Linear Algebra – Lesson 11

The four fundamental subspaces

1. Some motivation – the inverse matrix

Early in the course (lesson 4) we talked about the inverse matrix. For a real $n \times n$ matrix A we said that the inverse matrix A^{-1} exists if the following equivalent conditions are met:

- $\det(A) \neq 0$
- The columns of A are linearly independent
- $\dim\big(Im(A)\big) = n$
- $rank(A) = n$
- A is full-rank
- $\ker(A) = \{0\}$ (the nullspace of A has only the zero vector)

But what if our matrix is $m \times n$?

Today we will see if we can define an inverse matrix for an $m \times n$ matrix (spoiler alert: we can't, but we will define a left-inverse or a right-inverse).

2. Rectangular $\boldsymbol{m} \times \boldsymbol{n}$ matrices

Watch the footnote video in the 3blue1brown series, "Nonsquare matrices as transformations between dimensions".

Let A be an $m \times n$ matrix. For example:

$$A = \begin{pmatrix} 1 & 2 \\ 2 & 0 \\ 0 & 3 \end{pmatrix}$$

This matrix takes as input a vector $\bar{v} \in \mathbb{R}^2$ and gives as output a vector $(A\bar{v}) \in \mathbb{R}^3$.

Given a matrix, we immediately think of its input space $(\mathbb{R}^n)$ and its output space $(\mathbb{R}^m)$.

If we think about it geometrically, we can see that A takes all vectors in the 2D space $\mathbb{R}^2$, and maps them to the 3D space $\mathbb{R}^3$. Since any result $A\bar{v}$ is a linear combination of the columns of A, the result is some plane in $\mathbb{R}^3$. This plane is $Im(A)$.

In fact, A does more than that. It divides each of these vector spaces into two subspaces. We say that A induces four vector subspaces, which we call the four fundamental subspaces. In this lesson, we will spend some time on understanding the different subspaces and the relations between them.

3. The four fundamental subspaces

In class, we built together this schematic diagram of the four fundamental subspaces:

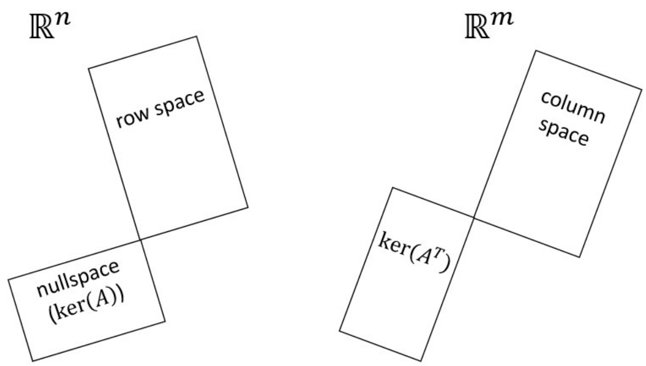

The four fundamental subspaces. $rowsp(A)$ and $ker(A)$ are orthogonal subspaces, and together (by taking all possible linear combinations of their vectors) they make up the entire input space $\mathbb{R}^n$. Similarly, $colsp(A)$ and $ket(A^T)$ are orthogonal subspaces, and together they make up the entire output space $\mathbb{R}^m$.

3.1 Column space

The column space, $colsp(A) = Im(A)$, is the set of all linear combinations of the columns of A. Since each column has m entries, $colsp(A) \subseteq \mathbb{R}^m$.

3.2 Row space

Similarly, the row space, $rowsp(A)$ is the set of all linear combinations of the rows of A. Since each row has n entries, $rowsp(A) \subseteq \mathbb{R}^n$.

3.3 Kernel, null space

The kernel of A, $\ker(A)$, also called the null space of A, is the set of all input vectors that A maps to zero. Therefore, $\ker(A) \subseteq \mathbb{R}^n$.

In one of the exercises, you proved that the null space and row space are orthogonal to each other. In fact, they are called **orthogonal complements**, because together they span the entire $\mathbb{R}^n$ space (what is the single vector that they share?)[1].

3.4 The left null space

The left null space consists of all vectors $\bar{y}$ such that:

$$\bar{y}^T A = 0$$

It is also called the kernel of A^T, since taking the transpose of both sides gives:

$$A^T \bar{y} = 0$$

Since each column of A (and therefore each row of A^T) has m entries, $\ker(A^T) \subseteq \mathbb{R}^m$.

