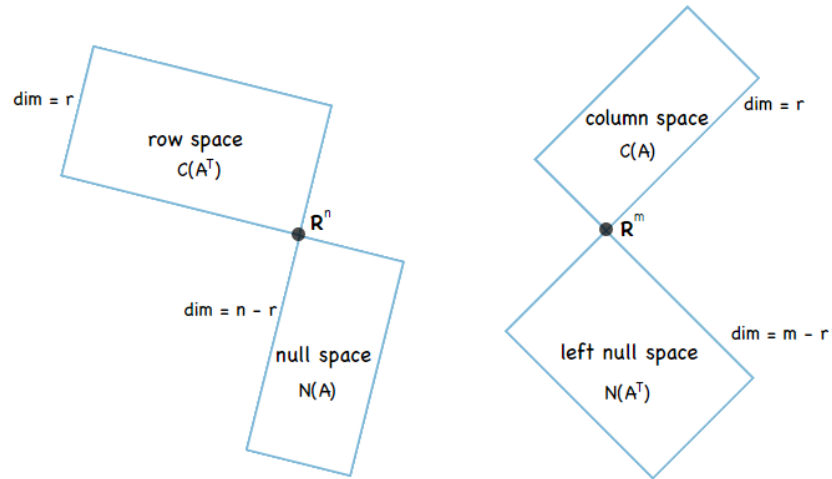# Four fundamental subspaces

https://www.youtube.com/watch?v=uQhTuRlWMxw

Inverse of a matrix $\mathbf{A}$: $\mathbf{A}^{-1}$
Applying the matrix and its inverse in succession restore the original vector, so $\mathbf{A}^{-1}\mathbf{A} = \mathbf{I}$
There is no inverse if $\det(\mathbf{A}) = 0$

$$\mathbf{A}\mathbf{v} = \mathbf{x}$$
$$\mathbf{A}^{-1}\mathbf{A}\mathbf{x} = \mathbf{A}^{-1}\mathbf{v}$$
$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{v}$$

1. Column space $C(\mathbf{A})$ is all possible linear combinations of column vectors, zero vector is always in the column space. It is the span of the columns of the matrix.

$C(\mathbf{A}) = \text{span}\{\mathbf{v_1}, \mathbf{v_2}, \ldots \mathbf{v_n}\} = \{x_1\mathbf{v_1} + x_2\mathbf{v_2} + \ldots + x_n\mathbf{v_n} \mid x_1, x_2, \ldots, x_n \in R, \mathbf{v_i} \in R^n\}$
If $\mathbf{v_1}\ \mathbf{v_2} \ldots \mathbf{v_n}$ are linearly independent set, they are valid basis of column space

Rank r = Dim($C(\mathbf{A})$): number of linearly independent column vectors, it equals to the dimensions in the output

2. Null space $N(\mathbf{A}) = \{\mathbf{x} \in R^n \mid \mathbf{A}\mathbf{x} = \mathbf{0}\}$, also termed 'Kernel', gives all of the possible solutions for the equation $\mathbf{A}\mathbf{x} = \mathbf{0}$

$N(\mathbf{A}) = N(\text{rref}(\mathbf{A}))$

Nullity = Dim($N(\mathbf{A})$): Number of free variables in rref(A)

Column vectors of A are linearly independent $\Leftrightarrow$ $N(A) = \{0\}$
Because $x_1v_1 + x_2v_2 + \ldots + x_nv_n = 0$, $x_1, x_2, \ldots, x_n = 0$

3. <u>Row space $R(A)$ or $C(A^T)$</u> is all possible linear combinations of row vectors, zero vector is always in the row space; equal to column space of A^T

4. <u>Left null space $N(A^T)$</u> is the null space of A^T

The big picture of linear algebra

MIT lecture by Prof. Gilbert Strang on the <u>orthogonality</u> and <u>dimensionality</u> of the four fundamental subspaces:
https://www.youtube.com/watch?v=ggWYkes-n6E

Let **A** be a matrix with m rows and n columns, then there are 4 fundamental subspaces:

- Row space $R(A) = C(A^T)$, it contains all linear combinations of the rows of **A**, or columns of A^T
- Column space $C(A)$, it contains all linear combinations of the columns of **A**
- Null space $N(A)$, it contains all solutions to the system $Ax = 0$
- Left null space $N(A^T)$, it contains all solutions to the system $A^Ty = 0$

