

INTELLIGENCE UPDATE

AI in data: sorting reality from hallucination



Dr. Rand Talib

4 Dec 2025

Many people do not use the term artificial intelligence correctly: vendors, investors, and even some operators label everything from basic automation scripts to deep learning controllers as AI. This inflation of the term has commercial and strategic motives. AI branding helps attract funding, creates differentiation in the market, and positions traditional analytics as cutting-edge solutions.

However, this broad usage also breeds confusion and skepticism. Data center operators, uncertain about the level of autonomy or risk they face, often hesitate to implement even safe, deterministic systems.

Many operators remain hesitant to implement AI in their data centers, often citing fears of hallucination — the risk that an AI system might generate false or invented information. Yet not all AI behaves this way, and the term is frequently misapplied. By clarifying the different types of AI, how they vary in capability and reliability, and which pose genuine hallucination risks, operators can better distinguish dependable automation from the marketing-driven “AI washing” that fuels confusion and obscures real risk.

The spectrum of AI in data centers

AI in data centers spans a broad continuum, from deterministic, data-driven algorithms to advanced systems capable of adaptive or autonomous decision-making. Treating these technologies as a single category obscures important differences in capability, reliability and operational risk. Understanding this spectrum is critical for evaluating what each system can — and cannot — safely automate.

Table 1 compares the different types of AI used in modern data centers.

Table 1 AI types used in data centers

AI type	Core role in data centers	Key applications	Autonomy level
Rule-based/ symbolic AI	Automated decision-making using explicit rules.	Optimizing cooling and power systems, enforcing operational policies, safety interlocks.	Moderate (within strict boundaries)
Machine learning (ML)	Learns statistical patterns from operational data.	Predictive maintenance, forecasting, anomaly detection, energy and workload optimization.	Low/moderate (assisted)
Deep learning/ neural network	Detects complex, nonlinear patterns across sensors and telemetry.	Vision analytics, thermal pattern recognition, security anomaly detection.	Low/moderate (assisted)
Large language model (LLM)	Interprets and generates natural language from logs, documents and tickets.	Chatbots, log summarization, incident triage support, documentation.	Low (advisory)
Generative AI	Produces new content or synthetic data based on learned distributions.	Report drafting, configuration templates, synthetic data creation.	Low (advisory)
Neurosymbolic AI	Combines ML with symbolic reasoning for transparent and rule-aligned decisions.	Compliance support, explainable recommendations, energy/cost trade-off modeling.	Low/moderate (decision-support)
Neuromorphic AI	Uses brain-inspired computing for real-time, low-power inference.	Edge monitoring, local fault detection in dark or remote sites.	Low (embedded inference)
Agentic AI	Autonomous, goal-driven orchestration and decision-making.	Workload balancing, self-healing systems, adaptive resource management.	High (proactive)

Mislabeling AI: the root of hallucination fears

Across the tech sector, and within data center operations in particular, everything from basic regression models to large transformer networks is labeled as AI. This conflation blurs the operational reality:

Predictive and optimization models (ML, neural networks) rely on measurable data and statistical learning. They rarely improvise.

Generative and language models (LLMs, GenAI) produce content probabilistically, often without grounding in external data, which creates a risk of fabrication.

Agentic AI orchestrates systems and can call on other models (including LLMs) to plan or communicate, but its reliability depends on which components it uses.

This terminological blur feeds operator anxiety. A predictive control loop that tunes chillers based on real-time feedback is not at risk of hallucination, yet many operators equate it with the

behavior of chatbots and generative systems. In practice, hallucination is a property of generative AI, not of deterministic automation or data-driven control.

Understanding which AI types can hallucinate, and why, is essential for evaluating their operational reliability. **Table 2** below clarifies the differences across major AI categories used in data centers.