We have previously shown that $\ker(A^T)$ is orthogonal to $colsp(A)$. In fact, they are also orthogonal complements, because together they span the entire $\mathbb{R}^m$ space.

[1] What does it mean that "together they span the entire $\mathbb{R}^n$ space? This is not to say that any vector is either in $rowsp(A)$ or in $\ker(A)$. The idea is that any vector in $\mathbb{R}^n$ can be decomposed into a component in $rowsp(A)$ and an orthogonal component in $\ker(A)$.

A note about dimensions

This rank-nullity theorem states that the dimension of both the column space and the row space is equal to the rank of the matrix:

$$\dim\big(colsp(A)\big) = \dim\big(rowsp(A)\big) = r$$

Therefore, you can now see that the dimension of the null spaces must be:

$$\dim(\ker(A)) = n - r$$
$$\dim(\ker(A^T)) = m - r$$

3.5 The abstract picture and the geometric picture (of the four fundamental subspaces)

Our first example above was a 3×2 matrix. A schematic representation of what this matrix does is shown on the right – it maps the 2D plane to some plane in 3D space.

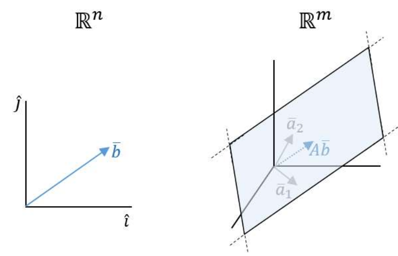

Question: Where is each one of the spaces in this geometric picture?

Answer:

- The row space is the entire input space $\mathbb{R}^2$.
- $\ker(A)$ is only the zero vector in $\mathbb{R}^2$. No other vector is mapped to zero, because the columns of A are linearly independent (and this is why we end up with a plane in $\mathbb{R}^3$).
- $colsp(A)$ is $Im(A)$, and that is the plane illustrated on the right (notice that it must include the origin, as it must include the zero vector).
- $\ker(A^T)$, the left null space, is all the vectors in $\mathbb{R}^m$ that are orthogonal to $Im(A)$. This is the line of vectors that goes through the origin, and is orthogonal to the blue plane.

Question: Is the transformation A from the previous illustration invertible?

Answer: First, remember that an inverse matrix exists only for square matrices. However, let's see if the transformation is invertible in some sense. Let an output vector $A\bar{v}$ be a vector in $Im(A)$, in the blue plane. Since each original 2D vector was mapped to a unique vector in $\mathbb{R}^3$, we can "undo" this transformation, and recover $\bar{v}$. In other words, A has a left inverse:

$$A_{left}^{-1} A\bar{v} = \bar{v}$$

However, strictly speaking, A is not invertible since it does not have a right-inverse. See a note about the existence of right- and left- inverse below (4.2).

Conclusions from Ex11. Why is I-P a projection matrix?

4. Left-inverse and right-inverse

4.1 Left and right inverse

Back to the inverse matrix. The inverse matrix A^{-1} exists if and only if:

$$r = m = n$$

It is also called a two-sided inverse matrix, since: $AA^{-1} = A^{-1}A = I$.

In this case, A is called a **full-rank matrix**. This is the special case in which $\ker(A) = \{0\}$ and also $\ker(A^T) = \{0\}$. No vector is mapped to 0 (except for the zero vector), and there exists an inverse matrix.

A rectangular matrix cannot have a left-inverse **and** a right-inverse. Why?
Because once you have more rows (or more columns), then they are necessarily linearly dependent. Therefore, they induce some null space, and this null space "spoils" everything and doesn't allow to go back to the original vector[2].

4.2 $A^T A$ and the four fundamental subspaces

We'll take a short detour talking about $A^T A$ and that will help us in learning about the left inverse and right inverse in the next section.
We have seen the matrix $A^T A$ come out before, for example when we discussed the least squares solution to an overdetermined system $A\bar{x} = \bar{b}$, and reached the following equation:

$$A^T A\bar{x} = A^T \bar{b}$$

We now prove some properties of $A^T A$.

Theorem I: The null space of $A^T A$ is the same as the null space of A.

