AI in chemical engineering: From promise to practice

Downloaded from: https://research.chalmers.se, 2026-07-25 08:27 UTC

Citation for the original published paper (version of record):
Chew, J., Andersson, R., Bierweiler, T. et al (2026). AI in chemical engineering: From promise to practice. AICHE Journal, 72(7). http://dx.doi.org/10.1002/aic.70358

N.B. When citing this work, cite the original published paper.

research.chalmers.se offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all kind of research output: articles, dissertations, conference papers, reports etc. since 2004. research.chalmers.se is administrated and maintained by Chalmers Library

(article starts on next page)

Check for updates

Received: 18 December 2025 | Revised: 25 February 2026 | Accepted: 6 March 2026

DOI: 10.1002/aic.70358

AIChE JOURNAL

REVIEW ARTICLE

Process Systems Engineering

AI in chemical engineering: From promise to practice

Jia Wei Chew[1] | Ronnie Andersson[1] | Thomas Bierweiler[2] | Patrik Ryttestal[3] | Torsten Wik[4]

[1]Department of Chemistry and Chemical Engineering, Chalmers University of Technology, Gothenburg, Sweden

[2]Siemens AG, Digital Industries, Process Automation, Karlsruhe, Germany

[3]Siemens AB, Digital Industries, Molndal, Sweden

[4]Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

Correspondence
Jia Wei Chew, Department of Chemistry and Chemical Engineering, Chalmers University of Technology, 412 96, Gothenburg, Sweden.
Email: jia.chew@chalmers.se

Funding information
Chalmers Gender Initiative for Excellence (Genie)

Abstract

Artificial intelligence (AI) in chemical engineering has moved from promise to practice: physics-aware (gray-box) models are gaining traction, reinforcement learning complements model predictive control (MPC), and generative AI powers documentation, digitization, and safety workflows. Near-term value arises where AI augments, rather than replaces, process system engineering (PSE) practice (e.g., through soft sensing and surrogate models), while autonomous operations, fully automated hazard and operability (HAZOP) analysis, and large-scale mechanistic discovery remain largely at the research stage. The decisive bottleneck is reliable deployment: AI models must be treated like any other engineered system, with validation, monitoring, and governance aligned with emerging frameworks such as the EU AI Act and NIST risk management framework (RMF). With incubator labs, open benchmarks, and retooled education pipelines, AI can become a safe, reliable, and sustainable co-worker in the process industries within years.

KEYWORDS

AI governance and assurance, generative AI, gray-box modeling, industrial AI, physics-aware models, reinforcement learning for process control

1 | INTRODUCTION: WHY A FOLLOW-UP NOW

Artificial intelligence (AI) has transitioned from a predominantly research-driven topic in academia to an increasingly deployed technology within chemical engineering practice. Over the past decade, advances in machine learning, physics-aware (gray-box) modeling, reinforcement learning, and generative AI have expanded the toolkits available to process engineers. Concurrently, the deployment context has shifted: AI systems are now expected to meet industrial standards for validation, robustness, governance, and lifecycle management.

A noteworthy reference point is the 2019 perspective that posed the question—"is it here, finally?"—framing a pivotal moment for artificial intelligence (AI) in chemical engineering (ChE).[1] That perspective article traced ChE's AI journey from its process systems engineering roots (PSE) in the 1960s, through the expert-systems wave of the 1980s and the rise of neural networks in the 1990s, to the deep-learning and data-science era that began in the mid-2000s. That article captured a moment when data availability, computing infrastructure, and learning algorithms were beginning to converge, with capabilities like soft sensing, monitoring, knowledge-data hybrid modeling (i.e., gray-box models), and decision support becoming practical possibilities, but large-scale industrial deployment remained limited. The present study does not seek to extend or update that perspective directly; rather, 2019 is referenced as a pivotal timepoint from which subsequent technical and organizational developments can be assessed.

Historically, ChE advanced in an Edisonian (trial-and-error), largely stochastic fashion—iterating empirically toward workable recipes and

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.
© 2026 The Author(s). *AIChE Journal* published by Wiley Periodicals LLC on behalf of American Institute of Chemical Engineers.

equipment designs when theory and/or data were scarce. From the 1960s through the 1990s, however, the discipline transitioned toward a more deterministic paradigm: PSE and digital computation formalized process design, optimization, and control, while transport phenomena unified momentum, heat, and mass transfer into a first-principles foundation that still anchors today's curricula and simulation models.

The 2000s surge of AI and machine learning (ML) appeared to mark a swing back to stochastic, data-driven approaches. Two enduring takeaways from the 2019 article continue to shape the discussion[1]: (i) ML— rather than AI— was then the practical frontier across design, control, diagnosis, and safety; and (ii) greater impact would depend on judicious integration with domain knowledge and ensuring model explainability to impose scientific rigor. It was hypothesized that data, computing architectures, and algorithms were finally converging to meaningfully impact the discipline.

Since then, several watershed moments have redefined expectations for AI's capability and momentum. AlphaFold's breakthrough in protein-structure prediction[2]—recognized by the 2024 Nobel Prize in Chemistry—demonstrated that learning-enabled models can bridge long-standing scientific gaps. The late-2022 launch and rapid adoption of ChatGPT brought general-purpose AI assistants into everyday life, accelerating the translation of AI from labs to mainstream usage.

Against this backdrop, in line with the call for integrating physical principles into AI,[1] 2019 marked a key inflection point with the emergence of physics-informed neural networks (PINNs),[3] which bridged stochastic black-box learning and mechanistic white-box governance to enhance accuracy, generalizability, and reliability. Since then, physics-aware (gray-box) approaches have evolved from conceptual ideas to practical toolkits.[4] The field has matured into ChE practice, with recent ChE-focused tutorials providing detailed, step-by-step hybrid modeling workflows.[5,6]

In parallel, generative AI (GenAI) has moved from molecule/text generation to practical engineering utilities, including flowsheet/P&ID digitization and enhancement,[7] documentation assistance, and early HAZOP support[8]—spanning scales from molecular to process.[9] Concurrently, foundation models (i.e., large AI models trained on vast, diverse datasets to perform a wide range of tasks, serving as a versatile base for specialized applications) and mature open-source stacks have lowered the prototype barrier, while hybrid (gray-box) modeling libraries/workflows have narrowed the gap between first-principles models and ML.[5] Reinforcement learning (RL) for process control likewise matured, with growing consensus on its role alongside classical model predictive control (MPC).[10] With respect to data infrastructure, tighter historian/distributed control system (DCS) integration[11] and modern data-engineering practices (e.g., reliable data archiving, ingestion and preprocessing of heterogeneous process/equipment data)[12,13] now underpin robust pipelines for soft sensors, monitoring, and control.[14] Since 2023, time-series foundation models (TSFMs) trained on multiple-domain datasets have emerged, offering scalable approaches to address data scarcity in narrow, domain-specific AI applications.[15,16] TSFMs support a variety of use cases during commissioning and operations. Time-series pre-trained transformers (TPTs; i.e., industrial-grade generative AI models trained on vast amounts of industrial data to enable expert-level autonomous reasoning and optimization) have been described in vendor-reported case materials as being applied across chlor-alkali, thermal power, and petrochemical sectors, with claimed improvements in operational efficiency, productivity, quality, and safety.[17] This indicates emerging industrial activity, though independent peer-reviewed validation remains limited.

Also, AI governance and MLOps (machine learning operations) shifted from optional to expected. The EU AI Act, in force since 1 August 2024,[18,19] introduces obligations in stages between 2025 and 2027—starting with bans on prohibited practices in early 2025 and requirements for general-purpose AI models in mid-2025 to key high-risk system rules by 2026–2027—while the NIST AI Risk Management Framework (RMF) complements with voluntary, sector-agnostic guidance. Together, these frameworks, which have spurred organizations to formalize validation, change control, and use post-market monitoring, regulate industry-grade AI.

Recent reviews have been valuable in providing insights on various facets of the AI-ChE landscape in depth: GenAI in ChE spanning molecular to process scales,[9] responsible use of GenAI in practice,[20] physics-informed machine learning (PIML) applications,[21] and step-by-step tutorials.[5,6] In process control, a focused review position RL as a complement to classical MPC, addressing sim-to-real transfer and safety issues but primarily at the algorithmic level.[22] In process safety, evaluations of AI-assisted HAZOP emphasize augmentation (e.g., pre-population, retrieval) rather than full automation.[8] These contributions provide valuable depth at the method or application level and have helped consolidate best practices within individual subdomains. However, the emphasis has been on algorithmic developments, case studies, or disciplinary advances, with comparatively limited attention to deployment maturity, governance, lifecycle management, and regulatory alignment in industrial settings.

In contrast to prior reviews, the distinctive contribution here is an industry-centric synthesis that ties recent AI advances with the practical requirements for safe, reliable, and scalable deployment in chemical process industries. Rather than reviewing algorithms in isolation, the developments in generative AI, physics-aware (gray-box) modeling, and reinforcement learning are integrated with evidence standards, MLOps practices, governance frameworks, and emerging regulation (e.g., EU AI Act, NIST AI RMF). By explicitly distinguishing research-grade capabilities from industry-ready applications, and by anchoring discussion to real engineering workflows such as process control, safety analysis, sustainability assessment, and education, this review aims to address the chief question: What constitutes "industry-grade" AI in ChE practice today?

Recent surveys of industry indicate that implementation gaps remain: 40% of respondents regard AI as untrustworthy, 92% report a shortage of AI-skilled experts, and only 16% have achieved their AI-related goal.[23] Supporting this, a 2021 Google Cloud White Paper found that 72% of organizations that initiated AI pilots before 2019 had not deployed a single application in production, and 55% of

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

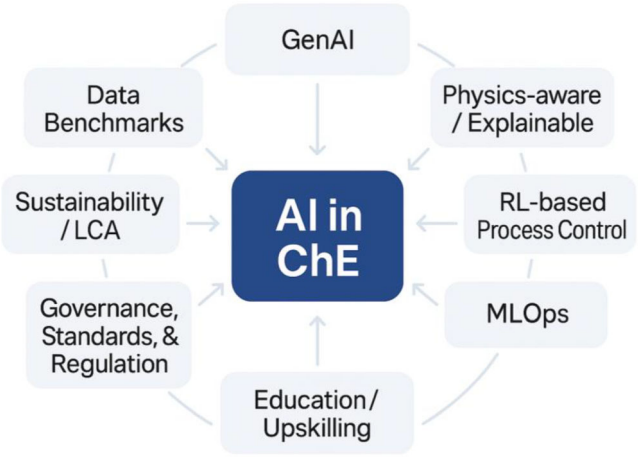

FIGURE 1 Schematic overview of key technical and organizational domains shaping AI adoption in chemical engineering (ChE).

companies subscribing to enterprise ML platforms had yet to operationalize an ML model.[24]

Clearly, deployment maturity lags technological advances. A reassessment of AI in ChE is therefore both necessary and timely to address persistent gaps in AI adoption by chemical process industries, as framed schematically in Figure 1. This study anchors recommendations to the EU AI Act timeline and NIST AI RMF/GenAI Profile (covering risk roles, technical files, post-market monitoring), and foregrounding verification harnesses (e.g., golden sets, stress tests, citation-first retrieval-augmented generation (RAG)) for GenAI and hybrids in safety-critical workflows (e.g., AI-assisted P&ID digitization for HAZOP preparation and analysis). In essence, this review is structured as an industry-focused review and perspective on AI in ChE, emphasizing deployment, assurance, and workforce readiness alongside technical progress. Commercial examples cited throughout the manuscript are intended to illustrate current industry capabilities and deployment patterns and should not be interpreted as endorsements or as privileging specific vendors.

2 | 2019 PREDICTION SCORECARD

The 2019 Perspective outlined the most plausible early landing zones for AI in ChE[1]—soft sensing, monitoring and diagnosis, hybrid (gray-box) modeling, and decision support—and, by and large, that roadmap has held true.

