The Paradigm Shift Toward AI-Native Process Simulation: DWSIM Open-Source Developments and a Vision for 2030

Daniel Wagner Oliveira de Medeiros[a], Nicolas Spogis[b,c]*

[a]*DWSIM Developer, Mossoró-RN, Brazil*
[b]*School of Chemical Engineering, University of Campinas, Campinas-SP, Brazil*
[c]*AI4Tech, Campinas-SP, Brazil*
**nicolas.spogis@gmail.com*

ABSTRACT

The integration of artificial intelligence with process systems engineering is transforming the design, simulation, and optimization of chemical processes. Although commercial simulators have begun incorporating AI as auxiliary modules, open-source platforms present a distinct opportunity to embed intelligence natively across every layer of the simulation workflow. This paper reports recent developments in DWSIM version 10, including a context-aware large language model (LLM) assistant, an artificial neural network (ANN)-based convergence enhancer for thermodynamic calculations, expanded dynamic simulation capabilities, and four community-developed unit operations that demonstrate the platform's extensible architecture. Additionally, the AI4Tech Suite—a companion web-based platform—is introduced, implementing a unified pipeline from design of experiments and surrogate modeling through multi-objective optimization and operability analysis. A forward-looking perspective for 2030 is presented, examining the trajectory toward autonomous simulation agents, self-calibrating digital twins, and broader democratization of advanced process engineering methods.

Keywords: Process simulation; Artificial intelligence; Open-source software; Surrogate modeling; DWSIM

1. Introduction

Process Systems Engineering (PSE) has historically relied on rigorous first-principles models implemented within commercial simulation environments such as Aspen Plus, AVEVA Process Simulation, and Honeywell UniSim. While these tools provide robust thermodynamic and unit-operation modeling capabilities, they impose substantial barriers to broader adoption: high licensing costs limit access for universities and small enterprises, closed-source architectures constrain reproducibility and customization, and their monolithic designs resist integration with modern data-driven workflows (Grossmann and Harjunkoski, 2019). As artificial intelligence rapidly advances across adjacent engineering disciplines, a critical question emerges for the PSE community: can process simulation evolve beyond its established paradigm toward a truly AI-native architecture?

An AI-native simulator differs fundamentally from the incremental addition of AI features to legacy software. Rather than treating intelligence as an auxiliary module, the AI-native approach embeds computational intelligence at every layer of the simulation workflow: neural network models provide high-quality initial estimates that accelerate thermodynamic convergence; natural language interfaces enable conversational interaction with process flowsheets; and optimization pipelines connect rigorous first-principles models with data-driven surrogates. Recent advances in large language models (LLMs), surrogate modeling, and automated machine learning have rendered this architecture technically feasible (Schweidtmann et al., 2021).

DWSIM is an open-source chemical process simulator with more than two decades of continuous development and a user base spanning academia and industry in over 190 countries. Its version 10 release incorporates a set of AI-native features that constitute the primary technical contributions of this work. Concurrently, the AI4Tech Suite—a companion platform developed to extend DWSIM's capabilities toward data-driven decision-making—implements a complete PSE workflow encompassing design of experiments, surrogate modeling, multi-objective optimization, and operability assessment.

This paper makes three interconnected contributions: (i) a description of the AI-native features introduced in DWSIM 10, including the LLM assistant, ANN-based convergence enhancer, and new dynamic simulation models; (ii) a presentation of four community-developed unit operations that illustrate the platform's extensible plugin architecture; and (iii) an introduction to the AI4Tech Suite as an integrated framework for intelligent process engineering. The paper concludes with a forward-looking discussion on how these developments may reshape the PSE landscape by 2030.

