

The Jonesi Consultancy | Houston, TX

AI LITERACY · RESOURCE HANDOUT FOR LEADERSHIP

25 AI TERMS FOR NON-TECHNICAL LEADERS

Plain Definitions of AI Terminology and How Each One is
Relevant to Executives

*Leaders who know the vocabulary make better calls, for both the
business and the people in it.*

HOW TO USE THIS

Enough vocabulary to ask a good question, hold a vendor to an answer, and read a proposal critically.

Nobody leading an organization needs to understand how these systems are built. What is genuinely useful is knowing what the words mean when they appear in a vendor proposal, an IT recommendation, or a board paper, because that is where the vocabulary is doing work.

Each entry below gives a plain definition and a line on why it matters to you specifically. The second line is the reason the entry is here.

1 Foundations

01 ARTIFICIAL INTELLIGENCE

A broad label for software that performs tasks normally requiring human judgment, such as recognizing images, understanding language, or making predictions.

WHY IT MATTERS: The term is so broad that it is nearly meaningless in a proposal. Ask what the system actually does rather than accepting the category.

02 MACHINE LEARNING

Software that improves at a task by finding patterns in data rather than by being given explicit rules.

WHY IT MATTERS: Explains why these systems reflect whatever was in their data, including its errors and its assumptions.

03 LARGE LANGUAGE MODEL

A system trained on very large amounts of text to predict what text should come next. This is what sits behind ChatGPT, Claude, Gemini, and Copilot.

WHY IT MATTERS: Predicting plausible text is not the same as knowing facts. Almost every practical limitation follows from this one sentence.

04 GENERATIVE AI

AI that produces new content: text, images, audio, video, or code, rather than classifying or predicting from existing content.

WHY IT MATTERS: This is the category that has changed knowledge work, and the category your policy needs to address.

05 TRAINING DATA

The material a model learned from, gathered before the model was released.

WHY IT MATTERS: Determines what a model knows, what it is good at, and which biases it carries. Your organization is almost certainly not in it.

06 KNOWLEDGE CUTOFF

The date after which a model has no built-in information, because its training data ended there.

WHY IT MATTERS: Anything more recent has to be supplied by you or looked up by the tool. Ask about recency before trusting an answer about current conditions.

07 MULTIMODAL

A model that handles more than one kind of input or output, such as reading an image and describing it in words.

WHY IT MATTERS: Useful in practice: photographs of documents, charts, and screenshots become workable inputs.

2 How They Behave

08 HALLUCINATION

When a model produces information that is fabricated but presented with the same confidence as anything else it says.

WHY IT MATTERS: The most important term on this page. It is a property of how these systems work rather than a defect awaiting a fix, which is why human verification is a permanent requirement.

09 NONDETERMINISM

Asking the same question twice can produce different answers.

WHY IT MATTERS: Matters enormously for consistency. If a process needs the same output every time, a language model alone is the wrong tool for that step.

10 CONTEXT WINDOW

The amount of text a model can consider at once, covering both what you send and what it returns.

WHY IT MATTERS: Explains why very long documents get truncated or summarized, and why quality can degrade in a long conversation.

11 TOKEN

The unit these systems process text in, roughly three quarters of a word.

WHY IT MATTERS: Usage and API pricing are measured in tokens, so this is the unit your invoice is built from.

12 BIAS

Systematic skew in output that reflects patterns in the training data or in how the system was tuned.

WHY IT MATTERS: Carries real legal and ethical weight anywhere a model touches hiring, lending, admissions, or service eligibility.

13 MODEL DRIFT

A model's behavior changing over time, either because the vendor updated it or because the conditions it operates in have shifted.

WHY IT MATTERS: A prompt or workflow that worked reliably last quarter can quietly stop working. Build in a periodic check rather than assuming stability.

THE ONE TO REMEMBER

Hallucination explains why verification is permanent rather than temporary. A fabricated answer arrives in the same confident voice as a correct one, so the check has to sit outside the system.

3 Working with Them

14 PROMPT

The instruction you give the model.

WHY IT MATTERS: Quality of input drives quality of output more than any other controllable factor.

15 PROMPT ENGINEERING

The practice of writing instructions that produce reliable results.