Several predictions have materialized. In monitoring and diagnosis, data-driven soft sensors and fault detection and diagnosis (FDD) have moved from promising demonstrations to routine engineering tools. Recent reviews and applications describe adaptive and just-in-time soft sensors as well as deep-learning-based models that sustain performance under changing operating regimes.[14,25] In design and optimization, surrogate models have become standard accelerators—ranging from unit-operation replacements embedded in superstructure optimization to plant-scale case studies—with 2024–2025 tutorials laying out step-by-step hybrid and surrogate modeling workflows for ChE practitioners.[6,26] In operations optimization, learning-augmented scheduling has progressed beyond proofs-of-concept into a growing body of studies, including those using RL for production scheduling, that explicitly addresses real-world constraints.[27]

Clear progress is also evident in process control, trending toward autonomy-leaning control architectures in which RL augments or supervises MPC. Recent studies define benchmarks, explore safety and sim-to-real transfer issues, and introduce control-informed RL, while noting that large-scale industrial deployment remains limited.[10,22] These approaches couple the disturbance-rejection and set-point-tracking strengths of traditional proportional-integral-derivative (PID) control with the nonlinear modeling capacity of deep RL. On generalization across plants or units, transfer-learning, robustness techniques, and knowledge-guided or hybrid models are reducing retraining burden and improving extrapolation.[28] For decision support and operator trust, explainable AI (XAI) has become increasingly essential for transparency and governance. Although studies, such as explainable RL for process safety,[29] demonstrate progress, standardized playbooks for process operations are still emerging.[29]

Two developments have notably exceeded the 2019 expectations. First, the outsized, rapid impact of GenAI: large language models (LLMs) and multimodal foundation models have become practically useful far sooner than anticipated, with a recent ChE-focused review charting applications from molecular and materials discovery to process documentation and reasoning aids.[30] Second, hybrid (gray-box) models have outpaced purely black-box approaches: in data-scarce and constraint-heavy settings typical of process industries, physics-aware models have delivered superior generalization and safer extrapolation than deep learning alone, shifting the momentum decisively toward mechanism-infused ML and AI.[5,6]

3 | GenAI

GenAI in ChE spans multiple model classes, including molecular and materials generation, graph-based generative models for process flows, and text and multimodal foundation models (i.e., models trained jointly on text and visual inputs such as diagrams and figures). While recent industrial adoption has been driven largely by document- and diagram-centric workflows, these represent one subset of a broader GenAI portfolio relevant to chemical engineering design and analysis. For legacy piping and instrumentation diagrams (P&IDs) and process flow diagrams (PFDs), vision-LLM and computer-vision (CV) methods can now detect symbols, text, lines, and connectivity to produce outputs compatible with DEXPI/ISO 15926, enabling direct transfer to engineering software, simulation environments, computerized maintenance management system (CMMS), and HAZOP preparation. Several commercial solutions now offer AI-assisted P&ID/PFD digitalization capabilities, with examples including DigiCo Digitization Companion, DeepIQ P&ID Digitization, iDrawings Intelligent Project Solutions,

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

SymphonyAI P&ID Ingestion, PIDFlow, PIDVision.AI, and Siemens Industrial Copilot. Recent studies demonstrate end-to-end P&ID graph extraction and DEXPI-compliant output, while transformer-based architectures (e.g., neural networks that employ self-attention to weigh relationships among all input elements in parallel[31]) achieve joint object-relation extraction, outperforming conventional, piecemeal optical character recognition (OCR) pipelines.[32,33] On the receiving end, DEXPI (ISO 15926-based) standard is well-documented and supported by major vendors (e.g., AVEVA P&ID or Siemens COMOS provide ISO 15926/DEXPI import–export), making it a practical interoperability target for GenAI/CV-derived P&ID data.[34] Within such digitization flows, tag extraction and normalization combine OCR with contextual cues from layouts and symbols, often pairing neural detectors with rule-based mapping to plant conventions before integration into historians or asset-management systems.[35]

Beyond document interpretation, graph-based generative models provide a natural representation for ChE systems. Process flow diagrams (PFDs), flowsheets, and networked process configurations can be represented as attributed graphs, enabling generative and constraint-aware graph models to propose, modify, or validate process structures. Recent studies demonstrate the potential of graph-based representations for process flowsheet synthesis and design. Transformer-based models have been proposed to learn flowsheet grammar and provide interactive synthesis recommendations using SFILES (Simplified Flowsheet Input-Line Entry-System) notation during early-stage design.[36] Graph neural networks (GNNs) have also been used within hierarchical reinforcement learning frameworks to generate flowsheets from graph representations.[37] Unlike traditional artificial neural networks (ANNs), which operate on fixed-size vector or grid-like inputs, GNNs are designed to learn from graph-structured data. This makes them well suited for representing topologies such as process flowsheets, molecular graphs, or equipment networks.

Additionally, emerging generative approaches such as Graph-to-SFILES demonstrate how generative models can take flowsheet topologies as graph input and produce control-extended process diagrams.[7] While most current industrial deployments focus on diagram digitization rather than automated synthesis, these graph-centric GenAI methods represent a growing research direction with potential future impact on conceptual design workflows.

GenAI has also become credible for document copilots. In software engineering, randomized controlled trials have shown large productivity gains from AI coding assistants—capability that translates naturally to ChE code development and data analysis, and to retrieval-augmented generation (RAG) over standard operating procedures (SOPs), vendor manuals, and standards.[38] Industrial case studies report RAG improving technical questions and answers (Q&A), and troubleshooting in manufacturing and maintenance, including multimodal RAG systems capable of reasoning jointly over text and figures from equipment manuals.[31]

At the molecular scale, generative models such as variational autoencoders, diffusion models, and transformer architectures have enabled inverse design of molecules, catalysts, solvents, and materials subject to target property constraints. Although often developed within computational chemistry and materials science, these capabilities increasingly feed upstream into process design, sustainability assessment, and digital twin workflows in chemical engineering. Commercial interest in generative molecular design is also emerging, an example of which is Evogene's AI-driven platform that applies proprietary generative models to design novel small molecules optimized across multiple performance criteria. For equipment and component design, commercial generative design and design-space exploration tools (e.g., Dassault SIMULIA Isight, Altair Inspire (Altair was acquired by Siemens in 2025), Ansys optiSLang, Autodesk Fusion 360 Generative Design, Siemens Simcenter HEEDS) enable automated, multi-objective search under constraints, and interfaces with all commercial CAD and CAE environments, thereby bridging process-level calculations with detailed geometric and thermo-structural optimization. These methods rely on algorithmic design-space exploration, not on text-to-CAD generation.

HAZOP augmentation is emerging but not as a replacement for expert-led efforts. Studies highlight value in pre-populating guideword-deviation tables, surfacing similar incidents, and drafting preliminary cause-consequence relationships, while emphasizing gaps in knowledge coverage and the need for expert verification.[8,39] Plant-tuned RAG systems that index SOPs, manuals, and incident corpora (e.g., U.S. Chemical Safety Board (CSB) reports[40]) help link deviations to equipment, safeguards, and prior events, supporting HAZOP, Layer of Protection Analysis (LOPA), and management of change (MOC) workflows.

Given these expanding capabilities, governance guardrails are essential. The EU AI Act (Article 14: Human Oversight) mandates human oversight and sign-off particularly for high-risk systems, while NIST AI RMF 1.0 calls for rigorous testing, evaluation, verification, validation (TEVV) and continuous monitoring. Recent RAG studies recommend adoption of golden-set evaluation (i.e., curated, fixed test suites that report both retrieval and generation metrics) to ensure reproducibility and regression testing.[41] Standardizations are also needed: target DEXPI/ISO 15926 for diagram semantics; verify toolchain compatibility (e.g., AVEVA or Siemens COMOS import/export) before scaling up; and treat OCR/CV pipelines as evidence-generating systems, retaining logs and test harnesses for auditability.[42]

Across these GenAI modalities, common challenges for industrial deployment include constraint satisfaction, validation, traceability, integration with existing engineering standards and workflows, and governance and assurance considerations.

4 | MECHANISTIC AND EXPLAINABLE AI IN ChE

Progress in data integration, coupled with the rapid emergence of mature AI/ML frameworks, open-source libraries, and production-ready workflows, has moved AI in ChE from promise to practice. These enablers have lowered barriers to deployment and accelerated adoption. However, purely black-box data-driven models offer limited value for design, control, optimization, and mechanistic interpretability

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

in ChE applications, where physical consistency, adherence to constraints, and explainability are essential. Moreover, the scarcity of data in many ChE applications, along with concerns about their accuracy and reliability in representing the true state of the system, further limits the effectiveness of purely data-driven models. Specifically, industrial process data are often sparse, noisy, and unevenly distributed across operating regimes. Common challenges include sensor drift, calibration errors, unlogged manual interventions, and limited fault data. Also, extreme conditions or process upsets, precisely where predictive robustness is most needed, are rarely well represented in historical datasets. Thus, purely data-driven models may overfit or extrapolate unreliably.

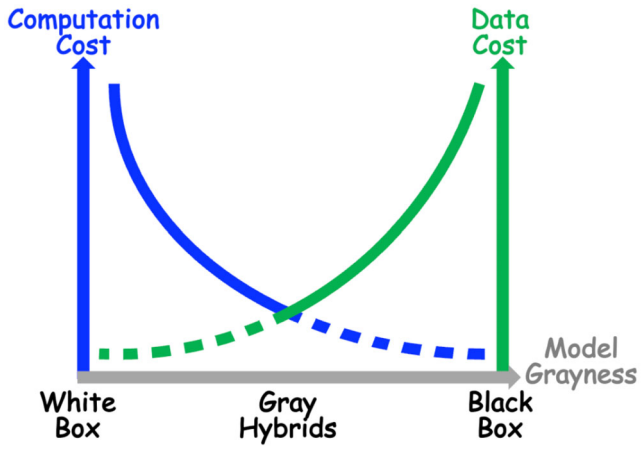

FIGURE 2 Knowledge-data hybrid modeling: qualitative trade-off between computation cost and data requirement as model "grayness" varies from white-box (physics-based) to black-box (data-driven). Model grayness denotes the degree of physics-based governance (i.e., mechanistic structure and constraints) vs. data-driven learned components. Dashed curves indicate limiting endpoint characteristics: the low data requirement of pure white-box models and the low computational cost of pure black-box models.

In response, a family of approaches has emerged that incorporate physical principles into learning systems, either by enforcing governing equations and constraints, preserving conservation laws and thermodynamic consistency, or modeling only the residuals around established first-principles models (e.g., constrained learners, physics-informed architectures, hybrid gray-box/residual models). *Physics-informed* approaches typically incorporate physical laws as soft penalties in the loss function, while *physics-constrained* ones enforce hard constraints by construction (e.g., via model architecture or constrained optimization). By reducing effective degrees of freedom and leveraging prior process knowledge, such models can improve robustness, reduce data requirements, and enhance extrapolation beyond observed operating windows. This section surveys these approaches and examines how ChE domain knowledge can be embedded to yield models that are both accurate and epistemically transparent (i.e., whose internal representations and outputs align with the underlying physical mechanisms).[4,43]

Such gray-box models balance the trade-off between the high computation costs of white-box models and the data costs of black-box models (Figure 2). Figure 3A shows how mechanistic accuracy generally decreases as model grayness increases. However, in industry, reliance has predominantly been on white-box models, with some implementations on pure black box models (e.g., Predictive Emissions Monitoring System specified in Performance Specification PS 16 (USA); EU: CEN/TS 17198; VDI-EE 3952:2024–04 Stationary source emissions (Germany); soft sensors), while gray-box models remain least mature in industrial deployment. Figure 3B contrasts ChE processes with facial recognition in terms of mechanistic governance and transferability. The former is governed by mechanistic principles, and so transferability enhances with the infusion of more mechanistic governance. As the extent of mechanistic governance decreases, the amount of data needed for the same extent of transferability increases. On the other hand, phenomenological systems like facial recognition are described using observed phenomena rather than

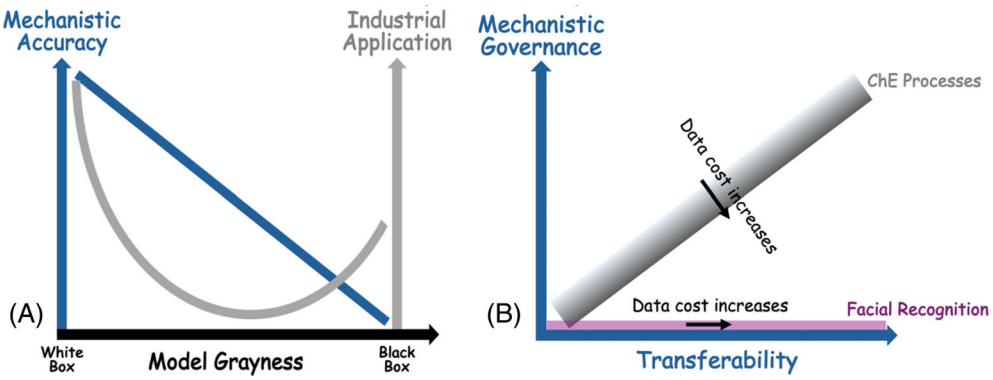

FIGURE 3 (A) Conceptual relationship between model grayness, mechanistic accuracy, and representative industrial applicability across model types. Model grayness refers to the relative contribution of physics-based governance (white-box) vs. data-driven learned components (black-box). (B) Conceptual relationship between mechanistic governance and generalizability/transferability across units and scales. Purely phenomenological models (e.g., image or speech recognition) rely primarily on data diversity to generalize, whereas chemical engineering applications often benefit from mechanistic structure to transfer reliably across operating conditions. Accordingly, for a given transferability target, lower mechanistic insight generally requires broader and more representative data to close the gap.