2. DWSIM 10: An AI-Native Process Simulator

The principal innovation introduced in DWSIM 10 is a context-aware AI assistant that operates natively within the simulation environment, as illustrated in Figure 1. The assistant maintains full awareness of the active flowsheet—including stream compositions, operating conditions, and process topology—enabling engineers to query process states, request result summaries, and receive technically grounded improvement suggestions through natural language interaction. The system supports multiple LLM backends: local models deployed via Ollama for data-sensitive industrial settings, and cloud-based models such as OpenAI GPT and Anthropic Claude for enhanced reasoning. Communication occurs exclusively through a local HTTP API, ensuring that proprietary process data remain within the user's own infrastructure.

To enrich the assistant's contextual understanding, DWSIM 10 introduces structured flowsheet metadata fields comprising process-type classification, detailed process descriptions, and key compound identification across feed streams, reactants, and products. These structured inputs enable the model to generate responses that are not merely linguistically coherent but technically grounded in the specific simulation context, addressing a critical limitation of general-purpose LLM deployments in engineering environments.

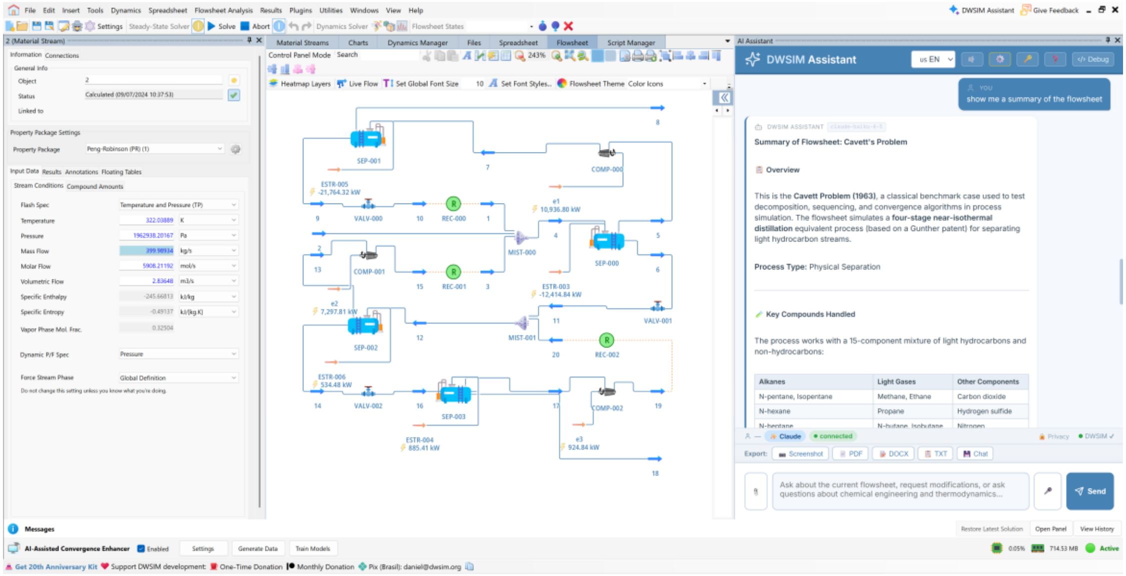

2 (Material Stream)

Information Connections

General Info
Object: 2
Status: Calculated (09/07/2024 10:37:53)
Linked to

Property Package Settings
Property Package: Peng-Robinson (PR) (1)

Input Data Results Annotations Floating Tables

Stream Conditions Compound Amounts

Flash Spec	Temperature and Pressure (TP)	
Temperature	322.03889	K
Pressure	1962938.20167	Pa
Mass Flow	399.98934	kg/s
Molar Flow	5908.21192	mol/s
Volumetric Flow	2.83648	m3/s
Specific Enthalpy	-245.66813	kJ/kg
Specific Entropy	-0.49137	kJ/[kg.K]
Vapor Phase Mol. Frac.	0.32504	
Dynamic P/F Spec	Pressure	
Force Stream Phase	Global Definition	

Do not change this setting unless you know what you're doing.

