

Neuro-Symbolic Artificial Intelligence for Carbon-Constrained Chemical Process Synthesis: A Review of Advances, Challenges, and Future Directions

Johnson Gani¹, Salihu Muhammad Aliyu², Barnabas Emobwetsah Micah³, Ansar Bilyaminu Adam⁴, Musa Yahaya Abubakar^{5*}

¹Department of Veterinary Biochemistry and Physiology, Federal University of Agriculture Zuru

²Department of Industrial Chemistry, Federal University Wukari, Taraba State

³Department of Chemistry, Federal University Wukari, Taraba State

⁴Department of Electrical/ Electronic Engineering, Federal University of Technology Minna, Niger State.

*Corresponding author: yahayaabubakarmusa2015@gmail.com

Received: 10th Jun 2026; Revised: 18th Jun 2026; Accepted: 25th Jun 2026

Abstract: In recent years, the application of artificial intelligence (AI) in chemical process engineering has gained momentum, driven by the rising demand for sustainable transformation in the industry, especially in the field of carbon constrained-constrained process optimization and low emission-constrained-emission industrial processes. The review delves into the evolving importance of the role of neuro-symbolic artificial intelligence (AI) in enhancing sustainable chemical processes by combining data-driven learning and first-principles scientific reasoning. The study investigates carbon constrained-emission-constrained optimization methods such as energy minimization strategies, emission-aware process synthesis methods, hybrid thermodynamic-AI modeling, and multi-objective optimization methods, -constrained methods, which are instrumental for net-zero industrial transitions. Additionally, the review addresses various uses of neuro-symbolic AI for reaction pathway optimization, catalysis design, intelligent process monitoring, process fault diagnosis, and emission reduction strategies for industrial processes. A conceptual model is presented that describes the relationship between process data, neural architectures, symbolic reasoning engines, thermodynamic constraints, and optimized low carbon constraints,-carbon outputs. The review also sheds light on the increasing significance of explainable AI systems that preserve physical consistency while enhancing the predictive capabilities in complex industrial setups. However, the current state of progress comes with a number of drawbacks, such as computational requirements, limited application in the industry, availability of data, the absence of a uniform set of carbon-aware AI frameworks, and difficulties in integrating with existing process infrastructure. Overall, the study shows that this neuro-symbolic AI is a promising direction in the quest for intelligent, explainable, and sustainable chemical manufacturing systems to help meet the future decarbonization targets and circular industrial targets.

Keywords: Neuro-symbolic artificial intelligence, Carbon-constrained process optimization, Sustainable chemical engineering, Hybrid AI-thermodynamic modeling, industrial decarbonization

How to cite this article: Gani J, Aliyu SM, Micah BE, Adam AB, Abubakar MY. Neuro-symbolic artificial intelligence for carbon-constrained chemical process synthesis: a review of advances, challenges, and future directions. NextGen Chem Discov. 2026;2(2):117-133.

Conflict of interest: None

Introduction

The chemical process industry contributes a significant part of the industrial carbon footprint and GHG emissions while at the same time contributing to the global economic development. With the advent of international climate agreements and net zero requirements, the systems for chemical manufacturing are under greater pressure to shift towards low-emission and sustainable manufacturing processes (Biswas, 2026). The transition must not only involve incremental gains in process efficiency, but also a rethinking of the fundamentals of chemical process design, optimization efficiency but optimization, and control under a carbon constraint.

Digital transformation in chemical engineering has grown to be a major tool for this paradigm shift. The incorporation of computational tools, sophisticated analytics optimization, analytics, and automation technologies have greatly improved process monitoring, simulation precision has precision, and flexibility. These advances have enabled the transition of traditional steady state precision, state design methods to “data”-driven methods for making decisions in systems with multiple variables and complexity. Artificial Intelligence (AI) is playing an increasing role in the optimization of processes and has proven to be a game changer in pattern recognition, predictive modeling, and system-level optimization.

Although traditional process design methods are based on heuristic rules and mechanistic models, they often fail to cope with the trade-offs that apply to the high-dimensional optimization problems of modern carbon constrained-state-constrained processes (Fózer *et al.*, 2023). Mechanistic methods are physically interpretable, but constrained interpretable but can be inflexible, while data-driven AI methods can be very predictive, but interpretable but predictive but can often lack physical consistency. In chemical engineering systems, where the thermodynamic laws and conservation principles must be followed, this limitation is quite significant.

In response to this challenge, a hybrid of two paradigms, neuro-symbolic Artificial Intelligence predictive but artificial intelligence (AI) have artificial intelligence, has emerged as a promising paradigm. It combines symbolic reasoning with explicit physical laws and constraints with neural network-based learning systems that are able to discover patterns from large sets of data. This synergy allows for more powerful and interpretable, an, has interpretable and physically consistent, model interpretable and consistent models that can more readily inform complex process optimization challenges (Mann *et al.*, 2025; Srinivasan *et al.*, 2024).

In spite of the high achievements in AI chemical engineering, the chemical process industry still has much to overcome in terms of inefficiencies consistent models inefficiencies in design, operation inefficiencies operation, and optimization strategies to operation, strategies the tribute to a high strategies that carbon footprint. However, when it comes to large-scale industrial systems (Biswas, 2026), there are existing approaches that fail to meet both high process efficiency and strict carbon emission constraints at the same time.

Current process design methodologies are mainly based on a deterministic or semi-empirical process model that is not easily adaptable to varying operating conditions. Such approaches are not useful to solve

the nonlinear and multi-objective problems common to today's sustainability-focused optimization problems. Consequently, they are not always able to offer the best solutions in carbon-constrained situations (Fózer *et al.*, 2023).