- The null space $N(A)$, is perpendicular to/orthogonal complement of its row space $R(A)$
- The left null space $N(A^T)$, is perpendicular to/orthogonal complement of its column space $C(A)$

- $r = \text{rank}(A) = \dim(C(A)) = \dim(R(A))$
- $\dim(N(A)) = n - r$
- $\dim(N(A^T)) = m - r$
- $\text{rank}(A^T) + \text{nullity}(A^T) = m$
- $\text{rank}(A) + \text{nullity}(A) = n$

Change of basis

Let V be a vector space and let $B = \{ \mathbf{v}_1, \mathbf{v}_2, \ldots , \mathbf{v}_n \}$ be a set of vectors in V.

Recall that B forms a <u>basis</u> for V if the following two conditions hold:
1. $\mathbf{v}_1, \mathbf{v}_2, \ldots , \mathbf{v}_n$ are linearly independent
2. B spans V

If $B = \{ \mathbf{v}_1, \mathbf{v}_2, \ldots , \mathbf{v}_n \}$ is a basis for V, then every $\mathbf{v} \in V$ can be expressed uniquely as a linear combination of $\mathbf{v}_1, \mathbf{v}_2, \ldots , \mathbf{v}_n$:
$$\mathbf{v} = c_1 \mathbf{v_1} + c_2 \mathbf{v_2} + \ldots + c_n \mathbf{v_n}$$

$c_1, c_2, \ldots c_n$ are just the <u>coordinates</u> of $\mathbf{v}$ relative to basis B. If V has a dimension of n, then every set of n linearly independent vectors in V forms a basis for V.

We have a nice coordinate representation of our vector $\mathbf{v}$, which we can construct by projecting our vector onto each basis vector individually. Now, what happens if we don't find the standard basis particularly convenient for our problem, and we would like to look at $\mathbf{v}$ from a different perspective?
We have a choice as to what basis to use!
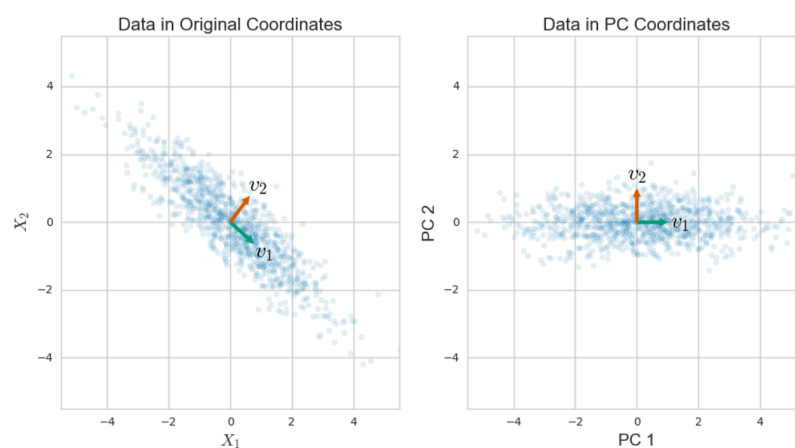
Here we will focus on vectors in R^2, although all of this generalizes to R^n. The standard basis in R^2 is $\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \}$

Let us specify other bases with reference to this rectangular coordinate system

Let $\mathbf{B} = \{\mathbf{x}, \mathbf{y}\}$ and $\mathbf{B'} = \{\mathbf{x'}, \mathbf{y'}\}$ be two bases for R^2. For a vector $\mathbf{v} \in V$, given its coordinates $[\mathbf{v}]_B$ in basis B, we would like to be able to express $\mathbf{v}$ in terms of its coordinates $[\mathbf{v}]_{B'}$ in basis $\mathbf{B'}$, and vice versa.