YOU
show me a summary of the flowsheet

DWSIM ASSISTANT claude-haiku-4-5

Summary of Flowsheet: Cavett's Problem

Overview

This is the **Cavett Problem (1963)**, a classical benchmark case used to test decomposition, sequencing, and convergence algorithms in process simulation. The flowsheet simulates a **four-stage near-isothermal distillation** equivalent process (based on a Gunther patent) for separating light hydrocarbon streams.

Process Type: Physical Separation

Key Compounds Handled

The process works with a 15-component mixture of light hydrocarbons and non-hydrocarbons:

Alkanes	Light Gases	Other Components
N-pentane, Isopentane	Methane, Ethane	Carbon dioxide
N-hexane	Propane	Hydrogen sulfide
N-heptane	N-butane, Isobutane	Nitrogen

Claude • connected Privacy • DWSIM ✓

Export: Screenshot PDF DOCX TXT Chat

Ask about the current flowsheet, request modifications, or ask questions about chemical engineering and thermodynamics...

Send

AI-Assisted Convergence Enhancer ☑ Enabled Settings Generate Data Train Models

Get 20th Anniversary Kit ♥ Support DWSIM development: One-Time Donation Monthly Donation Pix (Brasil): daniel@dwsim.org

Figure 1. The DWSIM 10 AI Assistant interface, showing a natural language interaction with a process flowsheet.

Beyond conversational interaction, DWSIM 10 introduces a dedicated module to address one of the most persistent computational challenges in process simulation: the convergence of thermodynamic equilibrium calculations. Flash algorithms, Gibbs energy minimization reactors, and phase envelope generation are iterative by nature and frequently fail to converge for complex multicomponent systems at extreme conditions.

The AI Convergence Enhancer addresses this challenge by employing pre-trained ANN models to generate high-quality initial estimates for thermodynamic calculations. When standard iterative algorithms fail, the module provides approximate fallback solutions that allow the simulation to proceed. The system operates at three configurable assistance levels: (1) an estimation mode that supplies ANN-generated initial guesses to accelerate standard convergence; (2) a reactive mode that activates ANN assistance only upon algorithm failure; and (3) a comprehensive mode that combines initial estimates with fallback approximations for persistently non-converging cases.

An opt-in data-sharing mechanism complements the convergence enhancer: users may upload anonymized thermodynamic calculation results to a central server, enabling iterative refinement of the ANN prediction models. This federated approach improves model accuracy collectively across the user base while preserving the confidentiality of individual process data.

DWSIM's dynamic simulation capabilities are substantially extended in version 10, an area where open-source tools have historically lagged behind commercial alternatives. Three new dynamic unit operation models are introduced: (1) the Pipe Segment, implementing transient hydraulic and thermal calculations; (2) the Absorption Column, enabling dynamic gas treatment analysis; and (3) the Safety Relief Valve, designed for overpressure scenario studies. An enhanced heat balance model for the Separator Vessel further supports rigorous thermal calculations during depressurization. Together, these models establish a coherent framework for dynamic safety analysis within an open-source environment.

In combination, these dynamic models enable the simulation of complete pressure relief sequences—from initial overpressure through valve actuation to final blowdown—a workflow previously confined to specialized commercial tools. Thermodynamic capabilities are further extended through integration with the ThermoPack library (Hammer et al., 2020), which provides access to advanced equations of state including PC-SAFT, SPC-SAFT, SAFT-VR(Q) Mie, and CPA variants, as well as multi-parameter reference equations such as NIST-MEOS and MBWR. A new immiscible liquids flash option improves the treatment of liquid–liquid equilibrium systems prevalent in biochemical and extraction processes.

3. Community-Driven Extensibility: New Unit Operations

A key advantage of open-source simulation platforms is their capacity for community-driven extensibility: domain experts may develop and distribute specialized models that integrate directly with the platform's thermodynamic engine and graphical interface. DWSIM's plugin architecture enables this through external library distribution, and this section describes four unit operations recently developed to address persistent gaps in the open-source simulation landscape.

