MODULES 1–4

Comprehensive Glossary

Every term used across the four modules, A–Z

AI for Emergency Management Education Program
AIDE Initiative · Markle Foundation · Aspen Digital · RAND

HOW TO USE THIS GLOSSARY

This glossary defines the terms used across all four modules in a single alphabetical list. Each term carries the same meaning in every module. Where a term has specific implications for emergency management, an **In EM** note describes them. **See also** lists related terms. Definitions draw on the most complete description provided in the module materials.

A

A-C-R-S

The four-lens output check: Accuracy, Completeness, Relevance, Soundness. The four checks you run before anything leaves your desk. Then make the call: accept, verify, revise, or reject. It is a tool for reaching a decision, not a grade.

In EM: Most real outputs land in verify or revise — and a person owns the final call.

See also: Hallucination · Human in the loop · Prompt framework · Stakes tier

Access model

How you pay for a product. Common access models include enterprise license, per-user, purchase outright, pay-per-use, and open source.

In EM: How you pay determines what happens when usage spikes during an incident.

See also: Pay-per-use · Enterprise license · Open source · Deployment model

Agent

An LLM you let do things, not only answer. You give a goal and it carries out the steps, then reports back. A human still sets the goal and stays responsible.

See also: Agentic AI · Tool / add-on

Agentic AI

AI that doesn't just answer a question — it plans a sequence of actions and executes them, often with limited human oversight. More proactive than generative AI: it works toward a goal and iterates. A human still sets the goal and stays responsible.

See also: Agent · Generative AI · Human in command

AI (Artificial Intelligence)

Technology enabled by computational systems that replicates or augments human cognitive functions — such as reasoning, prediction, problem solving, learning, planning, perception, and communication.

See also: AI Effect · AI-enabled product

AI Effect

The tendency for people to stop calling a capability "AI" once it becomes reliable and familiar. Spam filtering, GPS, and voice recognition were all once cutting-edge AI.

In EM: The AI Effect is why "AI-powered" is a starting point for questions, not evidence.

See also: AI · AI-enabled product

AI-enabled product

A named, purchasable tool with documented features, commercial terms, and repeatable deployment. It is not a one-off prototype, a research paper, or a raw capability — it is a finished offering you can buy, license, or deploy.

See also: Purpose-built for EM · Functional ladder · Deployment model

AI governance

The set of policies and practices to ensure that AI is built, deployed, and used in alignment with organizational priorities and to manage risks. Includes model validation, documentation of limitations, accountability structures, and protections for people whose data the system uses.

See also: NIST AI RMF · ISO/IEC 42001 · Human in the loop

API (Application Programming Interface)

A technical connector that allows your IT team to connect a product into your own systems using code. More powerful and flexible than SaaS, but requires technical staff and regular maintenance.

See also: SaaS · Integration requirement · Deployment model

Author of record

You — the person whose name is on the final product. The tool never is.

In EM: On anything that leaves your desk, you own the output. Check it before it goes out.

See also: Human in command · A-C-R-S

Automation bias

The human tendency to accept an AI recommendation without cross-checking.

In EM: EOC staff accepting an AI resource recommendation without verifying with field teams who know conditions on the ground.

See also: Human in the loop · A-C-R-S

B

Base model

The underlying AI engine that does the "thinking" (e.g., a GPT, Claude, or Gemini model). The "brain" behind the app.

See also: Chatbot · Model size · Foundation model

Bias / bias and blind spots

When AI reflects slants in training data or leaves groups out. Can skew tone, recommendations, and outputs — and quietly exclude whole populations.

In EM: A flood-risk model underestimating impact in historically underserved areas with sparser historical data.

See also: Training data · A-C-R-S · AI governance

Billboard test

A quick check before pasting anything into a public AI tool: would you be OK seeing it on a billboard by the highway? If not, it doesn't go in — no personal data, no sensitive or controlled information, nothing you don't own.

See also: PII (Personally Identifiable Information) · Privacy risk

C

Chain-of-thought prompting

Asking the model to show its reasoning steps — "make it think out loud." Documented gains are strongest on multi-step reasoning (arithmetic, logic, sequencing). Surfaces assumptions but does not verify them.

See also: Zero-shot prompting · Few-shot prompting · Role prompting

Chatbot

The app or website you actually talk to (ChatGPT, Claude, Gemini, Copilot…). Think of the chatbot as the car — login, history, and chat window — and the model as the engine under the hood that produces the answer.

See also: Base model · Enterprise tool

Closed AI system

An AI system that answers only from a defined set of information rather than the open internet.