However, data-driven machine learning models are highly effective at making predictions but may be lacking in some fundamental aspects of mechanistic interpretability and have certain constraints when applied to extrapolations beyond their training data (Venkatasubramanian & Chakraborty, 2024). This means that reliability, safety, and compliance with regulatory standards are issues of concern in industrial applications.

Another hurdle is the separation between data-driven AI models and physics-based process models. Existing frameworks lack integration between the two domains and, consequently, are not able to develop strong, explainable, and scalable solutions to carbon-constrained process synthesis (Chowdhury *et al.*, 2025).

This review aims at critically examining the potential of using neuro-symbolic artificial intelligence to overcome the limitations of data-driven learning strategies and physics-based modeling frameworks in carbon-constrained chemical process synthesis, especially in its ability to link the two. In particular, the review aims to critically analyze the current state of the art in the field of neuro-symbolic AI applications for chemical process optimization, evaluate the current state of the art in carbon-aware process design, and review the main limitations and challenges in implementing hybrid AI systems in real-world industrial systems and will highlight future research directions to advance hybrid AI systems in sustainable chemical engineering. Moreover, the review examines the potential of hybrid AI architectures to improve the interpretability, reliability, and decision-making in intricate process settings, particularly in situations where strict emission reduction targets and sustainability constraints are being taken into account during process design and operational optimization.




Figure 1: Conceptual Framework of AI-Driven Carbon-Constrained Process Synthesis

Method

A Systematic Literature Review (SLR) method was used to critically analyze and synthesize the existing literature on the use of AI in chemical process engineering, specifically in carbon-constrained and sustainable industrial processes. Academic reliability was ensured, and bias was reduced by using a rigorous and transparent review selection process with predefined inclusion and exclusion criteria. In addition, a narrative synthesis approach was used to qualitatively combine the results of the studies that were too heterogeneous to be quantitatively aggregated. To ensure comprehensive coverage of the chemical engineering, artificial intelligence, and sustainability research, relevant literature was obtained from the major scientific databases such as Scopus, Web of Science, IEEE Xplore, ScienceDirect, and SpringerLink. The search strategy involved using keywords like “neuro-symbolic AI,” “chemical process synthesis,” “carbon-constrained optimization,” and “AI-driven process engineering” and employing Boolean operators to refine the search and make it more precise and relevant. Selected studies were analyzed systematically using structured data extraction, thematic coding, comparative analysis, and trend mapping to identify methodological developments, new techniques for artificial intelligence application used, application fields, and gaps for sustainable chemical process optimization.

Developments in AI for Chemical Process Engineering

The evolution of artificial intelligence (AI) in chemical process engineering has evolved over the years in distinct stages, each corresponding to advances in computational power, data access, and complexity of the process systems in industry. This evolution is a move away from the rigid knowledge-driven systems to adaptive hybrid systems that can combine data-driven learning with physical/mechanical constraints.

Rule-Based Expert Systems

The first application of AI in chemical process engineering was rule-based expert systems, which used prewritten logical rules based on human knowledge. These systems are intended to emulate the decision-making process in areas including process control, fault diagnosis, and simple optimization problems. They operated by using “if-then” rules under certain operating conditions and were therefore very interpretable and transparent in their reasoning.

But, although they were well-structured, rule-based systems were also very rigid. They were not resilient to change in the process and were not effective at dealing with uncertainty or the nonlinear behavior of the system. Therefore, they were only suitable for use in modern chemical plants where the operating conditions were variable.

Machine Learning-Based Models

Machine learning (ML) models constituted a major step forward in chemical process engineering using AI methodology. While rule-based systems rely on abstract rules, ML systems learn directly from past and experimental data, and so they are able to uncover hidden patterns and relationships in complex data. This made the prediction, optimization, and monitoring of the processes much easier.

For various applications, such as process fault detection, yield prediction, and parameter optimization, machine learning methods like regression models, support vector machines, and decision trees gained wide traction. While the benefits are significant, there are still major challenges with ML models. They tend to be uninterpretable and perform poorly on distributions different from the one used for training. This limitation is concerning for chemical engineering systems, where the laws of the real world must be followed and reliability and safety of the model are important concerns (Sivaram & Venkatasubramanian, 2022; Thon *et al.*, 2021).

Deep Learning Approaches

Deep learning (DL) was another development that took AI capabilities in chemical process engineering to another level. Deep neural networks consist of many hidden layers and are able to model very nonlinear and complex system behavior. These models have shown good performance in several applications like reaction kinetics modeling, process monitoring, soft sensing, and advanced optimization applications.

Though powerful, deep learning models are often said to be black box models. They need a massive amount of data to learn, and they don't offer much insight into how the prediction is made. In safety-critical chemical processes, where knowing the "how" and "why" is vital to ensure safety, this opacity is a significant disadvantage. Furthermore, deep learning models often don't guarantee physical consistency for data outside their range of observation (Kiss & Grievink, 2026; Venkatasubramanian & Chakraborty, 2024).

Transition toward Neuro-Symbolic AI Systems

Today, there is a new wave of AI in chemical process engineering with the development of neuro-symbolic AI systems that combine neural networks and symbolic reasoning. This is what this hybrid paradigm tries to do – merge the strengths of both: the learning capability of neural networks and the structured reasoning of symbolic systems.

This shift is a major change from data-driven modeling to hybrid approaches, where domain knowledge and physical laws are integrated into AI systems. This then allows neuro-symbolic AI to be both physically consistent and have the advantages of large-scale data-driven learning. This trend is part of a larger transition from being solely heuristic to being hybrid with both predictive and mechanistic considerations (Sivaram & Venkatasubramanian, 2022; Thon *et al.*, 2021).

