

Math S-21b – Lecture #4 Notes

In this lecture we define and study subspaces of \mathbf{R}^n , the span of a collection of vectors, and what it means for a collection of vectors to be linearly independent. In particular, we'll focus on the kernel and image of a linear transformation – subspaces of the domain and codomain, respectively – to motivate the definitions and methods. We'll trim down a redundant spanning set for a given subspace to get a minimal spanning set or **basis for a subspace**, and we prove that the number of spanning vectors in any basis for a given subspace is always the same – the **dimension of the subspace**. We then define the **rank** and **nullity** of a matrix (and of the linear transformation that it represents).

Subspaces of \mathbf{R}^n

Definition: A subspace V of \mathbf{R}^n is a subset that is closed under vector addition and scalar multiplication. That is, for any vectors $\mathbf{v}_1, \mathbf{v}_2 \in V$ and scalars c_1, c_2 , it must be the case that $c_1\mathbf{v}_1 + c_2\mathbf{v}_2 \in V$. This extends to all linear combinations of vectors in the subspace V .

Proposition: The zero vector $\mathbf{0}$ must be in any subspace.

Proof: If $\mathbf{v} \in V$ is any vector, then $0\mathbf{v} = \mathbf{0} \in V$.

Note: Subspaces can be simply visualized as “flat things through the origin”. If $\mathbf{v} \in V$ is any vector, the all scalar multiples of \mathbf{v} must also be in V , i.e. a line passing through the origin. If $\mathbf{v}_1, \mathbf{v}_2 \in V$ are nonparallel vectors, then all linear combinations of the form $c_1\mathbf{v}_1 + c_2\mathbf{v}_2 \in V$, i.e.

a plane through the origin. In higher dimensions, we continue to understand subspaces of \mathbf{R}^n to be lines, planes, and higher-dimensional “flat things through the origin”. We use the term “affine” to refer to parallel objects that do not pass through the origin.

Span of a collection of vectors

Definition: Given a collection of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\} \in \mathbf{R}^n$, we define the span of these vectors to be the set of all linear combinations of these vectors, i.e.

$$\text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\} = \{c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k \text{ where } c_1, \dots, c_k \text{ are scalars}\}.$$

By its definition, the span of any collection of vectors is automatically a subspace. That is, for appropriate scalars,

$$\alpha(c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k) + \beta(d_1\mathbf{v}_1 + \dots + d_k\mathbf{v}_k) = (\alpha c_1 + \beta d_1)\mathbf{v}_1 + \dots + (\alpha c_k + \beta d_k)\mathbf{v}_k.$$

Though we often define subspaces by conditions, we usually specify a subspace by producing a collection of vectors that span the subspace. For example, we describe a line through the origin as the span of a single vector, and a plane through the origin as the span of two nonparallel vectors. It’s important to note, however, that we could also describe a plane as the span of more than two vectors that all lie in that plane. Eliminating such redundancy will motivate the concept of a basis for a subspace.

Image and kernel of a linear transformation

Suppose \mathbf{A} is an $m \times n$ matrix that represents a linear transformation $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ by $T(\mathbf{x}) = \mathbf{Ax}$. Its domain is \mathbf{R}^n and its codomain is \mathbf{R}^m . We define:

$$\text{image}(T) = \text{image}(\mathbf{A}) = \text{im}(\mathbf{A}) = \{ \mathbf{Ax} : \mathbf{x} \in \mathbf{R}^n \} \subset \text{codomain}(\mathbf{A})$$

$$\text{kernel}(T) = \text{kernel}(\mathbf{A}) = \text{ker}(\mathbf{A}) = \{ \mathbf{x} \in \mathbf{R}^n : \mathbf{Ax} = \mathbf{0} \} \subset \text{domain}(\mathbf{A})$$

Proposition: (1) $\text{im}(\mathbf{A})$ is a subspace of the codomain \mathbf{R}^m ; and
(2) $\text{ker}(\mathbf{A})$ is a subspace of the domain \mathbf{R}^n .