Suppose the basis vectors $\mathbf{x'}$ and $\mathbf{y'}$ for $\mathbf{B'}$ have the following coordinates relative to the basis $\mathbf{B}$: $[\mathbf{x'}]_B = \begin{bmatrix} a \\ b \end{bmatrix}$ $[\mathbf{y'}]_B = \begin{bmatrix} c \\ d \end{bmatrix}$

The change of coordinate matrix from $\mathbf{B'}$ to $\mathbf{B}$: $\mathbf{C} = \begin{bmatrix} a & c \\ b & d \end{bmatrix}$, meaning

$\mathbf{x} = a\mathbf{x'} + b\mathbf{y'}$
$\mathbf{y} = c\mathbf{x'} + d\mathbf{y'}$

Matrix $\mathbf{C}$ governs the change of coordinates of $\mathbf{v} \in V$:

$$[\mathbf{v}]_B = \mathbf{C}[\mathbf{v}]_{B'}$$
$$[\mathbf{v}]_{B'} = \mathbf{C}^{-1}[\mathbf{v}]_B$$

Suppose we know the coordinates $[\mathbf{v}]_{B'}$ in the new basis $\mathbf{B'}$, and we can find the coordinate of $[\mathbf{v}]_B$ in the old basis $\mathbf{B}$. The underlying mapping remains the same when matrix $\mathbf{M}$ is used in the original basis and the matrix $\mathbf{M'}$ is used in the new basis:

$$\mathbf{M'} = \mathbf{C}^{-1}\mathbf{MC}$$
$$\mathbf{M} = \mathbf{CM'C}^{-1}$$

The big take away: there are arbitrary different bases we can use to represent R^n, so we can have different matrices to represent the same linear transformation under different coordinate systems!

Exercise1: We have a vector $\mathbf{d} = \begin{bmatrix} 8 \\ -6 \\ 2 \end{bmatrix}$, with the change of basis matrix from standard coordinates: $\begin{bmatrix} 1 & 1 \\ 2 & 0 \\ 3 & 1 \end{bmatrix}$, find its new coordinates in the new basis.

Hint: equivalent to solving a linear system of equations: $\begin{bmatrix} 1 & 1 \\ 2 & 0 \\ 3 & 1 \end{bmatrix}\begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} 8 \\ -6 \\ 2 \end{bmatrix}$

a = -3, b = 11

Exercise2: Let us alternate R^2 basis to $\mathbf{B} = \{\begin{pmatrix}1\\2\end{pmatrix} \begin{pmatrix}2\\1\end{pmatrix}\}$, what is the representation of the transformation matrix $\mathbf{M} = \begin{bmatrix}3 & -2\\2 & -2\end{bmatrix}$ in the new basis?

The change of basis matrix for $\mathbf{B}$ would be $\mathbf{C} = \begin{bmatrix}1 & 2\\2 & 1\end{bmatrix}$

$\det(\mathbf{C})$ = -3

$\mathbf{C}^{-1}$ = -1/3$\begin{bmatrix}1 & -2\\-2 & 1\end{bmatrix}$

$\mathbf{M}' = \mathbf{C}^{-1}\mathbf{M}\mathbf{C} = \begin{bmatrix}-1 & 0\\0 & 2\end{bmatrix}$ It is a diagonal matrix!

When we transfer to a new basis, all of a sudden the transformation matrix is much simpler: we are only taking scaling factors of the corresponding terms during matrix multiplication. This is a neat result! It is super easy to multiply, to invert, to calculate the determinant, etc.

Linear algebra is the art of choosing the right basis!

Later we will learn eigenvectors and eigenvalues, we will see that eigenvectors make for good bases vectors/coordinate systems.

Eigenvectors and Eigenvalues

https://www.youtube.com/watch?v=PFDu9oVAE-g

In general, vectors will change direction as well as length when multiplied by a matrix. However, there will be vectors that might change length, but not direction. In other words, for these vectors, multiplication by a matrix is no different than multiplication by a simple scalar: $\mathbf{Av} = \lambda\mathbf{v}$

$\mathbf{A}$ therefore acts by stretching $\mathbf{v}$ (by a scalar factor of λ), but not changing its direction. We say that $\mathbf{v}$ is an eigenvector, and λ is the corresponding eigenvalue which determines how much $\mathbf{v}$ is shortened or lengthened by a linear transformation.

Example: the shear transformation we saw last time, $\mathbf{A} = \begin{bmatrix}1 & 1\\0 & 1\end{bmatrix}$, $\lambda = 1$

L.T. knocks off vectors from its original span, i.e. rotate them, but eigenvectors remain on their own span after transformation, the corresponding eigenvalue is the factor it gets stretched or squished after L.T.

*Any vector in eigenvector direction is still an eigenvector with the same corresponding eigenvalue, though we usually consider eigenvectors as unit vectors

Q: What does it mean when an eigenvalue is negative?
The vector gets flipped after L.T., but stays on the same line it spans out

Intuitive example: Consider spinning a globe, every location faces a new direction, except the north and south pole.

