

ELECTRIC COOPERATIVE DIRECTORS RESOURCE

AI & Machine Learning

Jargon Reference Guide

A plain-English guide to common AI and machine learning terms.



Highley Holdings, LLC
www.highley.ai

Version 1.2 | October 2026

Purpose

This guide is a quick reference for electric cooperative directors and employees who want plain-English explanations of common artificial intelligence (AI) and machine learning (ML) terms.

Every definition is written in plain english. Each term also includes a short "Why it matters" note that connects the concept back to cooperative governance and operations, plus a link to a free encyclopedia article for readers who want to go deeper. Some items include context relevant examples.

How to Use This Guide

Terms are grouped into six categories, from basic concepts through AI agents and governance risks. Use the table of contents to jump to a category, or the alphabetical index at the back to find a specific term. This guide is a starting point, not a technical manual or legal reference.

Notes on Links

Reference links point primarily to Wikipedia, a free and openly editable encyclopedia, rather than to any single company's website. This keeps the guide vendor-neutral and useful as a general education resource.

What Is Not Included

This draft focuses on the most common terms directors are likely to encounter. It does not attempt to cover every AI or data-science term, specific vendor products, or legal/regulatory detail. Suggestions for additions or corrections are welcome.

Questions or suggestions: gary@highley.ai

Table of Contents

[The Basics](#) (5 terms)

[AI Models & Language Tools](#) (6 terms)

[How AI Learns and Trains](#) (7 terms)

[AI Agents and Agentic AI](#) (10 terms)

[Data & Infrastructure](#) (7 terms)

[Risks & Oversight](#) (7 terms)

[Alphabetical Index](#)

The Basics

Artificial Intelligence (AI)

Computer systems built to do tasks that normally need human thinking, such as understanding language, spotting patterns, or making decisions.

Why it matters: AI already touches many aspects of electric cooperatives, outage prediction, member service, and grid operations, but each may require a different approach or tool.

[Learn more →](#)

Machine Learning (ML)

A type of AI where the computer learns from examples instead of being told an exact rule for every situation.

Why it matters: ML has been around for many years and powers everyday tools, from spam filters to load forecasting and outage prediction.

[Learn more →](#)

Algorithm

A step-by-step set of instructions a computer follows to solve a problem or make a decision.

Why it matters: Every AI tool runs on an algorithm; understanding the various algorithms helps us ask better questions about how a tool actually works.

[Learn more →](#)

Deep Learning

A more advanced form of machine learning that uses many layered steps, loosely modeled on the human brain, to find complex patterns in data.

Why it matters: Deep learning is behind most recent AI breakthroughs, including image recognition and chatbots.

[Learn more →](#)

Neural Network

A computer model made of connected pieces that pass signals to one another, loosely the way brain cells do, to recognize patterns.

Why it matters: Neural networks are the building blocks behind deep learning and most modern AI tools.

[Learn more →](#)

AI Models & Language Tools

Natural Language Processing (NLP)

The branch of AI focused on helping computers understand and generate human language.

Why it matters: *NLP is what lets a chatbot read a member's question and answer back in plain English.*

[Learn more →](#)

Large Language Model (LLM)

A very large AI system trained on huge amounts of text so it can understand and produce human-like writing.

Why it matters: *LLMs power tools like ChatGPT, Claude, Gemini, Grok, CoPilot, and others and will likely show up in co-op chatbots, internal search, and document review.*

[Learn more →](#)

Generative AI (Gen AI)

AI that creates brand-new content, based on it's training, such as text, images, or summaries, instead of only analyzing existing data.

Why it matters: *Gen AI can draft documents, policies or summarize board packets, but its output always needs a human check for accuracy.*

[Learn more →](#)

Foundation Model

A large, general-purpose AI model trained on broad data that can then be adapted for many specific tasks.

Why it matters: *Most commercial AI tools a co-op might buy are built on top of a foundation model, not created from scratch.*

[Learn more →](#)

Transformer

The technical design behind most modern language AI. It lets a model weigh which words in a sentence matter most to understand meaning.

Why it matters: *You don't need the math, but "transformer" is a term you will see often in vendor descriptions of AI products. Think of a Transformer as a type of machine learning algorithm.*

[Learn more →](#)

Chatbot

A computer program designed to hold a conversation with a person, usually through typed or spoken text.