Proof: (1) Any two vectors in $\text{im}(\mathbf{A})$ must be of the form $\mathbf{Ax}_1, \mathbf{Ax}_2$ for some vectors $\mathbf{x}_1, \mathbf{x}_2$ in the domain. Therefore $c_1\mathbf{Ax}_1 + c_2\mathbf{Ax}_2 = \mathbf{A}(c_1\mathbf{x}_1 + c_2\mathbf{x}_2) \in \text{im}(\mathbf{A})$ by linearity.

(2) If $\mathbf{x}_1, \mathbf{x}_2 \in \text{ker}(\mathbf{A})$, then $\mathbf{Ax}_1 = \mathbf{0}$ and $\mathbf{Ax}_2 = \mathbf{0}$. So

$$\mathbf{A}(c_1\mathbf{x}_1 + c_2\mathbf{x}_2) = c_1\mathbf{Ax}_1 + c_2\mathbf{Ax}_2 = c_1\mathbf{0} + c_2\mathbf{0} = \mathbf{0}, \text{ so } c_1\mathbf{x}_1 + c_2\mathbf{x}_2 \in \text{ker}(\mathbf{A}).$$

Special case of an $n \times n$ (square) matrix \mathbf{A}

When \mathbf{A} is a square matrix, the image and kernel give us a new way of characterizing when a matrix is invertible.

Proposition: Let \mathbf{A} be an $n \times n$ (square) matrix. Then the following statements are equivalent:

(1) \mathbf{A} is invertible

(4) $\text{rank}(\mathbf{A}) = n$ (full rank)

(2) The system $\mathbf{Ax} = \mathbf{b}$ has a unique solution \mathbf{x}
for all $\mathbf{b} \in \mathbf{R}^n$

(5) $\text{im}(\mathbf{A}) = \mathbf{R}^n$

(3) $\text{rref}(\mathbf{A}) = \mathbf{I}_n$

(6) $\text{ker}(\mathbf{A}) = \{\mathbf{0}\}$ (the zero subspace)

The proof of the equivalence of these statements is left to your observations and knowledge of the definitions.

Calculation of the image and kernel of an $m \times n$ matrix \mathbf{A}

Proposition: The image of any matrix \mathbf{A} is the span of its column vectors.

Proof: If the matrix \mathbf{A} is expressed in terms of its columns as $\mathbf{A} = \begin{bmatrix} \uparrow & & \uparrow \\ \mathbf{v}_1 & \cdots & \mathbf{v}_n \\ \downarrow & & \downarrow \end{bmatrix}$, we know that:

$\mathbf{Ae}_1 = \mathbf{v}_1$, $\mathbf{Ae}_2 = \mathbf{v}_2$, ..., $\mathbf{Ae}_n = \mathbf{v}_n$, so the column vectors are clearly all in the image of \mathbf{A} .

However, any vector \mathbf{x} in the domain can be written as $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} = x_1\mathbf{e}_1 + \cdots + x_n\mathbf{e}_n$, so

$\mathbf{Ax} = \mathbf{A}(x_1\mathbf{e}_1 + \cdots + x_n\mathbf{e}_n) = x_1\mathbf{Ae}_1 + \cdots + x_n\mathbf{Ae}_n = x_1\mathbf{v}_1 + \cdots + x_n\mathbf{v}_n$, i.e. $\text{im}(\mathbf{A}) = \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$.

This is why the image of a matrix \mathbf{A} is also referred to as the “**column space**” of \mathbf{A} .

Note: Though it's true that $\text{im}(\mathbf{A}) = \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, it is not necessarily the case that all of these column vectors are necessary to span the image. There may be some redundancy. Eliminating this redundancy can be accomplished by understanding the kernel of the matrix. Indeed, every vector in the kernel of a matrix will give a linear interdependency of the columns of the matrix.

Example #1: Determine the image and kernel of the matrix $\mathbf{A} = \begin{bmatrix} 2 & 1 & 0 \\ 1 & -1 & 4 \\ 0 & 3 & -8 \end{bmatrix}$ by providing a spanning set of vector for each of these subspaces.

