

Use of AI in Chemical Manufacturing

Artificial intelligence (AI) seems to be everywhere in industry, but the nuances of its application are not generally understood completely. Presented here are descriptions of the different subtypes of AI and their main applications currently in chemical processing

The term “artificial intelligence” (AI) has become the buzzword of the year in the chemical process industries (CPI), along with almost every other sector of industry and daily life. According to IoT Analytics, over 39% of CEOs mentioned AI during their Q1 2025 earnings calls, second only to the word “tariffs” [1]. Ever since OpenAI (San Francisco, Calif.; www.openai.com) unveiled its large language model ChatGPT in November 2022, attention on AI tools, along with their possibilities for raising productivity and spurring innovation, has grown dramatically.

But while awareness of AI is at an all-time high, a satisfying definition for the term remains elusive. For now, we will refrain from asking ChatGPT for a definition of AI, and instead refer to Webster’s Dictionary, which defines AI as follows:

“1. The capability of computer systems or algorithms to imitate intelligent human behavior. 2. A branch of computer science dealing with the simulation of intelligent human behavior by computers [2].”

This definition seems too broad and too generic to be of any great use, and makes it difficult to definitively determine what falls under this umbrella and what does not. This difficulty was famously called “the AI effect” by the American computer scientist and AI pioneer John McCarthy. The AI effect is an attempt to capture the idea that as soon as an AI system becomes proficient at a task or function, that task or function is often no longer viewed as being part of AI, and the goalposts for what qualifies as “AI” are moved accordingly [3]. The Australian robotics scientist and Massachusetts Institute of Technology (MIT; Cambridge, Mass.; www.mit.edu) professor Rodney Brooks states the idea as this: “Every time we figure out a piece of [AI], it stops being magical [and] we say, ‘Oh, that’s just a computation’” [4].

This article is designed to provide a description of the different AI technologies currently being employed in chemical manufacturing processes, and an overview of the applications for which they are being used.

**Jonathan
Alexander**
Albemarle Corp.

IN BRIEF

AI: A BRIEF HISTORY

THREE CATEGORIES
OF AI

TYPES OF MACHINE
LEARNING

CHOOSING ML
ALGORITHMS

PRACTICAL
APPLICATIONS OF ML

GENERATIVE AI IN
MANUFACTURING

FUTURE OF AI IN THE CPI




FIGURE 1. The roots of AI trace back centuries, to the first attempts at automated computation (photo: Wikimedia Commons)

AI: a brief history

Beliefs surrounding the “AI Effect” landscape can vary, but some suggest that the journey of AI within the manufacturing sector is rooted in centuries of innovation. In 1642, French mathematician and physicist Blaise Pascal built a mechanical calculator to automate arithmetic. Because the term “AI” wouldn’t be coined until a few centuries later, some may have thought of Pascal’s system as black magic. Just under 200 years later, the English polymath and inventor Charles Babbage designed the “Analytical Engine,” a programmable mechanical computer that predated digital electronics by over a century (Figure 1) [5]. These and other developments

laid the groundwork for the computational logic that would eventually underpin AI.

In more modern times, British mathematician Alan Turing, considered to be the father of AI, posed a profound question. To open his 1950 paper “Computing Machinery and Intelligence,” Turing wrote “I propose to consider the question . . . can machines think?” He then led the development of what is now called the “Turing Test,” in which a human assessor examines a written dialogue between a human and a machine, with the objective of identifying which participant is the machine. A machine is deemed to have “passed” the test if the human evaluator is unable to reliably distinguish its responses from those of the human.

The field of AI has been evolving at an accelerating pace ever since. The 1960s saw the development of ELIZA, an early natural language processing program, 50 years before the introduction of Amazon’s Alexa or Apple’s Siri.

AI has made a mark on pop culture as well. Widely seen films like *Star Wars* and *Back to the Future* brought AI ideas into the public imagination. But unfortunately, the world found itself losing hope when personal communication droids like C-3PO and time travel failed to materialize. The period from the 1970s to the 1990s is referred to as “the AI Winters,” and was characterized by a massive reduction in research and public interest in artificial intelligence [6].

The late 20th century witnessed significant milestones, as computers and technology were finally able to live up to people’s expectations. IBM’s (Armonk, N.Y.; www.ibm.com) Deep Blue supercomputer defeated chess champion Garry Kasparov in 1997, and IBM’s AI tool Watson won the American gameshow *Jeopardy!* in 2011. The introduction of consumer products like Roomba and iPhone further demonstrated AI’s practical applications. This brings us to the current time, where generative AI is dramatically surpassing most people’s wildest imaginations, and AI is garnering a great deal of funding. Significant advancements in AI have recently come from OpenAI, Google, Microsoft, Anthropic, xAI, Deepseek

and a host of other companies and research organizations.

