LINEAR SYSTEM SOLUTION

SYNOPSIS. The traditional problem within linear algebra is to find the scaling coefficients of a linear combination to exactly represent some given vector. Methods with a long history of hand computation have been developed for this purpose, and can still offer insight into properties of linear mappings and their associated matrices.

1. Orthogonal projectors and linear systems

Consider the linear system $\boldsymbol{A}\boldsymbol{x}=\boldsymbol{b}$ with $\boldsymbol{A}\in\mathbb{R}^{m\times n}$, $\boldsymbol{b}\in\mathbb{R}^m$ given. The scaling coefficients $\boldsymbol{x}\in\mathbb{R}^n$ are sought and are said to be a solution of the linear system when the equation $\boldsymbol{A}\boldsymbol{x}=\boldsymbol{b}$ is satisfied. Orthogonal projectors and knowledge of the four fundamental matrix subspaces allows us to succintly express whether there exist no solutions, a single solution of an infinite number of solutions:

- Consider the factorization $\boldsymbol{Q}\boldsymbol{R}=\boldsymbol{A}$, the orthogonal projector $\boldsymbol{P}=\boldsymbol{Q}\boldsymbol{Q}^T$, and the complementary orthogonal projector $\boldsymbol{I}-\boldsymbol{P}$
- If $\|(\boldsymbol{I}-\boldsymbol{P})\boldsymbol{b}\|\neq 0$, then $\boldsymbol{b}$ has a component outside the column space of $\boldsymbol{A}$, and $\boldsymbol{A}\boldsymbol{x}=\boldsymbol{b}$ has no solution
- If $\|(\boldsymbol{I}-\boldsymbol{P})\boldsymbol{b}\|=0$, then $\boldsymbol{b}\in C(\boldsymbol{Q})=C(\boldsymbol{A})$ and the system has at least one solution
- If $N(\boldsymbol{A})=\{\boldsymbol{0}\}$ (null space only contains the zero vector, i.e., null space of dimension 0) the system has a unique solution
- If $\dim N(\boldsymbol{A})=n-r>0$, then a vector $\boldsymbol{y}\in N(\boldsymbol{A})$ in the null space is written as

  $$\boldsymbol{y}=c_1\boldsymbol{z}_1+\ldots+c_{n-r}\boldsymbol{z}_{n-r}$$

  and if $\boldsymbol{x}$ is a solution of $\boldsymbol{A}\boldsymbol{x}=\boldsymbol{b}$, so is $\boldsymbol{x}+\boldsymbol{y}$, since

  $$\boldsymbol{A}(\boldsymbol{x}+\boldsymbol{y})=\boldsymbol{A}\boldsymbol{x}+c_1\boldsymbol{A}\boldsymbol{z}_1+\ldots+c_{n-r}\boldsymbol{A}\boldsymbol{z}_{n-r}=\boldsymbol{b}+\boldsymbol{0}+\ldots+\boldsymbol{0}=\boldsymbol{b}$$

  The linear system has an $(n-r)$-parameter family of solutions

If a solution exists, it can be found by backsubstitution solution of $\boldsymbol{R}\boldsymbol{x}=\boldsymbol{Q}^T\boldsymbol{b}$. If multiple solutions exist, an orthonormal basis $\boldsymbol{Z}$ is found for the null space and the family of solutions is $\boldsymbol{x}+\boldsymbol{Z}\boldsymbol{y}$.

2. Gaussian elimination and row echelon reduction

Suppose now that $\boldsymbol{A}\,\boldsymbol{x}=\boldsymbol{b}$ admits a unique solution. The QR factorization approach of reducing the problem to $\boldsymbol{R}\boldsymbol{x}=\boldsymbol{Q}^T\boldsymbol{b}$ is one procedure to compute the solution. It has the benefit of working with the orthonormal $\boldsymbol{Q}$ matrix. Finding the orthonormal $\boldsymbol{Q}$ matrix is however a computational expense. Recall that orthogonality implied linear independence. Other approaches might exist that only impose linear independence, without orthogonality. Gaussian elimination is the main such approach. Consider the system

$$\left\{\begin{array}{rcl} x_1+2x_2-x_3 & = & 2\\ 2x_1-x_2+x_3 & = & 2\\ 3x_1-x_2-x_3 & = & 1 \end{array}\right.$$

The idea is to combine equations such that we have one fewer unknown in each equation. Ask: with what number should the first equation be multiplied in order to eliminate x_1 from sum of equation 1 and equation 2? This number is called a Gaussian multiplier, and is in this case -2. Repeat the question for eliminating x_1 from third equation, with multiplier -3.

