HOW DOES AI ACTUALLY PRODUCE AN ANSWER?

From your words to tokens, patterns, predictions, and a response — with you in charge of checking it.



You do NOT need an AI account

Paper and pencil are enough for every core challenge.

Useful for any experience level

Explorer • Builder • Engineer / Researcher

Public, self-guided lesson • 15–20 min core • activities + optional deeper path

How to use this public lesson

Choose your depth. Complete the lesson alone, in a group, or while watching the replay.

EXPLORER

Understand the big idea.
Use familiar examples.
Explain it in your own words.

BUILDER

Try the paper lab.
Write a better prompt.
Design a simple test.

ENGINEER / RESEARCHER

Inspect probabilities, sampling, context, evaluation, and reliability.

STOP / DRAW / TEST whenever you see a challenge. A correct question is progress.

Six weeks, one connected story

Every week increases the resolution of an everyday experience.

01	Technology is a system, not a screen
02	One tap: sensor → app → response
03	Why technology feels fast, slow, or unstable
04	What is inside a phone or computer
05	How information travels across networks
06	**How AI produces an answer**

A normal moment: ask a helpful question

“Explain why a rainbow appears, in language a ten-year-old could understand.”

WHAT YOU SEE

You type a question.
A moment later, the screen shows a fluent explanation.

WHAT WE WILL INVESTIGATE

How did the system process your text?
Why did those words appear?
How would you know if they were right?

Guiding question: does a well-written answer automatically mean a true answer?

AI is a broad family — not one single tool

This lesson focuses on generative language models, which produce text.

AI

Broad field: systems that recognize patterns, predict, plan, classify, or generate.

MACHINE LEARNING

Many AI systems learn statistical patterns from data instead of using only hand-written rules.

GENERATIVE AI

Produces content such as text, images, audio, or code. Language models are one example.

A spam filter and a chatbot may both involve AI, but they solve different kinds of problems.

Zoom in: a generative language model

A language model estimates what tokens are likely to follow the tokens in its context.

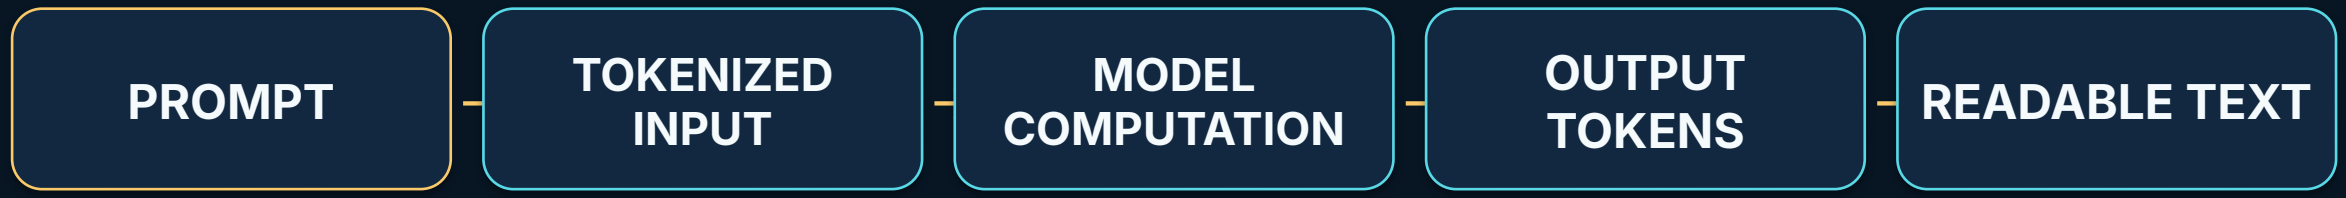

A powerful distinction

The model can generate plausible explanations without independently checking every factual claim. Some AI products add search, calculators, or other tools.

Step 1: the words become tokens

Tokens can be words, parts of words, punctuation, or other pieces of text.

“A sunflower grows toward the light.”

A	sun	flower	grows	toward	the	light.

ILLUSTRATION ONLY: real token boundaries depend on the tokenizer and may differ.

Try predicting the next word yourself

Think like a simple language predictor for one sentence.

“The ice cream melted because the day was ____.”

A • HOT

Would this fit the sentence?

B • BLUE

Would this fit the sentence?

C • SILENT

Would this fit the sentence?

Pause. Choose a continuation and explain what clues in the sentence you used.

Context changes the continuation

A useful prediction depends on what came before.