Why it matters: *Many co-ops are piloting chatbots for member service and technical support.*

[Learn more →](#)

How AI Learns and Trains

Supervised Learning

Training an AI by giving it many examples that are already labeled with the correct answer.

Why it matters: *This is the most common, predictable training method, used for tasks like fraud detection.*

Example: Training a model on historical meter data already labeled “confirmed theft” vs. “normal usage” so it can flag likely energy diversion cases for investigation.

[Learn more →](#)

Unsupervised Learning

Training an AI on data with no labels, so it finds patterns and groupings on its own.

Why it matters: *Useful for spotting unusual patterns or patterns difficult for the human brain to uncover in large data, such as unexpected outage clusters, without telling the system what to look for.*

Example: Feeding years of outage management system (OMS) data into a model with no predefined categories, letting it discover on its own that certain equipment ages combined with specific weather patterns or animal migration tend to precede failures.

[Learn more →](#)

Self-Supervised Learning

A training method where the AI generates its own practice questions from raw data, essentially teaching itself without people labeling anything.

Why it matters: *This is how most large language models learn from vast amounts of text; it needs huge data but little human labeling effort.*

Example: This step is usually done by the AI vendor, not the co-op. The AI reads years of unlabeled maintenance logs and manuals and teaches itself the language of utility terminology, with no human labeling involved.

[Learn more →](#)

Reinforcement Learning

A training method where an AI learns by trial and error, earning a reward for good decisions and a penalty for bad ones, similar to training a pet.

Why it matters: *Reinforcement learning fine-tunes AI behavior and is central to robotics and optimization tools.*

Example: A system that adjusts voltage regulation or demand-response scheduling, testing different settings and earning a reward when it reduces losses or cost, gradually learning the best strategy through trial and error.

[Learn more →](#)

Reinforcement Learning from Human Feedback (RLHF)

A training step where real people rate an AI's answers, and the model adjusts to give more of the answers people prefer.

Why it matters: *RLHF is a big reason modern chatbots sound polite and helpful instead of raw and unfiltered.*

Example: Staff rate a chatbot's draft answers to member questions ("too technical," "too pushy"), and the vendor uses those ratings to shape future responses to be more appropriate in tone.

[Learn more →](#)

Fine-Tuning

Taking an already-trained AI model and giving it extra, focused training on a smaller, specific set of data.

Why it matters: *A co-op might fine-tune a model on its own data, but this carries data-security risks unless done in a secure environment.*

Example: Taking a general-purpose chatbot and giving it focused training on your co-op's specific broadband packages, pricing, and sales approach, so a broadband sales agent chatbot sells knowledgeably instead of generically.

[Learn more →](#)

Prompt Engineering

The skill of writing clear, well-structured instructions or questions to get better answers from an AI tool.

Why it matters: *A practical staff-training topic; well-crafted prompts noticeably improve output quality and reduce mistakes.*

Example: Asking "Summarize this 40-page rate case filing into 3 bullet points focused on member rate impact" instead of just "summarize this," for a far more useful result with no retraining needed.

[Learn more →](#)

AI Agents and Agentic AI

AI Agent

A software program that can sense information, make decisions, and take actions on its own to reach a goal, with limited step-by-step direction from a person.

Why it matters: *AI agents go beyond chatbots. They can look things up, use other software tools, and complete multi-step tasks.*

[Learn more →](#)

Agentic AI

A newer category of AI systems that act as agents: pursuing goals, making a series of decisions, and adapting their actions over time with little ongoing human input.

Why it matters: *This is the direction AI is heading. Boards should ask what decisions an agentic system may make without a human sign-off.*

[Learn more →](#)

Context Window

The amount of text or information an AI model can "see" and consider at one time when forming a response.

Why it matters: A larger context window lets an AI review a longer document or conversation at once; a small one means it may lose track of earlier details.

[Learn more →](#)

Chunking

Breaking a large document into smaller pieces before an AI stores or searches it, since most AI tools can only work with one bite-sized piece at a time.

Why it matters: This happens behind the scenes whenever a co-op's policies, contracts, or manuals are loaded into an AI search tool; how the pieces are split affects how well the AI can answer questions about them.