$$\left\{\begin{array}{rcr} x_1+2x_2-x_3 & = & 2 \\ 2x_1-x_2+x_3 & = & 2 \\ 3x_1-x_2-x_3 & = & 1 \end{array}\right. \Rightarrow \left\{\begin{array}{rcr} x_1+2x_2-x_3 & = & 2 \\ -5x_2+3x_3 & = & -2 \\ -7x_2+2x_3 & = & -5 \end{array}\right.$$

Now, ask: with what number should the second equation be multiplied to eliminate x_2 from sum of second and third equations. The multiplier is in this case $-7/5$.

$$\left\{\begin{array}{rcr} x_1+2x_2-x_3 & = & 2 \\ -5x_2+3x_3 & = & -2 \\ -7x_2+2x_3 & = & -5 \end{array}\right. \Rightarrow \left\{\begin{array}{rcr} x_1+2x_2-x_3 & = & 2 \\ -5x_2+3x_3 & = & -2 \\ -\frac{11}{5}x_3 & = & -\frac{11}{5} \end{array}\right.$$

Starting from the last equation we can now find $x_3=1$, replace in the second to obtain $-5x_2=-5$, hence $x_2=1$, and finally replace in the first equation to obtain $x_1=1$.

The above operations only involve coefficients. A more compact notation is therefore to work with what is known as the "bordered matrix" and work with coefficients arising in rows

$$[\,\boldsymbol{A}\;\;\boldsymbol{b}\,]=\left[\begin{array}{rrr|r} \boxed{1} & 2 & -1 & 2 \\ 2 & -1 & 1 & 2 \\ 3 & -1 & -1 & 1 \end{array}\right] \sim [\,\boldsymbol{A}_1\;\;\boldsymbol{b}_1\,]=\left[\begin{array}{rrr|r} 1 & 2 & -1 & 2 \\ 0 & \boxed{-5} & 3 & -2 \\ 0 & -7 & 2 & -5 \end{array}\right] \sim [\,\boldsymbol{A}_2\;\;\boldsymbol{b}_2\,]=\left[\begin{array}{rrr|r} 1 & 2 & -1 & 2 \\ 0 & -5 & 3 & -2 \\ 0 & 0 & \boxed{-\frac{11}{5}} & -\frac{11}{5} \end{array}\right].$$

In Julia the above operations would be carried out as

```
∴ A=[1. 2 -1 2; 2 -1 1 2; 3 -1 -1 1]; A[2,:]=A[2,:]-2*A[1,:]; A[3,:]=A[3,:]-3*A[1,:];
```

```
∴ A
```

$$\left[\begin{array}{rrrr} 1.0 & 2.0 & -1.0 & 2.0 \\ 0.0 & -5.0 & 3.0 & -2.0 \\ 0.0 & -7.0 & 2.0 & -5.0 \end{array}\right] \tag{1}$$

```
∴ A[3,:]=A[3,:]-(7/5)*A[2,:]; A
```

$$\left[\begin{array}{rrrr} 1.0 & 2.0 & -1.0 & 2.0 \\ 0.0 & -5.0 & 3.0 & -2.0 \\ 0.0 & 0.0 & -2.1999999999999993 & -2.2 \end{array}\right] \tag{2}$$

```
∴
```

Once the above *triangular* form has been obtained, the solution is found by back substitution, in which we seek to form the identity matrix in the first 3 columns, and the solution is obtained in the last column.

$$\left[\begin{array}{rrrr} 1 & 2 & -1 & 2 \\ 0 & -5 & 3 & -2 \\ 0 & 0 & -\frac{11}{5} & -\frac{11}{5} \end{array}\right] \sim \left[\begin{array}{rrrr} 1 & 2 & -1 & 2 \\ 0 & -5 & 3 & -2 \\ 0 & 0 & 1 & 1 \end{array}\right] \sim \left[\begin{array}{rrrr} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{array}\right].$$

The operations arising in Gaussian elimination are successive linear combinations of rows that maintain the solution of the linear system. This idea is useful in identifying the fundamental subspaces associated with a matrix. The matrices arising at successive stages of the procedure are said to be similar to one another

$$\boldsymbol{A} \sim \boldsymbol{A}_1 \sim \boldsymbol{A}_2,$$

and since $\boldsymbol{A}_k$ is obtained by linear combination of the rows of $\boldsymbol{A}_{k-1}$, the row space is not changed

$$C(\boldsymbol{A}^T) = C(\boldsymbol{A}_1^T) = C(\boldsymbol{A}_2^T) = \cdots.$$