In EM: Utah's AI system for the non-emergency line answers from a defined set of information and hands anything that turns out to be a real emergency to a person.

See also: Grounding · Human in the loop

Computer vision

The AI subfield focused on interpreting images and video.

In EM: Damage assessment from aerial photos, or drone-based person detection.

See also: Multimodal · Machine learning / predictive modeling

Connectivity

Whether a product works without internet. Three levels: **offline capable** (functions fully without internet), **intermittent tolerant** (keeps working during brief outages and syncs when connection returns), and **requires online** (stops when internet goes down).

In EM: Approximately 78% of EM products require continuous connectivity — for response tools this can be non-negotiable.

See also: Edge (processing) · On-premises · SaaS

Context gaps

A limitation where the model lacks awareness of real-world changes not reflected in its training data.

In EM: Not knowing a highway overpass is under construction, and routing evacuees through it.

See also: Training cutoff · Grounding · Hallucination

Context window

The total amount of information the model can consider at once — your prompt, the chat so far, files, and its own replies. Measured in tokens. Picture a desk of fixed size: pile on too much and older papers slide off the edge.

See also: Token · Memory (in LLMs)

Cost justification

Comparing measurable value — time saved, throughput, consistency, fewer errors, reduced risk — against the full burden of using a tool: license or usage fees, setup and integration, data preparation, training and review time, governance, and the cost of mistakes.

In EM: Before scaling up, establish a baseline, pilot the product, compare value with burden, then decide whether to adopt.

See also: Access model · Pay-per-use

D

Data input requirements

The ongoing burden on your staff — not just at setup, but every day or every use. Categories: **vendor maintained** (vendor handles the data), **user input at use** (someone provides data each time), and **ongoing user input** (a staff member has a recurring responsibility to keep the system current).

See also: Integration requirement · Access model

Deep learning

A subset of machine learning using neural networks with many layers that can learn features on their own — nobody writes a rule for what a face looks like. Unlocked practical perception capabilities (seeing, hearing, reading) in the 2010s.

See also: Neural network · Machine learning / predictive modeling

Deployment model

Where a product physically runs. Options include SaaS (browser-based, vendor-hosted), API, installed software, on-premises (your servers), hardware (physical devices), and edge (processing happens locally).

In EM: Shapes whether you need IT staff, whether it works offline, and whether data leaves your building.

See also: SaaS · On-premises · Edge (processing) · Connectivity

E

Edge (processing)

When AI processing happens locally at your site rather than in the vendor's cloud. Useful for connectivity-constrained environments.

See also: Connectivity · On-premises · Deployment model

Enterprise license

An access model where your organization negotiates a contract for ongoing use across the entire organization, often multi-year.

See also: Access model · Pay-per-use

Enterprise tool

An AI tool your organization has approved under a formal agreement, with defined data-handling rules. Inputs are usually **not** used to train the model.

In EM: The distinction that decides what you may paste — never put incident data or PII into an unapproved tool.

See also: PII · Chatbot · Privacy risk

F

False alarms and misses

A limitation where AI incorrectly flags something (false positive) or fails to detect something (false negative).

In EM: A debris-detection model flagging intact structures as destroyed, triggering unnecessary deployment.

See also: Variability · Bias · Automation bias

FedRAMP

Federal Risk and Authorization Management Program. Verifies that a cloud product meets federal security standards for government data. A meaningful certification, but not designed to address AI-specific risks.

See also: SOC 2 · ISO 27001 · ISO/IEC 42001

Few-shot prompting

Adding one, two, or three examples so the model learns from them — "show it examples." The model learns formatting and style from your examples, not the reasoning behind them.

See also: Zero-shot prompting · Role prompting · Chain-of-thought prompting

Foundation model

One large model that is adaptable to many tasks; typically multimodal (works across text, images, and audio together) and may allow agentic actions. The dominant paradigm of the current AI era.

See also: Base model · Multimodal · Agentic AI

Fragmentation

When AI products cover only a slice of a task area's full scope. **Sequential fragmentation:** products cover some stages but not others. **Domain fragmentation:** products deliver one data stream when the task requires many.

See also: Stacking / stacking problem · Functional ladder

Functional ladder

A framework for assessing what level of contribution an AI product makes — from the bottom (perception/detection) through prediction, aggregation, summarization, narrow optimization, up to cross-domain reasoning at the top.

In EM: Most EM products operate on the bottom rungs.

See also: AI-enabled product · Fragmentation

G

Generative AI

AI that creates something new from a prompt — text, images, audio, video, code. Systems designed to produce content by learning and reproducing patterns in existing content. LLMs are the branch specialized for language.

