $x_{k+1} = \cdots = x_n = 0$, i.e., $x = \begin{bmatrix} x_1 \\ \vdots \\ x_{k-1} \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$. Then equations

$k, k+1, \ldots, n$ are satisfied. We can find $x_1, \ldots, x_{k-1}$ such that equations $1, \ldots, k-1$ are satisfied by using back substitution:

equation $k-1$ yields a unique x_{k-1} (since $u_{k-1,k-1} \neq 0$),..., equation 1 yields a unique x_1 (since $u_{11} \neq 0$). We found a nonzero vector x such that $Ax = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$, hence A is singular.

(2.) $\Longrightarrow$ (3.): Let $\tilde{A} := \begin{bmatrix} \text{row } p_1 \text{of } A \\ \vdots \\ \text{row } p_n \text{of } A \end{bmatrix}$. Then applying Gaussian elimination WITHOUT pivoting to the matrix $\tilde{A}$ yields exactly the same L, U. Hence we have $\tilde{A} = LU$.

(3.) $\Longrightarrow$ (1.) The linear system with the matrix $Ax = b$ is equivalent to solving a linear system $Ux = y$. Since U is nonsingular the matrix A is nonsingular.

Quadratic forms

A **symmetric matrix** $A \in \mathbb{R}^{n \times n}$ defines a quadratic function $q(x_1, \ldots, x_n)$:
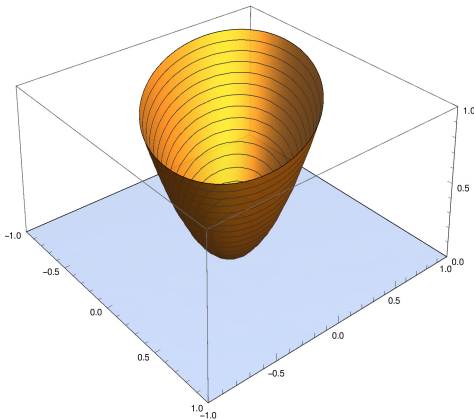
$$q(x_1, \ldots x_n) = \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} x_i x_j = x^\top A x$$

Such a function is called a **quadratic form**. For $x = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$ we obviously have $q(x) = 0$. What happens for $x \neq \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$?

Examples:

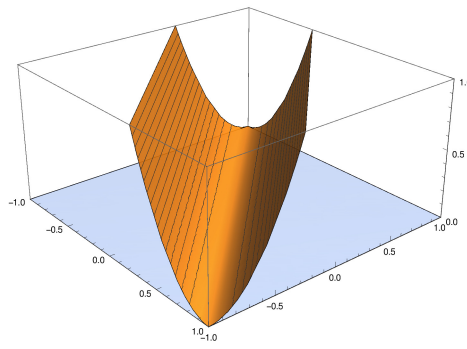$$A = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$$
$$q(x_1, x_2) = 2x_1^2 + 2x_1x_2 + 2x_2^2$$

$$A = \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix}$$
$$q(x_1, x_2) = 2x_1^2 + 4x_1x_2 + 2x_2^2$$
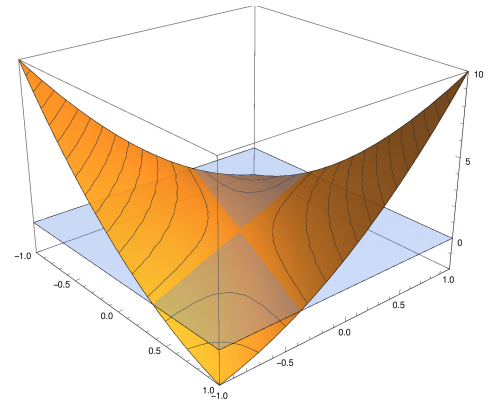
$$A = \begin{bmatrix} 2 & 3 \\ 3 & 2 \end{bmatrix}$$
$$q(x_1, x_2) = 2x_1^2 + 6x_1x_2 + 2x_2^2$$

for $x \neq \begin{bmatrix} 0 \\ 0 \end{bmatrix}$: $q(x)$ is positive

eigenvalues of A are 3,1

for $x \neq \begin{bmatrix} 0 \\ 0 \end{bmatrix}$: $q(x)$ is positive or zero

eigenvalues of A are 4, 0

for $x \neq \begin{bmatrix} 0 \\ 0 \end{bmatrix}$: $q(x)$ is positive or zero or negative

eigenvalues of A are $5, -1$

If $q(x) = x^\top A x > 0$ for all $x \neq \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$ we call the matrix A **positive definite**.