During the procedure a pivot element is identified in the diagonal position, as shown bordered above. If a zero value is encountered rows are permuted to bring a non-zero element to the pivot position. If a non-zero pivot value cannot be found by row permutation, one is sought by column permutations also. If a non-zero pivot cannot be found by either row or column permutations, the matrix is rank-deficient $r = \operatorname{rank}(\boldsymbol{A}) < \min(m, n)$ and has a non-trivial null space as in the following examples

$$\boldsymbol{A} = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 1 & 1 \\ 1 & 2 & 3 \end{bmatrix} \in \mathbb{R}^{3\times 3}, \boldsymbol{b} = \begin{bmatrix} 3 \\ 1 \\ 3 \end{bmatrix} \in \mathbb{R}^3, \boldsymbol{c} = \begin{bmatrix} 3 \\ 1 \\ 4 \end{bmatrix} \in \mathbb{R}^3.$$

$$[\, \boldsymbol{A} \;\; \boldsymbol{b} \,] = \left[\begin{array}{ccc|c} 1 & 2 & 3 & 3 \\ 0 & 1 & 1 & 1 \\ 1 & 2 & 3 & 3 \end{array}\right] \sim [\, \boldsymbol{A}_1 \;\; \boldsymbol{b}_1 \,] = \left[\begin{array}{ccc|c} 1 & 2 & 3 & 3 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{array}\right] \Longleftrightarrow \begin{cases} x_1 + 2x_2 + 3x_3 & = 3 \\ x_2 + x_3 & = 1 \\ 0 & = 0 \end{cases},$$

$$[\, \boldsymbol{A} \;\; \boldsymbol{c} \,] = \left[\begin{array}{ccc|c} 1 & 2 & 3 & 3 \\ 0 & 1 & 1 & 1 \\ 1 & 2 & 3 & 4 \end{array}\right] \sim [\, \boldsymbol{A}_1 \;\; \boldsymbol{c}_1 \,] = \left[\begin{array}{ccc|c} 1 & 2 & 3 & 3 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{array}\right] \Longleftrightarrow \begin{cases} x_1 + 2x_2 + 3x_3 & = 3 \\ x_2 + x_3 & = 1 \\ 0 & = 1 \end{cases}.$$

The $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ has an infinite number of solutions, while the $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{c}$ system has no solutions. Note that $\boldsymbol{A}_1$ has a row of zeros, hence the rows must be linearly dependent and $N(\boldsymbol{A}) \neq \{\boldsymbol{0}\}$. By the FTLA when $\boldsymbol{b} \in C(\boldsymbol{A})$ an infinite number of solutions is obtained, and for $\boldsymbol{c} \notin C(\boldsymbol{A})$ no solutions are obtained.

The rows with non-zero pivot elements are linearly independent, and reduction to the above row-echelon form is useful to identify the rank of a matrix. The first non-zero entry on a row is called either a pivot or a leading entry. A matrix is said to be brought to reduced row-echelon form when:

- all zero rows are below non-zero rows;
- in each non-zero row, the leading entry is to the left of lower leading entries;
- each leading entry equals 1 and is the only non-zero entry in its column.

In contrast to the Gram-Schmidt procedure, Gaussian elimination does not impose orthogonality between rows, nor that a row have unit norm. This leads to fewer computations, and is therefore well-suited to hand computation of small-dimensional matrices.

The steps in Gaussian elimination can be precisely specified in a format suitable for direct computer coding.

Algorithm Gauss elimination without pivoting

```
for s = 1 to m − 1
   for i = s + 1 to m
      t = −a_is / a_ss
      for j = s + 1 to m
         a_ij = a_ij + t · a_sj
      b_i = b_i + t · b_s

for s = m downto 1
   x_s = b_s / a_ss
   for i = 1 to s − 1
      b_i = b_i − a_is · x_s

return x
```

The variant of the above algorithm that accounts for possible zeros arising in a diagonal position is known as Gauss elimination with pivoting.

Algorithm Gauss elimination with partial pivoting

```
p = 1: m (initialize row permutation vector)
for s = 1 to m − 1
   piv = abs(a_{p(s),s})
   for i = s + 1 to m
      mag = abs(a_{p(i),s})
      if mag > piv then
         piv = mag; k = p(s); p(s) = p(i); p(i) = k
      if piv < ϵ then break("Singular matrix")
      t = −a_{p(i)s} / a_{p(s)s}
      for j = s + 1 to m
         a_{p(i)j} = a_{p(i)j} + t · a_{p(s)j}
      b_{p(i)} = b_{p(i)} + t · b_{p(s)}


for s = m downto 1
   x_s = b_{p(s)} / a_{p(s)s}
   for i = 1 to s − 1
      b_{p(i)} = b_{p(i)} − a_{p(i)s} · x_s

return x
```