Proof: First, we reformulate the question. We want to prove that if $A^T A\bar{x} = 0$ then also $A\bar{x} = 0$ (the other direction is trivial – can you see why?). So:

$$A^T A\bar{x} = 0$$

Multiply from the left by $\bar{x}^T$:

$$\bar{x}^T A^T A\bar{x} = 0$$

But this is just the inner product between a vector and itself:

$$\langle A\bar{x}, A\bar{x}\rangle$$

And the norm of a vector is 0 if only if it is the zero vector, so:

$$A\bar{x} = 0$$

Theorem II: $A^T A$ is invertible if and only if A has independent columns.

Proof[3]

Direction 1

We know that $A^T A$ is invertible, so $\ker(A^T A) = \{0\}$.
According to theorem I, also $\ker(A) = \{0\}$. Hence, the columns of A are independent.

Direction 2

Given: The columns of A are independent (i.e., $\ker(A) = \{0\}$).
To prove that $A^T A$ is invertible, we will show that $\ker(A^T A) = \{0\}$.
Assume that for some $\bar{x}$: $A^T A\bar{x} = 0$. We will show that necessarily $\bar{x} = 0$.
Multiply by $\bar{x}^T$ from the left:

$$\bar{x}^T A^T A\bar{x} = 0$$

Which implies (just like in the proof for Theorem I):

$$A\bar{x} = 0$$

---

[2] If thinking about the right inverse confuses you, wait until you read the end of 4.2.2, and then come back here.
[3] In a proof of the form "x is true if and only id y is true", we have to divide the proof in two: (1) assume x and prove y, and (2) assume y and prove x.

And according to what is given, this forces

$$\bar{x} = 0 \blacksquare$$

4.2.1 Left inverse (full column-rank)

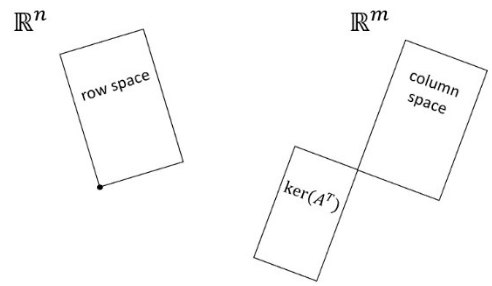

In order to have a left inverse, we must require that A sends no vector to zero except for the zero vector itself. In other words, $\ker(A) = \{0\}$.
Such a matrix is called full column-rank. The columns of A are independent, and we get:

$$r = n < m$$

How many solutions does $A\bar{x} = \bar{b}$ have? This was the case of a unique solution or no solution at all. You can see the reason for this in the illustration. If there is a solution, it is unique. Here are some justifications to why this is true:

- $\bar{b} \in \mathbb{R}^m$. Having a solution $\bar{x}$ means that $\bar{b} \in colsp(A)$. In other words, $\bar{b}$ is a linear combination of the columns of A. Since the columns are independent, we have no freedom in choosing the appropriate linear combination that gives $\bar{b}$. Hence, there is only one such combination.
- When there are infinite solutions, each solution can be written as the sum of a particular solution plus some solution of the homogeneous system $A\bar{x} = 0$:
  $$\bar{x} = \bar{x}_p + \bar{x}_{\mathcal{N}}$$
  In other words, they require that $\ker(A)$ will have more vectors in it, and not only the zero vector like we have here.

On the other hand, if $\bar{b} \notin colsp(A)$, we have no solution at all.

Now, how can we find the left inverse A_{left}^{-1}? This is where $A^T A$ comes in. Theorem II above tells us that in this case $A^T A$ is invertible. We can use this to obtain a formula for A_{left}^{-1}:

$$A_{left}^{-1} = (A^T A)^{-1} A^T$$

Let's check that this formula really gives us the left inverse:

$$A_{left}^{-1} A = (A^T A)^{-1} A^T A = I_{n \times n}$$

Notice how Theorem II ("if A is full column-rank, then $(A^T A)^{-1}$ exists") is necessary here.
What about a right inverse for A in this case?

A friend from the past

What happens if we try to use A_{left}^{-1} as a right inverse?

$$A A_{left}^{-1} = A(A^T A)^{-1} A^T$$

This is exactly the projection matrix onto $Im(A)$, which we have seen before.