From the proposition above, we know that if $\mathbf{v}_1 = \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix}$, $\mathbf{v}_2 = \begin{bmatrix} 1 \\ -1 \\ 3 \end{bmatrix}$, and $\mathbf{v}_3 = \begin{bmatrix} 0 \\ 4 \\ -8 \end{bmatrix}$ are the columns of \mathbf{A} , then $\text{im}(\mathbf{A}) = \text{span}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$. That's all there is to it.

For the kernel, we solve the homogeneous system $\mathbf{A}\mathbf{x} = \mathbf{0}$. This is most easily done using row reduction:

$$\left[\begin{array}{ccc|c} 2 & 1 & 0 & 0 \\ 1 & -1 & 4 & 0 \\ 0 & 3 & -8 & 0 \end{array} \right] \rightarrow \left[\begin{array}{ccc|c} 1 & -1 & 4 & 0 \\ 2 & 1 & 0 & 0 \\ 0 & 3 & -8 & 0 \end{array} \right] \rightarrow \left[\begin{array}{ccc|c} 1 & -1 & 4 & 0 \\ 0 & 3 & -8 & 0 \\ 0 & 3 & -8 & 0 \end{array} \right] \rightarrow \left[\begin{array}{ccc|c} 1 & -1 & 4 & 0 \\ 0 & 1 & -8/3 & 0 \\ 0 & 0 & 0 & 0 \end{array} \right] \rightarrow \left[\begin{array}{ccc|c} 1 & 0 & 4/3 & 0 \\ 0 & 1 & -8/3 & 0 \\ 0 & 0 & 0 & 0 \end{array} \right]$$

If we let $x_3 = 3t$, we get that $\left\{ \begin{matrix} x_1 = -4t \\ x_2 = 8t \\ x_3 = 3t \end{matrix} \right\}_{t \in \mathbf{R}}$ or $\mathbf{x} = t \begin{bmatrix} -4 \\ 8 \\ 3 \end{bmatrix}$.

This says simply that $\ker(\mathbf{A}) = \text{span} \left\{ \begin{bmatrix} -4 \\ 8 \\ 3 \end{bmatrix} \right\}$, so we have our spanning set.

Note that because $\begin{bmatrix} -4 \\ 8 \\ 3 \end{bmatrix} \in \ker(\mathbf{A})$, this means that

$$\mathbf{A} \begin{bmatrix} -4 \\ 8 \\ 3 \end{bmatrix} = \begin{bmatrix} \uparrow & \uparrow & \uparrow \\ \mathbf{v}_1 & \mathbf{v}_2 & \mathbf{v}_3 \\ \downarrow & \downarrow & \downarrow \end{bmatrix} \begin{bmatrix} -4 \\ 8 \\ 3 \end{bmatrix} = -4\mathbf{v}_1 + 8\mathbf{v}_2 + 3\mathbf{v}_3 = \mathbf{0}, \text{ so there's a } \mathbf{linear\ dependency} \text{ among}$$

these vectors. In particular, this means that we can solve for any one of these vectors in terms of the remaining vectors. For example, $\mathbf{v}_3 = \frac{4}{3}\mathbf{v}_1 - \frac{8}{3}\mathbf{v}_2$. This means that $\mathbf{v}_3 \in \text{span}\{\mathbf{v}_1, \mathbf{v}_2\}$, so we don't need it in our spanning set for the image. We could, of course, have eliminated any one of these three vectors in this manner, but it's good standard practice to solve for the later vectors in terms of its predecessors. There are no other linear interdependencies that can be used to eliminate redundancy, so this is the best we can do. That is, $\boxed{\text{im}(\mathbf{A}) = \text{span}\{\mathbf{v}_1, \mathbf{v}_2\}}$.

Note, in particular, that we retained in our spanning set for the image exactly those column vectors of the original matrix \mathbf{A} that eventually yielded Leading 1's in the reduced row-echelon form of this matrix. We can always choose our spanning set for the image such that this is the case. In order to understand this better, we need one more very important definition.

Definition: A set of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\} \in \mathbf{R}^n$ is called **linearly independent** if given any linear combination of the form $c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k = \mathbf{0}$, this implies that $c_1 = \dots = c_k = 0$. That is, there is no nontrivial way to combine these vectors to yield the zero vector.