Three categories of AI

While generative AI and large language models (LLMs) dominate the headlines today, they represent only the most visible tip of a much broader and more mature spectrum of AI capabilities. The foundation of modern AI is built upon several interrelated subfields, each contributing critical functionality to intelligent systems used across industries. The functions include predictive maintenance in manufacturing to autonomous vehicles and personalized healthcare (Figure 2). Here are descriptions of the main categories of AI:

Machine learning (ML). At its core, machine learning is the science of enabling computers to learn from data and improve their performance over time without being explicitly programmed for every task. ML encompasses a wide variety of algorithms that can recognize patterns, classify data, predict outcomes or cluster information into meaningful groups. In manufacturing, ML powers everything from yield prediction and process optimization to supply-chain forecasting and energy management.

Artificial neural networks (ANNs). Artificial neural networks are a class of ML models loosely inspired by the structure of the human brain. Composed of interconnected “neurons” or nodes arranged in layers, ANNs can model complex, nonlinear relationships within large datasets. They are particularly effective in tasks like pattern recognition, sensor fusion and fault detection.

Deep learning (DL). Deep learning is a specialized subset of machine learning that utilizes neural networks with multiple hidden layers, giving rise to the term “deep.” These architectures allow DL models to learn hierarchical representations of data, making them exceptionally good at processing high-dimensional inputs, such as images, audio and unstructured text. Deep learning is what powers advanced computer vision systems for quality inspection, autonomous navigation and speech recognition.

Generative AI (GenAI). Generative AI refers to models capable of

creating new content, such as text, images, audio, programming code, or even molecular structures, based on patterns learned from vast training datasets. The most well-known subset of GenAI today is LLMs, such as OpenAI’s GPT, Google’s Gemini and Meta’s LLaMA. These models are trained on massive amounts of text data and can generate human-like responses to prompts, summarize documents, translate languages, write computer code and more. In industrial contexts, GenAI is beginning to transform how operators interact with systems, enabling natural-language interfaces for data queries, generating shift reports from operational data, drafting standard operating procedures (SOPs) and even simulating training scenarios.

Types of machine learning

If you tried to catalog every machine-learning algorithm ever developed, you would need an entire library, and maybe a few backup hard drives. There are literally hundreds of distinct models, each with their own quirks, strengths and ideal use cases. The sheer variety of ML algorithms can feel overwhelming, especially for those starting out in the field.

Machine learning is typically organized into three core categories (Figure 3). Think of these not as competing techniques, but as different tools in an AI toolbox, with each area suited to a particular type of problem. The categories are described here:

Unsupervised learning. Unsupervised learning is targeted toward discovering hidden patterns in data, without any labeled outcomes to guide the way. It is considered “unsupervised” because the model is not told what to look for. Instead, the ML model sifts through data to find natural groupings, anomalies or correlations on its own. This is used for pattern recognition, outlier and anomaly detection, and finding shifts and clusters in operating data. The key takeaway here is that there is no “prediction” — the model is simply identifying complex correlations and patterns in data.

Supervised learning. In contrast, supervised learning is designed to make predictions based on histori-

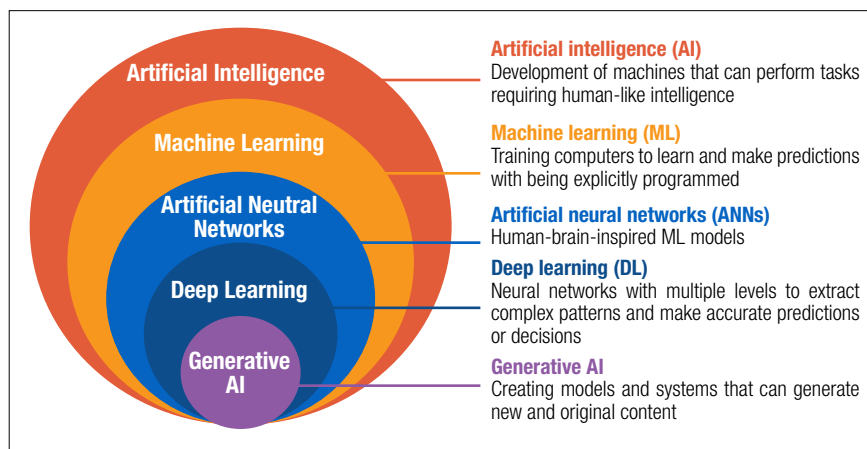


FIGURE 2. Modern AI is built upon a several interrelated subfields, each contributing critical functionality to intelligent systems used in industrial operations

cal outcomes. These models train on labeled data, where we already know the answer, and learn to generalize those patterns to new situations. If you've ever built a regression model with a Y variable, then you've already dabbled in supervised learning. In fact, regression models use the same logic behind most image-recognition systems. Feed the model thousands of photos labeled "cat" or "dog," for example, and it learns to classify future images with surprising accuracy. Supervised learning is a powerful ML technique, but it is only as good as the data-labels fed into it.

