I&EC research
Industrial & Engineering Chemistry Research

pubs.acs.org/IECR
Perspective

The Importance of First-Principles in the Era of Artificial Intelligence: A Perspective on Hybrid Modeling and Industrial Case Studies

Published as part of Industrial & Engineering Chemistry Research special issue "Importance of First-Principles Modeling in the Era of Growing AI".

Devin Griffith,* Saud Abdulaziz Al-Ghwainem, and Giampaolo Serra

Cite This: *Ind. Eng. Chem. Res.* 2026, 65, 16960–16969

ABSTRACT: In this work, we will give a perspective on the importance of first-principles in the era of AI and ML, as well as showcase real case studies with industrial collaborators that make the value of a hybrid approach apparent. The first case study shows multiple examples of hybrid modeling applied to polymer production processes. The second case study focuses on the modeling of a fluid catalytic cracking (FCC) process for the purposes of production planning optimization.

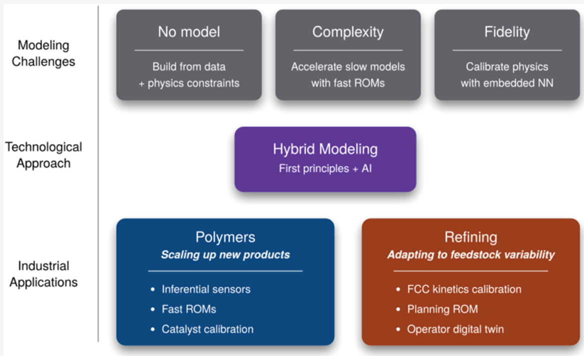

1. INTRODUCTION

Over the past decade, artificial intelligence (AI) and machine learning (ML) in process modeling have exploded in popularity, dominating in academic research, industrial applications, and startup companies. This trend has completely shifted the landscape of innovation, as well as the skillsets of new engineers joining the workforce. The previously niche chemometrics departments of established industrial companies have exponentially grown in importance, encompassing entire divisions dedicated to data science and AI.

However, in practice, there is often still a gap between the latest innovations and what is actually implemented in the plant. This is due to a myriad of practical, technical, and interpersonal limitations. However, recent innovations allow engineers to overcome these limitations by intelligently combining modern advances in data science and machine learning with decades of history and expertise in first-principles modeling. This allows us to overcome the drawbacks and leverage the strengths of each approach in order to provide synergistic and fit-for-purpose solutions.

Historically speaking, chemical engineers have been combining first-principles models with insights from data for many decades. In the past, these insights were called "empirical correlations" and often employed when rigorously deriving a relationship was impractical. Examples range from thermodynamic property correlations to capital cost models, often used in conjunction with rigorous first-principles mass and energy balance models. Today, the possibilities of such a "hybrid" approach have expanded greatly with the incorporation of advances from AI and ML, and the rate of innovation has exploded with previously unsolvable problems being addressed at an impressive rate.

In this work, we will elaborate on our perspective of the importance of first-principles in the era of AI, including use cases, strategies, and workflows. The advantages of this approach will be showcased via the application of hybrid modeling with two industrial collaborators. The information contained in this paper is meant to be accessible and understandable by any industrial practitioner and will therefore forego equations and in-depth technical descriptions in favor of focusing on philosophical approaches and value statements. Although there is now an impressive and growing span of academic literature on this topic,[1−4] here we will focus on what has been shown to be impactful in an industrial setting.

Received: December 9, 2025
Revised: May 27, 2026
Accepted: June 15, 2026
Published: August 6, 2026

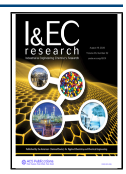

https://doi.org/10.1021/acs.iecr.5c04353
Ind. Eng. Chem. Res. 2026, 65, 16960–16969

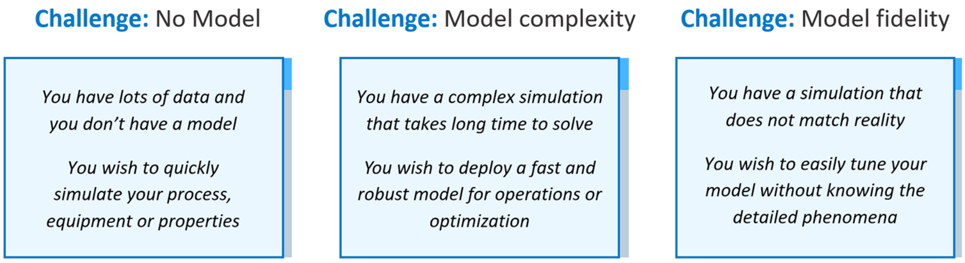

Figure 1. Common challenges in process modeling.