Note that this definition also means that it's impossible to solve for any one of these vectors in terms of the others. This is the essential quality of linear independence – there is no redundancy among a linearly independent set of vectors.

Test for linear independence

Given a collection of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\} \in \mathbf{R}^n$, if we write $\mathbf{A} = \begin{bmatrix} \uparrow & & \uparrow \\ \mathbf{v}_1 & \cdots & \mathbf{v}_k \\ \downarrow & & \downarrow \end{bmatrix}$, then the

expression $c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k = \mathbf{0}$ can be expressed as $\mathbf{Ac} = \mathbf{0}$ where $\mathbf{c} = \begin{bmatrix} c_1 \\ \vdots \\ c_k \end{bmatrix}$.

So the statement that $c_1\mathbf{v}_1 + \cdots + c_k\mathbf{v}_k = \mathbf{0}$ implies that $c_1 = \cdots = c_k = 0$ can be restated very succinctly as $\mathbf{A}\mathbf{c} = \mathbf{0}$ implies $\mathbf{c} = \mathbf{0}$. That is, a collection of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ will be linearly independent if and only if $\ker(\mathbf{A}) = \{\mathbf{0}\}$.

In our example above, we found that for $\mathbf{A} = \begin{bmatrix} \uparrow & \uparrow & \uparrow \\ \mathbf{v}_1 & \mathbf{v}_2 & \mathbf{v}_3 \\ \downarrow & \downarrow & \downarrow \end{bmatrix}$, $\ker(\mathbf{A}) = \text{span} \left\{ \begin{bmatrix} -4 \\ 8 \\ 3 \end{bmatrix} \right\}$. Therefore these column vectors were not linearly independent.

Definition: Given a subspace V of \mathbf{R}^n , a collection of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\} \in V$ is called a **basis** of V if $\text{Span}\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\} = V$ and $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ are linearly independent.

A basis is a minimal spanning set, and it's important to note that any given subspace can have many different bases.

Note: \mathbf{R}^n is itself a subspace of \mathbf{R}^n (the whole space). The standard basis $\mathcal{E} = \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$ is familiar to us, but, in fact, any linearly independent collection of n vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ in \mathbf{R}^n would provide an alternate basis for \mathbf{R}^n .

Example #2: Determine bases for the image and kernel of the matrix $\mathbf{A} = \begin{bmatrix} 1 & -1 & -1 & 1 & 1 \\ -1 & 1 & 0 & -2 & 2 \\ 1 & -1 & -2 & 0 & 3 \\ 2 & -2 & -1 & 3 & 4 \end{bmatrix}$.

We begin by noting that $\text{im}(\mathbf{A}) = \text{span}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4, \mathbf{v}_5\}$ where $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4, \mathbf{v}_5\}$ are the columns vectors of \mathbf{A} .

The kernel of \mathbf{A} is found by row reduction:
$$\left[\begin{array}{ccccc|c} 1 & -1 & -1 & 1 & 1 & 0 \\ -1 & 1 & 0 & -2 & 2 & 0 \\ 1 & -1 & -2 & 0 & 3 & 0 \\ 2 & -2 & -1 & 3 & 4 & 0 \end{array} \right] \rightarrow \left[\begin{array}{ccccc|c} \boxed{1} & -1 & 0 & 2 & 0 & 0 \\ 0 & 0 & \boxed{1} & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & \boxed{1} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right]$$

where the Leading 1's are highlighted. We must introduce two parameters s, t to describe all solutions, namely:

$$\left. \begin{cases} x_1 = s - 2t \\ x_2 = s \\ x_3 = -t \\ x_4 = t \\ x_5 = 0 \end{cases} \right\}_{s,t \in \mathbf{R}} \Rightarrow \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} = s \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + t \begin{bmatrix} -2 \\ 0 \\ -1 \\ 1 \\ 0 \end{bmatrix} \Rightarrow \ker(\mathbf{A}) = \text{Span} \left\{ \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -2 \\ 0 \\ -1 \\ 1 \\ 0 \end{bmatrix} \right\} \text{ and these form}$$