Reinforcement learning. Reinforcement learning (RL) takes a different approach. RL models learn by trial and error, receiving feedback (called "rewards") as they explore and act within an environment. Over time, the models learn which actions maximize their cumulative reward. It is analogous to training a dog, only with math instead of treats. RL is most famously used in strategy board games like chess and the Chinese game go, where the model simulates millions of matches, learning from both its wins and losses. Increasingly, RL is making its way into industrial settings, where it is being applied to dynamic control problems, optimization tasks and autonomous operations.

In essence, unsupervised learning finds hidden structure, while supervised learning learns from labeled examples, and reinforcement learning learns by doing. Together, the three forms of ML constitute the backbone of modern AI systems.

Choosing ML algorithms

With hundreds of machine-learning algorithms to choose from, and nearly all of them having made at least a cameo appearance in chemical manufacturing, the real question is not as much about whether the algorithms work, but which ones consistently deliver value. Over time, a handful of workhorse algorithms have emerged as being especially well-suited to handling the messy, high-dimensional, sensor-rich world of process manufacturing. The four algorithms described here are the ones on which the author has relied when building industrial AI solutions in the past. These four have all shown themselves to be effective across thousands of real-world use cases.

Principal component analysis (PCA). PCA is an unsupervised learning algorithm used for dimensionality reduction and clustering. In plain terms, PCA takes complex, high-dimensional process data and transforms it into a smaller set of orthogonal components, each capturing a slice of the variance in the data. What does that mean in practice? It means PCA helps users see patterns and correlations that would likely not be noticed by just staring at raw variables. PCA has been successful in the following functions:

- Fault detection
- Golden batch analysis
- Real-time process monitoring

If this author was forced to make due with only one ML algorithm, I would choose PCA.

Partial least squares (PLS). PLS can be thought of as PCA's more

predictive sibling. While PCA is geared toward uncovering structure among data, PLS uses that same mathematical foundation to model relationships between your input variables (X) and a target output (Y).

PLS shines in situations where data are noisy, or where variables are highly collinear, which basically describes data from every chemical plant, ever. PLS is especially useful in the following areas:

- Chemometrics
- Spectroscopy
- Real-time soft sensors for predicting product qualities like viscosity, pH or melt index

Succinctly, PCA shows users what's going on, while PLS tells them what is likely to happen next.

Neural networks. Neural networks are the "Swiss Army knives" of ML. Inspired by the human brain, they are flexible, nonlinear function approximators that thrive in high-dimensional, dynamic environments.

Neural networks are particularly effective for the following tasks:

- Predictive maintenance using equipment sensor data
- Image-based defect detection with convolutional neural networks
- Dynamic system modeling and closed-loop simulations

When used with deep-learning architectures, neural networks can model highly nonlinear, multistep relationships that traditional models would not be able to see.

XGBoost. Extreme gradient boosting (XGBoost) is a workhorse technique for when more complex analysis and predictions are needed, especially with time-series data. XGBoost is a high-performance, supervised-learning algorithm that builds a series of optimized decision trees. XGBoost can be thought of as a more powerful version of the "random forest" algorithm. It is faster, more accurate and often better at teasing out subtle interactions between variables.

XGBoost algorithms excel when used for several tasks in particular. These include yield prediction, batch classification and root-cause analysis from historical process data.