See also: Large language model (LLM) · Predictive AI / traditional ML · Agentic AI

Grounding

Giving the tool the source material so it answers from your document rather than from memory. The closest thing there is to a fix for invention.

See also: Hallucination · Pretrained knowledge · Training cutoff

Guardrail bypass

An AI-specific cybersecurity risk where an attacker finds ways to circumvent the safety restrictions built into an AI system. Traditional certifications were not designed to address this.

See also: Prompt injection · Model manipulation · SOC 2

H

Hallucination

A confident, fluent output that is wrong or invented. Not a malfunction — a consequence of how these systems work. The model is predicting plausible text, and sometimes that text is factually incorrect.

In EM: An AI-drafted evacuation notice referencing a road that doesn't exist, or a shelter already at capacity.

See also: Grounding · A-C-R-S · Sycophancy

Human in command

A qualified person does the reasoning and owns the decision — or you do not use the tool at all. Applies to high-stakes tasks: protective actions, survivor benefits, legal obligations, resource decisions.

In EM: In protective-action decisions, the authorized official remains the decision authority and decision-of-record.

See also: Human in the loop · Human on the loop · Stakes tier · Author of record

Human in the loop

A person reviews and approves every output before it is used. The authorized official remains the decision authority. Applies to medium-stakes tasks.

See also: Human in command · Human on the loop · Automation bias · Stakes tier

Human on the loop

The system acts while a person monitors and can intervene. Applies to low-stakes tasks where spot-checking and sampling is appropriate.

See also: Human in the loop · Agentic AI · Stakes tier

I

Integration requirement

How much work your team needs to do before a product becomes functional. **Turnkey:** ready to go quickly (~15% of products). **Light:** configuration without writing code (~21%). **Requires technical integration:** an IT professional must connect the product to existing systems (~62%).

See also: API · Data input requirements · Deployment model

ISO 27001

An international standard that verifies an information security management system exists and is being maintained. A meaningful certification, but not designed to address AI-specific risks like bias or model accuracy.

See also: SOC 2 · FedRAMP · ISO/IEC 42001

ISO/IEC 42001

An AI management system standard published in December 2023. The closest thing to an AI governance certification, but adoption is extremely rare — less than 1% of EM products claim alignment.

See also: AI governance · NIST AI RMF

K

Knowledge representation and reasoning

The AI subfield focused on encoding facts and rules so a system can draw conclusions — the old expert-systems tradition. The foundation of rule-based AI from the 1950s–1980s.

See also: Rules-based automation

L

Large language model (LLM)

The branch of generative AI specialized for language. A computer system trained on huge amounts of text that generates responses by predicting what words (or word-pieces) are likely to come next — think "very advanced autocomplete." Every LLM is generative AI; not all generative AI is an LLM.

See also: Generative AI · Token · Natural language processing (NLP)

M

Machine learning / predictive modeling

The AI subfield focused on learning patterns from examples to classify, score, detect, recommend, and forecast. The root of what vendors call predictive analytics. Quality depends on the training examples and whether new cases resemble them closely enough.

See also: Supervised learning · Training data · Predictive AI / traditional ML

Memory (in LLMs)

Unless a "memory" feature is on, each new chat starts blank. Context lives in the conversation, not in a stored profile — that's why you re-paste background information.

See also: Context window · Grounding

Model manipulation

An AI-specific cybersecurity risk where an attacker tampers with or exploits the AI model itself. Traditional security certifications were not designed to address this.

See also: Prompt injection · Guardrail bypass

Model size

Small/fast vs. large/advanced "brains" on a menu. Match the model to the job: quick summary vs. hard reasoning. Bigger models cost more per token. Think scooter, sedan, or semi-truck: small models for quick rewrites, mid-size as the daily driver, large for careful reasoning.

See also: Base model · Token

Multimodal

Able to work with more than one type of input or output — text, documents, images, audio, or video, depending on the tool. Foundation models are typically multimodal.

In EM: Multimodal prompting widens what you can hand a tool — and widens the privacy exposure with it.

See also: Computer vision · Foundation model · PII

N

Natural language processing (NLP)

The AI subfield focused on reading and producing human language — the family that LLMs belong to.

See also: Large language model (LLM) · Generative AI

Neural network

A computational architecture modeled loosely after the human brain, with many interconnected layers. Training adjusts many small weights until outputs improve. Powerful pattern matching — not thinking. The overlapping layers find complex patterns but are harder to interpret. Information enters an input layer, is processed across hidden layers, and leaves through an output layer.