3. LU-factorization

The operations arising in Gaussian elimination correspond to a matrix factorization, analogous to how the Gram-Schmidt procedure can be stated as the QR factorization. Revisting the previous example

$$\left\{ \begin{array}{rcl} x_1 + 2x_2 - x_3 & = & 2 \\ 2x_1 - x_2 + x_3 & = & 2 \\ 3x_1 - x_2 - x_3 & = & 1 \end{array} \right. \Longleftrightarrow \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}, \boldsymbol{A} = \left[\begin{array}{rrr} 1 & 2 & -1 \\ 2 & -1 & 1 \\ 3 & -1 & -1 \end{array} \right], \boldsymbol{b} = \left[\begin{array}{r} 2 \\ 2 \\ 1 \end{array} \right],$$

the idea is to express linear combinations of rows as a matrix multiplication. Recall that $\boldsymbol{A}\boldsymbol{x}$ is a linear combination of columns, and $\boldsymbol{A}\boldsymbol{X}$ expresses multiple column linear combinations. Linear combinations of columns are expressed as products in which the first factor contains the columns and the second contains the scaling coefficients. Analogously linear combinations of rows are expressed by products $\boldsymbol{L}\boldsymbol{A}$ where now the left factor contains the scaling coefficients entering into a linear combination of the rows of $\boldsymbol{A}$. For example, the first stage of Gaussian elimination for the above system can be expressed as

$$\boldsymbol{L}_1\boldsymbol{A} = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ -3 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & -1 \\ 2 & -1 & 1 \\ 3 & -1 & -1 \end{bmatrix} = \begin{bmatrix} 1 & 2 & -1 \\ 0 & -5 & 3 \\ 0 & -7 & 2 \end{bmatrix}.$$

The next stage is also expressed as a matrix multiplication, after which an upper triangular matrix $\boldsymbol{U}$ is obtained

$$\boldsymbol{L}_2\boldsymbol{L}_1\boldsymbol{A} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -7/5 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & -1 \\ 0 & -5 & 3 \\ 0 & -7 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 2 & -1 \\ 0 & -5 & 3 \\ 0 & 0 & -11/5 \end{bmatrix} = \boldsymbol{U}.$$

For a general matrix $\boldsymbol{A} \in \mathbb{R}^{m\times m}$ the sequence of operations is

$$\boldsymbol{L}_{m-1} \dots \boldsymbol{L}_2\boldsymbol{L}_1\boldsymbol{A} = \boldsymbol{U}.$$

DEFINITION. *The matrix*

$$\boldsymbol{L}_k = \begin{pmatrix} 1 & \dots & 0 & \dots & 0 \\ 0 & \ddots & 0 & \dots & 0 \\ 0 & \dots & 1 & \dots & 0 \\ 0 & \dots & -l_{k+1,k} & \dots & 0 \\ 0 & \dots & -l_{k+2,k} & \dots & 0 \\ \vdots & \dots & \vdots & \ddots & \vdots \\ 0 & \dots & -l_{m,k} & \dots & 1 \end{pmatrix}$$

with $l_{i,k} = a_{i,k}^{(k)} / a_{k,k}^{(k)}$, and $\boldsymbol{A}^{(k)} = (a_{i,j}^{(k)})$ the matrix obtained after step k of row echelon reduction (or, equivalently, Gaussian elimination) is called a Gaussian multiplier matrix.

The inverse of a Gaussian multiplier is

$$\boldsymbol{L}_k^{-1} = \begin{pmatrix} 1 & \dots & 0 & \dots & 0 \\ 0 & \ddots & 0 & \dots & 0 \\ 0 & \dots & 1 & \dots & 0 \\ 0 & \dots & l_{k+1,k} & \dots & 0 \\ 0 & \dots & l_{k+2,k} & \dots & 0 \\ \vdots & \dots & \vdots & \ddots & \vdots \\ 0 & \dots & l_{m,k} & \dots & 1 \end{pmatrix} = \boldsymbol{I} - (\boldsymbol{L}_k - \boldsymbol{I}).$$

From $(\boldsymbol{L}_{m-1}\boldsymbol{L}_{m-2}\dots\boldsymbol{L}_2\boldsymbol{L}_1)\boldsymbol{A} = \boldsymbol{U}$ obtain

$$\boldsymbol{A} = (\boldsymbol{L}_{m-1}\boldsymbol{L}_{m-2}\dots\boldsymbol{L}_2\boldsymbol{L}_1)^{-1}\boldsymbol{U} = \boldsymbol{L}_1^{-1}\boldsymbol{L}_2^{-1}\cdot\ldots\cdot\boldsymbol{L}_{m-1}^{-1}\boldsymbol{U} = \boldsymbol{L}\boldsymbol{U}.$$