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

underlying physics, so no mechanistic governance is needed, and generalizability or transferability depends on the amount of data used for training.

In Figure 4, the six dimensions provide a holistic view of model performance across the three modeling strategies, namely, mechanistic, hybrid (gray-box), and black-box. Data efficiency captures how quickly a model achieves high accuracy with limited data, favoring approaches that leverage prior knowledge. Computational efficiency reflects normalized inference speed, highlighting differences between solver-heavy mechanistic models and typically fast black-box models. Generalization (including transferability) measures robustness to unseen conditions and domain shifts, a critical factor for real-world reliability. Interpretability refers to how easily the model's logic can be understood. Uncertainty quantification (UQ) assesses how well a model estimates the confidence or reliability of its predictions, while UQ rigor describes the degree to which a model systematically and transparently quantifies uncertainty. Although classical statistical models (e.g., polynomial regressions and autoregressive models) exhibit high UQ rigor by providing analytically derived confidence and prediction intervals under explicit assumptions, many AI/ML black-box models (e.g., neural networks) lack inherent uncertainty handling unless augmented with Bayesian or ensemble methods, contributing to lower UQ rigor in their default formulations. Finally, industrial maturity captures the extent to which each modeling approach is standardized and adopted in industrial practice, reflecting its readiness for deployment in ChE environments.

The focus has transitioned from black-box, data-driven prediction toward mechanism-aware learning and explainable decision support. In practice, that translates to selecting hybrid models when data are scarce or constraints are stringent; using operator learners for fast, physically faithful partial differential equation (PDE) surrogates; preferring inherently interpretable glass-box models where possible; adding post-hoc explanations where possible; and coupling models with causal and sensitivity analyses (e.g., PCMCI (a causal discovery and time-series analysis algorithm)[44]) along with explicit verification and validation plans.

4.1 | Mechanism-aware modeling

Across monitoring, control, and scale-up tasks, significant gains have been realized by hybridizing learning with physical laws.[45] Physics-informed neural networks (PINNs) are emblematic of this shift from black-box to mechanism-aware learning,[3,4] which is achieved by augmenting the empirical loss with a physics-residual term that penalizes deviations from governing equations and associated boundary or initial conditions. This formulation improves data efficiency, enforces physical consistency, and enhances robustness under extrapolation. PINNs support both forward problems (predicting system states given model parameters and inputs) and inverse problems (where unknown model parameters or material properties are inferred from data). Extensions address multi-physics coupling, partial or noisy observations, and adaptive sampling.[46,47]

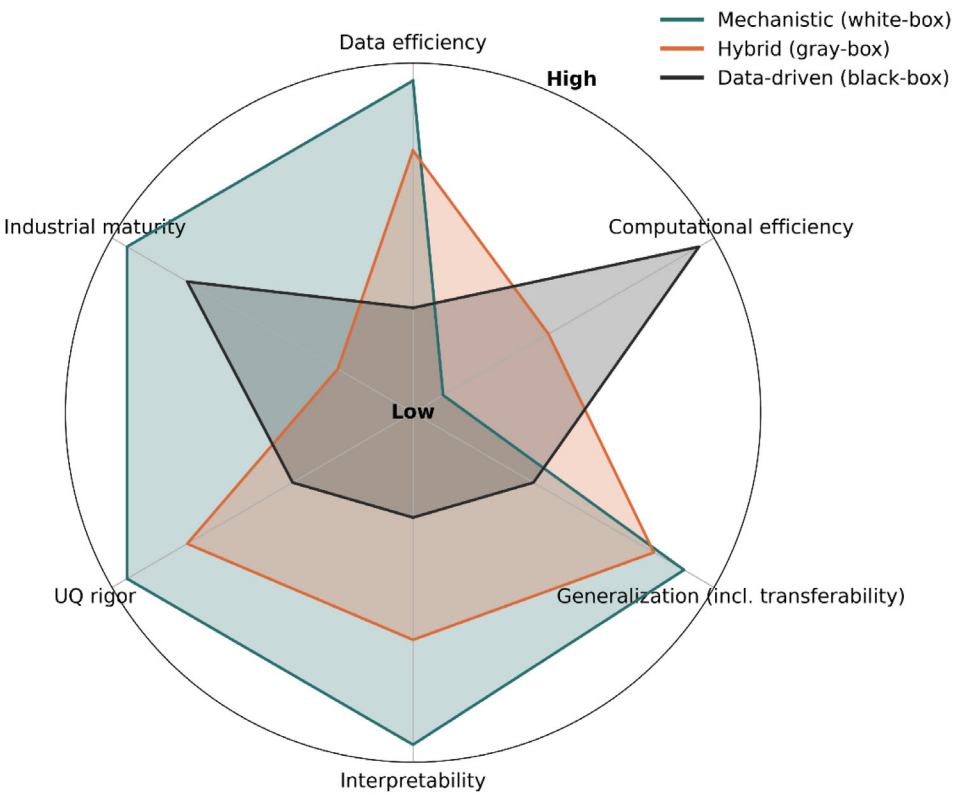

FIGURE 4 Model comparison: mechanistic (white-box), hybrid (gray-box), and data-driven (black-box).

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

The physics-informed ML toolbox has matured rapidly since the seminal PINN paper in 2019,[3] with subsequent comprehensive reviews consolidating advances, capabilities and limits.[21,48,49] While PINNs are often associated with time- and space-dependent PDE systems, physics-aware machine learning has also been applied to enforce thermodynamic consistency, embed equation-of-state relationships, incorporate phase-equilibrium constraints, and preserve energy and entropy balances within surrogate models. Moreover, beyond neural networks for the data-driven structures, Gaussian processes, kernel methods, ensemble regressors, and symbolic regression approaches have been employed.[50–52] In ChE specifically, including biochemical processes,[53] hybrid (gray-box) models that fuse mechanistic structure with data-driven components have become mainstream in process systems engineering,[5,6] and recent tutorials translate these concepts into practical, implementable workflows.[6] However, standard PINNs are typically trained for specific boundary and initial conditions and do not generalize to new scenarios without retraining, which restricts their applicability in tasks such as control system design and what-if analyses. To overcome this, recent extensions enable PINNs to handle multiple boundary or initial conditions, improving their utility in optimization and control contexts.[54]

Recent approaches replace iterative solvers in complex partial or ordinary differential-algebraic equations (PDAE/ODAE)-based ChE processes with feedforward neural networks, yielding orders-of-magnitude speedups while extrapolation to a broader input space.[55] Another approach, when rapid what-if analysis, controller-in-the-loop simulation,[56] or design-space exploration[57,58] are needed, are neural operators that learn solution operators of PDEs and generalize across meshes and discretizations.[59,60] The Fourier Neural Operator (FNO)[61] and DeepONet[54] families deliver orders-of-magnitude speedups relative to classical solvers while retaining physical structure, making them attractive drop-in accelerators for computational fluid dynamics (CFD), transport and reactor models. To improve robustness under previously unseen conditions in complex or emerging ChE applications, neural-operator and related methods have been successfully applied to stiff chemical-kinetics surrogates,[62] reacting-flow predictions,[63] and constrained variants that enforce conservation laws.[64,65] Neural and universal differential equation (NDE/UDE) approaches embed differential-equation structure into the learning process, enabling models to learn the governing system dynamics rather than merely fit observed time-series data, as is typical for recurrent neural networks. This inductive bias enhances the ability to predict previously unseen trajectories, supporting reliable what-if analysis and design exploration. However, as the model learns the system dynamics only within the confines of the state variables, its ability to extrapolate to entirely new regimes remains limited.

Mechanistic structure learning includes sparse regression-based equation discovery methods, such as sparse identification of nonlinear dynamics (SINDy), which identify governing equations by selecting candidate functional forms that satisfy both data and physical constraints. These methods yield compact, closed-form expressions that enable interpretation and model reduction. Symbolic regression is often implemented as a post-hoc distillation step following UDE training, where a flexible neural model first captures dynamics, but it can also be used directly on experimental or simulated datasets, particularly for low-dimensional systems or when interpretability is prioritized from the outset.[66]

4.2 | Explainability and model assurance

Explainability in ChE AI earns its place when it transforms physics-constrained models into auditable evidence that supports design, control, and scale-up, while remaining valid as data, equipment, and operating envelopes evolve. Achieving this requires more than post-hoc narratives. Interpretability layers such as model-agnostic attributions (e.g., SHapley Additive exPlanations (SHAP)[67]) and glass-box surrogates (Generalized Additive Models (GAMs)/Explainable Boosting Machines (EBMs); InterpretML[68]) should be embedded alongside rigorous physics checks (e.g., residual norms, conservation, feasibility), global sensitivity and identifiability,[69,70] and operator-oriented what-if probes. These elements make model behavior inspectable and ensure that explanations respect balances, units, and admissible domains, rather than purely statistical data-driven abstractions. In safety- or reliability-critical applications, this concept can be extended further: instead of imposing physical constraints via loss penalties (as in soft-constraint PINNs), certain models enforce them directly into the architecture as hard-constraint formulations. This ensures that fundamental laws (e.g., conservation or thermodynamic) are satisfied by design, mitigating the risk of physically invalid outputs and further strengthening baseline trustworthiness. The payoff is twofold. First, because mechanism-aware training (i.e., training that integrates mechanistic knowledge) narrows the hypothesis space to physically admissible behaviors, explanations gain mechanistic plausibility: global trends can be cross-checked against kinetics and transport logics, and counterfactuals are restricted to feasible manipulations, reflecting how process engineers reason. Second, when paired with calibration and out-of-distribution (OOD) safeguards, explainability becomes actionable. Predictive intervals can provide meaningful coverage on unseen campaigns, and OOD sentinels prevent recommendations when the system leaves the validated envelope. Explanation thus becomes a testable property of the model and its defined scope.

However, limitations remain: (i) most attribution methods are correlational—agreement with known mechanisms signals plausibility, not proof; (ii) local perturbation-based probes can mislead if perturbations stray into unphysical regions; (iii) common explainability tools are sensitive to collinearity and modeling choices, which can destabilize explanations in complex process spaces. These concerns argue for prioritizing inherently interpretable formulations, where feasible, and for treating explanations as hypotheses to be tested rather than definitive proofs. In short, explainability delivers value in ChE only when welded to physics checks, calibrated uncertainty, and explicit operating envelopes; used carelessly, it risks becoming a veneer of understanding without reliability.

Mechanism-aware modeling approaches, while promising, are not without limitations. PINNs can suffer from convergence and stability

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

issues, particularly for stiff or multiscale dynamics, under-constrained settings, sparse or noisy data, or when loss terms are poorly weighted. A range of remedies has been presented, including adaptive collocation, dynamic loss weighting, and alternative loss formulations (e.g., maximum norm loss, L_∞, which penalizes the worst-case error rather than the average squared error) that have been shown to improve stability for certain classes of PDEs. Symbolic regression and structure discovery methods may yield overly complex or non-generalizable expressions if regularization is insufficient or the function library is poorly specified. These methods are also sensitive to hyperparameter tuning and robustness may be compromised in noisy or high-dimensional settings. Continued research in training strategies, regularization, and model selection, integrated with rigorous validation and stress testing, are needed to improve the reliability and scalability of these approaches in practical applications.

5 | PROCESS CONTROL IN RL ERA

PID controllers are still used in most applications today. MPCs are usually set on top of traditional PID controllers rather than directly controlling actuators. Though methods for determining PID parameters are well-established, disturbances to normal operation must usually be applied, which may be undesired, especially for reconfigurable plants with fast-changing production targets. In addition to traditional first-principles modeling for MPC, recent research explores leveraging different kinds of recurrent neural networks to model system dynamics within MPC frameworks, enabling efficient data-driven system identification and enhanced control performance.[71] To maintain connection to first principles, physics-informed neural nets for control (PINC), which extends PINNs with control inputs and an autoregressive evaluation over short intervals, has been shown to be compatible with MPC-type formulations by enabling simulations over variable time horizons and facilitating control tasks through efficient inference.[72]

RL increasingly complements rather than replaces the classical MPC (Table 1). In practice, RL is most effective for policy tuning, disturbance forecasting, and supervisory setpoint optimization, where MPC or PID controllers remain the primary workhorse in industrial environments.[22] A pragmatic middle ground has emerged in the form of control-informed RL, which embeds classical control structure and knowledge (e.g., PID/MPC priors, stability certificates) within RL algorithms to enhance stability and sample efficiency.[10]

Contemporary safe-RL practice emphasizes formulations and tools well-suited to chemical processes: constrained Markov decision processes (CMDPs), shielded policies, Lyapunov and barrier-function methods, and offline or model-based RL trained explicitly to remain within safety envelopes.[73] Sector-specific reviews (e.g., power and energy systems[74]) document mature safety strategies, such as risk-aware rewards, constraint critics, and certified fallback policies that are readily transferable to chemical process industries with minimal adaptation.