a basis for the kernel. These two vectors also give that $\mathbf{v}_1 + \mathbf{v}_2 = \mathbf{0}$ and $-2\mathbf{v}_1 - \mathbf{v}_3 + \mathbf{v}_4 = \mathbf{0}$, so we can solve for $\mathbf{v}_2 = -\mathbf{v}_1$ and $\mathbf{v}_4 = 2\mathbf{v}_1 + \mathbf{v}_3$. Therefore $\text{im}(\mathbf{A}) = \text{span}\{\mathbf{v}_1, \mathbf{v}_3, \mathbf{v}_5\}$ and these form a basis for the image.

We can easily see that these vectors are now linearly independent because if we create a new matrix with just these vectors as its columns, our test for linear independence gives

$$\left[\begin{array}{ccc|c} 1 & -1 & 1 & 0 \\ -1 & 0 & 2 & 0 \\ 1 & -2 & 3 & 0 \\ 2 & -1 & 4 & 0 \end{array} \right] \rightarrow \left[\begin{array}{ccc|c} \boxed{1} & 0 & 0 & 0 \\ 0 & \boxed{1} & 0 & 0 \\ 0 & 0 & \boxed{1} & 0 \\ 0 & 0 & 0 & 0 \end{array} \right].$$

This follows exactly the same steps as the original matrix, only now the Leading 1's occupy all columns (except the column corresponding to the right-hand-sides of the corresponding homogeneous linear equations). This example illustrates how *the rank of a matrix coincides with the minimum number of vectors necessary to span the image of the matrix.*

Dimension of a subspace

Intuitively, we would expect to have just one vector to form a basis for a line through the origin, and two vectors to form a basis for a plane through the origin. This is the subject of an important theorem:

Theorem: Given any two bases for a subspace V of \mathbf{R}^n , the number of vectors in both bases must be the same. This uniquely determined integer is called the **dimension** of V (written $\dim(V)$).

Proof: Suppose the subspace $V \subseteq \mathbf{R}^n$ has two bases $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p\}$ and $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_q\}$. We need to show that $p = q$. Start by noting that because $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_q\}$ is a basis, it spans V , so we can express each of the vectors in the first basis in terms of these vectors in the second basis,

$$\text{i.e. } \left\{ \begin{array}{l} \mathbf{v}_1 = c_{11}\mathbf{w}_1 + c_{12}\mathbf{w}_2 + \dots + c_{1q}\mathbf{w}_q \\ \vdots \\ \mathbf{v}_p = c_{p1}\mathbf{w}_1 + c_{p2}\mathbf{w}_2 + \dots + c_{pq}\mathbf{w}_q \end{array} \right\}.$$

If we assemble each of these vectors as the columns of a matrix and use our definition of the product of a matrix and a vector, we can write these p vector equations as a single matrix equation:

$$\underbrace{\begin{bmatrix} \uparrow & & \uparrow \\ \mathbf{w}_1 & \cdots & \mathbf{w}_q \\ \downarrow & & \downarrow \end{bmatrix}}_{\mathbf{M}} \underbrace{\begin{bmatrix} c_{11} & c_{21} & \cdots & c_{p1} \\ \vdots & \vdots & \cdots & \vdots \\ c_{1q} & c_{2q} & \cdots & c_{pq} \end{bmatrix}}_{\mathbf{A}} = \underbrace{\begin{bmatrix} \uparrow & & \uparrow \\ \mathbf{v}_1 & \cdots & \mathbf{v}_p \\ \downarrow & & \downarrow \end{bmatrix}}_{\mathbf{N}}$$

Note that \mathbf{M} is an $n \times q$ matrix, \mathbf{A} is an $q \times p$ matrix, and \mathbf{N} is an $n \times p$ matrix. We can think of each of these matrices as representing linear transformations, $\mathbf{R}^p \xrightarrow{\mathbf{A}} \mathbf{R}^q \xrightarrow{\mathbf{M}} \mathbf{R}^n$ and the composition $\mathbf{R}^p \xrightarrow{\mathbf{N}} \mathbf{R}^n$ with $\mathbf{MA} = \mathbf{N}$.

