



AI: under the hood

5 notions of how it works, their practical consequences — and 3 experiments to verify them.

THE 5 NOTIONS

1

A trained generator, not a database

Billions of dials tuned during training: it reconstructs, it doesn't copy — hence facts restated wrong.

2

It reads tokens, not words

Text chopped into fragments: letter-counting fails, rare terms derail. Stuck? Rephrase with common words.

3

Data sets the ceiling — and the biases

Abundant = mastered, rare = hallucinations (same confidence). The “standard” answer reflects the data's majority, not a neutral truth.

4

The context window explains the “forgetting”

The whole conversation is reread every turn, within a limit. One topic = one conversation; restate the essential context.

5

A probability machine

Deliberate randomness: same question, different answers. Variance signals uncertainty — where it varies, you verify.

3 experiments to see it yourself — no code

1. Long conversation: “remind me of my first question” — the window at work.
2. Same question ×3 then rephrased: the stable is solid, the variable gets verified.
3. Mainstream topic vs your niche: the confidence of the tone doesn't change — verify both.

Understanding the mechanism means predicting the errors.

The levers that follow

Context is your #1 lever: formulation, section-by-section work, fresh conversations.

Rephrasing and re-asking is a legitimate strategy — not a whim.

The basic reflexes (verification, confidentiality): sheet 09 on domaformalis.com