In what is perhaps the best feature of all, XGBoost is highly effective

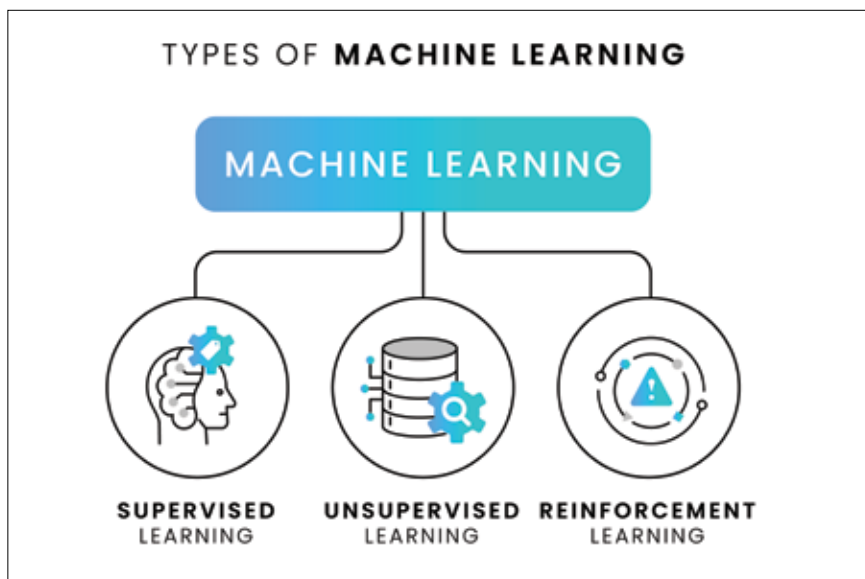


FIGURE 3. The field of machine learning can be divided into different types, each with its own strengths

tive at handling noisy, irregular process data, such as those coming from temperature sensors, pressure transmitters, flowmeters and other process instruments.

Practical applications for ML

ML techniques have moved beyond the research-and-development and proof-of-concept stages to become serious, scalable enablers across the entire chemical manufacturing life-cycle. What used to be niche experiments at chemical manufacturing sites are now systems that drive real, repeatable value, such as boosting uptime, improving product quality, optimizing resource use and advancing sustainability goals. Whether a facility is running batch reactors or continuous distillation columns, AI is now starting to feel less like a buzzword and more like part of the plant's nervous system. The following are nine areas in which ML technologies are now being applied in chemical processing environments.

Open-loop advisory systems. Open-loop advisory systems provide real-time decision support to operators, without directly interfacing with the control layer. These tools analyze historical and live process data using techniques like PCA, clustering and gradient boosting (such as XGBoost), and then surface actionable insights.

In a pulp mill case study, a mill-wide optimization (MWO) tool was run in advisory mode. Operators

reviewed the AI's suggested targets and applied them as setpoints, thereby using the solution as an "open-loop advisory system" to drive optimal production [7].

Modeling of "golden batches."

Golden-batch modeling uses unsupervised learning techniques like PCA, k-means clustering (grouping similar data points into clusters) and dynamic time warping (DTW; comparing unsynchronized temporal sequences) to create a statistical fingerprint of the best production runs at a site. Once that fingerprint is locked in, users can monitor live batch data in real time and catch deviations before they become yield losses or material requiring rework.

In practice, this means comparing each batch's data to the golden-batch standard and adjusting parameters. For instance, experts from Siemens AG (Munich, Germany; www.siemens.com) note that with AI, plants can pinpoint what defines the golden batch and catch negative trends early so that batches can be saved instead of discarded [8].

Predictive maintenance. ML models are proving to be game changers in asset reliability programs, especially when it comes to predicting mechanical or electrical failures using vibration, temperature, pressure and acoustic data.

Algorithms like random forests, support vector machines (SVMs) and deep neural networks can identify

degradation patterns long before humans or traditional rule-based systems would be able to detect. McKinsey & Co. (New York, N.Y.; www.mckinsey.com) reports that AI image recognition and pattern analysis help identify anomalies and predict equipment failures in the field, enabling preventive maintenance and improved uptime [9].

Predictive analytics for quality and yield. Predictive analytics lets users forecast product quality, yield, or emissions hours (or even days) before laboratory testing results come in. Models like PLS, neural networks and XGBoost are used to correlate live process data with critical quality attributes (CQAs).

For example, one petrochemical site built a neural net model that predicted sulfur content with 94% accuracy, allowing operators to adjust blend ratios in real time and avoid off-specification product [10]. Similar tools are now used across the industry for polymer grade prediction, cycle time optimization and compliance forecasting.

Soft sensors and virtual instrumentation. Soft sensors estimate hard-to-measure process parameters using correlated surrogate variables and ML models. These are especially useful when laboratory samples are slow to analyze or online analyzers are too expensive or unreliable.

Common techniques include PLS regression, multivariate adaptive regression splines (MARS), and LSTM (long-short-term memory) neural networks. Researchers developed ML-based soft sensors to estimate key water quality parameters, such as chemical oxidation demand (COD), total suspended solids (TSS) and *E. coli* levels, in real time using easily measurable in-line data from an onsite wastewater treatment system. Despite limited datasets, the models accurately predicted COD ($R^2 = 0.96$) and TSS ($R^2 = 0.99$), demonstrating the potential of soft sensors for cost-effective, real-time performance monitoring [11].