A landmark industrial demonstration of control-informed RL is Yokogawa and JSR's deployment of the Factorial Kernel Dynamic Policy Programming (FKDPP) algorithm in a full-scale chemical plant.[75,76] The RL controller operated autonomously for 35 days, maintaining product quality, yield, and energy efficiency while handling disturbances previously handled manually via PID/advanced process control (APC). FKDPP complements PID and MPC by autonomously optimizing multivariable processes under conflicting objectives, achieving faster stabilization, safe integration with DCS interlocks, and digital-twin compatibility, illustrating control-informed RL in real-world industrial practice.

To bridge the sim-to-real gap, controller transfer workflows typically combine domain randomization (e.g., vary kinetics, dead times, noise), explicit disturbance modeling, and stress testing against unseen upsets and actuator limits before deployment.[77–79] In process control contexts, recent studies design sim-to-real pipelines around process-specific models, demonstrating how simulation-trained agents can be transferred with plant-realistic perturbations and validation steps.[78,80] A recent review echo this trend toward plant-tailored RL transfer in process industries.[22]

On the safety and governance front, robust deployments align controller behavior with existing operating procedures and trip logic by incorporating runtime safety layers. These layers include shielding mechanisms that block unsafe actions and Lyapunov/barrier function

TABLE 1 Comparison of traditional vs. AI/ML-based control methods.

Aspect	Mechanistic	Current practice	AI/ML-based
Basis	MPC[a]	PID	Data-driven models (e.g., policy trained by RL, forecast model as internal controller model of MPC)
Models	First-principles, linear or nonlinear	Linear, analytic	Nonlinear, learned
Complexity handling	Limited for unmodeled dynamics	Limited	Excellent
Explainability	High	High	Low/medium
Adaptivity	Low	Low	High (can learn online)
Data requirement	Low/moderate	Low	High
Trust and safety	Mature	Mature, reliable	Requires safeguards
Best-case scenario	Known physics	Stable, well-known processes	Complex, nonlinear, (slowly) changing processes

[a]MPC is also common for multivariable plants where the input–output relations are experimentally determined, for example, by step responses.

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

or CMDP-based formulations that enforce hard constraints—patterns that map readily to interlocks and shutdown thresholds traditionally used in industrial plants.[81]

Digital twins (DTs), which are real-time, adaptive representations of physical plants, promise continuous monitoring, prediction (e.g., future conditions, predictive maintenance), optimization, and control. Yet full implementation in industry remains challenging.[82] RL is emerging as a key AI-in-the-loop enabler for DTs, especially where first-principles models are incomplete or control problems are too complex for conventional optimization solvers. Advances in dynamic hybrid modeling reinforce this trajectory. For instance, Caspari et al.[83] introduced an incremental identification framework coupling mechanistic and ML models for dynamic process control. Their hybrid MPC demonstrations exemplify the "control-informed AI" philosophy, leveraging physics-based structure for safe, data-efficient learning and reliable DT deployment.

Beyond tracking setpoints, modern control objectives can be formulated in terms of profit, energy, or emissions. The economic MPC (EMPC) framework specifically targets economic optimization. Foundational EMPC studies establish stability and performance guarantees via dissipativity theory and Lyapunov-based analyses, offering practical design recipes.[84] In deployment, EMPC typically resides within a two-layer hierarchy, where a real-time optimization (RTO) or supervisory layer optimizes economics and a regulatory layer enforces constraints.[85] Building on this foundation, two-layer RL-EMPC hybrids are emerging.[86] In these systems, an RL agent in the upper layer learns to optimize economic objectives or setpoints under uncertainty and dynamic pricing, while EMPC (or MPC) in the lower layer handles multivariable constraints and actuator realities. Recent studies describe RL for real-time optimization (RL-RTO) and setpoint selection on top of MPC/EMPC, as well as control-informed RL frameworks that embed classical structure to ensure stability and sample efficiency.[87] Parallel developments in RL-based EMPC and MPC-for-RL approaches demonstrate practical integrations where learning augments economic objectives yet defers constraint enforcement to MPC, aligning with established industrial control hierarchies and thus ease adoption.[88]

6 | MLOps (MACHINE LEARNING OPERATIONS) FOR PROCESS INDUSTRIES

In the context of chemical engineering, MLOps can be viewed as the lifecycle management discipline for machine learning models. Just as pumps, reactors, and control loops require commissioning, validation, documentation, and periodic requalification, ML models require structured governance once deployed. MLOps formalizes this lifecycle (including development, testing, deployment, monitoring, retraining, and retirement) to ensure models remain reliable under changing plant conditions.

Table 2 puts MLOps in perspective relative to familiar product lifecycle management (PLM) paradigms. Unlike traditional software systems, ML systems are probabilistic and data-dependent, which fundamentally changes how they must be tested, validated, and maintained. Compared to traditional PLM, MLOps can be highly beneficial to process industries because both continuous and batch operations are subject to drift, non-idealities, equipment changes, and evolving operating procedures. As a result, models must be managed as engineering lifecycles, requiring commissioning, monitoring, drift detection, re-qualification, and redeployment, with documented tests and sign-offs. A recent bioprocess study exemplifies this end-to-end loop: an industrial fed-batch fermentation soft-sensor pipeline automating development, deployment, online monitoring, and maintenance within an integrated MLOps framework.[89]

Effective industrial MLOps platforms combine several capabilities: version control for both data and models, reproducible training pipelines, automated deployment workflows (CI/CD), continuous performance monitoring, and documented approval processes.[24] These mechanisms ensure that any model update can be traced, tested, and, if necessary, rolled back, which is similar to established change-control procedures in plants. Practitioner frameworks from large-scale ML operations describe recurring patterns (e.g., continuous training, champion–challenger evaluation, rollback), while industrial ML guidance emphasizes data engineering, monitoring, logging/auditing, and business key performance indicators (KPIs) as integral platform features.

TABLE 2 Comparison of traditional PLM, DevOps, and MLOps.

Aspect	Traditional PLM	DevOps	MLOps
Primary focus	Executable delivery (apps/services)	Frequent deployment of updates and enhancements	Model lifecycle, data, code
Artifacts	Apps, services	Apps, services	Trained ML models
Determinism	High	High	Probabilistic
Data dependency	Low	Low	High
Testing	Functional and integration	Continuous integration and test	Data, model performance, fairness
Monitoring	System metrics	Health status	Drift, data quality, predictions
Tooling	Executable	Containers	CI/CD + ML-specific tooling
Retraining	Not needed	Not needed	Essential
Team	Developers, Architect, UX-Expert	Developers, Architect, UX-Expert, Ops	Data scientist + ML engineer + Data Engineer + Ops

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Integration of ML with process-control and data infrastructure is needed. MLOps in process plants must interface seamlessly with historians, DCS, and manufacturing execution system (MES) layers via robust, secure messaging protocols. In the classical manufacturing model (ISA-95), Enterprise Resource Planning (ERP; Level 4) handles business planning (e.g., orders, inventory), while the Manufacturing Execution System (MES; Level 3) translates plans into executable recipes and schedules. At Levels 1–2, supervisory control and data acquisition (SCADA) systems, distributed control systems (DCS), and programmable logic controllers (PLCs) perform real-time control and monitoring. Data flows top-down (from plans to setpoints) and bottom-up (maps status/quality to KPIs), typically via ISA-95/ISA-88 models and OPC UA interfaces. Using ISA-95 models ensures clean handoffs (ERP↔MES↔SCADA/PLC) and, with OPC UA, provides secure, interoperable telemetry, which is essential for low-latency scoring, reliable fallbacks, and traceable root-cause analysis.

A monitoring-first discipline is vital for production ML. Systems should continuously check for data, target, and feature drift, enforce service-level agreements (SLAs) for model performance, and trigger operator actions (e.g., revert, retrain). Studies describe practical drift detection and remediation patterns for production ML, particularly important in safety- and quality-critical contexts.[90] To promote robustness, engineering checklists such as Hidden Technical Debt (designing for maintainability) and the ML Test Score (28 actionable tests for production readiness) can be adapted to process-safety contexts (e.g., stress-testing worst-case scenarios, sensor dropouts, and actuator saturation).[91] Every model should be treated like an instrument: shipped with Model Cards (intended use, limits, slices) and Datasheets for Datasets (provenance, sampling, consent/ethics where relevant).[92] Such documentation accelerates audits, change control, and knowledge transfer across shifts and sites.

For structured risk management, several emerging standards provide guidance. The NIST AI Risk Management Framework (RMF 1.0; Govern/Map/Measure/Manage) offers a voluntary framework for identifying, measuring, and managing AI-related risks. ISO/IEC 23894:2023 (AI risk-management guidance) and ISO/IEC 42001:2023 (AI—Management system) extend this guidance into organizational risk-management and governance systems. In the EU, the AI Act (effective since 1 August 2024)[18,19] introduces regulatory obligations for high-risk AI systems, including post-market monitoring, incident logging, traceability, and technical documentation, requirements of which are conceptually analogous to safety case documentation in process industries. Additionally, ISO/IEC 23053:2022 defines a conceptual framework for AI systems employing ML, detailing components, functions, and lifecycle phases, including development, deployment, operation and monitoring. ISO/IEC 42001:2023 provides a structural approach to managing the risks and opportunities associated with AI within organizations.

A structured operational protocol for industrial MLOps is emerging.[24] The recommended sequence begins with commissioning in "shadow" mode, where models are validated using historian data and accompanied by a published Model Card describing intended use and limitations. Next, the model is deployed behind safety layers, employing canary gating and operator overrides, with telemetry connected to continuous monitoring services. Once deployed, ML systems must be treated as continuously operating assets. During operation, systems should monitor for drift (e.g., data input, target drift), constraint violations, and business KPIs, integrating these insights into established plant workflows. Requalification is activated by defined triggers (e.g., drift, MOC event, equipment change), followed by re-execution of acceptance tests (e.g., ML Test Score items, plant-specific stress suites). Finally, models are (re)deployed via CI/CD pipelines with versioned artifacts and auditable approvals, archiving all runs and decisions. This protocol merges modern MLOps discipline with the assurance and change-control culture that has long underpinned reliability in process industries.

While most plant AI systems are advisory rather than safety-instrumented systems (SIS), they should adopt comparable lifecycle rigor. Such systems should coexist with existing safety interlocks (e.g., IEC 61511) and cybersecurity standards that protect Safety Controls/Alarms/Interlocks (e.g., ISA-TR84.00.09, ISA/IEC 62443), preserving fail-safe behavior and formal change-control. The current industrial consensus is to keep ML systems advisory or supervisory, preserve existing interlocks, and maintain explicit change-control and fail-safe behavior, as described for example in the NAMUR Open Architecture (i.e., NAMUR Open Architecture NOA Concept (NE 175), NAMUR Open Architecture NOA Verification of Request (NE 178)).

The bottom line is that industrial MLOps is not about another toolchain, but about institutionalizing a lifecycle familiar to plants, such as requirements, verification, handoffs, and monitoring, augmented for data-driven systems and aligned with emerging AI governance. With these practices in place, ML can evolve from lab prototypes to reliable, industry-grade assets.

7 | SUSTAINABILITY DECISIONS AND PROSPECTIVE LIFE-CYCLE ANALYSIS (LCA)

At the design stage, structure-based ML surrogates now enable orders-of-magnitude faster screening of chemicals' life-cycle impacts compared with conventional LCAs (Table 3). These models support early portfolio triage and rapid iteration during route and process selection. Recent studies demonstrate that ANN- and GNN-based impact predictors can estimate life-cycle indicators directly from molecular descriptors or product structure.[93,94] A key requirement for forward-looking sustainability analysis is prospective LCA (pLCA), which accounts for future conditions (e.g., energy mix, supply chains, technology learning). The Premise tool addresses this by projecting the ecoinvent database into future years using integrated assessment model (IAM) scenarios (e.g., REMIND, IMAGE), ensuring LCAs are evaluated against scenario-consistent grid mix, supply chains, and technology pathways rather than today's conditions.[95] Recently, TIAM-UCL scenarios have also been incorporated in Premise to enhance scenario diversity for pLCA.[96]

Methodological progress in ML-LCA has been well mapped. Approaches span QSPR/QSTR molecular descriptors, graph neural

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

TABLE 3 Comparison of traditional vs. AI/ML-based sustainability decisions and prospective LCA.

