

Demystifying AI

what it is, what it is not, and how to use it well



A plain-English workshop for everyday users — no math, no marketing, no hype.

Key takeaway: It's a predictive pattern engine, not a thinking machine. Knowing the difference changes how you use it.

WORKSHOP ROADMAP

What You'll Walk Away With

Three things every smart user of these tools should understand — and most people don't.

How it actually works

A look under the hood at how these systems turn your question into an answer. It's mechanical, not magical.

Where it fails

Why these systems confidently produce wrong answers — and how to spot it before it costs you.

How to use it well

The handful of habits that separate people who get real value from those who get burned.

Net: By the end, you'll know what's actually happening when you type into a chatbot — and why that matters.



MEET YOUR FACILITATOR

Dan Bond

Founder & Principal Advisor at LumenForge Advisors — and your guide for this workshop.

26+ years at Microsoft

Most recently Director of AI Strategy & Transformation for the Americas, turning ambitious AI strategy into measurable, real-world adoption.

Built to teach

Years in Learning & Development at Microsoft — developing and delivering curriculum that turns complex technical topics into content people can actually use.

Why I built this

Founder of LumenForge Advisors, bringing enterprise-grade enablement to everyday teams — distilled here into plain-English, hands-on practice.

Net: I've spent a career helping people learn to work with new technology. This series is me doing that with you.

WORKSHOP AGENDA

What We'll Cover Today

Six stops, one through-line — from how these tools work to using them well.

01

6 MIN

The Mental Shift

Why these tools predict rather than think — the one idea everything else builds on.

02

8 MIN

The Vocabulary

AI, ML, neural networks, and GenAI — what each term means and how they nest.

03

14 MIN

Under the Hood

How a chatbot answers you: the five-step tokenize, embed, predict pipeline.

04

10 MIN

Why It Makes Things Up

Hallucinations, knowledge cutoffs, and why every output still needs a human check.

05

8 MIN

Where It Wins

Summarizing and reshaping text — the tasks these tools are genuinely great at.

06

12 MIN

Use It Safely & Well

Failure modes, what not to paste, and five habits to take with you.

Net: About sixty minutes, no math and no jargon — just a mental model you can use the same afternoon.



It's Pattern-Matching, Not Thinking

“Calling it ‘artificial intelligence’ is a marketing decision, not a technical one. These systems don’t reason or understand — they predict what text statistically comes next.”

100%

of output is statistical prediction. The system never “looks up” a fact — it generates text that fits the pattern.

0

facts the model itself verifies before answering. It produces plausible-sounding text and stops.

Takeaway: this is the single most important shift in how you think about these tools. Everything else flows from it.



FOUR WORDS PEOPLE MIX UP

AI, ML, Neural Networks, GenAI — The Nesting Doll

These four terms get used interchangeably in news headlines. They're not the same. Each one sits inside the next, like nesting dolls — and being precise helps you read coverage clearly.

AI — the outer box

The biggest, oldest, fuzziest term. Any computer system doing something we'd call "smart." A 1970s chess program counts. Your spam filter counts.

Machine Learning — inside AI

A specific approach: instead of programming rules, you show the system examples and let it find patterns. Recommends your Netflix shows. Reads checks at the ATM.

Neural Networks — inside ML

One specific style of machine learning, loosely inspired by brain cells. Big stacked grids of math. The thing powering ChatGPT, image generators, and voice assistants.


Generative AI — inside neural nets

The newest layer — and the one everyone's met. Neural nets trained to create text, images, and code. ChatGPT, Copilot, and image generators all live here.

HOW IT WORKS · THE 5-STEP PIPELINE

How a Chatbot Actually Answers You

Every reply follows the same five-step pipeline. The next few slides zoom in on the middle three — that's where the “intelligence” lives, or doesn't.



Outcome: what looks like a thoughtful answer is a long sequence of educated guesses, stitched together.

STEP 1 · TOKENIZATION

Your Words Become Numbers

Think of it like a translator that turns text into a numbered code book. From here on, the system never sees actual letters again — just numbers.

Tokens, not words

THE UNIT

A token is usually a word fragment, not a whole word. “Unbelievable” might become three tokens: un / believ / able.

Each gets a number

THE ID

Every token in the system’s vocabulary has a unique number. “The” might be 464. “Cat” might be 8415.

~100K tokens total

THE VOCABULARY

Modern systems use vocabularies of roughly 50,000 to 200,000 tokens, covering most languages plus code symbols.


Letters disappear

THE TRADE-OFF

From here on the system works only with numbers. That’s why it miscounts letters or struggles with rhymes — it can’t see the characters anymore.

STEP 2 · EMBEDDINGS

Numbers Get Coordinates



MEANING-SPACE

Each token gets placed in a huge multi-dimensional map of meaning.

Think of a giant 3D map — except with thousands of dimensions, not three. Words that mean similar things land near each other. “Dog” sits near “puppy.” “King” minus “man” plus “woman” lands near “queen.” The system doesn’t know what a king is — it just knows its position.

Net: this is how the system gets a sense of relationships between words, without ever “understanding” them.

STEP 3 · PREDICTION

It Guesses the Next Chunk — Over and Over

This is the core trick. The system looks at everything you wrote plus everything it's written so far, and asks one question: what token statistically comes next? It picks one, adds it, then asks again.