SHORT CONTEXT

“The student brought a ____ to class.”

Many objects could fit.

MORE CONTEXT

“The student brought a ____ to class because today they measured angles.”

A protractor becomes more plausible.

What extra information changed your guess? How might misleading context change it?

Not one guaranteed answer: a distribution

Illustrative toy scores, NOT measurements from a real model.

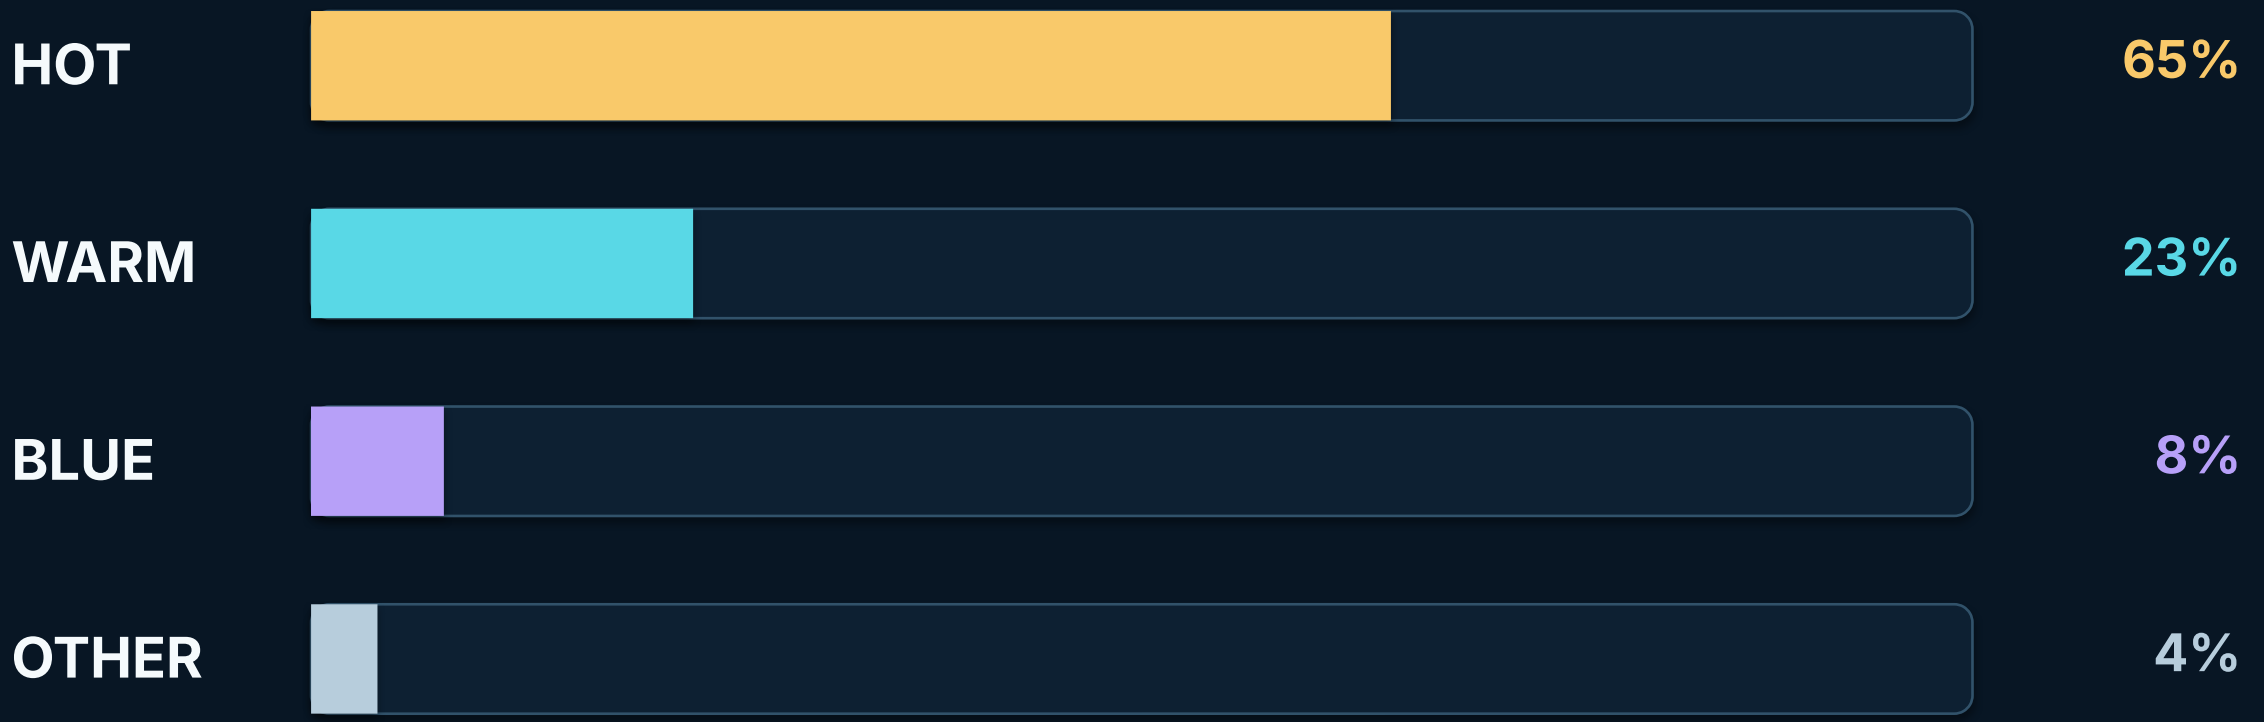

A system could select the highest-scored option or sample among options; settings and model design matter.

Training and answering are different activities

During training, parameters are updated; during ordinary inference, the model uses learned parameters.

TRAINING • LEARNING PATTERNS

Examples / data → model predicts → error or feedback is evaluated → parameters adjusted (many repeats).

INFERENCE • PRODUCING AN ANSWER

Your prompt → model computation using learned parameters → generated tokens → displayed response.

A normal user prompt does not mean the base model retrains itself on the spot.

A simplified training loop

A prediction is compared with training feedback, and parameters are adjusted.



Repeat across many examples and iterations.
Different model families and training stages use different objectives.

The arrows are a teaching simplification, not a complete model-training recipe.

What is stored in the model?

Learned parameters help transform input patterns into output probabilities.

PARAMETERS / “WEIGHTS”

Numbers adjusted during training that shape the model’s behavior. They are not a neatly organized encyclopedia.

TRAINING DATA

Text and other data help establish patterns. How much is memorized or reproduced varies; avoid assuming every reply is a database lookup.

A model may reproduce a known fact, infer something new, or generate a convincing error.

A response unfolds one step at a time

For common autoregressive text models, each new token becomes part of the context for the next.

PROMPT

“A rainbow has”

NEXT TOKEN

“many”

NEXT TOKEN

“colors”

FINISH

“.”