One of the strengths of a neuro-symbolic system is its capability to overcome the shortcomings of the purely data-driven models, which lack physical validity and capacity for extrapolation in unseen scenarios. These systems include symbolic reasoning to guarantee that predictions are aligned with thermodynamic principles and process constraints (Kiss & Grievink, 2026; Venkatasubramanian & Chakraborty, 2024). Furthermore, symbolic regression can be used to uncover meaningful mathematical relationships from intricate data sets, thus promoting transparency and interpretability of models.

Moreover, this hybrid method fills in the long-standing gap between high-fidelity neural models and traditional first-principles-based engineering models. It allows the construction of efficient and physically relevant surrogate models and uncertainty quantification of multi-scale process simulations (Kumar *et al.*, 2025; Wehinger *et al.*, 2022; Neveux *et al.*, 2024). This feature is especially valuable when dealing with robustness for industrial applications.

The move towards neuro-symbolic AI systems is a major milestone in chemical process engineering. It boosts the reliability of decision-making, increases the interpretability of the processes, and enables adherence to new safety and governance requirements in industrial settings where AI systems are subject to strict regulatory frameworks (Chowdhury *et al.*, 2025; Chew *et al.*, 2026).

Carbon-Constrained Process Optimization Techniques

The transition toward low-carbon industrial systems has intensified the development of advanced process optimization techniques capable of balancing productivity, profitability, and environmental sustainability. Carbon-constrained optimization frameworks are increasingly being adopted to minimize greenhouse gas emissions while maintaining operational efficiency across industrial networks. These methodologies integrate process systems engineering, thermodynamic modeling, artificial intelligence, and mathematical programming to support sustainable industrial transformation. Modern optimization strategies emphasize energy efficiency, carbon minimization, and intelligent decision-making under strict environmental

regulations and economic limitations. These methodologies integrate multi-scale modeling and non-linear programming to address the increasingly complex requirements of net-zero industrial transitions (Biegler, 2024; Lye *et al.*, 2024). By incorporating life cycle assessment and mathematical programming, these frameworks ensure a rigorous balance between economic viability and stringent environmental constraints (Pistikopoulos & Tian, 2024; Sachin, 2025).

Energy Minimization-Based Approaches

Energy minimization remains one of the most fundamental strategies in carbon-constrained process optimization. Industrial systems consume substantial thermal and electrical energy during production processes, resulting in significant carbon emissions. Consequently, optimization frameworks now focus on reducing unnecessary energy consumption through systematic integration of heat recovery systems, utility optimization, and thermodynamic efficiency enhancement.

Pinch analysis has emerged as a major technique for identifying minimum energy requirements within industrial plants. By analyzing temperature intervals and heat exchange opportunities, pinch-based methods reduce fuel consumption and improve heat integration efficiency. The combination of pinch analysis with optimization algorithms enables industries to simultaneously minimize operational costs and carbon emissions. Furthermore, the application of multi-criteria decision-making allows industries to systematically prioritize decarbonization strategies that align with specific operational goals and stakeholder interests (Sun *et al.*, 2023).

Additionally, modern energy minimization models incorporate dynamic operational variables such as fluctuating energy prices, renewable energy integration, and carbon taxation policies. These systems frequently employ Mixed-Integer Linear Programming (MILP) and Non-Linear Programming (NLP) techniques to determine optimal operating conditions under carbon constraints. Specifically, these approaches utilize Multi-Objective Mixed-Integer Linear Programming to evaluate Pareto-optimal frontiers that reconcile throughput maximization with significant reductions in carbon intensity (N, 2026a, 2026b).

Table 1: Major Energy Minimization Techniques in Carbon-Constrained Industries

Technique	Primary Objective	Key Benefit
Pinch Analysis	Heat recovery optimization	Reduced utility consumption
MILP Optimization	Optimal process scheduling	Lower carbon intensity
Heat Integration Networks	Energy reuse enhancement	Improved thermal efficiency
Utility System Optimization	Efficient power distribution	Reduced operational costs
Renewable Energy Coupling	Low-carbon energy substitution	Emission reduction

Emission-Aware Process Design Frameworks

Emission-aware process design frameworks incorporate environmental performance directly into process synthesis and operational planning. Unlike conventional process design approaches that prioritize economic performance alone, emission-aware systems treat carbon emissions as a core design constraint. These frameworks integrate life cycle assessment, carbon footprint analysis, and environmental impact indicators during early-stage process development.

Advanced emission-aware frameworks evaluate alternative process pathways based on both environmental and economic performance metrics. Such systems enable industries to identify technologies capable of achieving low-emission production targets while maintaining competitive productivity levels. Furthermore, the integration of graphical pinch analysis with these mathematical models enables the

visualization of complex trade-offs, providing decision-makers with a transparent pathway for selecting sustainable technology portfolios (Akmal *et al.*, 2026).

Another important feature of emission-aware design is the incorporation of carbon pricing and environmental regulations into optimization objectives. Industries increasingly face strict emission reduction targets, making regulatory compliance an essential component of process optimization. To address the non-convex nature of industrial optimization problems, computational frameworks often apply piecewise linear approximations to simplify capital expenditure modeling and accelerate convergence in large-scale industrial systems (Weber *et al.*, 2021).

Hybrid AI-Thermodynamic Modeling Systems

The integration of artificial intelligence with thermodynamic modeling has significantly enhanced carbon-constrained process optimization. Hybrid AI-thermodynamic systems combine first-principles process models with machine learning algorithms to improve predictive accuracy, computational speed, and adaptive optimization capabilities. These systems are particularly valuable in large-scale industrial operations where process variables exhibit strong non-linear interactions.

Machine learning techniques such as artificial neural networks, reinforcement learning, and genetic algorithms are increasingly employed to predict process behavior, optimize operating conditions, and reduce carbon-intensive inefficiencies. AI-driven optimization frameworks can rapidly analyze extensive operational datasets, identify hidden process relationships, and recommend energy-efficient operating strategies.