WHY IT MATTERS: Worth learning and worth keeping in perspective. Techniques change with each model generation, while judgment about when to use these tools does not.

16 SYSTEM PROMPT

Standing instructions that sit behind every conversation in a configured tool, setting its role, tone, and limits.

WHY IT MATTERS: How an organization enforces consistency without depending on every individual writing a good prompt.

17 FEW-SHOT EXAMPLES

Including two or three examples of what you want inside the instruction.

WHY IT MATTERS: The cheapest reliable improvement available. Examples communicate format and tone better than any description of them.

18 RETRIEVAL AUGMENTED GENERATION

Connecting a model to your own documents so it answers from them rather than from memory. Usually shortened to RAG.

WHY IT MATTERS: The standard approach for making a general tool useful on organization-specific knowledge, and the phrase most likely to appear in a vendor proposal you are asked to approve.

19 FINE-TUNING

Additional training on your own material to change how a model behaves.

WHY IT MATTERS: Expensive, slower to change, and less often necessary than vendors suggest. Ask whether retrieval would solve the problem first.

20 AI AGENT

A system that takes multiple steps toward a goal on its own, using tools or taking actions rather than simply answering.

WHY IT MATTERS: Raises the stakes considerably. When software acts rather than advises, the accountability and oversight questions have to be settled in advance.

ASK RETRIEVAL BEFORE FINE-TUNING

Most requests to “train a model on our data” are really requests for retrieval. It is cheaper, easier to update, and easier to audit.

4 Risk and Governance

21 SHADOW AI

Employees using unapproved tools for work, generally on personal accounts.

WHY IT MATTERS: Present in almost every organization that has not published an approved list. Prohibition does not remove it, it removes your visibility of it.

22 HUMAN-IN-THE-LOOP

A required point where a person reviews and approves before output is used or released.

WHY IT MATTERS: The single control that actually catches confident wrong answers. Accountability stays with the person who signed off.

23 ENTERPRISE ACCOUNT

A business subscription, as distinct from a consumer one, typically with administrative controls and different data handling terms.

WHY IT MATTERS: The distinction matters more than the brand. The same product can carry very different data commitments depending on which tier you are on.

24 DATA RETENTION AND TRAINING OPT-OUT

Whether a vendor stores what you submit, for how long, and whether it may be used to improve their models.

WHY IT MATTERS: Get the answer in the contract rather than from a settings menu. Settings can be changed by someone who does not know what they are changing.

25 AI GOVERNANCE

The structure that decides who approves tools, who reviews use, and how exceptions and incidents are handled.

WHY IT MATTERS: Without it, a policy is accurate on the day it is signed and out of date within two quarters.

The 3 Terms that Carry the Most Weight

If you retain nothing else from this page, retain these.

HALLUCINATION

Explains why verification is permanent rather than temporary.

KNOWLEDGE CUTOFF

Explains why a confident answer about current conditions may be describing a world that has moved on.

HUMAN-IN-THE-LOOP

The control that addresses both, and the only one that has held up so far.

5 Five Phrases Worth Pushing Back On

These appear in proposals and presentations and sound more precise than they are.

"AI-powered"

Describes nothing. Ask what the system does, what it is trained on, and what happens when it is wrong.

"It learns from your data"

Could mean retrieval, fine-tuning, or nothing at all. Ask which, then ask where that data lives.

"Fully automated"

Ask where the human checkpoint is. If the honest answer is that there is not one, that is a decision requiring approval rather than a feature.

"Accuracy of ninety-five percent"

Measured on what, by whom, against what benchmark? Accuracy figures without a stated test are marketing.

"Enterprise-grade security"

Ask the specific questions instead: where the data is held, for how long, whether it trains their models, and what happens on termination.

VOCABULARY IS NOT THE POINT

Being able to ask a second question is. Most poor AI decisions inside organizations are made not because leadership did not understand the technology, but because nobody in the room felt equipped to ask what a particular sentence actually meant.

© 2026 The Jonesi Consultancy. All rights reserved.

Building the shared vocabulary a leadership team needs.

Our sessions are designed for people who will never write a line of code and still have to decide.