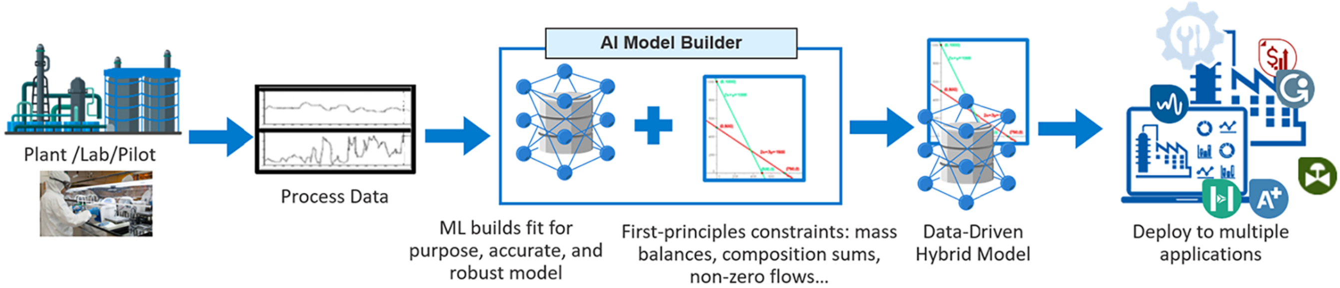

Figure 2. Data driven modeling workflow, AKA AI driven hybrid modeling.

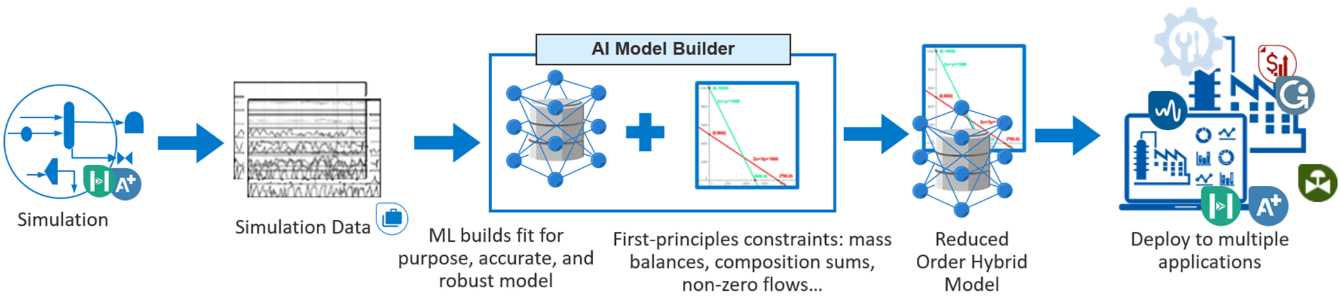

Figure 3. Reduced order hybrid modeling workflow.

2. CHALLENGES ADDRESSED BY HYBRID MODELING

There are many challenges associated with the application of simulation technology in the process industries. Creating a fit-for-purpose model that accurately represents a complex real-world process is often a nontrivial task involving technological, operational, and data constraints. In this section, we will give a broad and general overview of the types of challenges that are more and more commonly being addressed by hybrid modeling in the industry. A very brief summary of these challenges is given in Figure 1.

Note that, throughout this work, there will be limited discussion of any particular regression methods. The regression method should be chosen to appropriately fit a given application, and this choice tends to matter less than having a well-formulated problem statement and clear business motivation. Although advanced regression methods such as neural networks have a great deal of value in many cases, methods such as the least absolute shrinkage and selection operator (LASSO[5]) are often preferred when the system exhibits moderate nonlinearities for reasons of simplicity and understandability.

2.1. Model Creation

The first challenge is the most obvious, as it aligns with the typical data science or machine learning problem statement. In this case, it is desirable to have a prediction for something important in the plant, such as a critical KPI related to product quality, energy usage, or emissions. Suppose that a first-principles model for a particular phenomenon is not available, but process data is available. Therefore, it is desirable to create a model from data. This is possible using a wide variety of established regression techniques in conjunction with first-principles integrated as equality constraints, inequality constraints, or by simply embedding the regression model within a broader first-principles model. This workflow in the AspenTech toolset is summarized in Figure 2.