One token at a time

It doesn't plan ahead. It doesn't draft and revise. It generates left-to-right, one chunk at a time, never looking back to change its mind.

Probabilities, not certainty

At each step, thousands of possible next tokens compete with different probabilities. The system picks one with some randomness — so it doesn't sound robotic.

Patterns, not facts

What “statistically comes next” was learned from huge amounts of training text. The output reflects patterns in that text, not a verified understanding of reality.

PUTTING IT TOGETHER

What It THINKS You Asked



The system doesn't understand your question — it pattern-matches it.

Your question becomes coordinates in meaning-space, compared against patterns from billions of training examples. The most-likely-fitting response gets generated, token by token.

“Hallucinations” happen when the patterns match but the facts don’t.



WHY HALLUCINATIONS HAPPEN

It's Not a Bug — It's How It Works

“The system has no internal way to know whether what it’s generating is true. It only knows whether the text is plausible — and plausible is what it produces.”

Same

tone, right or wrong. The system sounds equally confident whether facts are real or invented.

Zero

built-in “am I making this up?” check. The model can’t flag uncertainty it doesn’t have.

Takeaway: this is the most dangerous failure mode. Confident wrongness reads exactly like confident rightness.



Why You Always Need a Person in the Loop


100%

of AI output that informs a decision needs human verification — every claim, every number, every name.

Not because the system is bad at it. Because it has no way to know when it's wrong. That job is yours.



Where AI Genuinely Shines #1: Summarizing



BEST USE CASE

Boiling long text down to short text is the most reliable thing these systems do.

Hand it a 40-page report, a meeting transcript, a long email thread, or a dense article — it gives you a clean summary in seconds. Why this works: the source material is right in front of it. The system isn't inventing — it's compressing. Much harder to hallucinate when the answer is already on the page.

Net: the more your task looks like “here’s the information, just compress it,” the safer and more useful AI gets.

Where AI Shines #2: Drafting & Reformatting

First-draft writing

STARTING POINT

Emails, memos, social posts, recaps of meetings you led. Easier to edit a 70% draft than to stare at a blank page.

Reformat & restyle

MECHANICAL WORK

Turn a list into a table. Convert formal to casual. Bulletize a paragraph. Translate. Pure pattern-shaping — exactly where it excels.

Brainstorm options

VOLUME OVER QUALITY

Ask for 20 names, 10 angles, 15 subject lines. The system produces volume; you bring the taste and pick the best one.

Plain-language explainer

TRANSLATOR

Make a legal document readable. Turn jargon into clear English. Doesn't need to invent — just rephrase what's in front of it.

Where It Falls Down

Facts & citations

MAKES THEM UP

Book titles, court cases, research papers, statistics — confidently fabricated. Always verify before using.

Math & counting

GETS WRONG OFTEN

Letters in a word, items in a list, arithmetic, dates. Use a calculator or spreadsheet for numbers, not a chatbot.

Anything recent

DOESN'T KNOW IT

Training has a cutoff months or years ago. Recent events, news, prices — unreliable unless the tool is web-connected.

Real reasoning

PATTERN-MATCHES IT

Looks like reasoning. Isn't. Fails on novel problems, long multi-step logic, anything outside training patterns.

Don't Paste That In

When you paste text into a chatbot, you're sending it to a company's servers. Treat it the way you'd treat sending an email to a stranger.

GENERALLY SAFE

- ▶ General questions and public information
- ▶ Drafts you'd also send in plain email
- ▶ Brainstorming, reformatting, summaries of public docs
- ▶ Anything you'd be fine seeing on a billboard

DO NOT PASTE

- ▶ Passwords, account numbers, SSN, credit cards
- ▶ Confidential work docs (unless your employer approved the tool)
- ▶ Health records, legal matters, anything privacy-protected
- ▶ Other people's personal information

Net: if you wouldn't put it on a postcard, don't paste it into a chatbot.

Five Habits That Make You Better At This

01**Verify before you use**

Treat every output as a confident-sounding first draft. Check facts, names, and numbers yourself.

02**Give it context**

Specific questions get better predictions. The more detail you give it, the better the output.

03**Use it for shape, not truth**

Drafting, summarizing, brainstorming — yes. Facts, math, medical or legal advice — find a better source.

04**Don't paste sensitive info**

If you wouldn't put it on a postcard, it doesn't go in. No passwords, no health records, no others' data.

05**Stay the expert**

You are the human in the loop. Your judgment, your responsibility. The tool assists; it doesn't replace you.

These five habits cover most of the “how do I use this well?” question.

WRAP

What You'll Remember From This



AI is a predictive pattern engine, not a thinking machine.

It's genuinely useful for shaping text — drafting, summarizing, brainstorming, reformatting. It's genuinely unreliable for facts, math, recent events, and real reasoning.

Use it like a fast assistant. Stay the expert. Verify before you trust.

Where This Goes Next

You've seen what AI actually is. Now let's talk about what it means for you.

Your next step: Book a free 30-minute AI Reality Check Call at lumenforge.ai — no pitch, just an honest read on where AI fits your business.

lumenforge.ai • Book Your AI Reality Check Call