See also: Deep learning · Training data · Weights

NIST AI RMF (AI Risk Management Framework)

A voluntary federal framework for managing AI risk. A vendor can claim alignment, but nobody can independently verify that claim — there is no audit or certificate.

See also: AI governance · ISO/IEC 42001

O

On-premises

A deployment model where the product runs on servers you physically own. Your IT team handles maintenance. Data stays in your building.

See also: SaaS · Edge (processing) · Deployment model

Open-ended vs. closed-ended prompts

Open-ended prompts explore and widen your options; closed-ended prompts pin down one specific, structured output. Neither is better. A good habit: open first to see the landscape, then narrow to act — and watch for a clean, narrow answer that hides what you didn't think to ask.

See also: Prompt · Prompt framework

Open source

Software with no license fee, often free to access, but someone still has to set it up, secure it, and maintain it. Free to access does not mean free from licensing terms — verify conditions before implementation.

See also: Access model · Integration requirement

P

Parent product

A product that another product (an add-on) requires in order to function.

In EM: About 1 in 5 AI products in the EM landscape cannot function without another product — and marketing doesn't always say so.

See also: Tool / add-on · Stacking / stacking problem · Vendor lock-in

Pay-per-use

An access model where you are charged per message, image, query, or other unit of consumption.

In EM: Costs may balloon during high-tempo disaster operations.

See also: Access model · Enterprise license

PII (Personally Identifiable Information)

Data that can identify a specific person. Never paste it into a public tool. Never enter it into unapproved tools.

See also: Privacy risk · Enterprise tool

Planning, scheduling, and optimization

The AI subfield focused on finding the best sequence or allocation.

In EM: Routing crews, staging resources, evacuation route planning.

See also: Rules-based automation · Agentic AI

Predictive AI / traditional ML

Writes nothing new. Scores, classifies, sorts, ranks, or flags things that already exist. It learned from past examples, and tells you what's likely based on historical patterns.

See also: Machine learning / predictive modeling · Generative AI

Predictive analytics

Using historical and current information to estimate a future or unknown outcome — a forecast, probability, quantity, or expected timeline. Often uses machine learning.

In EM: Forecasting flood conditions, or estimating shelter and hospital needs, to inform human decisions.

See also: Machine learning / predictive modeling · Predictive AI / traditional ML

Pretrained knowledge

What the model absorbed during training, stored as learned patterns — not a lookup table of verified facts. The model doesn't "know" things; it has statistical associations.

See also: Training cutoff · Grounding · Hallucination

Privacy risk

The condition where a product collects or processes PII, geolocation data, health data, law enforcement-sensitive information, or communications content.

In EM: Approximately 94% of AI products in the EM landscape present a potential privacy risk.

See also: PII · AI governance · Enterprise tool

Prompt

The instruction you give the AI — your input. Everything you type is the prompt.

See also: Prompt engineering · Context window

Prompt engineering

The practice of adjusting *how* you ask to improve *what* you get back. No coding required — just clearer requests.

See also: Zero-shot prompting · Few-shot prompting · Role prompting · Chain-of-thought prompting

Prompt framework

Six moves for building a prompt: Task (one clear job), Context (what it needs to know), Reference (the source or example to follow), Format (what the answer should look like), Constraints (the rules and limits), and Iterate (how you improve the first answer). A role is optional. Use more of the parts as the task gets more important, detailed, or repeatable.

In EM: A weak answer usually points to the missing part — invented details mean no reference; the wrong length or shape means no format.

See also: Prompt · Role prompting · A-C-R-S

Prompt injection

An AI-specific cybersecurity risk where an attacker crafts inputs designed to override the system's instructions or safety measures. Traditional certifications were not designed to address this.

See also: Guardrail bypass · Model manipulation

PSAP (Public Safety Answering Point)

The call center that answers 911 calls — and, often, non-emergency calls.

In EM: An estimated 65% of Utah's PSAP workload is non-emergency calls, which is why the state is using AI to answer its non-emergency line.

See also: Closed AI system

Purpose-built for EM

One of three product types in the EM landscape: products explicitly designed for one or more emergency management task areas. Contrasts with "evidence of EM use" (built for another sector but with documented EM adoption) and "may be useful to EM" (general-purpose, configurable, but not EM-designed).

See also: AI-enabled product · Functional ladder

R

Randomness

See Stochasticity. The reason the same input can give different answers — useful for brainstorming, risky when you need the same result twice.

See also: Stochasticity · Temperature · Variability

Robotics

The AI subfield focused on perception and action in the physical world — including drones and autonomous physical systems.

See also: Computer vision · Edge (processing)

Role prompting

Telling the model who to write as (e.g., "write as our county public information officer"). Sets voice, not accuracy.