Aspect	Traditional	AI/ML-based
Foundation	Engineering models, mechanistic equations, expert assumptions, ISO 14040/44 processes	Data-driven models, learning patterns from large datasets, hybrid process-aware models
Forecasting	Scenario-based, deterministic or linear projections	Predictive ML models, nonlinear, data-adaptive forecasts
Uncertainty	Expert judgment, sensitivity analysis	Probabilistic ML, uncertainty quantification, real-time adaptation
Decision support	Static frameworks (e.g., ISO LCA); some dynamic models and multi-objective optimization	Dynamic, predictive, multi-objective optimization, recommendations
Data requirements	Limited datasets, literature values, heuristics	Big data, sensor data, real-time environmental data
Granularity	Coarse system boundaries, broad assumptions	Fine-grained insights from high-resolution data
Interpretability	Highly interpretable but sometimes oversimplified; dependent on modeler expertise	May be less interpretable, but explainable AI (e.g., SHAP) supports transparency
Adaptivity over time	Changes require full reassessment; limited responsiveness to new data	Continual learning and model updates as new data streams become available
Development vs. execution effort	High upfront development effort driven by deterministic modeling and calibration; low to moderate execution cost per scenario due to repeated deterministic model evaluation	High ongoing training and retraining effort associated with data curation and model updates; very low execution cost per scenario via trained-model inference
Application	Baseline assessments, compliance reporting, high-level policy evaluation	Real-time strategy optimization, early-stage sustainability forecasting, circularity prediction, dynamic eco-design

features, and process-aware regressors that fill life-cycle inventory (LCI) and life-cycle impact assessment (LCIA) gaps. Several comprehensive reviews (2022–2025) summarize these options, emphasizing best practices in method selection, uncertainty handling, and validation.[97–100] Beyond molecular-level surrogates, process-specific predictive LCA frameworks combine ML with automated process design to estimate impacts early from reaction stoichiometry and molecular structure, which are useful when full mass and energy balances are unavailable.[101]

Beyond direct prediction of life-cycle indicators, ML models can also predict process-relevant quantities directly from molecular structure. Structure-based approaches (e.g., QSPR/QSTR methods, graph neural networks (GNNs), and physics-aware regressors) can estimate thermophysical properties (e.g., vapor pressure, solubility, heat capacity), reaction kinetics, transport coefficients, and separation-relevant parameters from molecular descriptors or graph representations.[102] These predicted properties can then serve as inputs to process simulators or surrogate models, enabling early-stage process screening when experimental data are limited. In this way, molecular-scale predictions propagate into mass and energy balances, equipment sizing, and ultimately sustainability metrics. However, this multi-scale coupling can amplify uncertainty: prediction errors at the molecular level may propagate nonlinearly through process and LCA models. Best practice therefore includes explicit uncertainty quantification and validation at both molecular and process scales when using structure-based ML predictions for sustainability decision-making. Recent efforts are beginning to integrate structure-based ML property predictors directly into automated process design and optimization frameworks, enabling closed-loop evaluation of candidate molecules under combined process and sustainability constraints. For example, integrated ML models have been coupled with computer-aided molecular and process design (CAMPD) frameworks to simultaneously optimize molecular selection and downstream process performance, demonstrating improvements in energy efficiency and process performance.[103]

In parallel, ML is used to fill LCI and LCIA data gaps, including prediction of ecotoxicity characterization factors (e.g., USEtox characterization factors via ML on physico-chemical properties) and estimation of missing inventory elements using neural networks or gradient-boosted models, enhancing generalizability to novel chemicals and emerging processes.[104,105] One study formalizes ex-ante, ML-assisted LCA for bio-based processes, combining ANN and decision-tree models with mixed inventory sources to enable environmental-driven screening before detailed design.[106] In practice, pipelines are trending toward GNN-embedded LCA surrogates that leverage molecular-graph structure, outperform conventional descriptor-based models, and scale effectively to large chemical libraries for early screening.[94]

To translate screening into actionable decisions, ML-LCA outputs can be incorporated as additional objectives or constraints (e.g., global warming potential (GWP), ecotoxicity) in multi-objective process design, scheduling, or control frameworks, thereby coupling economics and sustainability in closed-loop studies. A recent review maps ML's roles within LCA (e.g., predicting LCIA outcomes, filling LCI gaps, clustering, early-stage screening) and notes that ANNs and Random Forests dominate current practice, with calls to explore Bayesian and other underused models.[100] The same review highlights two cross-cutting cautions: (i) data quality and data shift, given static inventories or sparse/mismatched datasets; and (ii) uncertainty propagation, where ML predictive uncertainty must be carried through to LCA outcomes. Recommended practice include coupling ML with explicit UQ (ensembles/Bayesian regressors or quantile models), integrating predictive variance into Monte-Carlo LCA, reporting both scenario and model uncertainty instead of point estimates alone, and preference for dynamic or perspective LCAs so that background changes (e.g., grid mix) are correctly reflected.

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Finally, reporting and assurance should follow ISO 14040/14044 principles with respect to documenting dataset provenance and scenario assumptions, validating predictions against held-out LCAs (i.e., independent LCA datasets reserved for validation), and reporting uncertainty (intervals, sensitivity). These practices are increasingly emphasized across recent ML-LCA[99] and pLCA[107] studies, providing a foundation for credible, decision-ready sustainability assessments.

An additional sustainability consideration concerns the energy and carbon footprint associated with training and operating AI models themselves. Training large foundation models can require substantial computational resources, and the associated environmental impact depends on model size, hardware efficiency, and the carbon intensity of the datacenter energy mix.[108] From an LCA perspective, this raises an important system-boundary question: should the energy and emissions associated with model development and deployment be included in sustainability assessments of AI-enabled decision frameworks? In many industrial settings, the one-time cost of model training or fine-tuning may be small relative to long-term operational gains, such as reduced energy use, optimized feedstock consumption, or avoided waste. Transparent reporting of computational demand and carbon intensity is also consistent with ISO 14040/44 principles of completeness and system-boundary clarity. Emerging "green AI" practices (e.g., energy-efficient model architectures, hardware-aware optimization, and carbon accounting of compute workloads) offer pathways to ensure that AI-enabled sustainability tools do not inadvertently undermine the environmental objectives they aim to support.[109] Incorporating computational energy use into prospective LCAs of AI-enabled workflows may become increasingly relevant as models scale and as sustainability reporting standards evolve.

8 | GOVERNANCE, STANDARDS, AND REGULATION

AI systems that inform setpoints, alarms, release predictions, maintenance, and operational decisions introduce obligations in functional safety, cybersecurity, data stewardship, validation, and regulatory compliance. "Responsible" AI must be lawful, ethical, robust, and auditable throughout its lifecycle, rather than just being accurate on benchmark datasets.[20,110]

8.1 | Regulatory landscape interfacing with chemical plants

Since 1 August 2024, the EU AI Act has been in force,[18,19] introducing a risk-based regulatory regime for AI. Requirements for general-purpose AI (GPAI) and high-risk systems are phased in progressively through 2025–2026, covering conformity assessment, technical documentation, incident reporting, and post-market monitoring. These obligations align closely with long-standing governance expectations in the process industries.

Sector-specific safety and cyber standards by the International Electrotechnical Commission (IEC) remain binding safeguards. For example, IEC 61511 governs the safety instrumented systems (SIS) for process industries. If AI influences safety functions or related decision support, the overall safety performance must still satisfy IEC 61511. Similarly, IEC 62443 provides a comprehensive cybersecurity framework for industrial automation and control systems (IACS) and is applicable to modern architectures including components of digital-twin and cloud-enabled systems.

In parallel, the NIST AI Risk Management Framework (AI RMF 1.0)[111] and its companion Generative AI Profile (2024) serve as voluntary but widely adopted frameworks to assist organizations in systematically governing, mapping, measuring, and managing AI-related risks across the lifecycle of AI systems.

8.2 | Validation, verification, and assurance

Lifecycle-oriented validation and verification (V&V) that integrates domain knowledge and ML-specific considerations (i.e., gray-box models), with explicit uncertainty quantification, drift detection, and change control, is essential to augment confidence.[30,112]

Safety-engineering research, notably the Assurance of Machine Learning for use in Autonomous Systems (AMLAS) methodology, provides structured guidance for developing evidence-based safety cases for ML components in autonomous or safety-critical systems.[113] AMLAS outlines six stages: safety-assurance scoping (linking system hazards to the ML component), ML safety requirements elicitation, data-management assurance, model-learning assurance, model verification, and deployment/operational assurance including monitoring and change control. These frameworks are directly applicable when AI informs safety-related decisions in chemical plants.

8.3 | Safety and human-factors governance

When AI contributes to alarms, interlocks, abnormal-situation management, or operator decision support, organizations must demonstrate that Safety Instrumented System (SIS) and basic process control system (BPCS/PCS) performance under AI-assisted operation continues to meet all IEC 61511 safety-integrity requirements. Human-factors provisions of the standard also apply: operators must maintain situational awareness, manageable workload, and clear escalation paths.

For cyber-safety, plants can apply IEC 62443 cybersecurity principles, originally defined for industrial automation and control systems, to adjacent components such as data pipelines, model repositories, and runtime endpoints, implementing defense-in-depth across the AI lifecycle where these elements interact with the operational technology (OT; includes DCS, SCADA, PLC, SIS, etc.) environment.

8.4 | Practical governance

Operationalizing responsible AI in ChE begins with mapping legal and regulatory exposure, ensuring that AI-enabled applications comply

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

with the EU AI Act and align with existing product and process safety frameworks. This should be coupled with the NIST AI RMF roles and functions. Internally, project gates must explicitly reference established industrial standards, such as IEC 61511 for SIS and IEC 62443 for cybersecurity, particularly when AI influences control, alarms, or automation. Data governance should be equally rigorous, enforcing FAIR principles for provenance, versioning, retention, and access control,[114] while requiring Model Cards and Datasheets to accompany every deployed model.[92] On the assurance side, gray-box modeling should be prioritized where feasible, complemented by formal V&V plans, drift and incident logging, structured change control and management-of-change (MOC), and operator training that incorporates HAZOP-style analyses of AI failure modes. Finally, organizations must institutionalize audit and monitoring practices: post-market surveillance as mandated by the EU AI Act, periodic refresh of safety cases, stress-testing of models, and cyber-security measures in line with IEC 62443. Together, these measures transition AI from research prototypes into industry-ready, auditable systems that integrate safely and transparently into the operational fabric of chemical plants.

9 | DATA, BENCHMARKS, AND OPEN RESOURCES

Data (and, more critically, data integrity) remain defining bottlenecks for scaling and deploying AI in ChE. Plant-level records are often inconsistent: setpoints may not reflect actual PLC execution, manual interventions tend to be insufficiently logged, and gaps or corrupted entries can erode both reproducibility and model trustworthiness. Scale introduces an additional challenge: laboratory and pilot-plant datasets are typically the richest and most carefully measured, yet reliably transferring these insights to industry-scale operations demands robust protocols for hierarchical data fusion, uncertainty quantification, and validation across technology readiness levels (TRLs).

Addressing these challenges requires robust data standardization. Several established standards support interoperability across design, operations, and enterprise systems. ISO 15926 and DEXPI provide structured representations of process equipment and P&ID semantics, enabling consistent exchange across tools and vendors. ISA-95 and ISA-88 define hierarchical integration between enterprise resource planning (ERP) systems (business planning such as orders and inventory), manufacturing execution systems (MES; production scheduling, execution, and traceability), and plant control layers (SCADA/DCS/PLCs), forming the backbone of industrial data architecture. OPC UA (Open Platform Communications Unified Architecture) enables secure, interoperable machine-to-machine communication for real-time telemetry between industrial control systems and higher-level analytics platforms. Emerging Industry 4.0 frameworks, such as the Asset Administration Shell (AAS), further promote standardized asset metadata to support digital twin integration. Alignment with these standards enhances data portability, traceability, governance, and long-term maintainability of AI-enhanced systems in the process industries.

To accelerate progress, the community must also leverage curated and openly accessible resources spanning materials design, thermophysical properties, kinetics, process operations, and safety records. The growing ecosystem of design and discovery databases, thermophysical property archives, reaction repositories, AI-driven catalyst and polymer discovery platforms, benchmarked process datasets, open-source control toolboxes and LCA toolkits, and accident and safety databases offers essential building blocks for reproducible research and reliable AI development. When paired with canonical benchmarks and rigorous validation protocols, these resources can help bridge the persistent gap between demonstration-scale AI workflows and robust industrial deployment.

Some examples of such databases are listed in the Supporting Information Appendix.

