The **minimal polynomial** $m_A(x)$ is the **monic polynomial of smallest degree** such that

$$m_A(A) = 0.$$

- "Monic" = leading coefficient is 1
- "Smallest degree" = no lower-degree polynomial annihilates A

Why does it exist?

Because:

- The vector space of all $n \times n$ matrices has dimension n^2
- The set $\{I, A, A^2, \dots, A^{n^2}\}$ is linearly dependent
- So there exists a nonzero polynomial $p(x)$ with $p(A) = 0$

Among all such polynomials, choose the monic one of least degree → **minimal polynomial**

2. Minimal polynomial vs characteristic polynomial

Let:

- $\chi_A(x) = \det(xI - A)$ (characteristic polynomial)
- $m_A(x)$ = minimal polynomial

Fundamental relationships

1. **Every eigenvalue of A is a root of $m_A(x)$**
2. $m_A(x)$ **divides** $\chi_A(x)$
3. They have **the same distinct roots**, but possibly different powers

Example:

$$\chi_A(x) = (x - 2)^3(x + 1)^2$$

Possible minimal polynomials:

$$(x-2)(x+1),\ (x-2)^2(x+1),\ (x-2)^3(x+1)^2$$

3. How the minimal polynomial detects Jordan blocks

This is the key insight.

Theorem

> The exponent of $(x - \lambda)$ in $m_A(x)$ equals the size of the **largest Jordan block** for eigenvalue λ.

So:

- If $m_A(x)$ has **no repeated roots**, all Jordan blocks are size 1 → diagonalizable
- Repeated roots ⇒ nontrivial Jordan blocks

4. Diagonalizability criterion (central theorem)

Theorem (Minimal Polynomial Criterion)

A matrix A is diagonalizable over $\mathbb{F}$ **iff**

$$m_A(x) = \prod_{i=1}^{k}(x - \lambda_i) \quad \text{(all roots simple)}$$

Why?

- Each eigenvalue has only 1×1 Jordan blocks
- Algebraic multiplicity = geometric multiplicity

1. What is a Jordan block?

For eigenvalue λ, a Jordan block of size k is:

$$J_k(\lambda) = \begin{pmatrix} \lambda & 1 & 0 & \cdots & 0 \\ 0 & \lambda & 1 & \cdots & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \lambda & 1 \\ 0 & \cdots & 0 & 0 & \lambda \end{pmatrix}$$

- One eigenvector
- $k - 1$ **generalized eigenvectors**

2. Algebraic vs geometric multiplicity

For eigenvalue λ:

- **Algebraic multiplicity (AM)**
  = total size of *all* Jordan blocks for λ
- **Geometric multiplicity (GM)**
  = number of Jordan blocks for λ

Thus:

$$\boxed{\text{AM} = \sum \text{(sizes of Jordan blocks)}}$$

$$\boxed{\text{GM} = \text{number of Jordan blocks}}$$

3. What does the block *size* mean?

Key fact:

> A Jordan block of size k contributes exactly **1 eigenvector**, but counts as k algebraic multiplicity.

So:

- If all blocks are size 1 ⇒ AM = GM ⇒ diagonalizable
- If a block has size $k > 1$ ⇒ GM < AM

4. Is the block size "the difference" between AM and GM?

Not exactly — but it explains the difference.

For one eigenvalue:

$$\text{AM} - \text{GM} = \sum(\text{block size} - 1)$$

So:

- Each block of size k contributes $k - 1$ to the deficit
- The *largest* block determines how "bad" the non-diagonalizability is

5. Connection to minimal polynomial

For eigenvalue λ:

- Largest Jordan block size = power of $(x - \lambda)$ in minimal polynomial

Example:

$$m_A(x) = (x - \lambda)^3 \Rightarrow \text{largest block is size 3}$$

6. Geometric meaning (important intuition)

Jordan blocks tell you how many times you must apply:

$$(A - \lambda I)$$

before a generalized eigenvector becomes an actual eigenvector.