A symmetric matrix has real eigenvalues $\lambda_1, \ldots, \lambda_n$ and an orthonormal basis of eigenvectors $v^{(1)}, \ldots, v^{(n)}$. If we write the vector x in terms of eigenvectors we obtain

$$x = c_1 v^{(1)} + \cdots + c_n v^{(n)}, \qquad q(x) = [c_1, \ldots, c_n] \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \begin{bmatrix} c_1 \\ \vdots \\ c_n \end{bmatrix} = \lambda_1 c_1^2 + \cdots + \lambda_n c_n^2$$

We see: **a symmetric matrix is positive definite $\Longleftrightarrow$ all eigenvalues are positive**.

Cholesky decomposition

Many application problems give a linear system with a **symmetric matrix** $A \in \mathbb{R}^{n \times n}$, i.e., $a_{ij} = a_{ji}$ for all i, j.

Example 1: Consider the matrix $A = \begin{bmatrix} 9 & -6 & 6 \\ -6 & 5 & -1 \\ 6 & -1 & 15 \end{bmatrix}$. We perform Gaussian elimination WITHOUT pivoting and obtain $A = LU$:

$$\begin{bmatrix} 9 & -6 & 6 \\ -6 & 5 & -1 \\ 6 & -1 & 15 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ -\frac{2}{3} & 1 & 0 \\ \frac{2}{3} & 3 & 1 \end{bmatrix} \begin{bmatrix} 9 & -6 & 6 \\ 0 & 1 & 3 \\ 0 & 0 & 2 \end{bmatrix}$$

We can then use this to solve a linear system $Ax = b$: First solve $Ly = b$ using forward substitution, then solve $Ux = b$ using back substitution.

We would like to exploit the fact that A is symmetric and use a decomposition which reflects this:

> we want to find an upper triangular matrix $C \in \mathbb{R}^{n \times n}$ with $A = C^\top C$

$$\begin{bmatrix} 9 & -6 & 6 \\ -6 & 5 & -1 \\ 6 & -1 & 15 \end{bmatrix} = \begin{bmatrix} c_{11} & 0 & 0 \\ c_{12} & c_{22} & 0 \\ c_{13} & c_{23} & c_{33} \end{bmatrix} \begin{bmatrix} c_{11} & c_{12} & c_{13} \\ 0 & c_{22} & c_{23} \\ 0 & 0 & c_{33} \end{bmatrix}$$

$$= \begin{bmatrix} 3 & 0 & 0 \\ -2 & 1 & 0 \\ 2 & 3 & \sqrt{2} \end{bmatrix} \begin{bmatrix} 3 & -2 & 2 \\ 0 & 1 & 3 \\ 0 & 0 & \sqrt{2} \end{bmatrix}$$

Note:

- this decomposition $A = C^\top C$ is called **Cholesky decomposition**
- $A = C^\top C$ implies $A^\top = C^\top C = A$, i.e., such a decomposition can only be obtained for symmetric A.
- We can then use this decomposition to solve a linear system $Ax = b$: First solve $C^\top y = b$ using forward substitution, then solve $Cx = y$ using back substitution.

How can we find the matrix C?