This is an illustrative sequence, not the exact token boundaries or output of any real system.

Context helps the model connect words

Transformer-based models can use attention mechanisms to relate different parts of the input.

“Maya gave Jordan the map because Jordan was lost.”



Attention can help the model use relevant context. It does not make every interpretation correct.

The MODEL is not the whole AI PRODUCT

Tools, retrieval, filters, interfaces, and human review may surround the model.



WITHOUT RETRIEVAL

A generated response may use learned patterns but no fresh source lookup.

WITH AUTHORIZED TOOLS

The application may search, calculate, or read permitted data. Results still need interpretation.

Why can an AI answer sound right but be wrong?

Fluent language is not the same thing as verification.

PLAUSIBLE GENERATION

Model output can resemble a well-formed answer even when a name, date, citation, calculation, or explanation is incorrect.

CHECKABLE EVIDENCE

An independent reliable source, reproducible calculation, measurement, or trustworthy reference can support a claim.

NIST calls convincingly false generated content “confabulation” (often called hallucination).

Four habits that make AI more useful

Don’t ask only “Does this sound good?” Ask “Can I check it?”

1 DEFINE THE GOAL

Who is this for? What kind of answer do you need?

2 GIVE CONTEXT

Specify background, limits, format, and what to avoid.

3 VERIFY IMPORTANT CLAIMS

Check facts, citations, computations, and sensitive advice.

4 PROTECT PRIVATE DATA

Do not paste passwords, grades, medical details, or identifiers.

Good prompting is clear communication

A useful prompt often contains five parts.

GOAL

What do you want?

AUDIENCE

Who will use it?

CONTEXT

What must it know?

LIMITS

What should it avoid?

CHECK

How should uncertainty be handled?

Better instructions may improve relevance, but cannot guarantee correctness.

One request, two different prompts

Make the audience and constraints visible.

VAGUE PROMPT

“Explain fractions.”