Note that the lack of a first-principles model could be due to the lack of an available out-of-the-box model or that a first-principles model is otherwise costly to create for a given application. One example would be a heat exchanger model that predicts the rate of fouling into the future. Creating a first-principles model for fouling rate would likely require a complex CFD simulation, which is impractical to create for every heat exchanger in the plant. Another example would be hard-to-predict product properties such as porosity, yellowness index, or true vapor pressure.[6] An example related to polymers will be shown in Section 4.1.1.

2.2. Model Complexity

The second challenge is model complexity. In this case, suppose that a first-principles model has been developed that successfully accomplishes its intended purpose, making accurate predictions of something critical in the plant. However, in many cases, this model could be too computa-

https://doi.org/10.1021/acs.iecr.5c04353
Ind. Eng. Chem. Res. 2026, 65, 16960−16969

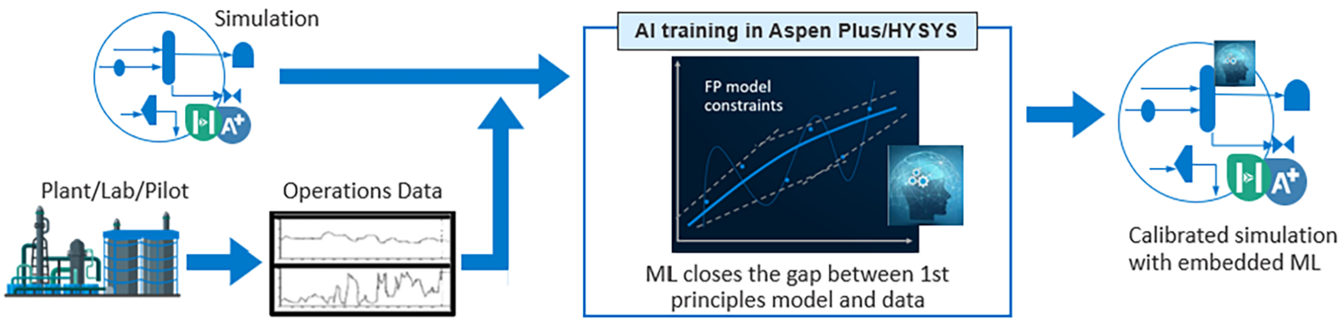

Figure 4. Model calibration workflow AKA first-principles driven hybrid modeling.

tionally complex for practical usage. To overcome this limitation, it is possible to create accurate reduced order models (ROMs, synonymous with surrogate models) from existing first-principles models.

In this case, data is generated from the first-principles model via some sampling strategy, which is then used by some regression method to create fast and accurate representation of the original model. By applying constraints within a hybrid modeling architecture, critical first-principles properties such as a mass balance that are not inherently guaranteed by general regression models can be preserved.

Note that sampling strategies are key here. Naive approaches such as a grid-based search can scale very poorly or give poor accuracy results in particular regions of the search space. The reduced order modeling workflow in the AspenTech toolset is summarized in Figure 3.

One common motivation for this workflow is that engineers need fast predictions in order to make better decisions in operations. Even if a first-principles model is perfectly accurate, if it takes a long time to solve, it may not be useful for real-time decision making. Thus, we deploy a ROM of the first-principles model as a process digital twin, enabling superior operational decisions with fast, on demand predictions. It is also possible to use a ROM to initialize a rigorous model for faster convergence. A polymers example is shown in Section 4.1.2.

Another critical use case is in deriving simplified representations of process simulations for use in higher level optimization problems, such as in production planning optimization. The ROM created in this case must be highly accurate, very reliable under optimization, and computationally tractable in the context of the broader scope problem. The most common applications are in refinery production planning,[7,8] including recent examples in biofuels.[9] We highlight an industrial application for a Fluid Catalytic Cracking reactor (FCC) in the context of production planning optimization in Section 4.2.2.

2.3. Model Fidelity

The third challenge occurs when a first-principles model is available that is very close to what exists in the real plant, but there is some error gap in the predictions. In this case, first-principles models are calibrated by embedding a regression model to predict key parameters that would otherwise be held constant. These parameters, such as kinetic rate constants or Murphree efficiencies, are then allowed to vary according to inputs to the first-principles model. The loss function for training is calculated from the first-principles model's outputs that are measurable in the plant, such as reaction or separation outlet compositions. Training itself becomes computationally more challenging, as backpropagation must occur in the context of the first-principles model with a regression model embedded. The workflow and model architecture in the AspenTech toolset are summarized in Figures 4 and 5.