The cooling tower unit operation is based on the classical Merkel method (Merkel, 1925), which formulates coupled heat and mass transfer through psychrometric properties. The model solves the enthalpy balance across fill segments while accounting for water evaporation losses and air–water contact conditions.

The multi-stream heat exchanger model addresses a notable limitation of open-source process simulation: the inability to rigorously model exchangers involving more than two process streams, which are essential in LNG liquefaction, cryogenic air separation, and integrated heat recovery networks. The implementation employs the segmented interval method with composite curve construction, consistent with approaches adopted in leading commercial simulators (Kamath et al., 2012; Watson et al., 2015). The model supports multiple specification modes—including outlet temperatures, minimum internal temperature approach (MITA), and overall UA specification—for both counterflow and co-current configurations.

The plate heat exchanger model incorporates geometry-specific heat transfer and pressure drop correlations for chevron-pattern plates, enabling more accurate simulation of compact exchangers used in food processing, pharmaceutical manufacturing, and HVAC applications, for which shell-and-tube correlations are not applicable.

The calandria evaporator model simulates vertical-tube, natural-circulation evaporators widely used in sugar refining, pulp and paper processing, and chemical concentration. The model couples steam-side condensation with process-side boiling calculations, accounting for boiling-point elevation due to dissolved solids and hydrostatic head effects.

All four unit operations are distributed as open-source packages and rely on DWSIM's native thermodynamic engine for physical property calculations, ensuring thermodynamic consistency with the broader simulation environment. This collaborative development model—in which academic researchers and industry practitioners contribute specialized equipment models within a shared framework—provides a scalable and sustainable mechanism for expanding open-source simulation capabilities without concentrating the development burden on a single maintainer.

4. AI4Tech Suite: Bridging Simulation and Intelligence

While DWSIM 10 embeds AI capabilities within the simulator itself, the AI4Tech Suite addresses the complementary challenge of connecting rigorous first-principles simulation with the broader data-driven workflow that modern process engineering requires. Implemented as a browser-accessible scientific computing platform, the suite integrates the complete PSE pipeline—from experimental design through simulation execution, statistical analysis, surrogate modeling, optimization, and operability assessment—into a unified environment with direct interfaces to DWSIM, Excel, Python scripts, and AVEVA simulation backends. Figure 2 illustrates the platform architecture and the data flow across its modules.

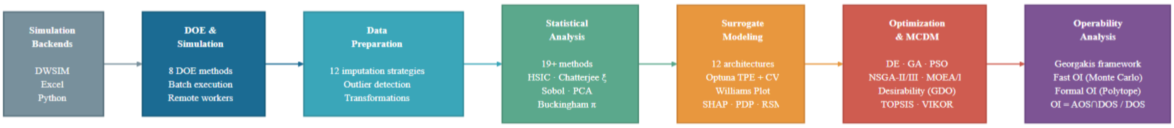

Figure 2. Architecture and workflow of the AI4Tech Suite.

The workflow begins with structured experimental design, offering eight sampling and design methods: Latin Hypercube Sampling, Full Factorial, Central Composite Design, Box-Behnken, Plackett-Burman, D-Optimal, Sobol sequences, and custom user-defined designs. These methods generate simulation campaigns across continuous, discrete, and fixed variable spaces. A distributed execution engine with remote worker nodes enables parallel simulation across multiple machines, substantially reducing the wall-clock time of large DOE campaigns. Once simulation data are collected, a data preparation pipeline applies 12 imputation strategies for missing values, detects outliers via Isolation Forest, Mahalanobis distance, and Z-score methods, and applies variable transformations to condition the dataset for surrogate modeling.

The statistical analysis module provides more than nineteen methods, organized in a progressive hierarchy from exploratory visualization to advanced dependence detection. Beyond classical correlation measures (Pearson, Spearman, and Kendall coefficients), the platform implements the Hilbert-Schmidt Independence Criterion (HSIC) for detecting nonlinear and non-monotonic statistical dependencies through kernel methods (Gretton et al., 2005), and Chatterjee's ξ correlation for the detection of asymmetric functional relationships (Chatterjee, 2021)—measures capable of identifying dependencies that classical linear coefficients may overlook.