A valid starting point, but it leaves the level, output, examples, and purpose unclear.

CLEARER PROMPT

“Explain equivalent fractions to an adult beginner in 5 short steps. Use a pizza drawing idea, one practice question, and an answer key. Flag anything uncertain.”

Pause: Which request is easier to evaluate? Why?

Your turn: become the language model

Use this paper activity with no phone or computer.

STEP A — GUESS

Complete “The workshop begins at ____.”

Write TWO plausible endings and why each might fit.

No right time has been given.

STEP B — ADD CONTEXT

Now read: “The flyer says the workshop begins at six.”

Would your continuation change? Why?

What you expect is not what you have verified. Context can contain evidence or misinformation.

Toy “probabilities” from pretend examples

These are invented counts for practicing arithmetic — not a model’s output.

In a made-up set of 10 completions for “The workshop begins at …”:

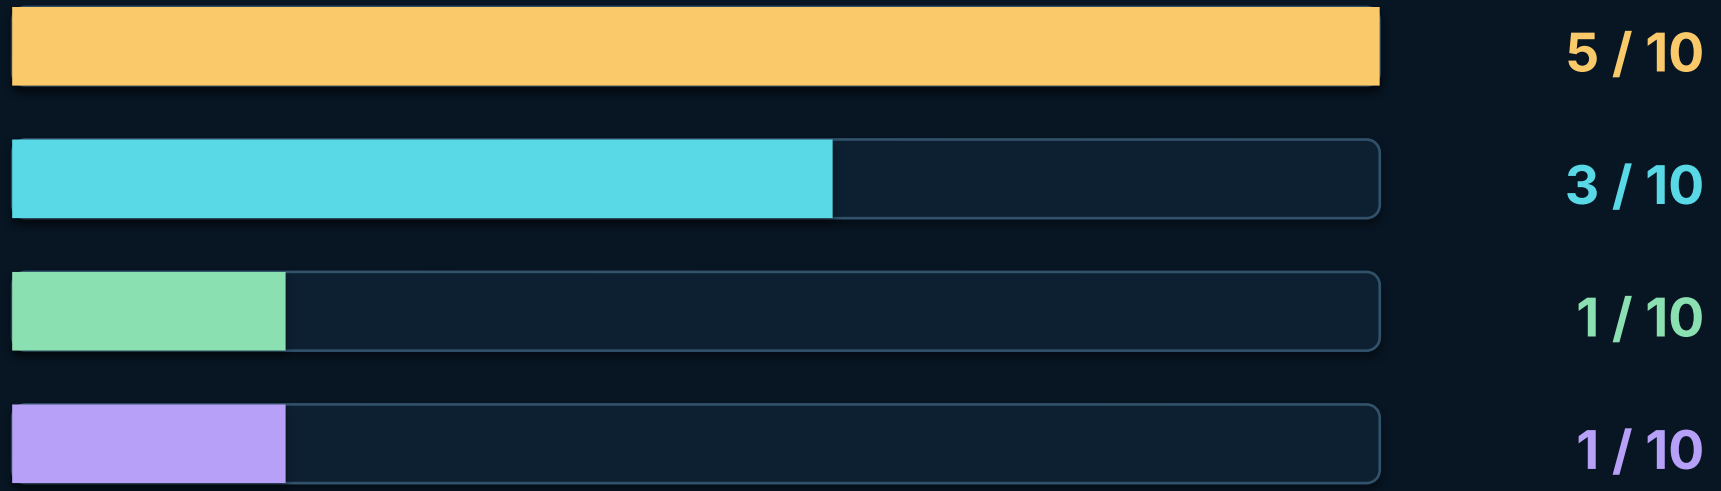

CHALLENGE: Find each percentage. Which option is most frequent? Is it necessarily true?

Worked solution: likely is not verified

Compare a best guess with a claim backed by a source.

TOY SCORES

SIX = 50%
FIVE = 30%
NOON = 10%
HOME = 10%

The most frequent is SIX.

WHAT IT DOES NOT PROVE

These fabricated counts cannot tell us the actual start time.

For a real event, check an official announcement or organizer.

General rule: a generated answer may be useful without being a source of truth.

Activity B: write a better AI request

Make a prompt someone else can understand and evaluate.

GOAL	What should the system do?
AUDIENCE	Who will read it?
CONTEXT	What information can be used?
LIMITS	What should it avoid?
CHECK	How should facts be verified?

Write one or two sentences using all five parts. Then underline what could be checked.