3D rotation of a tea pot, if we can find an eigenvector who remains on its own span, what does it mean? We find the axis of rotation (green vector)!

An example of calculating eigenvalues and eigenvectors of a 2x2 matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$

Exercise: look for eigenvalues and eigenvectors of matrix: $\mathbf{B} = \begin{bmatrix} 1 & 2 \\ 4 & 3 \end{bmatrix}$

1. Left matrix multiplication, right scalar multiplication
2. Rewrite the right hand side by some matrix multiplication:

$\mathbf{Av} = (\lambda\mathbf{I})\mathbf{v}$
$\mathbf{Av} - (\lambda\mathbf{I})\mathbf{v} = \mathbf{0}$
$(\mathbf{A} - \lambda\mathbf{I})\mathbf{v} = \mathbf{0}$

We want a non-zero eigenvector, a product of a matrix with a non-zero vector equals to zero, the transformation therefore squishes the space into a lower dimension, corresponding to a determinant of zero:

$\det(\mathbf{A} - \lambda\mathbf{I}) = 0$
$(a-\lambda)(d-\lambda) - bc = 0$
$(3-\lambda)(1-\lambda) - 8 = 0$

$\lambda_1 = 5, \lambda_2 = -1$

$\mathbf{v}_1 = \begin{bmatrix} 0.5 \\ 1 \end{bmatrix}, \mathbf{v}_2 = \begin{bmatrix} -1 \\ 1 \end{bmatrix}$

$\text{Eigenspace}_5 = \text{span}\{\begin{bmatrix} 0.5 \\ 1 \end{bmatrix}\}, \text{Eigenspace}_{-1} = \text{span}\{\begin{bmatrix} -1 \\ 1 \end{bmatrix}\}$

Eigenvectors and eigenvalues pop up in many areas such as:
Image compression: using SVD to find eigenvectors of the covariance matrix, and project data onto the eigenvectors of the largest eigenvalues, thereby throwing away the small eigenvalues.

Dimensionality reduction for machine learning and data analysis: using PCA to find the principal components that correspond to the largest eigenvalues of the covariance matrix. It allows us to understand where most of the variation in a data comes from.

Carsen will lecture on PCA, SVD in details in the following week.

Related to linear regression week:

Linear algebra view of least-squares regression

The goal of regression is to fit a mathematical model, i.e. a linear equation, to a set of observed points. We believe there is an underlying relationship that maps a to b, in the form of b = ca + d. Here, c and d are the regression coefficients we are looking for.

We can write the observed points as a simple linear system such as:
$b_1 = ca_1 + d$
$b_2 = ca_2 + d$
$b_3 = ca_3 + d$

We have a simple linear system and we only need to deal with vectors and matrices!
Here is our linear system in the matrix form: $\mathbf{Ax} = \mathbf{b}$

$$\begin{bmatrix} a1 & 1 \\ a2 & 1 \\ a3 & 1 \end{bmatrix} \begin{bmatrix} c \\ d \end{bmatrix} = \begin{bmatrix} b1 \\ b2 \\ b3 \end{bmatrix}$$

What this is saying is that we hope the vector **b** lies in the column space of **A**, i.e. C(**A**). In other word, we wish to find a linear combination of the columns of **A** that gives us our vector **b**.
But we know that most of the time **b** doesn't fit our model perfectly, meaning that it is outside the column space of **A**. We cannot simply solve the equation $\mathbf{Ax} = \mathbf{b}$ for vector **x**.

The linear regression answer is that we intend to swap out for **b** for another vector that is very close to it but fits in our model. Specifically, we want to pick a vector **p** that is in the column space of **A**, but is also as close as possible to **b**.

This is termed <u>least squares approximation</u> in linear algebra. We would like to find **p**, that could minimize the difference between observed **b** and the projected **p**. Here, **p** is the projection of **b** in C(**A**).

Geometry makes it pretty clear what's going on. We started with **b**, which doesn't fit the model, and then switched to **p**, which gives a good approximation and has the virtue of sitting in C(**A**).
$\mathbf{p} = \mathbf{Ax}^*$

Now solving the regression coefficients becomes solving for $\mathbf{x}^*$, which is the estimated regression coefficients c and d.

$\mathbf{x}^* = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{b}$