It's easy to see that $\ker(\mathbf{A}) \subseteq \ker(\mathbf{N})$. If $\mathbf{x} \in \ker(\mathbf{A})$, then $\mathbf{Ax} = \mathbf{0}$. So

$\mathbf{Nx} = \mathbf{MAx} = \mathbf{M}(\mathbf{Ax}) = \mathbf{M}(\mathbf{0}) = \mathbf{0}$. Therefore $\mathbf{x} \in \ker(\mathbf{N})$. However, because $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p\}$ is a basis and these vectors form the columns of \mathbf{N} , this matrix will have linearly independent columns. But this is the same as saying that $\ker(\mathbf{N}) = \{\mathbf{0}\}$. Therefore $\ker(\mathbf{A}) = \{\mathbf{0}\}$ as well. But this then tells us that the columns of \mathbf{A} are linearly independent. So $\text{rank}(\mathbf{A}) = p$. If we interpret this in terms of the number of rows (q) and columns (p) of the matrix \mathbf{A} and use the fact that the rank of a matrix can never be greater than the number of rows or the number of columns of the matrix, we get that $\text{rank}(\mathbf{A}) = p \leq \min\{p, q\}$. But this clearly implies that $p \leq q$.

If we were to now repeat this entire argument with the roles of the two bases interchanged, we would similarly conclude that $q \leq p$. Therefore $p = q$.

Now that the idea of dimension has meaning in the context of linear subspaces, a few definitions are in order:

Definition: Given an $m \times n$ matrix \mathbf{A} , we define:

$$\boxed{\text{rank}(\mathbf{A}) = \dim(\text{im } \mathbf{A})} \quad \text{and} \quad \boxed{\text{nullity}(\mathbf{A}) = \dim(\ker \mathbf{A})}.$$

Note: This new definition for the rank of a matrix does not contradict our previous definition in terms of the number of Leading 1's in the reduced row-echelon form of the matrix. As we saw in the example, a basis for the image of any matrix can be formed using only the columns of the matrix that ultimately produced Leading 1's in its reduced row-echelon form, and the number of such basis vectors will then be the rank of the matrix.

We might also note that the number of parameters necessary to describe all vectors in the kernel corresponds to those columns that do not yield Leading 1's. This is also the number of vectors in a basis for the kernel (based on previous observations), so this corresponds to the nullity of the matrix. Any $m \times n$ matrix \mathbf{A} has n columns (corresponding to its domain being \mathbf{R}^n). So if we combine all of these observations we get the following:

Rank-Nullity Theorem: Given any $m \times n$ matrix \mathbf{A} ,

$$\boxed{\text{rank}(\mathbf{A}) + \text{nullity}(\mathbf{A}) = n = \dim(\text{domain}(\mathbf{A}))}.$$

We can also state this as $\dim(\text{im}(\mathbf{A})) + \dim(\text{ker}(\mathbf{A})) = n = \dim(\text{domain}(\mathbf{A}))$.

Though all of these facts apply for any matrix, the special case of an $n \times n$ (square) matrix \mathbf{A} allows us to state the following proposition (proof left as an exercise) summarizing what it means for a matrix to be invertible.

Proposition: Let \mathbf{A} be an $n \times n$ (square) matrix. Then the following statements are equivalent:

- (1) \mathbf{A} is invertible
- (2) The system $\mathbf{Ax} = \mathbf{b}$ has a unique solution \mathbf{x} for all $\mathbf{b} \in \mathbf{R}^n$
- (3) $\text{rref}(\mathbf{A}) = \mathbf{I}_n$
- (4) $\text{rank}(\mathbf{A}) = n$ (full rank)
- (5) $\text{im}(\mathbf{A}) = \mathbf{R}^n$
- (6) $\text{ker}(\mathbf{A}) = \{\mathbf{0}\}$ (the zero subspace)
- (7) The columns of \mathbf{A} are linearly independent.
- (8) The columns of \mathbf{A} span \mathbf{R}^n .
- (9) The columns of \mathbf{A} form a basis for \mathbf{R}^n .

Notes by Robert Winters