Fault detection and root-cause analysis. Modern ML-based monitoring systems are redefining how process troubleshooting is accomplished. PCA, autoencoders and

decision trees are often deployed in combination with data historian and distributed control system (DCS) data to flag subtle anomalies and uncover complex root causes.

In one case, an unsupervised anomaly-detection model identified a pressure imbalance in a distillation tower. Engineers traced it back to a steam valve drift that had gone unnoticed by standard alarms [12]. ML helped cut down root-cause analysis time, prevented escalation and reduced the risk of repeat failures.

Inventory and supply chain optimization. Beyond the plant floor, ML tools are transforming how raw materials are managed, how production is planned and how supply-chain risks are reduced. Time-series forecasting, LSTM models and gradient boosting are being used to predict demand patterns, raw material usage and procurement lead times.

According to McKinsey & Co., AI-enabled systems can simulate various supply-chain scenarios and optimize inventory levels and costs. This can be done, for example, by evaluating different production and distribution plans to meet demand with minimal stock-outs or overstock [13]. This kind of optimization helps balance working capital with operational readiness.

Emissions monitoring and environmental compliance. Environmental monitoring is no longer just a compliance checkbox, it is becoming a core optimization lever. ML models like neural networks and regression-based predictors can forecast emissions such as NO_x, CO₂, and volatile organic compounds (VOCs) using real-time process data.

In a cement kiln application, artificial neural networks and support vector regression were used to develop soft sensors for real-time estimation of O₂, CO, NO and CH₄ concentrations in stack gas, providing a reliable alternative during analyzer downtime that is caused by interference from fine dust [14].

Autonomous closed-loop control systems. Autonomous, closed-loop control uses ML to do more than just advise. The ML takes action also. This is considered the bleeding edge of ML application at manufacturing sites

right now. Reinforcement learning, model predictive control (MPC) and digital twins are used to adjust control setpoints in real time to optimize for yield, emissions or energy use.

For instance, Imubit Inc. (Houston, Tex.; www.imubit.com) has implemented its Closed Loop Neural Network™ platform, which utilizes reinforcement learning to optimize complex industrial processes in real time. This technology has been deployed in over 90 applications across major industrial processing plants, leading to significant improvements in operational efficiency and energy consumption. Notably, in rotary kiln operations, the platform has achieved up to a 30% reduction in natural gas usage. These systems require robust simulation environments, safety constraints and rigorous validation, but they represent a clear step toward self-optimizing plants [15].

Machine learning is not a “plug-and-play” solution. It still requires clean, contextualized data, strong domain expertise and alignment with plant priorities. But when used strategically, it can become a powerful catalyst for making smarter decisions, improving the efficiency of operations and driving long-term competitiveness.

From predictive maintenance to autonomous control, the promise of AI in chemical manufacturing is no longer theoretical — it’s operational.

Generative AI in manufacturing

Because generative AI has been especially visible in the last two years, a logical question becomes: How exactly can GenAI be used in chemical processing plants?

In late 2022, at the time that ChatGPT was launched, my guess would have been that it would take 10 to 20 years before the kinds of capabilities exhibited by ChatGPT would be available to operators on the plant floor. As it turns out, the pace with which natural language AI models would be adopted at chemical plants was drastically underestimated.

The pace of advancement in GenAI is truly mind-blowing, and for once, the manufacturing community is not lagging behind — it is leaning in.

While traditional ML focuses on

prediction, classification and optimization, GenAI introduces something entirely new: content and knowledge creation. Trained on massive data sets from operator logs and technical manuals to equipment metadata, GenAI systems can synthesize text, generate recommendations, create documentation and even simulate future scenarios.

As the chemical manufacturing industries continue to digitize, GenAI is already starting to reshape knowledge management, operator enablement, engineering design and compliance workflows. Here are eight emerging, high-impact use cases for GenAI already gaining traction on the plant floor.

Standard operating procedure (SOP) generation and knowledge documentation. Generative AI can automatically draft SOPs, safety procedures and maintenance checklists by learning from existing documentation, historian data and control logic. Models such as GPT-4 or domain-tuned LLMs can parse engineering documents and synthesize coherent instructions customized to plant-specific operations.

For example, eGain’s AI Knowledge Hub uses generative AI to capture tacit knowledge from experienced employees and transform it into structured, usable guidance. This approach has been used in knowledge-intensive industries to reduce documentation time and address the looming risk of institutional knowledge loss due to retirements and workforce churn [16].