[Learn more →](#)

Memory (AI Agent Memory)

The ability of an AI system to remember information from earlier in a task, or across separate sessions, instead of starting fresh every time.

Why it matters: Memory lets an agent build on past interactions, but it also means sensitive information could be retained and must be governed carefully.

[Learn more →](#)

Connector

A tool or standard that lets an AI system securely reach outside data and software, such as email, a database, or a scheduling system.

Why it matters: Connectors are how AI agents get real work done; each one is also a potential access point that needs security review.

[Learn more →](#)

Model Context Protocol (MCP)

An open, shared set of rules that lets an AI system connect to outside tools, files, and data in a consistent way, instead of needing a custom-built connection for every app.

Why it matters: MCP is quickly becoming the common "plug" that AI agents use to reach a co-op's other systems, such as billing, GIS, or outage management; directors should ask what a vendor's MCP connections can access and how each one is secured.

[Learn more →](#)

A2A (Agent-to-Agent Protocol)

An open, shared set of rules that lets AI agents built by different companies find each other, talk, share information, and hand off tasks. Google started it in 2025, and it is now managed by the nonprofit Linux Foundation so no single company controls it.

Why it matters: MCP connects an agent to tools and data; A2A connects an agent to other agents. As vendors' agents start working together, directors should ask which outside agents a co-op's agent may talk to, what information it may share, and what it may agree to without a person signing off.

Example: During a major storm, the co-op's outage agent sends a work request straight to a mutual-aid contractor's scheduling agent and gets back a crew arrival time, much like crews from different utilities sharing a common radio channel, with no staff re-typing the details.

[Learn more →](#)

In-Context Learning

When an AI model adjusts its answers based only on examples or instructions given within the current conversation, without being retrained.

Why it matters: This is why giving an AI good examples in your prompt can improve results immediately, with no technical setup.

[Learn more →](#)

Retrieval-Augmented Generation (RAG)

A method where an AI looks up relevant information from an approved set of documents before writing its answer.

Why it matters: RAG is one of the safest ways for a co-op to let AI answer questions using its own policies or manuals without retraining a model.

[Learn more →](#)

Data & Infrastructure

Cloud Computing

Using computing power, storage, and software delivered over the internet from a remote provider instead of on local equipment.

Why it matters: Most AI tools run in the cloud; directors and staff should understand where co-op data physically goes when employees use these tools.

[Learn more →](#)

On-Premises AI (Self-Hosted)

Running an AI model on the co-op's own servers and equipment instead of sending data to an outside company over the internet.

Why it matters: Keeps sensitive data fully under the co-op's control, but requires the co-op to own the cost, hardware, and expertise to run and maintain it.

[Learn more →](#)

Open-Source vs. Proprietary Model

Whether an AI model's inner workings and settings are made publicly available for anyone to inspect and run, or kept locked inside a single company's product.

Why it matters: Affects how much a co-op can verify what a model is actually doing, and whether it could be run in-house instead of relying on one vendor.

[Learn more →](#)

Vector Database

A specialized database that stores information in a mathematical format so an AI can quickly find related content by meaning, not just by matching keywords.

Why it matters: Vector databases are the engine behind RAG and internal AI search tools.

[Learn more →](#)

API (Application Programming Interface)

A defined way for two different pieces of software to talk to each other and share data.

Why it matters: APIs are how an AI tool connects to a co-op's other systems, such as billing or outage management.

[Learn more →](#)

GPU (Graphics Processing Unit)

A computer chip built to handle many calculations at once, which makes it especially good at running AI models.

Why it matters: GPUs are why AI computing costs more than typical office computing; this affects vendor pricing and infrastructure decisions.

[Learn more →](#)

Data Governance

The rules and processes an organization uses to decide how data is collected, stored, protected, and who may access it.

Why it matters: Strong data governance should be in place before adopting any AI tool that touches member or operational data.

[Learn more →](#)

Risks & Oversight

Hallucination

When an AI confidently gives an answer that is factually wrong or entirely made up.

Why it matters: Directors and staff should insist that employees treat AI as an assistant, not a final authority, and verify important outputs.

[Learn more →](#)

Algorithmic Bias

When an AI system produces unfair or skewed results because of patterns in the data it was trained on.

Why it matters: Co-ops must watch for biased outcomes in areas like hiring, member communications, storm restoration priority, or program eligibility.