- **row 1 of A:**

  diagonal element: $a_{11} = c_{11} c_{11} \Longrightarrow \boxed{c_{11} = \sqrt{a_{11}}}$

  for $k = 2, \ldots n$: $a_{1k} = c_{1k} c_{11} \Longrightarrow \boxed{c_{1k} = a_{1k}/c_{11}}$

- **row 2 of A:**

  diagonal element: $a_{22} = c_{12}^2 + c_{22}^2 \Longrightarrow \boxed{c_{22} = \sqrt{a_{22} - c_{12}^2}}$

  for $k = 3, \ldots, n$: $a_{2k} = c_{1k} c_{12} + c_{2k} c_{22} \Longrightarrow \boxed{c_{2k} = (a_{2k} - c_{1k} c_{12})/c_{22}}$

  $\vdots$

Algorithm:

For $j = 1, \ldots, n$:

  $s := a_{jj} - \sum_{\ell=1}^{j-1} c_{\ell j}^2$

  If $s \le 0$: stop with error "matrix not positive definite"

  $c_{jj} = \sqrt{s}$

  For $k = j+1, \ldots, n$:

  $$c_{jk} = \left(a_{jk} - \sum_{\ell=1}^{j-1} c_{\ell j} c_{\ell k} \right)/c_{jj}$$

How to compute row $j = 3$ of C:

$$
A = \begin{bmatrix}
* & * & * & * & * & * \\
* & * & * & * & * & * \\
* & * & \boxed{*} & \boxed{* & * & *} \\
* & * & * & * & * & * \\
* & * & * & * & * & * \\
* & * & * & * & * & *
\end{bmatrix}
\qquad
C = \begin{bmatrix}
\circledast & * & \boxed{*} & \boxed{* & * & *} \\
\circledast & * & \boxed{*} & \boxed{* & * & *} \\
 & & & & & \\
 & & & & & \\
 & & & & & \\
 & & & & &
\end{bmatrix}
\longmapsto
\begin{bmatrix}
\circledast & * & * & * & * & * \\
\circledast & & * & * & * & * \\
 & & \sqrt{\square - v^\top v} & (\square - v^\top M)/c_{jj} & & \\
 & & & & & \\
 & & & & & \\
 & & & & &
\end{bmatrix}
$$

Here $v = \texttt{C(1:j-1,j)}$ and $M = \texttt{C(1:j-1,j+1:n)}$. We obtain the following Matlab code:

```
function C = cholesky(A)
n = size(A,1);
C = zeros(n,n);
for j=1:n
  v = C(1:j-1,j);
  s = A(j,j) - v'*v;
  if s<=0
    error('matrix is not positive definite')
  end
  C(j,j) = sqrt(s);
  C(j,j+1:n) = (A(j,j+1:n)-v'*C(1:j-1,j+1:n))/C(j,j);
end
```

Work of Cholesky decomposition: For row j:

- compute c_{jj}: $(j-1)$ multiplications and 1 square root

- compute c_{jk} for $k = j+1,\ldots,n$: $(n-j)$ times $((j-1)$ multiplications and 1 division $)$

If we count multiplications, divisions and square roots we have $j + (n-j)j = (n-j+1)j$ operations for row j.

The total is therefore $\sum_{j=1}^{n}(n-j+1)j = \dfrac{n(n+1)(n+2)}{6} = \dfrac{n^3}{6} + O(n^2)$ operations. This is half the work of the LU decomposition which was $\dfrac{n^3}{3}$ operations.

Example 2: Consider the matrix $A = \begin{bmatrix} 9 & -6 & 6 \\ -6 & 5 & -1 \\ 6 & -1 & 12 \end{bmatrix}$.

The algorithm breaks down when we try to compute c_{33}: $s = 12 - (2^2 + 3^2) = -1 \leq 0$:

$$
C^\top C = \begin{bmatrix} 3 & 0 & 0 \\ -2 & 1 & 0 \\ 2 & 3 & 0 \end{bmatrix}\begin{bmatrix} 3 & -2 & 2 \\ 0 & 1 & 3 \\ 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 9 & -6 & 6 \\ -6 & 5 & -1 \\ 6 & -1 & 13 \end{bmatrix}
$$

Therefore we have $s = 12 - 13 = -1 \leq 0$, and the Cholesky decomposition does not exist.

The matrix A is symmetric and nonsingular. Why is there no Cholesky decomposition?

Assme that a matrix A has a Cholesky decomposition $A = C^\top C$ where C has nonzero elements.

For a vector $x \in \mathbb{R}^n$ we consider the scalar $x^\top A x = \sum_{i,j=1}^n a_{ij}x_i x_j$ (this is called a "quadratic form").

For any vector $x \neq \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$ we have $y := Cx \neq \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$ since C is nonsingular (upper triangular with nonzero diagonal entries).