Natural language interfaces to complex systems. One of the most accessible applications of GenAI is its ability to translate human language into structured queries or engineering logic. Using LLMs fine-tuned with domain-specific data and tools like LangChain or Azure OpenAI Service, operators can query large datasets or historian systems in plain English.

For instance, Celanese Corp. (Irving, Tex.; www.celanese.com) has developed a generative AI-powered chatbot named “Celia” as part of its digital transformation efforts. Celia serves as a user-friendly interface to the company’s digital twin, allowing manufacturing personnel to ask

questions like, “What is the production status of the Clear Lake facility?” and receive real-time, contextualized responses. This approach has improved visibility, collaboration, and decision-making across Celanese’s manufacturing sites [17].

Automated root-cause analysis reports. After a process deviation or equipment failure, GenAI can compile multi-system data (from DCS, alarms, maintenance logs and so on) and create narrative reports that summarize likely root causes, based on previous incidents and plant knowledge bases.

For instance, Intel Corp. (Santa Clara, Calif.; www.intel.com) developed a manufacturing “incident assistant” using Articulate’s generative AI platform to diagnose and resolve manufacturing problems. This system ingests and analyzes structured and unstructured data from diverse sources, including historical data and real-time feeds from sensors and semiconductor manufacturing equipment [18].

Technical report and audit preparation. GenAI is being used to pre-fill environmental audit templates, emissions disclosures and performance reviews by pulling data from production reports, historian tags and manufacturing execution systems (MES). These models can also generate summaries in regulatory language aligned with U.S. Environmental Protection Agency (EPA; www.epa.gov), Occupational Safety and Health Administration (OSHA; www.osha.gov), or E.U. Registration, Evaluation, Authorization of Chemicals (REACH) frameworks.

For instance, SG Systems Global (Farmer’s Branch, Tex.; sgsystems-global.com) has developed a recipe management software tailored for pharmaceutical manufacturers, which integrates seamlessly with enterprise resource planning (ERP) systems to provide electronic batch records and ensure compliance with 21 CFR Part 11 requirements. This digital system enhances quality, reduces costs, and eliminates errors compared to traditional paper-based systems by offering features like guided procedures, automatic data logging and traceability of batch information [19].

Simulation and scenario narra-

tive generation. When coupled with digital twins or process simulation engines, GenAI can generate interpretive narratives of simulated outcomes. This helps plant personnel better understand the implications of parameter changes, equipment failures or operating strategies without needing to manually analyze time-series outputs.

BMW has developed a comprehensive digital twin of its manufacturing facilities using NVIDIA’s (Santa Clara, Calif.; www.nvidia.com) Omniverse platform. This virtual environment allows BMW to simulate the entire production process, including the movement of vehicles through assembly lines, interactions with machinery and worker ergonomics. By incorporating GenAI, BMW can analyze these simulations to identify potential bottlenecks, optimize workflows, and predict the outcomes of changes in the production process [20].

Code and logic generation for engineering applications. LLMs like Codex or GitHub Copilot are increasingly being used to accelerate development of engineering scripts, logic blocks and visualization dashboards. Engineers can describe the intent (for example, “plot batch cycle times over the last six months and flag anomalies”), and the GenAI tool produces code in Python, SQL, or even SCADA scripting languages to accomplish that function.

DXC Technology (Ashburn, Va.; www.dxc.com) has implemented LLM-powered tools to transform data exploration for their oil-and-gas customers. These tools use LLMs to generate Python code, execute it, and return the results, enabling users to interact with complex datasets through natural language queries. This approach simplifies data analysis tasks, such as filtering and aggregating data, without requiring users to write code manually [21].

Training and operator enablement. GenAI systems are capable of generating quizzes, flashcards, safety walkthroughs and interactive training content based on actual plant procedures and historical performance data. By tailoring this content to specific unit operations or employee roles, companies can

dramatically reduce onboarding time and reinforce critical knowledge.

Colossyan Inc. (London, U.K.; www.colossyan.com) is a technology company that leverages GenAI to create corporate training videos. Their platform enables companies to produce training videos without traditional filming equipment, utilizing text-to-speech technology and AI avatars to deliver content with realistic lip-syncing. This approach allows for the rapid creation of multilingual training materials, including safety walkthroughs and equipment tutorials, tailored to specific roles and procedures [22].