Question: How did that happen? Why does it make sense that we got the projection matrix?
Answer: We will think about this together in class. Try thinking about a good explanation yourself.

4.2.2 A rectangular matrix cannot have both a left- and a right-inverse

A rectangular matrix cannot have a two-sided inverse because either that matrix or its transpose have a nonzero kernel. The rank of an $m \times n$ matrix is at most the minimum between m and n. When $m > n$ (you have more rows than columns), the rank cannot be greater than n. This means that the rows are necessarily linearly dependent, which means that $\ker(\boldsymbol{A^T})$ is non-zero. Therefore a **right**-inverse does not exist.

Now, thinking about a right-inverse might seem strange: we are used to thinking about the left-inverse as the matrix that "undoes" the transformation of A, and we know that it doesn't exist if A has a non-zero kernel. This can help you understand why $\ker(A^T) \neq \{0\}$ implies that the right-inverse doesn't exist. If it did exist, then:

$$AA_{right}^{-1} = I_{m\times m}$$

But then, taking the transpose of both sides:

$$\left(A_{right}^{-1}\right)^T A^T = I_{m\times m}$$

However, since A^T has a non-zero kernel, it cannot have a left inverse, and therefore A_{right}^{-1} cannot exist.

5. Pseudo-inverse (and SVD)

We are almost ready to talk the pseudo-inverse. As we said, the null spaces are those that cause us problems when we try to invert a matrix. However, if we leave the null spaces aside and just talk about the row space and column space, we may notice a useful property:

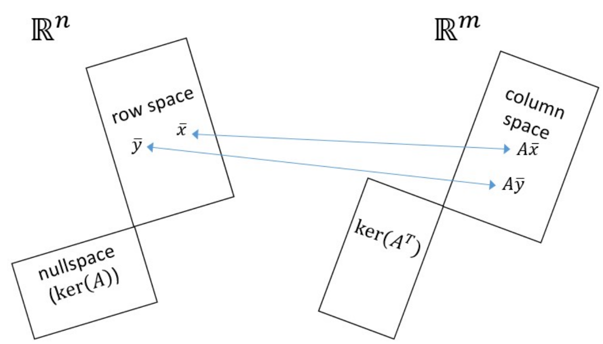

There is a one-to-one mapping[4] from the row space to the column space. This is equivalent to saying:

- For every vector $\bar{w} \in cols(A)$, there exists at most one vector $\bar{v} \in rowsp(A)$ such that $A\bar{v} = \bar{w}$.
- If $\bar{x}, \bar{y} \in rowsp(A)$ and $\bar{x} \neq \bar{y}$ then $A\bar{x} \neq A\bar{y}$. Can you prove it yourselves?
- Different inputs in $rowsp(A)$ give different outputs in $colsp(A)$.

Proof:

Let $\bar{x} \neq \bar{y}$ (both in $rowsp(A)$), and assume that $A\bar{x} = A\bar{y}$.

Then $A(\bar{x} - \bar{y}) = 0$. But this means that $(\bar{x} - \bar{y}) \in ker(A)$.

Now, $\bar{x} \in rowsp(A)$ and also $\bar{y} \in rowsp(A)$. Since the row space is a vector subspace, any linear combination of these is also in the row space, so:

$$(\bar{x} - \bar{y}) \in rowsp(A)$$

But we know that $\ker(A)$ and $rowsp(A)$ are orthogonal subspaces, and therefore $\bar{x} - \bar{y}$ can only be the zero vector, meaning:

$$\bar{x} = \bar{y}$$

In contradiction to our assumption. This finishes the proof.

[4] Do not confuse "one-to-one" with "onto". A transformation $A: \mathbb{R}^n \to \mathbb{R}^m$ is called "onto", if $Im(A) = \mathbb{R}^m$.

The pseudo-inverse

We are now satisfied that there is a one to one mapping between the row space and the column space. If $\bar{x} \in rowsp(A)$, then A maps it to $colsp(A)$.
The matrix that brings back a vector from the column space to the row space is called the pseudo-inverse of A (and is usually denoted by A^{+}).

We won't talk about it in this course, but know that it can be found using an extremely important concept called the Singular Value Decomposition (SVD).

Additional resources

A summary paper about the four fundamental subspaces, by Gilbert Strang.