10 | INTEGRATING AI IN ChE EDUCATION

The increasing centrality of AI to ChE practice makes education and workforce preparation an urgent priority. AI applications now permeate process modeling, control, optimization, and sustainability decision-making. This means future chemical engineers must be prepared not only to use these AI-enabled tools, but also to develop, scrutinize, and responsibly deploy them. This need has prompted calls for curriculum redesign, faculty upskilling, and expanded lifelong-learning programs for practicing engineers.[115] International guidance echoes these priorities: UNESCO's GenAI framework and the OECD-Education International guardrails emphasize human-centered design, transparency, teacher capacity, data protection, equity, and explainability as prerequisites for effective AI adoption in education.[116,117]

A recent article argues that AI and data science must be woven into the core of the ChE curriculum, rather than treated as electives or peripheral add-ons.[118] Program-level framing is critical: rather than "outsourcing" computational content to computer science departments, ChE programs are increasingly embedding data- and AI-driven approaches into transport phenomena, reaction engineering, process systems engineering, and design courses. Accreditation structures already permit this evolution: the Accreditation Board for Engineering and Technology (ABET) criteria mandate inclusion of appropriate mathematics and basic sciences while giving ChE programs the flexibility to integrate AI into engineering subjects and computational coursework, enabling modernization and substantial AI content without displacing core ChE fundamentals.[119] Several universities have begun piloting dedicated AI/ML courses tailored to chemical engineers at undergraduate and graduate levels, with feasible implementation pathways detailed in terms of topics, prerequisites, and assessment strategies.[120,121] These courses are increasingly complemented with hands-on labs and open educational resources, including datasets, coding exercises, and open-source toolkits, which facilitate reuse, adaptation, and cross-institutional collaboration.[122] Such materials provide essential practice in model development and validation while fostering reproducibility through shared benchmarks.

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Parallel efforts in industry and professional education reflect similar priorities. Industrial organizations have begun investing heavily in workforce upskilling to support digital transformation. Some companies have established internal AI academies, training programs in data engineering and MLOps, cross-functional digital teams, and partnerships with universities and technology providers. Shell, for example, partnered with Udacity to co-develop and deliver AI training to thousands of employees as part of its Shell.ai initiative,[123] while DuPont's Spark Digital Academy[124] provides company-wide instruction in data science and digital technologies through project-based learning. Professional organizations such as AIChE have introduced practitioner-oriented offerings, including the "AI for Chemical Engineers" course series.[125] Such initiatives are complemented by a growing number of AI courses offered by universities in different parts of the world. These range from programs accessible to working professionals alongside traditional students to more specialized offerings designed specifically for practicing chemical engineers. The latter are distinguished not only by content focused on relevant AI techniques, domain-specific case studies, and hands-on exercises, but also by delivery formats adapted to accommodate professional schedules. Collectively, these efforts reflect three key dimensions of recent developments in AI education for chemical engineers: (i) company-led upskilling initiatives, (ii) professional society offerings, and (iii) university-delivered continuing education courses. In this evolving landscape, it is imperative that academic institutions provide foundational training in process modeling, statistics, optimization, and transport phenomena, while industry-led programs emphasize deployment readiness, governance, toolchains, and organizational change management. Effective AI adoption in ChE will benefit from sustained collaboration between academia and industry, ensuring that engineers are both technically grounded and prepared for real-world implementation. In parallel, governments and public policy frameworks in multiple countries are increasingly emphasizing the strategic role of universities in lifelong learning and workforce upskilling. Through national AI and digital transformation initiatives, public funding and policy support are being directed toward academic institutions, enabling them to extend their mission beyond traditional degree programs and actively support continuing education, reskilling, and professional development throughout engineers' careers. GenAI brings both opportunities and challenges for teaching and learning.[126,127] Case studies highlight how large language models (LLMs) can act as tutoring, brainstorming, or coding assistants, aiding students in navigating complex workflows or debugging codes. At the same time, students must understand that responsible use requires clear policies on disclosure, academic integrity, and assessment redesign. Beyond pedagogy, hybrid models that combine first-principles knowledge with data-driven methods offer students exposure to the frontier of mechanism-informed AI. These approaches reinforce core ChE reasoning while engaging modern AI paradigms, bridging symbolic and numerical methodologies.[128] Such hybrid framing equips future chemical engineers with the mastery needed to convert AI's technological promise into safe, reliable industrial implementation.

11 | ROADMAP AND CALLS TO ACTION

The roadmap for AI in ChE must begin with education. Without proper training, engineers risk either treating AI as an unquestionable oracle or dismissing it altogether when misapplied. It is necessary not just to update the curriculum for undergraduates but also reskilling practicing chemical engineers to engage advantageously with AI systems. Without this cultural and educational grounding, technical advances may fail to translate into practice.

A priority is to build true two-way ML-ChE teams, not mere tool hand-offs. This requires funding interdisciplinary labs, joint appointments, and shared PhD training pipelines, with KPIs that combine process physics, optimization, and machine learning. Schweidtmann et al.[129] argue that ChE's transformation hinges on such joint efforts, pointing to critical gaps in physics–ML integration, data quality, and trust. Alongside, incubator labs that span proof-of-concept through pilot to plant can be created. These shared environments, combining datasets, simulators, and rigs, would allow vendors, plants, and universities to co-develop and validate AI workflows such as P&ID digitization, soft sensors, and RL-assisted control against common evidence packs. A recent study outlined such an incubator model for process industries, highlighting the need for shared testbeds and translational funding.[130]

Sector-specific roadmaps must also be published with phased value capture in mind. Each industry (e.g., bulk chemicals, specialty, pharma, energy) can identify near-term "fast lanes" like anomaly detection or hybrid surrogates; mid-term opportunities such as GenAI for documentation and HAZOP preparation or learning-enhanced MPC; and long-term visions including plant-wide coordination and autonomous operations. Bortz et al.[131] provide a useful template, mapping technical advances to adoption phases with associated risks and benefits. To underpin credibility, open benchmarks and evidence standards should be codified inside incubator environments. These should include canonical dataset splits, RL vs. NMPC baselines, constraint and safety metrics, and "golden-set" evaluation harnesses for GenAI/RAG tasks, with all releases tied to Model Cards and Datasheets. As Schweidtmann et al.[129] and Kockmann et al.[130] emphasize, trustworthy evaluation demands a community-endorsed benchmarking charter.

Finally, governance and education must advance in lockstep. Incubators can deliver governance playbooks aligned with both the EU AI Act and SIS lifecycles, including role-split templates (provider vs. deployer), technical-file checklists, post-market monitoring plans, and change-control procedures for retraining.[131] In parallel, education pipelines should be re-designed to mirror these roadmaps: required Python/ML foundations, modules on gray-box modeling, and plant-grade MLOps laboratories. Schweidtmann et al.[129] highlight the skills gap and interdisciplinary needs that justify such curriculum changes and industry-sponsored capstones. Together, these steps can align technical capability, regulatory assurance, and workforce readiness, ensuring that AI in ChE evolves from promise to sustainable practice.

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

11.1 | Challenges: Education

As AI systems increasingly generate technically plausible solutions in seconds, the traditional markers of chemical engineering competence—derivations, simulations, and point solutions (i.e., solutions that address a single, narrowly defined problem instance under fixed assumptions)—are no longer sufficient. The distinctive value of the chemical engineer now lies in interrogating assumptions, stress-testing automated outputs, and taking responsibility for safety, ethics, and real-world consequences that AI cannot own. This shift exposes a fundamental misalignment in ChE education: curricula and assessments continue to reward solution execution over judgment, creativity, and professional accountability. Without a deliberate reorientation toward these human-centered capabilities, AI risks amplifying errors rather than insight. As AI becomes deeply embedded in ChE practice, excellence will be defined not by computational speed, but by the engineer's ability to govern, contextualize, and responsibly deploy intelligent systems.

To meet this challenge, ChE programs must cultivate "AI-fluent engineers", professionals who can both harness and critique intelligent tools, bridging domain expertise with algorithmic literacy to ensure safe, transparent, and impactful applications. This redefinition of excellence must be reflected not only in undergraduate curricula but also in continuing professional education, ensuring that both new graduates and practicing chemical engineers are equipped to critically evaluate, validate, and integrate AI methodologies within the rigor and responsibilities of ChE practice.

11.2 | Challenges: Industry

From an industry perspective, the barriers to adopting AI are not merely technical but structural and cultural. While access to high-quality data, such as engineering documentation, standard operating procedures, and historical process data, is a critical prerequisite, a persistent challenge is often insufficient data structure and semantic consistency. In practice, inconsistent formats, missing metadata, and fragmented taxonomies often hinder integration, interoperability, and downstream AI readiness, even if large volumes of data are available.

While China's state-led data governance and national data infrastructure initiatives can lower barriers to data aggregation for AI development, Western industrial ecosystems remain highly fragmented, with data access constrained by firm-level incentives, liability concerns, and competitive defensiveness. Without deliberate contractual, legal, and governance mechanisms to enable trusted data collaboration, many AI initiatives will remain confined to small-scale pilots with limited impact.

Regulatory uncertainty compounds this inertia. Emerging frameworks such as the EU AI Act classify AI systems that may affect machinery as high-risk, shifting significant legal responsibility onto both developers and deployers. At the same time, legacy legislation, most notably the Machinery Directive (2006/42/EC), does not explicitly address risks arising from AI-enabled interactions between machines or adaptive, learning-based behavior. This regulatory misalignment creates ambiguity over accountability, certification, and lifecycle obligations, discouraging investment in advanced AI capabilities, particularly generative and learning-enabled control systems. Without clearer harmonization between AI-specific regulation and established industrial safety standards, generally risk-averse organizations are likely to default to inaction.

Finally, organizational readiness remains a decisive bottleneck. As highlighted in Section 11.1, the effective use of AI depends on engineers who can critically evaluate, govern, and responsibly deploy intelligent systems. Yet many organizations treat AI as a peripheral IT initiative rather than a core operational capability, failing to invest in workforce understanding, cross-functional ownership, and cultural change. Unless AI literacy, trust, and accountability are embedded across vendors, operators, engineers, and management, even technically robust solutions will struggle to scale. In this sense, the industrial adoption challenge is inseparable from the educational challenge: without aligned skills, incentives, and organizational commitment, AI risks becoming another unrealized promise rather than a transformative tool for ChE practice.

12 | CONCLUSIONS AND PERSPECTIVES

AI in chemical engineering has progressed from promise to practice, though with uneven adoption across domains. The arc sketched by the Perspective article in 2019[1]—renewed momentum across design, control, diagnosis, and safety—has largely materialized, though with a different center of gravity than anticipated. Physics-aware (gray-box) modeling have matured from niche curiosities into plausible toolkits; reinforcement learning (RL) has become a credible complement to model predictive control (MPC); and generative AI (GenAI) has vaulted from curiosity to utility in process documentation, digitization, and HAZOP support. These approaches embed mechanistic knowledge inside learning architectures, improving robustness, extrapolation, and operator trust.

What is usable in industry today is clearest where AI augments rather than replaces established PSE practice. Monitoring and diagnosis via soft sensors and fault detection/diagnosis (FDD) are now trustworthy thanks to curated benchmarks; surrogate and hybrid models accelerate design and optimization within structured frameworks; and LLM/RAG co-pilots support operations by parsing SOPs, vendor manuals, and code repositories. GenAI's gains also come from turning legacy P&IDs and PFDs into DEXPI/ISO 15926 graphs, extracting and normalizing tags, and wiring RAG over manuals and incident databases (e.g., CSB, eMARS) to augment confidence based on knowhow. Treated as document-engineering problems with verification harnesses (golden sets, regression tests, retained logs), these deliver measurable value while meeting governance expectations.

Still, key frontiers remain research-grade. Fully autonomous control beyond well-bounded units, HAZOP/LOPA automation beyond decision-support tools, and large-scale mechanistic discovery via operator learning or symbolic regression all demand enhanced reliability, validation, and uncertainty-quantification pipelines. The bottleneck is shifting from algorithmic capability to assurance, deployment discipline, and workforce readiness. Models must be treated as engineered instruments—commissioned in shadow, documented with Model Cards and Datasheets, monitored for drift and KPI impact, and re-qualified under change control. Moving forward, incubator labs, shared benchmarks, governance playbooks aligned to the EU AI Act

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

and NIST RMF, and upskilling pipelines that blend physics, data, and MLOps will be decisive. With these pieces in place, AI in ChE can progress from isolated successes toward broad, safe, and sustainable implementation within years.

AUTHOR CONTRIBUTIONS

Jia Wei Chew: Conceptualization; methodology; investigation; validation; funding acquisition; writing—original draft; writing—review and editing; formal analysis; project administration. **Ronnie Andersson**: Methodology; investigation; validation; writing—original draft; writing—review and editing; formal analysis. **Thomas Bierweiler**: Methodology; investigation; validation; writing—original draft; formal analysis. **Patrik Ryttestal**: Writing—review and editing. **Torsten Wik**: Writing—review and editing.

ACKNOWLEDGMENTS

Jia Wei Chew acknowledges funding from Chalmers Gender Initiative for Excellence (Genie).