Generative design and equipment layouts (emerging). While still early in their development, generative models are being trained to propose plant equipment layouts, process flow diagrams (PFDs) or even molecule structures based on desired outcomes or constraints. These tools can suggest design alternatives, compare tradeoffs and accelerate early-stage engineering work.

Airbus has collaborated with Autodesk to apply generative design techniques in reimagining structural aircraft components, such as the vertical tail plane of the A320. By utilizing generative design, Airbus developed lighter-weight parts that exceed performance and safety standards. Beyond aircraft components, Airbus is also exploring the use of generative design for the layout of adaptable, DGNB- (www.dgnb.de) and LEED-certified factories, aiming to streamline logistics and improve employee work conditions [23].

GenAI adds a new layer of intelligence to chemical manufacturing. It doesn’t just predict what’s going to happen. It actively creates knowledge, drafts documents, simulates scenarios and even writes programming code.

While these tools are still maturing, they are already accelerating documentation, training, analytics, and decision-making across the plant.

To deploy GenAI safely and effectively, manufacturers should:

- Use retrieval-augmented generation (RAG) to anchor outputs in verified plant knowledge
- Keep humans in the loop to validate content before execution

- Establish governance, audit trails and model traceability
- Fine-tune models with domain-specific data to reduce hallucinations and increase relevance

GenAI is not just another technology trend; it is a powerful augmentation layer for the modern chemical workforce, unlocking faster insights, leaner operations and a smarter, more intuitive interface between people and plant data.

Future of AI in the CPI

Given the pace of AI advancement over the past few years, I would not claim to be able to precisely predict how, when or where these technologies will ultimately take us. Just a decade ago, few imagined the possibilities that humans would be having fluent conversations with machines, or that engineers could generate reports from PDFs in seconds, or that they could use predictive models to fine-tune reactor conditions in real time. And yet, here we are.

In the world of chemical manufacturing, the promise of AI is not about replacing the deep domain knowledge of process engineers or the hard-earned intuition of control-room operators. It's about amplifying human expertise, unlocking insight buried deep in data and enabling faster, smarter decisions at every level — from the operator console to the executive boardroom.

The true power of AI does not lie with any one algorithm, model or software platform. It lies in the strategic fusion of three forces: the right data, the right problem and the right application. When aligned, these create a feedback loop of continuous learning and improvement — one that can be scaled across units, plants and global operations.

Imagine a future where the following scenarios are a reality:

- New operators train not by reading manuals, but by interacting with conversational AI tutors that speak in plant-specific language, referencing years of shift logs, batch reports and incident learnings.
- Every asset on the plant floor becomes its own historian and analyst, flagging anomalies, predicting failures and suggesting optimizations autonomously.

- Digital twins don't just simulate the process, they co-pilot it, learning and adapting to new feedstocks, environmental conditions and economic constraints in real time.
- Sustainability goals become engineering constraints, and AI navigates tradeoffs between efficiency, carbon footprint and cost at computing speeds no human could achieve.

This is not science-fiction. The building blocks for these capabilities are already available. AI will not revolutionize chemical manufacturing in a single “big bang.” It will do so quietly, relentlessly and exponentially, plant by plant, asset by asset and shift by shift.

The greatest threat is not AI itself — rather, it is doing nothing. Waiting too long to begin implementing AI tools is likely to result in witnessing competitors turn their predictive insights into lower costs, safer plants, smarter workforces and greener operations.

To succeed, manufacturers must stop thinking about AI as a project, and start thinking of it as a capability, a core pillar of operational excellence in the digital age. AI is no longer optional. It is the infrastructure of decision-making for the next generation of manufacturers.

The winners of tomorrow won't be the ones with the flashiest technology or the largest budgets. The winners will be the ones who deeply understand their process, embrace the messiness of real-world data, and treat AI as a multiplier of human talent, rather than a replacement.

So while I cannot and will not predict exactly what comes next, I will say that the plants that invest in AI today — those that experiment, learn, scale and lead from the front — will not only build smarter factories, they will build the future of manufacturing itself. And the future is already knocking. ■

Edited by Scott Jenkins

References

1. Wegner, P., What CEOs talked about in Q1 2025, IoT Analytics blog post, March 25, 2025, [iot-analytics.com](https://www.iot-analytics.com).
2. Merriam-Webster Inc., Online dictionary, accessed March 2025, www.merriam-webster.com.
3. Wikipedia, AI effect, accessed March 2025, en.wikipedia.org.
4. Kahn, J., It's Alive, *Wired* magazine, March 2002, www.wired.com.
5. Wikimedia Commons, Image of Analytical machine, Charles Babage, Upload by Mrjohncummings 2013-08-28 15:10, CC BY-SA 2.0., commons.wikimedia.org.
6. AI Newsletter.com, www.ainewsletter.com.