Thermodynamic models provide the physical and chemical constraints required to ensure realistic optimization results, while AI algorithms enhance predictive flexibility and adaptive learning capabilities. This hybridization improves process stability, minimizes waste generation, and supports real-time emission monitoring and control. As industrial systems continue to grow in complexity, hybrid AI-thermodynamic optimization frameworks are expected to play a critical role in enabling sustainable and carbon-neutral manufacturing systems.

Table 2: Components of Hybrid AI-Thermodynamic Optimization Systems

Component	Function	Contribution to Sustainability
Thermodynamic Models	Process constraint modeling	Accurate energy prediction
Artificial Neural Networks	Pattern recognition	Process efficiency improvement
Genetic Algorithms	Optimization search	Reduced carbon emissions
Reinforcement Learning	Adaptive process control	Real-time optimization
Data Analytics Platforms	Operational data integration	Improved decision-making

Applications of Neuro-Symbolic AI

Neuro-symbolic artificial intelligence (NSAI) represents a transformative advancement in industrial process optimization and sustainable manufacturing. By combining neural-network-based learning with symbolic reasoning and mechanistic process knowledge, neuro-symbolic systems provide both predictive intelligence and interpretable decision-making capabilities. Unlike purely statistical machine learning models, these hybrid systems integrate physical laws, engineering principles, and domain-specific constraints into computational learning architectures. By integrating mechanistic domain knowledge with data-driven machine learning, these frameworks overcome the reasoning limitations inherent in purely statistical models (Venkatasubramanian & Chakraborty, 2024). This hybrid paradigm facilitates the generation of interpretable, first-principles-based models that maintain physical consistency while

providing causal explanations for complex industrial processes (Kiss & Grievink, 2026; Sivaram & Venkatasubramanian, 2022).

The increasing complexity of industrial systems, particularly under carbon-constrained operational environments, has accelerated the adoption of neuro-symbolic AI across catalysis, process optimization, emission management, and intelligent manufacturing systems. These technologies enhance process transparency, improve fault diagnosis, and support autonomous industrial decision-making while maintaining compliance with engineering safety and sustainability constraints.

Reaction Pathway Optimization

Reaction pathway optimization is one of the most significant applications of neuro-symbolic AI in chemical and industrial engineering. Industrial reactions often involve highly nonlinear kinetics, competing pathways, and thermodynamic limitations that make conventional optimization approaches computationally expensive and difficult to interpret. Neuro-symbolic systems address these challenges by integrating symbolic reaction mechanisms with machine-learning-driven predictive analytics.

These frameworks can analyze extensive experimental and process datasets to identify optimal reaction conditions that maximize yield while minimizing energy consumption and unwanted by-products. Symbolic reasoning modules encode known chemical rules, stoichiometric relationships, and kinetic constraints, while neural learning components identify hidden nonlinear interactions among variables. This integration enables more accurate prediction of reaction pathways under varying operational conditions.

Furthermore, neuro-symbolic AI systems facilitate automated hypothesis generation for novel reaction mechanisms. By combining mechanistic constraints with adaptive learning, industries can accelerate process development while maintaining chemical feasibility and thermodynamic consistency. This approach allows for the derivation of sparse, interpretable mathematical expressions that bridge the gap between neural network nonlinearity and the rigorous constraints of physical laws ((63095) *et al.*, 2025; Kumar *et al.*, 2025).

Table 3: Neuro-Symbolic AI Functions in Reaction Pathway Optimization

NSAI Component	Function	Industrial Benefit
Symbolic Rule Engine	Encodes reaction mechanisms	Improved process interpretability
Neural Networks	Learns nonlinear process behavior	Enhanced predictive accuracy
Kinetic Constraint Models	Ensures thermodynamic feasibility	Reliable optimization
Automated Hypothesis Generation	Suggests alternative pathways	Accelerated process innovation
Process Data Analytics	Identifies operational patterns	Improved yield efficiency

Catalysis Design and Discovery

Catalyst development remains a critical factor in achieving sustainable industrial transformation. Traditional catalyst discovery methods often require extensive experimentation, high costs, and prolonged development cycles. Neuro-symbolic AI significantly accelerates catalyst design by integrating molecular-level scientific knowledge with data-driven predictive systems.

In catalyst optimization, symbolic AI modules encode physicochemical relationships such as adsorption energies, reaction kinetics, and electronic structure constraints, while neural networks analyze large-scale materials databases to identify promising catalyst candidates. This hybrid strategy enables efficient screening of catalytic materials with improved selectivity, stability, and energy efficiency.

Additionally, neuro-symbolic frameworks support inverse catalyst design, where desired reaction outcomes are specified first, and the AI system predicts catalyst structures capable of achieving those objectives. These systems improve catalyst performance prediction under industrial operating conditions and reduce experimental uncertainty. As a result, industries can rapidly develop catalysts that lower energy demand and reduce carbon-intensive process emissions.




Figure 2: Neuro-Symbolic AI Workflow in Catalyst Discovery

Intelligent Process Control Systems

Neuro-symbolic AI has become increasingly important in intelligent process control systems due to its ability to combine adaptive learning with explainable reasoning. Conventional control systems often struggle when dealing with uncertain operating conditions, nonlinear process dynamics, and incomplete datasets. Neuro-symbolic architectures address these limitations through real-time reasoning supported by process knowledge representation.

These systems continuously monitor industrial operations using sensor networks, process simulations, and historical datasets to predict deviations, detect anomalies, and optimize control actions. By constraining gradient descent trajectories to knowledge-consistent subspaces, these architectures enable efficient fault diagnosis and process monitoring even when faced with high-dimensional nonlinearities or limited labeled datasets (Yin *et al.*, 2026).