Global sensitivity analysis via Sobol variance-based indices quantifies the contribution of each input variable to output variance. Principal Component Analysis and K-Means clustering provide dimensionality reduction and pattern recognition capabilities. A Buckingham π theorem module automates dimensional analysis and identifies known dimensionless groups, supporting the derivation of physically interpretable correlations.

The surrogate modeling module offers twelve model architectures spanning the complexity–interpretability spectrum: polynomial regression; regularized linear models (Ridge, Lasso, Elastic Net); ensemble methods (Random Forest, XGBoost, LightGBM, CatBoost, and Stacking); and flexible nonlinear models (multilayer perceptron neural networks, Gaussian Process Regression, and Support Vector Regression).

Hyperparameter optimization is automated through Optuna's Tree-structured Parzen Estimator (TPE) with stratified k-fold cross-validation (Akiba et al., 2019), eliminating the need for manual tuning. The implementation emphasizes model validation and interpretability criteria relevant to engineering applications. The Williams Plot is used to assess applicability domain boundaries, flagging predictions that require extrapolation beyond the training data envelope (Tropsha et al., 2003). SHAP (SHapley Additive exPlanations) analysis provides a game-theoretic decomposition of feature importance, quantifying each input variable's contribution to individual model predictions (Lundberg and Lee, 2017). Partial dependence plots and three-dimensional response surface visualizations reveal marginal effects and interaction patterns among process variables, supporting physically informed model interpretation. Figure 3 illustrates representative outputs from the surrogate modeling and interpretation workflow.

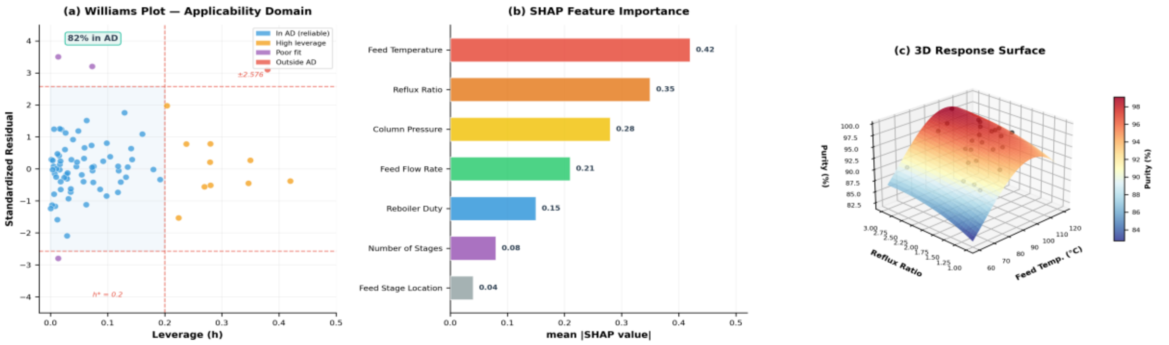

Figure 3. Representative outputs from the AI4Tech Suite: (a) Williams Plot showing the applicability domain of a trained surrogate model; (b) SHAP feature importance analysis; (c) 3D response surface visualization illustrating the interaction between two key process variables.

The optimization framework is structured hierarchically to match solver complexity to problem requirements. Single-objective problems are addressed using evolutionary algorithms (Differential Evolution, Genetic Algorithm, and Particle Swarm Optimization) and gradient-based methods (L-BFGS-B, SLSQP, and Basin Hopping). Multi-objective formulations employ NSGA-II, NSGA-III, and MOEA/D to generate Pareto-optimal trade-off fronts (Deb et al., 2002; Blank and Deb, 2020). A dedicated Goal-Driven Optimization module applies Derringer–Suich desirability functions for simultaneous multi-response optimization (Derringer and Suich, 1980), while a Process Change Optimization module minimizes operational disruption during transitions between steady states.