See also: Prompt engineering · Few-shot prompting

Rules-based automation

Follows logic a person wrote: if this, then that. Nothing learned, nothing invented. Transparent and reliable in narrow conditions, but brittle when reality changes.

See also: Knowledge representation and reasoning · Predictive AI / traditional ML

S

SaaS (Software as a Service)

Software you access through a browser. The vendor hosts it, maintains it, and updates it. Convenient, but your data leaves your building, it requires internet, and if the vendor's servers go down, your tool goes with them. The most common deployment model for EM AI products.

See also: On-premises · API · Connectivity

SOC 2

Service Organization Control Type 2. A cybersecurity certification that verifies security controls are in place and have been audited — covering data access, encryption, incident response, and monitoring. Meaningful, but not designed to verify AI-specific risks.

See also: ISO 27001 · FedRAMP · ISO/IEC 42001

Stacking / stacking problem

Assembling multiple products to cover portions of one task area's full scope. This is the norm, not a workaround. Costs include multiple procurements, products that don't communicate, staff learning multiple systems, accumulating license costs, and increased vendor lock-in risk.

See also: Fragmentation · Parent product · Vendor lock-in

Stakes tier

How much human oversight an AI-assisted task needs: low (human on the loop — spot-check and sample), medium (human in the loop — review every output), or high (human in command — a qualified person reasons and decides). Reach and consequence set the tier — not the topic, and not how routine the task feels.

In EM: Anything the public acts on for safety starts at high.

See also: Human on the loop · Human in the loop · Human in command

Stochasticity

The built-in randomness that lets the same prompt produce different answers. The model usually picks the most likely next word, but sometimes picks the 2nd or 3rd on purpose to sound natural. Great for generating options; risky when you need consistency. Also called randomness.

See also: Temperature · Variability · Randomness

Supervised learning

A machine learning approach where the model learns from labeled examples — data that has been tagged with the correct answer.

In EM: Used in FEMA applications such as predicting Public Assistance applicant quantities.

See also: Machine learning / predictive modeling · Training data

Sycophancy

The model's tilt toward telling you what you want to hear — agreeing, over-praising, or caving when you push back. You get agreement instead of truth; flawed premises go unchecked.

See also: Hallucination · A-C-R-S

T

Temperature

A “creativity dial” for model outputs. Higher = more varied and surprising; lower = more predictable. You rarely set it directly in a chatbot, but it's why “be creative” and “be precise” prompts behave differently.

See also: Stochasticity · Model size

Token

The small chunk of text the model reads and writes — a short word, part of a word, or punctuation. Roughly ¾ of a word each (~4 characters). Drives cost, speed, and context limits.

See also: Context window · Model size

Tool / add-on

An extra capability bolted onto a model or a parent product: web search, file analysis, a calculator/code tool, image or voice, memory, connectors. Some add-ons cannot function without a parent product. LLMs on their own struggle with arithmetic — be wary of trusting one to calculate without a math tool.

See also: Parent product · Web search / tools · Agent

Training cutoff

The date the model's built-in knowledge stops. It has no inherent knowledge of anything after that date.

In EM: It may quote outdated policies or plans with full confidence.

See also: Pretrained knowledge · Grounding · Web search / tools

Training data

The examples used to teach an AI model patterns. The quality of predictions depends entirely on the training examples and whether new cases resemble them closely enough to generalize.

See also: Bias · Supervised learning · Neural network

V

Variability

A limitation where the same input produces different outputs.

In EM: Two staff members running the same damage description through an AI tool and getting different resource estimates.

See also: Stochasticity · Temperature

Vendor lock-in

When integration effort or ecosystem dependencies make it costly or difficult to switch from one vendor to another — like buying tools that all use the same battery platform. The convenience of staying within one ecosystem means accepting inferior options when that vendor doesn't make what you need.

See also: Stacking / stacking problem · Parent product · Integration requirement

W

Web search / tools

Features that allow some models to fetch live information or run calculators, code, and similar. Shrinks the training-cutoff problem and can cite sources — but you must still verify the summary.

See also: Training cutoff · Tool / add-on · A-C-R-S

Weights

The strengths of the connections between layers in a neural network. Training adjusts the weights so outputs improve.

See also: Neural network · Training data

Z

Zero-shot prompting

Describing the task with no examples — "just ask." Effective when the task is common: press releases, summaries, plain-language rewrites, notes into action items.

See also: Few-shot prompting · Chain-of-thought prompting · Prompt engineering

> Some terms appear across multiple modules with slightly different emphases. The definitions above are synthesized to reflect the fullest description available across all source materials.