Another major advantage of neuro-symbolic control systems is their interpretability. Unlike black-box AI models, symbolic reasoning layers provide traceable explanations for operational decisions, thereby improving regulatory transparency and operator trust. Furthermore, this synergy mitigates the propensity for process-control hallucinations by subjecting model-generated causal narratives to rigorous verification against fundamental physical constraints and sensor data (Galitsky & Rybalov, 2026).

Table 4: Features of Neuro-Symbolic Intelligent Control Systems

Feature	Function	Operational Advantage
Real-Time Monitoring	Continuous process evaluation	Early fault detection
Symbolic Reasoning	Rule-based explanation generation	Improved transparency
Adaptive Neural Learning	Dynamic process optimization	Increased operational efficiency
Predictive Maintenance	Equipment failure prediction	Reduced downtime
Safety Constraint Enforcement	Regulatory compliance assurance	Enhanced industrial safety

Industrial-Scale Emission Reduction Strategies

Industrial decarbonization increasingly relies on intelligent systems capable of balancing production efficiency with stringent environmental targets. Neuro-symbolic AI contributes significantly to emission reduction by integrating process knowledge, carbon accounting frameworks, and machine-learning-driven optimization techniques.

These systems analyze plant-wide operational data to identify emission hotspots, optimize energy utilization, and recommend low-carbon operational strategies. Symbolic reasoning components encode environmental regulations, safety protocols, and sustainability targets, while neural learning modules identify complex emission-generation patterns across industrial systems.

Moreover, neuro-symbolic AI enables intelligent decision-support systems that align industrial processes with regulatory and sustainability frameworks. This integration bridges the persistent gap between human-centric process queries and the structured regulatory or safety ontologies required for reliable industrial decision-support systems (Shyalika *et al.*, 2026). Consequently, industries can improve carbon reporting accuracy, optimize renewable energy integration, and achieve substantial reductions in greenhouse gas emissions while maintaining economic competitiveness.

Table 5: AI Methods in Chemical Process Optimization

AI Approach	Strengths	Limitations	Carbon Reduction Potential
Machine Learning	Fast prediction, scalable	Low interpretability	Moderate
Deep Learning	High accuracy, nonlinear modeling	Data intensive, black-box nature	High
Physics-based Models	Strong physical consistency	Limited flexibility	Low-Moderate
Neuro-Symbolic AI	Combines reasoning + learning	Emerging, computational complexity	Very High

Conceptual Framework of Neuro-Symbolic AI in Chemical Processes

Neuro-Symbolic Artificial Intelligence (NSAI) enables a paradigm shift in designing sustainable and carbon-constrained chemical process systems. Neuro-symbolic systems combine symbolic reasoning and first-principles chemical knowledge with neural learning mechanisms, as opposed to only statistical learning that is the foundation of traditional artificial intelligence models. This integration allows industrial systems to be both mechanistically interpretable and predictively intelligent while respecting the physical and thermodynamic laws that exist in their systems.

The framework combines neural networks with symbolic reasoning to ensure that predictions based on data are always consistent with first-principles chemical and physical laws (Perumal *et al.*, 2026). This architecture alleviates the requirement for extensive experimental supervision, while simultaneously boosting the interpretability of the chemistry (Sivaram & Venkatasubramanian, 2022; Tayane & Gaber, 2025). Moreover, the symbolic constraints in combination with adaptive neural learning provide improved process transparency and reduce the number of non-physical predictions commonly seen in data-driven systems.

The conceptual framework of the neuro-symbolic AI in chemical processes is composed of four interrelated layers, namely the input layer, AI layer, constraint layer, and output layer. The layers all aid in the creation of optimized low-carbon process designs that balance productivity, sustainability, and regulatory compliance.

Input Layer

The first layer in the neuro-symbolic framework is the input layer. This layer gathers and collates the information related to processes that are generated during industrial operations, such as process variables, reaction conditions, energy utilization profiles, and emission constraints. High-dimensional datasets of plant operations are continually generated by sensor networks, industrial databases, or real-time monitoring systems.

In addition, reaction kinetics, catalyst performance parameters, and environmental regulations are included in the input stage of the framework. The various data sets provide the AI system with insight into the complexity of the operation and allow it to form connections between the process conditions and sustainability goals. The use of emission constraints at this stage guarantees that carbon reduction goals are built into the optimization procedure.

AI Layer

The AI layer is the computing intelligence backbone of the framework. It integrates artificial neural networks and symbolic reasoning engines to attain the learning capability and explainable reasoning. Large-scale industrial data are processed by neural networks to detect nonlinear patterns, optimize the operating conditions of the process, and predict its behavior under different operating conditions. At the same time, symbolic reasoning modules include chemical knowledge as rules, engineering principles as rules, and mechanistic process knowledge as logical structures that are comprehensible to humans.

The model automatically translates chemical equations into systems of ordinary differential equations, guaranteeing thermodynamic consistency and ensuring the prevention of non-physical concentration states (Tayane & Gaber, 2025a, 2025b). The combined use of neural architectures and symbolic frameworks also facilitates identifying implicit physical principles from complex data sets, enhancing model generalization and robustness in data-limited environments (Hajireza *et al.*, 2025; Tayane & Gaber, 2025).

Another important feature of this layer is symbolic regression that allows for the purely symbolic discovery of mathematical relationships governing reaction systems. In addition, this integration enables the independent discovery of reaction pathways by enabling the generation of interpretable mathematical expressions from time-resolved concentration data via symbolic regression (Ji & Deng, 2021; *et al.*, 2025). This allows the AI layer to provide improved predictive performance and insight into chemical processes.