7. Uniyal, R., Hultgren, M. and others, Empowering Productivity and Quality in the AI Era, Indian Pulp and Paper Technical Association (IPPTA), Quarterly Journal, Vol. 37(1), pp. 116-120, 2025, ippta.co.
8. Siemens AG, Going for the Golden Batch with AI-based Performance Analytics Process, Digitalization Tech talks, podcast, episode 22, 2022, podtail.com/tr/podcast/siemens/.
9. Mori, L., Macak, M., Perdur, M. and others, How AI Enables New Possibilities in Chemicals, McKinsey and Co., article, November 2024, www.mckinsey.com/industries/chemicals/.
10. Kazerooni, N.M., Adib, H. and others, Toward an intelligent approach to H₂S content and vapor pressure of sour condensate of south pars natural gas processing plant, *Journal of Natural Gas Science and Engineering*, vol. 28, pp. 216-221, 2016, www.sciencedirect.com/science/.
11. Shyu, H.Y., Castro, C., Bair, R. and others, Development of a soft sensor using machine learning algorithms for predicting the water quality of an onsite wastewater treatment system, *ACS Environmental Au*, vol. 3(5), 308-318, June 2023, pubs.acs.org.
12. Adebayo, O., Imitiaz, S. and Ahmed, S., Machine learning for early detection of distillation column flooding, *Chemical Engineering Science*, vol. 312, June 2025, www.sciencedirect.com/science/article/pii/S0009250925004269.
13. Mori, L., Macak, M., Perdur, M. and others, How AI Enables New Possibilities in Chemicals, McKinsey and Co., article, November 2024, www.mckinsey.com/industries/chemicals/.
14. Khosrozade, A. and Mehranbod, N., Comparison of support-vector regression and neural network-based soft sensors for cement-plant exhaust gas composition, *International Journal of Environmental Science and Technology*, vol. 17(5), October 2019, www.researchgate.net/publication/336865261.
15. Imubit Inc., Unlock new value, webpage, accessed March 2025, imubit.com/deep-learning-process-control-platform-for-hydro-carbon-processors/.
16. Seigel, E., Capturing tacit knowledge from the great retirement cohort using genAI, eGain blog post, April 2025, www.egain.com/blog/capturing-tacit-knowledge-from-the-great-retirement-cohort-using-genai/.
17. Pelletier, S., Building a Generative AI-powered digital twin, Manufacturing Leadership Council, ML Journal, August 2024, manufacturingleadershipcouncil.com/.
18. Articul8 Inc., Intel: Manufacturing Root-Cause Analysis, case study, Webpage, accessed March 2025, www.articul8.ai/case-studies/intel-manufacturing-root-cause-analysis.
19. SG Systems Global, Pharmaceutical Manufacturing Regulatory Compliance, slide presentation, January 2023, www.slideshare.net/slideshow/pharmaceutical-manufacturing-regulatory-compliancepdf/255243126.
20. Kobie, N., The dream of the metaverse is dying. Manufacturing is keeping it alive, *Wired* magazine, May 2025, www.wired.com/story/the-metaverse-is-here-and-its-industrial/.
21. Genevay, A., Charuvaka, A. and others, DXC transforms data exploration for their oil and gas customers with LLM-powered tools, Amazon Web Services (AWS) blog post, November 2024, aws.amazon.com/blogs/machine-learning/.
22. Colossyan Inc., corporate webpage, www.colossyan.com/, accessed April 2025.
23. Deplazes, R., Autodesk and Airbus demonstrate the impact of generative design on making and building, Autodesk blog post, November 2019, adsknews.autodesk.com/en/news/autodesk-airbus-generative-design-aerospace-factory.

Author



Jonathan Alexander is the manufacturing AI and advanced analytics manager within Albemarle's (4250 Congress St., #900, Charlotte, NC 28209; Email: jonathan.alexander@albemarle.com) global manufacturing excellence team. Alexander has held a variety of roles within manufacturing, including working within maintenance and operations for both Albemarle's hydroprocessing catalyst plant in Bayport, Tex., and the aluminum alkyl plants in Pasadena, Tex. He has also led a site-wide digital transformation program. In his current global role, he has been focused on empowering people by scaling AI. Alexander has also been working on changing Albemarle's culture around manufacturing data analytics, an integral part of their journey toward digital transformation. Alexander holds a bachelor's degree in electrical engineering from Louisiana State University, and an MBA from The University of Houston in Clear Lake.