Table 2 Hallucination behavior and risks across AI types

Model type	Can Hallucinate?	Reason/behavior on data	Example of problematic behavior in data centers
Rule-based/ symbolic AI	No	Deterministic logic grounded in explicitly defined rules and conditions; does not generate novel semantic content.	May respond incorrectly or fail safe when conditions fall outside its ruleset, but does not invent new states or explanations.
Machine learning (ML)	Rarely	Finds patterns statistically from operational data; can misclassify points or anomalies but does not invent semantic content.	Vulnerable to bad or shifted data (e.g., mislabels a developing fault as normal), potentially delaying maintenance.
Deep learning/ neural network (predictive)	Rarely	Can misclassify data points or anomalies, but does not invent semantic content.	Can misidentify a vibration or thermal anomaly, but does not invent non-existent equipment or sensors.
Large language model (LLM)	Yes	Predicts next tokens statistically without firm factual grounding; may generate plausible but false statements.	IT copilots can misreport system status or create fictitious error explanations or remediation steps.
Generative AI (non-LLM, e.g., template/text/ image/code generators)	Yes	Produces synthetic outputs from pattern completion rather than verified operational data.	Auto-generated maintenance reports or configuration templates can contain fabricated details.
Neurosymbolic AI	Very rarely	Combines neural networks with symbolic AI to add structure and constraints; residual errors can arise from mismatches between neural output and symbolic rules.	In a constrained decision-support system, an edge case produces an internally inconsistent recommendation that conflicts with encoded policies.
Neuromorphic AI	Occasionally	Performs low-power, real-time inference; may “fill in” or generalize from training data when faced with unfamiliar signals, but typically at the level of misclassification rather than rich semantic hallucination.	At the edge, misinterprets a novel sensor pattern and under- or overreports a local fault condition.
Agentic AI	Indirectly	Orchestration layer that can inherit hallucinations or errors from embedded LLM/generative components used for reasoning or communication.	An autonomous orchestration agent executes an unnecessary or sub-optimal workflow after being misinformed by its LLM-based advisory module.

Managing and mitigating hallucination risks

Operators can apply a focused set of safeguards that keep AI useful while limiting unsafe or fabricated outputs:

- Constrain generative models to verified, domain-specific sources such as

maintenance manuals, runbooks, building management system (BMS)/data center infrastructure management (DCIM) logs, incident records and approved knowledge articles.

- Use retrieval-augmented generation (RAG) so that models base responses on current operational data rather than general training alone.
- Adopt hybrid architectures that pair LLM copilots with deterministic rule engines or physics-based digital twins, which can verify or veto proposed actions before they affect live systems.
- Require human-in-the-loop validation before AI can change configurations, control physical systems, or execute high-impact runbook steps.
- Establish clear governance that makes a distinction between “assistive AI” (documentation, recommendations, analysis) and “operational AI” (any system that can directly change configurations or physical infrastructure).
- Apply strict scoping and access control so more powerful generative or agentic components start in read-only or advisory modes and follow least-privilege principles for credentials and APIs.

The Uptime Intelligence View

Much of the data center industry’s caution around AI appears to come from treating it as a single, generative technology rather than a stack of distinct capabilities. In real deployments, predictive models are typically aligned with control and optimization tasks; emerging agentic approaches support orchestrated, multi-step decision flows; and LLMs or other generative systems are best suited for documentation, reasoning support, and advisory use under governance constraints. When these distinctions are made explicit, AI can be a potential enabler of resilient, self-optimizing facilities and poses less risk of becoming a direct threat to uptime.

ABOUT THE AUTHOR



Dr. Rand Talib

Dr. Rand Talib is a Research Analyst at Uptime Institute with expertise in energy analysis, building performance modeling, and sustainability. Dr. Talib holds a Ph.D. in Civil Engineering with a concentration in building systems and energy efficiency. Her background blends academic research and real-world consulting, with a strong foundation in machine learning, energy audits, and high-performance infrastructure systems.

rtalib@uptimeinstitute.com

About Uptime Institute

Uptime Institute is the Global Digital Infrastructure Authority. Its Tier Standard is the IT industry's most trusted and adopted global standard for the proper design, construction, and operation of data centers – the backbone of the digital economy. For over 25 years, the company has served as the standard for data center reliability, sustainability, and efficiency, providing customers assurance that their digital infrastructure can perform at a level that is consistent with their business needs across a wide array of operating conditions.

With its data center Tier Standard & Certifications, Management & Operations reviews, broad range of related risk and performance assessments, and accredited educational curriculum completed by over 10,000 data center professionals, Uptime Institute has helped thousands of companies, in over 100 countries to optimize critical IT assets while managing costs, resources, and efficiency.