Constraint Layer

The constraint layer guarantees all optimization results are physically realistic and industrially feasible. This layer incorporates thermodynamic laws, mass conservation equations, energy balance relationships, and process safety constraints directly into the optimization architecture. The constraints help keep the number of unrealistic process states to a minimum and ensure consistency with accepted chemical engineering practices.

The first-principles constraints add a lot of reliability to the model in comparison to black-box AI systems. The framework ensures thermodynamic consistency, reducing prediction errors, boosting process stability, and aiding regulatory compliance in industrial applications.

Output Layer

The output layer creates optimized low-carbon chemical process designs that can meet sustainability targets while preserving operation performance. These outputs can be optimum reactor designs, energy-saving operating parameters, emission minimization methods, and optimized suggestions for process control.

Finally, the combination of neuro-symbolic AI enhances sustainable manufacturing by minimizing carbon emissions, optimizing energy usage, and facilitating adaptive process optimization. Within this framework, there is thus a viable roadmap towards intelligent, interpretable, and environmentally friendly chemical process synthesis.

Challenges and Limitations

In spite of the many significant advances made in the potential of neuro-symbolic artificial intelligence (NSAI) in chemical engineering and carbon-constrained industrial optimization, many technical, computational, and operational issues remain that hinder its broad industrial implementation. The synergy of data-driven learning and symbolic reasoning makes neuro-symbolic systems sufficiently powerful; however, the synergy also creates a significant amount of complexity in these systems, affecting their scalability, reliability, and deployment efficiency. The current challenges include the barrier to industrial-scale adoption, computational burden, lack of standardization, data quality, and legacy industrial infrastructure compatibility.

Limited Industrial-Scale Deployment of Neuro-Symbolic AI

A primary challenge with neuro-symbolic AI is that it has been used in industries at a fairly small scale. Although numerous NSAI models have been shown to yield good results in research settings, there are difficulties in implementing them on a large scale in industrial settings. Industrial systems have very dynamic operating conditions, complex interprocess relationships, and strict safety requirements, which make it difficult to implement a hybrid AI system.

One of the challenges comes from reliance on symbolic knowledge produced by experts. It is crucial to carefully design ontologies, logical rules, and process constraints that represent accurately the industrial operations in neuro-symbolic systems. For hybrid systems, however, there are no well-established metrics for measuring interpretability, and for domains without digitized domain expertise (DeLong *et al.*, 2024), expert-curated knowledge is a major hurdle. As a result, industries that have a less developed digital infrastructure or lack documentation of their processes can often find it difficult to effectively adopt a robust neuro-symbolic framework.

Furthermore, manual ontology construction and knowledge base maintenance can be time-consuming and prone to scalability issues, making it difficult to quickly adapt to the rapidly changing industrial contexts (Hakim *et al.*, 2025). However, as the industrial processes change, symbolic knowledge structures need to be updated on an ongoing basis, which is costly and challenging.

High Computational Complexity of Hybrid Systems

Neuro-symbolic AI systems typically integrate neural learning algorithms with symbolic reasoning engines, which significantly increases their compute requirements compared to traditional AI methods. Optimization problems in industry are often high-dimensional, with complex nonlinear process models, real-time constraints, and so on, which also makes for a more complex computational problem.

The integration of symbolic logic with deep neural networks often requires simultaneous optimization of continuous and discrete variables. This hybridization significantly increases model complexity and computational latency during training and inference. Additionally, the inherent overhead associated with balancing symbolic precision and neural flexibility often introduces latency that hinders real-time operational efficiency (Lu *et al.*, 2024).

In large-scale industrial settings such as petrochemical plants, refinery systems, and carbon capture facilities, computational delays may negatively affect process responsiveness and operational stability. Furthermore, the increasing size of industrial datasets intensifies memory requirements and prolongs model convergence times. These limitations restrict the practical use of NSAI in real-time industrial control systems where rapid decision-making is essential.

Table 6: Major Computational Challenges in Neuro-Symbolic AI Systems

Challenge	Description	Industrial Impact
High-dimensional data processing	Large-scale industrial datasets increase computational demand	Reduced processing efficiency
Symbolic-neural synchronization	Simultaneous reasoning and learning operations	Increased latency
Nonlinear optimization complexity	Complex process interactions require advanced optimization	Longer convergence times
Memory-intensive architectures	Large knowledge graphs and neural layers	High hardware requirements
Real-time inference limitations	Delayed process-response capability	Reduced operational stability

Inadequate Standardized Carbon-Aware AI Frameworks

A further major challenge is the lack of standardized models for carbon-aware neuro-symbolic AI systems. The current implementations are in many cases very domain-specific and cannot be interoperable between industrial sectors. It makes it difficult to benchmark, validate for regulatory purposes, and compare the performance of different optimization models.

Furthermore, there is a lack of strong, generic frameworks that require domain-specific optimizations to be performed, making integration with existing software engineering and manufacturing processes cumbersome, error-prone, and time-consuming (Mastropaolo & Poshyvanyk, 2025). This can lead to specialized AI architectures, which are often unique to each industry, and restrict scalability and widespread use.

Non-standardization of sustainability metrics/outcomes, carbon accounting approaches, and explainability assessment adds to the complexity of regulatory acceptance. If there are no standardized methods of validation, industries might struggle to certify the safety, clarity, and eco-friendliness of neuro-symbolic optimization systems.

Data Quality and Availability Constraints

Accurate, diverse, and high-quality industrial data are crucial for the effectiveness of neuro-symbolic AI. But in industrial settings, often sensors can take incomplete measurements, records can be inconsistent, and there may be few labels. Failing to get the data right has a huge impact on model accuracy and ability to predict.