CONFLICT OF INTEREST STATEMENT

Thomas Bierweiler and Patrik Ryttestal are employed by Siemens, which develops and commercializes digitalization and AI technologies for the process industries. These affiliations are disclosed for transparency. The views expressed in this article are those of the authors and do not constitute endorsement of any specific commercial product or service. This review is intended as a general, industry-focused perspective sharing broadly applicable engineering practices from the industry perspective.

DATA AVAILABILITY STATEMENT

Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

ORCID

Jia Wei Chew https://orcid.org/0000-0002-6603-1649

REFERENCES

1. Venkatasubramanian V. The promise of artificial intelligence in chemical engineering: is it here, finally? *AIChE J*. 2019;65(2):466-478.
2. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596(7873):583-589.
3. Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys*. 2019;378:686-707.
4. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. *Nat Rev Phys*. 2021;3(6):422-440.
5. Chakraborty A, Taskiran NP, Kottooru R, Mann V, Venkatasubramanian V. Building hybrid AI models in chemical engineering: a tutorial review. *Comput Chem Eng*. 2025;201:109236.
6. Shah P, Pahari S, Bhavsar R, Kwon JS-I. Hybrid modeling of first-principles and machine learning: a step-by-step tutorial review for practical implementation. *Comput Chem Eng*. 2025;194:108926.
7. Balhorn LS, Caballero M, Schweidtmann AM. Toward autocorrection of chemical process flowsheets using large language models. In: Manenti F, Reklaitis GV, eds. *Computer Aided Chemical Engineering*. Vol 53. Elsevier; 2024:3109-3114.
8. Elhosary E, Moselhi O, Bucur C. Evaluation of AI-assisted HAZOP Software Tools. *IChemE Hazards 34 Symposium Series 171*. the Institution of Chemical Engineers (IChemE); 2024.
9. Schweidtmann AM. Generative artificial intelligence in chemical engineering. *Nat Chem Eng*. 2024;1(3):193.
10. Bloor M, Ahmed A, Kotecha N, Mercangöz M, Tsay C, del Río-Chanona EA. Control-informed reinforcement learning for chemical processes. *Ind Eng Chem Res*. 2025;64(9):4966-4978.
11. Crompton J. Chapter 5—Data management from the DCS to the historian. In: Bangert P, ed. *Machine Learning and Data Science in the Oil and Gas Industry*. Gulf Professional Publishing; 2021:83-110.
12. Ismail A, Truong H-L, Kastner W. Manufacturing process data analysis pipelines: a requirements analysis and survey. *J Big Data*. 2019;6(1):1.
13. Siivonen J, Naakka K, Salminen K, Kärki T, Aho M. Industrial data pipelines for manufacturing applications. *Procedia CIRP*. 2023;120:69-74.
14. Yeo WS, Saptoro A, Kumar P, Kano M. Just-in-time based soft sensors for process industries: a status report and recommendations. *J Process Control*. 2023;128:103025.
15. Feng C, Huang L, Krompass D. Only the curve shape matters: training foundation models for zero-shot multivariate time series forecasting through next curve shape prediction. arXiv Preprint arXiv:240207570. 2024.
16. Feng C, Huang L, Krompass D. General Time Transformer: an Encoder-only Foundation Model for Zero-Shot Multivariate Time Series Forecasting. *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*. Association for Computing Machinery; 2024:3757-3761.
17. SUPCON. *One Model Runs Your Plants: Time-Series Pre-Trained Transformer*. 2025. Accessed February 21, 2026. https://global.supcon.com/resources/tpt-time-series-pre-trained-transformer
18. AI Act Enters Into Force. European Commission. 2024. Accessed February 21, 2026. https://commission.europa.eu/news-and-media/news/ai-act-enters-force-2024-08-01_en
19. EU AI Act: First Regulation on Artificial Intelligence. 2023. Accessed February 21, 2026. https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence
20. Daniel T, Xuan J. Responsible use of generative AI in chemical engineering. *Digit Chem Eng*. 2024;12:100168.
21. Wu Y, Sicard B, Gadsden SA. Physics-informed machine learning: a comprehensive review on applications in anomaly detection and condition monitoring. *Expert Syst Appl*. 2024;255:124678.
22. Park J, Jung H, Kim JW, Lee JM. Reinforcement learning for process control: review and benchmark problems. *Int J Control Autom Syst*. 2025;23(1):1-40.
23. Ryttestal P, Siemens. Digitalization in chemical process engineering. *Chalmers Chemical Process Engineering (CPE) Meeting*. Chalmers University of Technology; 2025.
24. Salama K, Kazmierczak J, Schut D. *Practitioners Guide to MLOps: A Framework for Continuous Delivery and Automation of Machine Learning*. Google Cloud; 2021.
25. Perera YS, Ratnaweera DAAC, Dasanayaka CH, Abeykoon C. The role of artificial intelligence-driven soft sensors in advanced sustainable process industries: a critical review. *Eng Appl Artif Intell*. 2023;121:105988.
26. Misener R, Biegler L. Formulating data-driven surrogate models for process optimization. *Comput Chem Eng*. 2023;179:108411.
27. Hubbs CD, Li C, Sahinidis NV, Grossmann IE, Wassick JM. A deep reinforcement learning approach for chemical production scheduling. *Comput Chem Eng*. 2020;141:106982.
28. Liu L, Zhou W, Guan K, et al. Knowledge-guided machine learning can improve carbon cycle quantification in agroecosystems. *Nat Commun*. 2024;15(1):357.

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

29. Szatmári K, Horváth G, Németh S, Bai W, Kummer A. Resilience-based explainable reinforcement learning in chemical process safety. *Comput Chem Eng*. 2024;191:108849.
30. Decardi-Nelson B, Alshehri AS, You F. Generative artificial intelligence in chemical engineering spans multiple scales. *Front Chem Eng*. 2024;6:1458156.
31. Heredia Álvaro JA, Barreda JG. An advanced retrieval-augmented generation system for manufacturing quality control. *Adv Eng Inform*. 2025;64:103007.
32. Kim BC, Kim H, Moon Y, Lee G, Mun D. End-to-end digitization of image format piping and instrumentation diagrams at an industrially applicable level. *J Comput Des Eng*. 2022;9(4):1298-1326.
33. Moon Y, Lee J, Mun D, Lim S. Deep learning-based method to recognize line objects and flow arrows from image-format piping and instrumentation diagrams for digitization. *Appl Sci*. 2021;11(21):10054.
34. Sierla S, Azangoo M, Fay A, Vyatkin V, Papakonstantinou N. Integrating 2D and 3D digital plant information towards automatic generation of digital twins. Paper presented at: IEEE 29th International Symposium on Industrial Electronics (ISIE); 17–19 June 2020. 2020.
35. Ji X, Gong F, Wang N, Zhao Y, Ma Y, Shi Z. Pixel-level semantic parsing in complex industrial scenarios using large vision-language models. *Inf Fusion*. 2025;116:102794.
36. Vogel G, Schulze Balhorn L, Schweidtmann AM. Learning from flowsheets: a generative transformer model for autocompletion of flowsheets. *Comput Chem Eng*. 2023;171:108162.
37. Stops L, Leenhouts R, Gao Q, Schweidtmann AM. Flowsheet generation through hierarchical reinforcement learning and graph neural networks. *AIChE J*. 2023;69(1):e17938.
38. Peng S, Kalliamvakou E, Cihon P, Demirer M. The impact of AI on developer productivity: evidence from Github Copilot. arXiv Preprint arXiv:2302.06590. 2023.
39. Elhosary E, Moselhi O. Automation for HAZOP study: a state-of-the-art review and future research directions. *J Inf Technol Construct*. 2024;29(33):750-777.
40. U.S. Chemical Safety Board Releases Volume 2 of Chemical Incident Reports. 2025. Accessed March 10, 2026. https://www.csb.gov/us-chemical-safety-board-releases-volume-2-of-chemical-incident-reports/
41. Yu H, Gan A, Zhang K, Tong S, Liu Q, Liu Z. *Evaluation of Retrieval-Augmented Generation: A Survey*. Springer Nature; 2025.
42. Cameron DB, Otten W, Temmen H, Hole M, Tolksdorf G. DEXPI process: standardizing interoperable information for process design and analysis. *Comput Chem Eng*. 2024;182:108564.
43. Yuan Y, Chaffart D, Wu T, Zhu J. Transparency: the missing link to boosting AI transformations in chemical engineering. *Engineering*. 2024;39:45-60.
44. Runge J, Nowack P, Kretschmer M, Flaxman S, Sejdinovic D. Detecting and quantifying causal associations in large nonlinear time series datasets. *Sci Adv*. 2019;5(11):eaau4996.
45. Polak J, von Stosch M, Sokolov M, et al. Hybrid modeling supported development of an industrial small-molecule flow chemistry process. *Comput Chem Eng*. 2023;170:108127.
46. Specht T, Nagda M, Fellenz S, Mandt S, Hasse H, Jirasek F. HANNA: hard-constraint neural network for consistent activity coefficient prediction. *Chem Sci*. 2024;15(47):19777-19786.
47. Rosenberger D, Barros K, Germann TC, Lubbers N. Machine learning of consistent thermodynamic models using automatic differentiation. *Phys Rev E*. 2022;105(4):045301.
48. Cuomo S, Di Cola VS, Giampaolo F, Rozza G, Raissi M, Piccialli F. Scientific machine learning through physics–informed neural networks: where we are and what's next. *J Sci Comput*. 2022;92(3):88.
49. Zhao C, Zhang F, Lou W, Wang X, Yang J. A comprehensive review of advances in physics-informed neural networks and their applications in complex fluid dynamics. *Phys Fluids*. 2024;36(10):101301.
50. Schweidtmann AM, Zhang D, von Stosch M. A review and perspective on hybrid modeling methodologies. *Digit Chem Eng*. 2024;10:100136.
51. Sharma N, Liu YA. A hybrid science-guided machine learning approach for modeling chemical processes: a review. *AIChE J*. 2022;68(5):e17609.
52. Zendehboudi S, Rezaei N, Lohi A. Applications of hybrid models in chemical, petroleum, and energy systems: a systematic review. *Appl Energy*. 2018;228:2539-2566.
53. Narayanan H, von Stosch M, Feidl F, Sokolov M, Morbidelli M, Butté A. Hybrid modeling for biopharmaceutical processes: advantages, opportunities, and implementation. *Front Chem Eng*. 2023;5:1157889.
54. Lu L, Jin P, Pang G, Zhang Z, Karniadakis GE. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nat Mach Intell*. 2021;3(3):218-229.
55. Huang Y, Zou C, Li Y, Wik T. MINN: learning the dynamics of differential-algebraic equations and application to battery modeling. *IEEE Trans Pattern Anal Mach Intell*. 2024;46(12):11331-11344.
56. Bhan L, Qin P, Krstic M, Shi Y. Neural operators for predictor feedback control of nonlinear delay systems. *Proc Mach Learn Res*. 2025;283:1-26.
57. Ramezankhani M, Deodhar A, Parekh RY, Birru D. An advanced physics-informed neural operator for comprehensive design optimization of highly-nonlinear systems: an aerospace composites processing case study. *Eng Appl Artif Intell*. 2025;142:109886.
58. Shukla K, Oommen V, Peyvan A, et al. Deep neural operators as accurate surrogates for shape optimization. *Eng Appl Artif Intell*. 2024;129:107615.
59. Kovachki N, Li Z, Liu B, et al. Neural operator: learning maps between function spaces with applications to PDEs. *J Mach Learn Res*. 2023;24:1-97.
60. Azizzadenesheli K, Kovachki N, Li Z, Liu-Schiaffini M, Kossaifi J, Anandkumar A. Neural operators for accelerating scientific simulations and design. *Nat Rev Phys*. 2024;6(5):320-328.
61. Li Z, Kovachki NB, Azizzadenesheli K, Bhattacharya K, Stuart A, Anandkumar A. Fourier neural operator for parametric partial differential equations. Paper presented at: International Conference on Learning Representations. 2020.
62. Weng Y, Li H, Zhang H, Chen ZX, Zhou D. Extended Fourier neural operators to learn stiff chemical kinetics under unseen conditions. *Combust Flame*. 2025;272:113847.
63. Zhang H, Weng Y, Zhao Z, Zhou D. Learning transient evolution of multidimensional reacting flows by multiscale Fourier neural operators. *Proc Combust Inst*. 2024;40(1):105714.
64. Wang T, Yi Y, Yao J, Xu Z-QJ, Zhang T, Chen Z. Enforcing physical conservation in neural network surrogate models for complex chemical kinetics. *Combust Flame*. 2025;275:114105.
65. Hajireza K, Eklund D, Andersson R. Enhancing generalization and training efficiency of neural ordinary differential equations for chemical reactor modeling. *Chem Eng J*. 2025;519:165183.
66. Makke N, Chawla S. Interpretable scientific discovery with symbolic regression: a review. *Artif Intell Rev*. 2024;57(1):2.
67. Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. *Proceedings of the 31st International Conference on Neural Information Processing Systems*. Curran Associates Inc.; 2017:4768-4777.
68. Nori H, Jenkins S, Koch P, Caruana R. InterpretML: a unified framework for machine learning interpretability. arXiv Preprint arXiv:190909223. 2019.
69. Saltelli A, Ratto M, Andres T, et al. *Global Sensitivity Analysis*. The Primer; 2008.
70. Sobol IM. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Math Comput Simul*. 2001;55(1):271-280.