Hence the quadratic form is positive:

$$
x^\top A x = x^\top C^\top C x = y^\top y > 0 \tag{1}
$$

Definition: A symmetric matrix $A \in \mathbb{R}^{n \times n}$ is called **positive definite** if for all $x \neq \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$ we have $x^\top A x > 0$.

E.g., the matrix $A = \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix}$ is positive definite since $x^\top A x = 2x_1^2 + 3x_2^2 > 0$ for $x \neq \begin{bmatrix} 0 \\ 0 \end{bmatrix}$.

On the other hand the matrix $A = \begin{bmatrix} 2 & 0 \\ 0 & -1 \end{bmatrix}$ is not positive definite since $[0,1]^\top A \begin{bmatrix} 0 \\ 1 \end{bmatrix} = -1 \leq 0$.

Note: A symmetric matrix A has real eigenvalues. The matrix is positive definite if and only if all eigenvalues are positive.

We can check whether a matrix is positive definite by trying to find the Cholesky decomposition. If C with nonzero diagonal elements exists, the matrix is positive definite because of (1).

If the algorithm breaks down we claim that the matrix is not positive definite.

E.g., the matrix A from Example 2 is not positive definite:

Pick $x_3 = 1$. Solve $Cx = [0,0,0]^\top$ using back substitution:

$$x = \begin{bmatrix} -8/3 \\ -3 \\ 1 \end{bmatrix}, \qquad x^\top A x = x^\top C^\top C x + x^\top \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix} x = -1$$

Theorem: For a symmetric matrix $A \in \mathbb{R}^{n \times n}$ the following three statements are equivalent:

1. The matrix is positive definite.

2. The Cholesky algorithm succeeds and gives $C \in \mathbb{R}^{n \times n}$ with nonzero diagonal elements.

3. There exists a decomposition $A = C^\top C$ where $C \in \mathbb{R}^{n \times n}$ is upper triangular with nonzero diagonal elements.

Proof:
(1.) $\implies$ (2.): Assume the algorithm breaks down in row j with $s \leq 0$. We want to show that A is not positive definite by finding a nonzero vector x with $x^\top A x \leq 0$: Let, e.g., $j = 3$. We have already computed the first two rows of C with $c_{11} \neq 0$, $c_{22} \neq 0$ and have $s = a_{33} - (c_{13}^2 + c_{23}^2) \leq 0$

$$C_{(3)}^\top C_{(3)} = \begin{bmatrix} c_{11} & 0 & 0 \\ c_{12} & c_{22} & 0 \\ c_{13} & c_{23} & 0 \end{bmatrix} \begin{bmatrix} c_{11} & c_{12} & c_{13} \\ 0 & c_{22} & c_{23} \\ 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{31} & c_{13}^2 + c_{23}^2 \end{bmatrix}$$

$$A_{(3)} = C_{(3)}^\top C_{(3)} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & s \end{bmatrix}$$

Now we construct a vector $\tilde{x} = \begin{bmatrix} x_1 \\ x_2 \\ 1 \end{bmatrix}$ with $\underbrace{\begin{bmatrix} c_{11} & c_{12} & c_{13} \\ 0 & c_{22} & c_{23} \\ 0 & 0 & 0 \end{bmatrix}}_{C_{(3)}} \begin{bmatrix} x_1 \\ x_2 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$: We use back substitution to find x_2 and

then x_1. This works since $c_{11}, c_{22} \neq 0$. Hence

$$\tilde{x}^\top A_{(3)} \tilde{x} = \tilde{x}^\top C_{(3)}^\top \underbrace{C_{(3)} \tilde{x}}_{\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}} + \tilde{x}^\top \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & s \end{bmatrix} \tilde{x} = 0 + s \leq 0$$

Now we define $x = \begin{bmatrix} x_1 \\ x_2 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$ and obtain $x^\top A x = \tilde{x}^\top A_{(3)} \tilde{x} = s \leq 0$.

(2.) $\Longrightarrow$ (3.): The algorithm gives C with $C^\top C = A$.

(3.) $\Longrightarrow$ (1.): as shown above in (1) $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$