Moreover, structured knowledge representations and semantic consistency are essential for symbolic reasoning systems. Data fragmentation and documentation of inconsistent processes are a problem for many industrial applications in developing reliable symbolic knowledge bases. Moreover, existing symbolic representation learning approaches have difficulty disambiguating between predicates that share similar meanings, resulting in less robust logical reasoning in high signal-to-noise ratio scenarios (Hakim *et al.*, 2025; Yu *et al.*, 2023).

Additionally, data privacy concerns and proprietary industrial information hinder data sharing among organizations, thereby decreasing chances of joint AI model creation and cross-industry learning. As a result, a number of neuro-symbolic systems are limited in both the amount of training they receive and the completeness of their domain knowledge.

Integration Difficulties with Legacy Industrial Systems

Many industrial sites still rely on old process-control systems that were not built to support advanced AI technologies. The deployment of a neuro-symbolic AI into these current systems has compatibility and operational issues.

Most legacy systems use old software architectures, small sensor networks, and automation protocols that are incompatible with current AI-optimized software. The retrofit of industrial plants using neuro-symbolic infrastructures can entail large-scale investments in hardware, software restructuring, and retraining of the plant staff.

In addition, industries might be hesitant to embrace systems with high levels of autonomy because of concerns about cybersecurity, transparency in operations, and system reliability. Therefore, careful processes of validation, regulatory approval, and strong mechanisms of human oversight need to be carefully considered when integrating NSAI into mission-critical industrial operations. While the ability of neuro-symbolic AI to drive sustainable changes in the industrial sector is significant, the key challenge lies in the barriers to widespread industrial adoption, such as the integration of AI systems.

Conclusion

The review illustrates that the potential of neuro-symbolic AI is great for revolutionizing chemical process engineering by leveraging the predictive power of machine learning and the interpretability and physical consistency of symbolic reasoning and first-principles modeling. By combining these technologies, sophisticated carbon-constrained optimization systems can be created to solve energy efficiency, emission reduction, process safety, and economic performance simultaneously. Industries can optimize process synthesis, discover reaction pathways, monitor processes with intelligence, and control processes adaptively with hybrid AI-thermodynamic approaches to enable sustainable manufacturing goals.

The modern optimization techniques such as multi-objective mathematical programming, life-cycle-based optimization, and process synthesis using AI-driven approaches are found to be effective approaches to achieve net-zero industrial operations, the study also shows. By combining neural networks with symbolic representations, these architectures provide a more transparent and explainable approach to decision-making that surpasses the shortcomings of data-driven black-box models, particularly in industrial applications. Utilizing the thermodynamic laws, kinetic equations, and symbolic constraints in a neural learning system improves the robustness, generalization, and mechanistic understanding in complex chemical environments.

However, there are some technical and operational hurdles that still need to be overcome to make the practical application of neuro-symbolic AI a reality. The high computational requirements, lack of good quality industrial datasets, requirement for expert knowledge bases, lack of standardized validation frameworks, and integration challenges with industrial infrastructures remain major challenges to

widespread adoption. Furthermore, the lack of a common carbon-aware AI standard makes it difficult to assess and compare the efficiency of different sustainable process optimization systems in different industries.

However, further development of explainable AI, multi-scale process modeling, generative AI, symbolic neural integration, and autonomous process systems will further speed up the industrialization of intelligent sustainable manufacturing technologies. Scalable hybrid architectures, standardized assessment criteria, interoperable industrial AI platforms, and comprehensive real-time deployment strategies are areas that should be explored further in the future. In conclusion, neuro-symbolic AI could prove to be a potent and scientifically sound approach towards fostering resilient, transparent, energy-efficient, and low-carbon chemical industries that can contribute to the sustainable future and decarbonization of the global economy.

References

- [1]. Akmal M, Foo DCY, Ooi J, Lim CH. A novel hybrid method combining mixed-integer linear programming and pinch analysis for technology selection. *Process Integr Optim Sustain*. 2026. doi:10.1007/s41660-025-00661-1.
- [2]. Biegler LT. Multi-level optimization strategies for large-scale nonlinear process systems. *Comput Chem Eng*. 2024;185:108657. doi:10.1016/j.compchemeng.2024.108657.
- [3]. Biswas S. Artificial intelligence in chemistry for future sustainability: An in-depth scientific review. *Int J Innov Sci Res Technol*. 2026;712. doi:10.38124/ijisrt/26mar158.
- [4]. Chew JW, Andersson R, Bierweiler T, Ryttestal P, Wik T. AI in chemical engineering: From promise to practice. *AIChE J*. 2026. doi:10.1002/aic.70358.
- [5]. Chowdhury SH, Ghosh A, Acharya S, Saha SC, Pal PK, Roy S, *et al*. Transforming chemical process engineering: The role of AI and machine learning in revolutionizing process systems. *Can J Chem Eng*. 2025;104(2):699-731. doi:10.1002/cjce.70032.
- [6]. Decardi-Nelson B, Alshehri AS, Ajagekar A, You F. Generative AI and process systems engineering: The next frontier. *Comput Chem Eng*. 2024;187:108723. doi:10.1016/j.compchemeng.2024.108723.
- [7]. Fózer D, Nimmegeers P, Tóth AJ, Varbanov PS, Klemeš JJ, Mizsey P, *et al*. Hybrid prediction-driven high-throughput sustainability screening for advancing waste-to-dimethyl ether valorization. *Environ Sci Technol*. 2023;57(36):13449-13462. doi:10.1021/acs.est.3c01892.
- [8]. Galitsky B, Rybalov A. Neuro-symbolic verification for preventing LLM hallucinations in process control. *Processes*. 2026;14(2):322. doi:10.3390/pr14020322.
- [9]. 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. doi:10.1016/j.cej.2025.165183.
- [10]. Ji W, Deng S. Autonomous discovery of unknown reaction pathways from data by chemical reaction neural network. *J Phys Chem A*. 2021;125(4):1082-1092. doi:10.1021/acs.jpca.0c09316.
- [11]. Kiss AA, Grievink J. Perspectives on the essential role of first-principles modeling in the age of AI. *Ind Eng Chem Res*. 2026. doi:10.1021/acs.iecr.5c04156.
- [12]. Kumar D, Dixit V, Ramteke M, Kodamana H. Learning system physics using symbolic neural integration (SynISM) with applications to chemical processes. *Ind Eng Chem Res*. 2025;64(23):11441-11458. doi:10.1021/acs.iecr.5c00107.