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

71. Wu Z, Christofides PD, Wu W, et al. A tutorial review of machine learning-based model predictive control methods. *Rev Chem Eng*. 2025;41(4):359-400.
72. Antonelo EA, Camponogara E, Seman LO, Jordanou JP, de Souza ER, Hübner JF. Physics-informed neural nets for control of dynamical systems. *Neurocomputing*. 2024;579:127419.
73. Inamdar R, Sundarr SK, Khandelwal D, Sahu VD, Katal N. A comprehensive review on safe reinforcement learning for autonomous vehicle control in dynamic environments. *e-Prime*. 2024;10:100810.
74. Bui V-H, Mohammadi S, Das S, Hussain A, Hollweg GV, Su W. A critical review of safe reinforcement learning strategies in power and energy systems. *Eng Appl Artif Intell*. 2025;143:110091.
75. Yokogawa. *In a World First, Yokogawa and JSR Use AI to Autonomously Control a Chemical Plant for 35 Consecutive Days*. Yokogawa Press Releases; 2022.
76. Cui Y, Zhu L, Fujisaki M, Kanokogi H, Matsubara T. Factorial kernel dynamic policy programming for vinyl acetate monomer plant model control. Paper presented at: 2018 IEEE 14th International Conference on Automation Science and Engineering (CASE); 20–24 August 2018. 2018.
77. Kadokawa Y, Zhu L, Tsurumine Y, Matsubara T. Cyclic policy distillation: sample-efficient sim-to-real reinforcement learning with domain randomization. *Robot Auton Syst*. 2023;165:104425.
78. Liang H, Xie J, Huang B, Li Y, Sun B, Yang C. A novel sim2real reinforcement learning algorithm for process control. *Reliab Eng Syst Saf*. 2025;254:110639.
79. Ju H, Juan R, Gomez R, Nakamura K, Li G. Transferring policy of deep reinforcement learning from simulation to reality for robotics. *Nat Mach Intell*. 2022;4(12):1077-1087.
80. Kubosawa S, Onishi T, Tsuruoka Y. Sim-to-real transfer in reinforcement learning-based, non-steady-state control for chemical plants. *SICE J Control Meas Syst Integr*. 2022;15(1):10-23.
81. Anand A, Seel K, Gjærum V, Håkansson A, Robinson H, Saad A. Safe learning for control using control Lyapunov functions and control barrier functions: a review. *Proc Comput Sci*. 2021;192:3987-3997.
82. Peterson L, Gosea IV, Benner P, Sundmacher K. Digital twins in process engineering: an overview on computational and numerical methods. *Comput Chem Eng*. 2025;193:108917.
83. Caspari A, Bierweiler T, Fadda S, et al. Dynamic hybrid modeling: incremental identification and model predictive control. *Comput Chem Eng*. 2026;204:109413.
84. Ellis M, Durand H, Christofides PD. A tutorial review of economic model predictive control methods. *J Process Control*. 2014;24(8):1156-1178.
85. Darby ML, Nikolaou M, Jones J, Nicholson D. RTO: an overview and assessment of current practice. *J Process Control*. 2011;21(6):874-884.
86. Zhang Z, Li S. Enhanced reinforcement learning in two-layer economic model predictive control for operation optimization in dynamic environment. *Chem Eng Res Des*. 2023;196:133-143.
87. Powell BKM, Machalek D, Quah T. Real-time optimization using reinforcement learning. *Comput Chem Eng*. 2020;143:107077.
88. Alhazmi K, Albalawi F, Sarathy SM. A reinforcement learning-based economic model predictive control framework for autonomous operation of chemical reactors. *Chem Eng J*. 2022;428:130993.
89. Metcalfe B, Acosta-Pavas JC, Robles-Rodriguez CE, et al. Towards a machine learning operations (MLOps) soft sensor for real-time predictions in industrial-scale fed-batch fermentation. *Comput Chem Eng*. 2025;194:108991.
90. Eck B, Kabakci-Zorlu D, Chen Y, Savard F, Bao X. A monitoring framework for deployed machine learning models with supply chain examples. Paper presented at: 2022 IEEE International Conference on Big Data (Big Data); 17–20 December 2022. 2022.
91. Sculley D, Holt G, Golovin D, et al. Hidden technical debt in machine learning systems. *Proceedings of the 29th International Conference on Neural Information Processing Systems*. Vol 2. MIT Press; 2015:2503-2511.
92. Mitchell M, Wu S, Zaldivar A, et al. Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery; 2019:220-229.
93. Zhao K, Peng X, Hu S, Wang X. Machine learning models for rapid prediction of chemicals' life-cycle environmental impacts: current status, challenges, and future directions. *Green Carbon*. 20253(4):433-442.
94. Gao Q, Schulze Balhorn L, Laera A, et al. Environmental impacts prediction using graph neural networks on molecular graphs. *Comput Chem Eng*. 2026;204:109362.
95. Sacchi R, Terlouw T, Siala K, et al. PRospective EnvironMental Impact asSEment (premise): a streamlined approach to producing databases for prospective life cycle assessment using integrated assessment models. *Renew Sust Energ Rev*. 2022;160:112311.
96. Šimaitis J, Butnar I, Sacchi R, Lupton R, Vagg C, Allen S. Expanding scenario diversity in prospective LCA: coupling the TIAM-UCL integrated assessment model with Premise and ecoinvent. *Renew Sust Energ Rev*. 2025;211:115298.
97. Neupane B, Belkadi F, Formentini M, et al. Machine learning algorithms for supporting life cycle assessment studies: an analytical review. *Sustainable Prod Consumption*. 2025;56:37-53.
98. Ghoroghi A, Rezgui Y, Petri I, Beach T. Advances in application of machine learning to life cycle assessment: a literature review. *Int J Life Cycle Assess*. 2022;27(3):433-456.
99. Romeiko XX, Zhang X, Pang Y, et al. A review of machine learning applications in life cycle assessment studies. *Sci Total Environ*. 2024;912:168969.
100. Salla JVE, de Almeida TA, Silva DAL. Integrating machine learning with life cycle assessment: a comprehensive review and guide for predicting environmental impacts. *Int J Life Cycle Assess*. 2025;30:2423-2445.
101. Kleinekorte J, Kleppich J, Fleitmann L, Beckert V, Blodau L, Bardow A. APPROPRIATE life cycle assessment: a process-specific, predictive impact assessment method for emerging chemical processes. *ACS Sustain Chem Eng*. 2023;11(25):9303-9319.
102. Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. *Nature*. 2018;559(7715):547-555.
103. Rittig JG, Dahmen M, Grohe M, Schwaller P, Mitsos A. Molecular machine learning in chemical process design. arXiv Preprint arXiv:2508.20527. 2025.
104. Hou P, Jolliet O, Zhu J, Xu M. Estimate ecotoxicity characterization factors for chemicals in life cycle assessment using machine learning models. *Environ Int*. 2020;135:105393.
105. Aboagye E, Longo J, Conway M, et al. Leveraging machine learning algorithms to predict life cycle inventory assessments (LCIA) to facilitate sustainable process design. *Comput Chem Eng*. 2025;201:109217.
106. Karka P, Papadokonstantakis S, Kokossis A. Digitizing sustainable process development: from ex-post to ex-ante LCA using machine-learning to evaluate bio-based process technologies ahead of detailed design. *Chem Eng Sci*. 2022;250:117339.
107. Blanco CF, Pauliks N, Donati F, Engberg N, Weber J. Machine learning to support prospective life cycle assessment of emerging chemical technologies. *Curr Opin Green Sustain Chem*. 2024;50:100979.
108. Strubell E, Ganesh A, McCallum A. Energy and policy considerations for modern deep learning research. *Proc AAAI Conf Artif Intell*. 2020;34(9):13693-13696.
109. Schwartz R, Dodge J, Smith NA, Etzioni O. Green AI. *Commun ACM*. 2020;63(12):54-63.
110. Xu W, Wang Y, Zhang D, et al. Transparent AI-assisted chemical engineering process: machine learning modeling and multi-objective optimization for integrating process data and molecular-level reaction mechanisms. *J Clean Prod*. 2024;448:141412.
111. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST. NIST AI 100-1. 2023.

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

112. Daoutidis P, Lee JH, Rangarajan S, et al. Machine learning in process systems engineering: challenges and opportunities. *Comput Chem Eng*. 2024;181:108523.
113. Paterson C, Hawkins R, Picardi C, Jia Y, Calinescu R, Habli I. Safety assurance of machine learning for autonomous systems. *Reliab Eng Syst Saf*. 2025;264:111311.
114. Wilkinson MD, Dumontier M, Aalbersberg IJ, et al. The FAIR guiding principles for scientific data management and stewardship. *Sci Data*. 2016;3(1):160018.
115. Achenie LEK, Paulson J, Mittal J, et al. ChE in context: chemical engineers embrace the frontiers of AI/ML. *Chem Eng Prog*. 2024.
116. Miao F, Holmes W. *Guidance for Generative AI in Education and Research*. UNESCO; 2023.
117. OECD-Education-International. *Opportunities, Guidelines and Guardrails on Effective and Equitable Use of AI in Education*. OECD-Education International; 2023.
118. Gong J, Shallcross DC, Jiao Y, Venkatasubramanian V, Davis R, Arges CG. Rethinking chemical engineering education. *Nat Chem Eng*. 2024;1(2):127-133.
119. 2024–2025 Criteria for Accrediting Engineering Programs. ABET Engineering Accreditation Commission. 2023.
120. Wu M, Di Caprio U, Vermeire F, et al. An artificial intelligence course for chemical engineers. *Educ Chem Eng*. 2023;45:141-150.
121. Pilario KE. Teaching classical machine learning as a graduate-level course in chemical engineering: an algorithmic approach. *Digit Chem Eng*. 2024;11:100163.
122. Lavor V, de Come F, dos Santos MT, Vianna AS. Machine learning in chemical engineering: hands-on activities. *Educ Chem Eng*. 2024;46:10-21.
123. Caminiti S. *Royal Dutch Shell Reskills Workers in Artificial Intelligence as Part of Huge Energy Transition*. 2020. Accessed February 10, 2026. https://www.cnbc.com/2020/04/03/royal-dutch-shell-reskills-workers-in-ai-part-of-energy-transition.html
124. NAM. *How DuPont Teaches Digital Skills In-House*. 2023. Accessed February 10, 2026. https://nam.org/how-dupont-teaches-digital-skills-in-house-2-25125/
125. AIChE A. *AI for Chemical Engineers With Live Instructor Q&A Sessions*. 2026. Accessed February 10, 2026. https://www.aiche.org/ili/academy/courses/ble277/ai-chemical-engineers-live-instructor-qa-sessions
126. Caccavale F, Gargalo CL, Gernaey KV, Krühne U. Towards Education 4.0: the role of large language models as virtual tutors in chemical engineering. *Educ Chem Eng*. 2024;49:1-11.
127. Keith M, Keiller E, Windows-Yule C, Kings I, Robbins P. Harnessing generative AI in chemical engineering education: implementation and evaluation of the large language model ChatGPT v3.5. *Educ Chem Eng*. 2025;51:20-33.
128. Venkatasubramanian V. Teaching artificial intelligence to chemical engineers: experience from a 35-year-old course. *Chem Eng Educ*. 2022;56(4):231-240.
129. Schweidtmann AM, Esche E, Fischer A, et al. Machine learning in chemical engineering: a perspective. *Chem Ing Tech*. 2021;93(12):2029-2039.
130. Kockmann N, Schindler T, Urbas L. AI in process industries—incubator labs and use cases. *Chem Ing Tech*. 2023;95(7):963.
131. Bortz M, Dadhe K, Engell S, et al. AI in process industries—current status and future prospects. *Chem Ing Tech*. 2023;95(7):975-988.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Chew JW, Andersson R, Bierweiler T, Ryttestal P, Wik T. AI in chemical engineering: From promise to practice. *AIChE J*. 2026;72(7):e70358. doi:10.1002/aic.70358

15475905, 2026, 7, Downloaded from https://aiche.onlinelibrary.wiley.com/doi/10.1002/aic.70358 by Chalmers University Of Technology, Wiley Online Library on [16/07/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License