Example: a prompt for community learning

A strong example shows its audience, output, limits, and verification rule.

"Teach a curious adult how Wi-Fi differs from internet access. Use plain language, a 5-box paper activity, and one self-check question. Do not request personal details or router passwords. Identify claims that should be checked against a trustworthy networking reference."

The prompt is still a request — the learner must inspect the output and verify important facts.

Activity C: decide what needs checking

Mark each item “check” or “low-risk idea.” Explain your choice.

A “Invent three names for my fictional robot.”

B “This scholarship deadline is tomorrow.”

C “This medication is safe with your prescription.”

D “Here is a metaphor for how RAM works.”

Why does the cost of being wrong matter?

Answer key: when accuracy matters

The harder the consequences, the stronger the verification.

LOWER-RISK IDEAS

A: Fictional robot names are creative suggestions.

D: RAM metaphor is a teaching aid, but confirm it is not misleading.

VERIFY INDEPENDENTLY

B: Deadline — check the scholarship's official site.

C: Medication safety — consult a qualified health professional and authoritative drug information.

A source link can be wrong or misrepresented. Open it and check the actual claim.

Optional PHONE experiment: compare responses

No phone required. Only use a service if you are permitted to and choose to.

TRY A SAFE, NONPERSONAL TOPIC

Ask an AI tool to explain “Why can Wi-Fi show connected while the web is unavailable?”

Then request a 10-year-old explanation and a college-level explanation.

COMPARE / VERIFY

What changed: wording, examples, or facts?

Check a specific claim against a networking reference.

Do not provide personal files, passwords, grades, or identifying details.

If no tool is available, write how YOU would adjust an explanation for two audiences.

Engineer / researcher: go below the surface

Optional investigation of what changes output quality and reliability.

SAMPLING

Compare choosing the highest-scoring token with selecting among candidates. What changes diversity and repeatability?

EVALUATION

Design a set of prompts with known answers. Track correctness, citation validity, consistency, and failure cases.

SYSTEM DESIGN

Consider retrieval, freshness, caching, permissions, privacy, latency, cost, and human review.

Advanced challenge: what measurement would distinguish “sounds plausible” from “is correct”?

Design a useful AI assistant for a public library

What should it help with, and what must it never pretend to know?

COMMUNITY GOAL

Help visitors find public learning resources and understand technology concepts.

Accessible language. Paper alternatives. Human support when needed.

SAFETY + EVIDENCE

No private patron records by default.
Never invent opening hours or policies.
Cite current official sources.
Offer a staff handoff when unsure.

Create a 5-box system map: USER → INTERFACE → MODEL → TRUSTED SOURCES → HUMAN CHECK.

Four questions to check your understanding

Pause and answer before moving to the next slide.

1. Is every artificial-intelligence system a text chatbot?

2. What is a token in a language model?

3. Does a fluent answer prove the claim is true?

4. How is training different from ordinary inference?

Write your answer in one sentence. The next slide provides a model answer.

Answers: explain, then check

Your wording can differ from these examples.

1 NO	AI also includes classification, recommendation, perception, and more.
2 A TOKEN	A piece of text (for example, a word, part of a word, or symbol).
3 NO	Plausibility is not proof. Verify important facts and citations.
4 TRAINING VS INFERENCE	Training adjusts parameters; inference generates responses using the available model.

If you can teach one idea back accurately, you have learned something useful.

From curious learner to responsible creator

Every learning path can produce an artifact worth sharing.

EXPLORER

Draw how a question becomes an answer. Add a “verify facts” step.

Share a question you still have.

BUILDER

Create a prompt card and a simple evaluation checklist for a useful community task.

ENGINEER / RESEARCHER

Design a small, consent-respecting test set; measure factual errors and usability.

Keep learner information private. No identifying records or credentials in a shared artifact.

Where to go next

Free learning resources; no paid account or special device required for the paper activities.

THURSDAY OPEN LAB	gunnchos.com/labs/thursday
gunnchOS3k	gunnchos.com
WAIKE LEARNING	github.com/gunnchOS3k/waike-research-ops
THE TECHNOLOGY LANDSCAPE	github.com/gunnchOS3k/gunnchos-technology-landscape
GOOGLE MACHINE LEARNING CRASH COURSE	developers.google.com/machine-learning/crash-course/llm
NIST GENERATIVE AI PROFILE	nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

NEXT THURSDAY • NOV 12 • HOW DO HACKERS AND DEFENDERS THINK?