The above is known as an LU factorization, short for lower-upper factorization. Solving a linear system by LU-factorization consists of the steps:

1. Find the factorization $LU = A$
2. Insert the factorization into $Ax = b$ to obtain $(L\ U)x = L(U\ x) = L\ y = b$, where the notation $y = U\ x$ has been introduced. The system

$$Ly = b$$

   is easy to solve by forward substitution to find y for given b
3. Finally find x by backward substitution solution of

$$Ux = y$$

The various procedures encountered so far to solve a linear system are described in the following table.

Given $\boldsymbol{A} \in \mathbb{R}^{m \times n}$		
Singular value decomposition	Gram-Schmidt	Lower-upper
Transformation of coordinates	$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$	
$\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T = \boldsymbol{A}$	$\boldsymbol{Q}\boldsymbol{R} = \boldsymbol{A}$	$\boldsymbol{L}\boldsymbol{U} = \boldsymbol{A}$
$(\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T)\boldsymbol{x} = \boldsymbol{b} \Longrightarrow \boldsymbol{U}\boldsymbol{y} = \boldsymbol{b} \Longrightarrow \boldsymbol{y} = \boldsymbol{U}^T\boldsymbol{b}$	$(\boldsymbol{Q}\boldsymbol{R})\boldsymbol{x} = \boldsymbol{b} \Longrightarrow \boldsymbol{Q}\boldsymbol{y} = \boldsymbol{b},\ \boldsymbol{y} = \boldsymbol{Q}^T\boldsymbol{b}$	$(\boldsymbol{L}\ \ \boldsymbol{U})\boldsymbol{x} = \boldsymbol{b} \Longrightarrow \boldsymbol{L}\ \ \boldsymbol{y} = \boldsymbol{b}$ (forward sub to find)**y**
$\boldsymbol{\Sigma}\boldsymbol{z} = \boldsymbol{y} \Longrightarrow \boldsymbol{z} = \boldsymbol{\Sigma}^+\boldsymbol{y}$	$\boldsymbol{R}\boldsymbol{x} = \boldsymbol{y}$ (back sub to find $\boldsymbol{x}$)	$\boldsymbol{U}\boldsymbol{x} = \boldsymbol{y}$ (back sub to find $\boldsymbol{x}$)
$\boldsymbol{V}^T\boldsymbol{x} = \boldsymbol{z} \Longrightarrow \boldsymbol{x} = \boldsymbol{V}\boldsymbol{z}$		

4. Matrix inverse

For $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ the pseudo-inverse $\boldsymbol{A}^+$ has been introduced based on the SVD, $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T$ as

$$\boldsymbol{A}^+ = \boldsymbol{V}\boldsymbol{\Sigma}^+\boldsymbol{U}^T.$$

When $\boldsymbol{A} \in \mathbb{R}^{m \times m}$ is square and of full rank the system $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ has a solution that can be stated as $\boldsymbol{x} = \boldsymbol{A}^{-1}\boldsymbol{b}$, where $\boldsymbol{A}^{-1}$ is the inverse of $\boldsymbol{A}$. The matrix $\boldsymbol{A}$ is said to be invertible $\boldsymbol{X} \in \mathbb{R}^{m \times m}$ such that

$$\boldsymbol{A}\boldsymbol{X} = \boldsymbol{X}\boldsymbol{A} = \boldsymbol{I},$$

and in this case $\boldsymbol{X} = \boldsymbol{A}^{-1}$ is the inverse of $\boldsymbol{A}$.

The inverse can be computed by extending Gauss elimination

$$[\ \boldsymbol{A}\ \mid\ \boldsymbol{I}\ ] \sim [\ \boldsymbol{I}\ \mid\ \boldsymbol{X}\ ],$$

a procedure known as the Gauss-Jordan algorithm.

A square matrix has an inverse only when it is of full rank. The following are equivalent statements:

a) $\boldsymbol{A}$ invertible

b) $\boldsymbol{Ax} = \boldsymbol{b}$ has a unique solution for all $\boldsymbol{b} \in \mathbb{R}^m$

c) $\boldsymbol{Ax} = \mathbf{0}$ has a unique solution

d) The reduced row echelon form of $\boldsymbol{A}$ is $\boldsymbol{I}$

e) $\boldsymbol{A}$ can be written as product of elementary matrices