- [13]. Lye YX, Chew YE, Foo DCY, How BS, Andiappan V. Carbon emission reduction strategy planning and scheduling for transitioning process plants towards net-zero emissions. *RSC Sustain.* 2024;3(2):929-945. doi:10.1039/d4su00636d.
- [14]. Mann V, Lu J, Venkatasubramanian V, Gani R. A perspective on artificial intelligence for process manufacturing. *Engineering.* 2025;52:60-67. doi:10.1016/j.eng.2025.01.014.
- [15]. N KR. Decarbonizing industrial operations: A multi-objective optimization approach for balancing energy efficiency and production throughput in carbon-constrained environments. *Zenodo.* 2026. doi:10.5281/zenodo.18710893.
- [16]. Neveux T, Commenge J, Vermeire FH. Editorial: Artificial intelligence-assisted design of sustainable processes. *Front Chem Eng.* 2024;6. doi:10.3389/fceng.2024.1391673.
- [17]. Perumal D, Vaithiyalingam M, YS P, Dharanitharan K, Dinesh B, Kumar MS, *et al.* Neuro-symbolic AI in drug discovery. In: *Advances in Computational Intelligence and Robotics.* Hershey (PA): IGI Global; 2026. p. 307-332. doi:10.4018/979-8-3373-6400-1.ch011.
- [18]. Pistikopoulos EN, Tian Y. Advanced modeling and optimization strategies for process synthesis. *Annu Rev Chem Biomol Eng.* 2024;15(1):81-103. doi:10.1146/annurev-chembioeng-100522-112139.
- [19]. Rupprecht S, Gao Q, Karia T, Schweidtmann AM. Multi-agent systems for chemical engineering: A review and perspective. *Curr Opin Chem Eng.* 2025;51:101209. doi:10.1016/j.coche.2025.101209.
- [20]. Sachin J. Design of sustainable chemical processes and supply chains via mathematical programming and hybrid modeling [dissertation]. Zürich: ETH Zürich Research Collection; 2025. doi:10.3929/ethz-c-000785582.
- [21]. Schweidtmann AM, Esche E, Fischer A, Kloft M, Repke J, Säger S, *et al.* Machine learning in chemical engineering: A perspective. *Chem Ing Tech.* 2021;93(12):2029-2039. doi:10.1002/cite.202100083.
- [22]. Shyalika C, Sharma A, Kalach FE, Jaimini U, Henson C, Harik R, *et al.* CausalTrace: A neurosymbolic causal analysis agent for smart manufacturing. *Proc AAAI Conf Artif Intell.* 2026;40(47):40432-40439. doi:10.1609/aaai.v40i47.41486.
- [23]. Sivaram A, Venkatasubramanian V. XAI-MEG: Combining symbolic AI and machine learning to generate first-principles models and causal explanations. *AIChE J.* 2022;68(6). doi:10.1002/aic.17687.
- [24]. Srinivasan KK, Puliyanda A, Thosar D, Bhakte A, Singh K, Addo P, *et al.* Artificial intelligence and machine learning at various stages and scales of process systems engineering. *Can J Chem Eng.* 2024;103(3):1004-1035. doi:10.1002/cjce.25525.
- [25]. Sun X, Durkin A, Guo M. Life cycle optimisation tool development for process systems and centralised supply chain design. *RSC Sustain.* 2023;1(9):2224-2240. doi:10.1039/d3su00251a.
- [26]. Tayane S, Gaber J. SymChemAI: A symbolically informed neural network for physicochemical modeling of reaction systems. *ChemRxiv.* 2025. doi:10.26434/chemrxiv-2025-2hp4c-v2.
- [27]. Thon C, Finke B, Kwade A, Schilde C. Artificial intelligence in process engineering. *Adv Intell Syst.* 2021;3(6). doi:10.1002/aisy.202000261.
- [28]. Venkatasubramanian V, Chakraborty A. Quo vadis ChatGPT? From large language models to large knowledge models. *Comput Chem Eng.* 2024;192:108895. doi:10.1016/j.compchemeng.2024.108895.
- [29]. Venkatasubramanian V, Chakraborty A. Quo vadis ChatGPT? From large language models to large knowledge models. *Comput Chem Eng.* 2024;192:108895. doi:10.1016/j.compchemeng.2024.108895.

- [30]. Weber JM, Guo Z, Zhang C, Schweidtmann AM, Lapkin AA. Chemical data intelligence for sustainable chemistry. *Chem Soc Rev.* 2021;50(21):12013-12036. doi:10.1039/d1cs00477h.
- [31]. Wehinger GD, Ambrosetti M, Cheula R, Ding Z, Isoz M, Kreitz B, *et al.* Quo vadis multiscale modeling in reaction engineering? A perspective. *Process Saf Environ Prot.* 2022;184:39-58. doi:10.1016/j.cherd.2022.05.030.
- [32]. Yin M, Wang Y, Yu W, Wilson D, Ma X. Industrial processes fault diagnosis method based on expert system-guided neural network decision-space sparsification. *IEEE Trans Syst Man Cybern Syst.* 2026:1-13. doi:10.1109/TSMC.2026.3673583.