Following the steady-state operability framework of Georgakis and co-workers (Georgakis et al., 2003; Lima et al., 2010), the operability analysis module evaluates whether a given process design can achieve desired output specifications across the feasible range of available inputs. Two complementary computational approaches are implemented: Monte Carlo sampling for rapid probabilistic screening and geometric polytope computation for exact deterministic results. The Operability Index—defined as the ratio of the intersection between the Achievable Output Space and the Desired Output Space to the total Desired Output Space—provides a single scalar metric quantifying the inherent flexibility of the design. This analysis closes the loop between optimization (identifying what is optimal) and operability (determining what is achievable under real operating constraints).

5. A Vision for 2030

The developments reported in this paper represent an early but substantive stage in the broader transformation of process simulation toward AI-native paradigms. Several convergent technical trends are expected to reshape the field over the coming years, and the open-source ecosystem described herein is well positioned to both contribute to and benefit from these developments.

The most consequential near-term trend is the emergence of autonomous simulation agents. Building upon the LLM assistant introduced in DWSIM 10, future agent systems are expected to autonomously construct process flowsheets from natural language specifications, select thermodynamically appropriate models for the chemical system under study, resolve convergence failures through reasoning-guided strategies, and propose design modifications targeting economic or environmental objectives. The coupling of LLMs with structured simulation APIs—in which the model both reasons about engineering principles and executes simulation commands—creates a foundation for human–AI collaboration in which the process engineer provides domain intent and critical judgment while the agent manages routine computational tasks.

A second trend involves self-calibrating digital twins. The combination of rigorous first-principles simulation with surrogate modeling and data-driven optimization, as implemented in DWSIM and the AI4Tech Suite, establishes a natural pathway toward process digital twins that continuously update model parameters based on real-time plant data. ANN-enhanced convergence modules enable simulation execution at speeds approaching real-time, while surrogate models provide the computational efficiency required for online optimization, closed-loop control, and predictive maintenance applications.

The federated data-sharing mechanism of the AI Convergence Enhancer illustrates a broader principle: community-driven improvement of thermodynamic model accuracy without requiring disclosure of proprietary process data. This approach is particularly relevant to the chemical industry, where competitive concerns have historically constrained data sharing and collaborative model development.

From a broader perspective, these developments contribute to the democratization of advanced PSE methods. By integrating surrogate modeling, multi-objective optimization, SHAP-based interpretability, and operability analysis within accessible browser-based interfaces, the tools described here substantially lower the technical and financial barriers for researchers, students, and small enterprises—particularly in developing economies—to apply methods that were, until recently, accessible only to organizations with substantial commercial software licenses and specialized expertise. The open-source licensing of this ecosystem ensures that these advances do not replicate existing technological inequalities but instead extend access to rigorous process engineering capabilities on a global scale.

6. Conclusions

This paper has reported recent advances in the DWSIM open-source ecosystem that collectively mark a transition toward AI-native process simulation. DWSIM 10 introduces two principal AI capabilities: a context-aware LLM assistant for natural language interaction with process flowsheets, and an ANN-based convergence enhancer for thermodynamic calculations. The release also expands dynamic simulation capabilities relevant to safety analysis and integrates the ThermoPack thermodynamic library for advanced equations of state. Four community-developed unit operations—a cooling tower, a multi-stream heat exchanger, a plate heat exchanger, and a calandria evaporator—demonstrate the viability and extensibility of the platform's plugin architecture.

The AI4Tech Suite complements the simulator with a complete data-driven workflow encompassing eight experimental design methods, twelve surrogate model architectures with automated hyperparameter optimization, single- and multi-objective optimization solvers, five MCDM methods, and operability analysis grounded in the Georgakis framework.

Collectively, these developments demonstrate that open-source platforms are not merely competitive with commercial alternatives in AI integration, but are in a position to pioneer new paradigms—offering transparency, extensibility, and broad accessibility that proprietary software cannot readily provide. As the PSE community advances toward 2030, the technical foundations for autonomous, intelligent, and democratized process engineering are being established within the open-source ecosystem.