Example: A model used to help prioritize storm restoration order could unintentionally deprioritize certain neighborhoods if those areas were historically restored last, baking a past inequity into future decisions instead of correcting it.

[Learn more →](#)

Explainable AI

AI that is designed or evaluated so that people can understand and trust how it reached a decision.

Why it matters: For high-stakes decisions, boards should ask whether a system's reasoning can be explained, not just its final answer.

[Learn more →](#)

Data Privacy

The right and practice of controlling how personal information is collected, used, and shared.

Why it matters: AI tools often send data to outside providers; directors should know what member data leaves the co-op's control and how it is protected.

[Learn more →](#)

Data Residency

The requirement or expectation that data stay stored within a specific location or country.

Why it matters: Matters when choosing an AI vendor, since some cloud AI services process or store data outside the United States.

[Learn more →](#)

Prompt Injection

A security trick where hidden text in a document, email, or webpage tricks an AI into ignoring its instructions and doing something it shouldn't.

Why it matters: Becomes a real risk the moment an AI agent starts reading member emails or outside documents on the co-op's behalf; a malicious file could hijack the AI's behavior.

[Learn more →](#)

Shadow AI

Employees using AI tools on their own, without IT's knowledge or approval, such as pasting sensitive co-op data into a free public chatbot.

Why it matters: A fast-growing, easy-to-overlook data security risk; boards should ask whether the co-op has a clear AI use policy and an approved tool list.

[Learn more →](#)

Alphabetical Index

Click any term to jump to its definition.

[A2A \(Agent-to-Agent Protocol\)](#) (AI Agents and Agentic AI)
[Agentic AI](#) (AI Agents and Agentic AI)
[AI Agent](#) (AI Agents and Agentic AI)
[Algorithm](#) (The Basics)
[Algorithmic Bias](#) (Risks & Oversight)
[API \(Application Programming Interface\)](#) (Data & Infrastructure)
[Artificial Intelligence \(AI\)](#) (The Basics)
[Chatbot](#) (AI Models & Language Tools)
[Chunking](#) (AI Agents and Agentic AI)
[Cloud Computing](#) (Data & Infrastructure)
[Connector](#) (AI Agents and Agentic AI)
[Context Window](#) (AI Agents and Agentic AI)
[Data Governance](#) (Data & Infrastructure)
[Data Privacy](#) (Risks & Oversight)
[Data Residency](#) (Risks & Oversight)
[Deep Learning](#) (The Basics)
[Explainable AI](#) (Risks & Oversight)
[Fine-Tuning](#) (How AI Learns and Trains)
[Foundation Model](#) (AI Models & Language Tools)
[Generative AI \(Gen AI\)](#) (AI Models & Language Tools)
[GPU \(Graphics Processing Unit\)](#) (Data & Infrastructure)
[Hallucination](#) (Risks & Oversight)
[In-Context Learning](#) (AI Agents and Agentic AI)
[Large Language Model \(LLM\)](#) (AI Models & Language Tools)
[Machine Learning \(ML\)](#) (The Basics)
[Memory \(AI Agent Memory\)](#) (AI Agents and Agentic AI)
[Model Context Protocol \(MCP\)](#) (AI Agents and Agentic AI)
[Natural Language Processing \(NLP\)](#) (AI Models & Language Tools)
[Neural Network](#) (The Basics)
[On-Premises AI \(Self-Hosted\)](#) (Data & Infrastructure)
[Open-Source vs. Proprietary Model](#) (Data & Infrastructure)
[Prompt Engineering](#) (How AI Learns and Trains)
[Prompt Injection](#) (Risks & Oversight)
[Reinforcement Learning](#) (How AI Learns and Trains)
[Reinforcement Learning from Human Feedback \(RLHF\)](#) (How AI Learns and Trains)
[Retrieval-Augmented Generation \(RAG\)](#) (AI Agents and Agentic AI)
[Self-Supervised Learning](#) (How AI Learns and Trains)
[Shadow AI](#) (Risks & Oversight)
[Supervised Learning](#) (How AI Learns and Trains)
[Transformer](#) (AI Models & Language Tools)
[Unsupervised Learning](#) (How AI Learns and Trains)
[Vector Database](#) (Data & Infrastructure)