Acknowledgments

The authors acknowledge the DWSIM open-source community and Patreon supporters whose contributions make continued development possible.

References

T. Akiba, S. Sano, T. Yanase, T. Ohta and M. Koyama: Optuna: A Next-generation Hyperparameter Optimization Framework, Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2623–2631, 2019. https://doi.org/10.1145/3292500.3330701

J. Blank and K. Deb: pymoo: Multi-Objective Optimization in Python, IEEE Access (8), 89497–89509, 2020. https://doi.org/10.1109/ACCESS.2020.2990567

S. Chatterjee: A New Coefficient of Correlation, Journal of the American Statistical Association (116), 2009–2022, 2021. https://doi.org/10.1080/01621459.2020.1758115

K. Deb, A. Pratap, S. Agarwal and T. Meyarivan: A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II, IEEE Transactions on Evolutionary Computation (6), 182–197, 2002. https://doi.org/10.1109/4235.996017

G. Derringer and R. Suich: Simultaneous Optimization of Several Response Variables, Journal of Quality Technology (12), 214–219, 1980. https://doi.org/10.1080/00224065.1980.11980968

C. Georgakis, D. Uztürk, S. Subramanian and D.R. Vinson: On the Operability of Continuous Processes, Control Engineering Practice (11), 859–869, 2003. https://doi.org/10.1016/S0967-0661(02)00217-4

A. Gretton, O. Bousquet, A. Smola and B. Schölkopf: Measuring Statistical Dependence with Hilbert-Schmidt Norms, Proceedings of the 16th International Conference on Algorithmic Learning Theory, 63–77, 2005. https://doi.org/10.1007/11564089_7

I.E. Grossmann and I. Harjunkoski: Process Systems Engineering: Academic and Industrial Perspectives, Computers & Chemical Engineering (126), 474–484, 2019. https://doi.org/10.1016/j.compchemeng.2019.04.028

M. Hammer, A. Aasen, Å. Ervik and Ø. Wilhelmsen: Choice of Reference, Influence of Non-Additivity and Present Challenges in Thermodynamic Perturbation Theory for Mixtures, Journal of Chemical Physics (152), 134106, 2020. https://doi.org/10.1063/1.5142771

R.S. Kamath, L.T. Biegler and I.E. Grossmann: Modeling Multistream Heat Exchangers with and without Phase Changes for Simultaneous Optimization and Heat Integration, AIChE Journal (58), 190–204, 2012. https://doi.org/10.1002/aic.12565

F.V. Lima, Z. Jia, M. Ierapetritou and C. Georgakis: Similarities and Differences Between the Concepts of Operability and Flexibility, AIChE Journal (56), 702–716, 2010. https://doi.org/10.1002/aic.12021

S.M. Lundberg and S.I. Lee: A Unified Approach to Interpreting Model Predictions, Advances in Neural Information Processing Systems (30), 4765–4774, 2017. https://papers.nips.cc/paper/7062

F. Merkel: Verdunstungskühlung, VDI-Zeitschrift (70), 123–128, 1925.

A.M. Schweidtmann, E. Esche, A. Fischer, M. Kloft, J.U. Repke, S. Sager and A. Mitsos: Machine Learning in Chemical Engineering: A Perspective, Chemie Ingenieur Technik (93), 2029–2039, 2021. https://doi.org/10.1002/cite.202100083

A. Tropsha, P. Gramatica and V.K. Gombar: The Importance of Being Earnest: Validation is the Absolute Essential for Successful Application of QSPR/QSAR Models, QSAR & Combinatorial Science (22), 69–77, 2003. https://doi.org/10.1002/qsar.200390007

H.A.J. Watson, K.A. Khan and P.I. Barton: Multistream Heat Exchanger Modeling and Design, AIChE Journal (61), 3390–3403, 2015. https://doi.org/10.1002/aic.14965