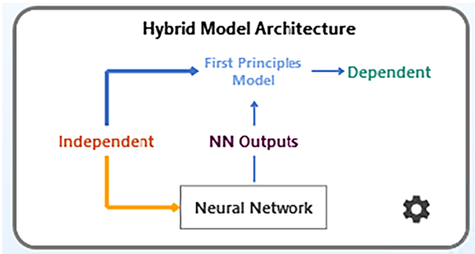

Figure 5. Model calibration architecture.

Estimating kinetic behavior is a common use case for the model calibration workflow. The exact approach employed in a given example depends on what is known and what is unknown. For example, stream conditions, components, and stoichiometry of a given reaction may be available, but not the kinetic parameters. Alternatively, stoichiometry may not be available, but yield data is. This conceptual workflow is flexible enough to accommodate many different practical situations.

This approach gives a model that is hybrid by default, since first-principles are enforced by the model that the regression model is embedded in. There is also a distinct advantage that much less data is needed than for the other two approaches, since a large portion of the behavior of the system is already captured in the first-principles model.

Applications include hybrid modeling for heat exchangers or steam reformers. An example in polymer reaction modeling is presented in Section 4.1.3, and an example in FCC reaction modeling is shown in Section 4.2.1.

3. WHY FIRST-PRINCIPLES MATTER

In this section, we will explain the top reasons that first-principles are still important in data-driven modeling applications. This is based on our experience in many real industrial applications.

1. **Accurate predictions with less data**

   Although it is well known that some regression models such as neural networks have beneficial properties as universal approximators,[10] these properties only describe the existence of such an approximator, and not necessarily how to find it. In reality, function approximation requires "enough" data. The quantification of "enough" is highly situation dependent. In our

https://doi.org/10.1021/acs.iecr.5c04353
Ind. Eng. Chem. Res. 2026, 65, 16960−16969

experience, the availability or generation of sufficient, high-quality data is always the primary limiting factor in data-driven modeling. One common challenge is that, since industrial processes tend to be operated stably at a setpoint, process historian data often does not contain much useful variation. Another challenge is that the most important process KPIs, such as product compositions or other properties, tend to be the hardest to measure, with data points coming from analyzers or lab analysis at a relatively infrequent rate. These challenges will remain relevant until data collection and historization methods improve, and some, such as the lack of variation inherent to stable operation, may persist regardless.

Since the availability of data is always a limiting factor, incorporating more information from first-principles is always beneficial. There is, in most cases, no practical reason to pose a data-driven modeling problem as learning everything from scratch. There is almost always a better starting point to be had by injecting mass and energy balances, known reaction kinetics, thermodynamic properties, or at least process structure. This thereby lessens the degree of information that must be learned from data, and thereby raises the likelihood of success for any given project.

2 **Better structure for optimization**

In the field of process modeling, it is often the goal to create a model that can then be optimized to find the best operating conditions for the plant. Historically, this is done by optimizing first-principles models that enforce physical constraints everywhere. If the model is rigorous and high-quality, then the end result will be a set of operating conditions that are physically consistent, obey process or safety bounds, and give superior process operation.

Given recent advances and trends, there is now desire to use ML models in optimization, either to replace a part or all of a first-principles process model. This has inherent risks, since ML models traditionally only enforce accuracy or other properties at the specific points in a data set, but not necessarily for unseen data. This means that an optimizer could drive the model to an operational condition unseen in the data set that violates known first-principles, e.g., a mass balance. In fact, the optimizer is incentivized to "cheat" by finding points in the ML model that generate mass for free. There are also other considerations such as the optimizer's ability to exploit nonphysical relationships apparent in data (such as a downstream independent variable impacting an upstream dependent variable), or noisy derivatives from potentially overfit regression models causing difficulties with convergence. These difficulties are largely overcome by enforcing known first-principles everywhere in the operational range, not just for given data points. This requires advanced and domain-specific techniques that are not possible with off-the-shelf ML methods.

3 **Improved extrapolation**

In process modeling and optimization, it is often desirable to find new operating conditions not previously seen in the plant. This is a very common issue, especially considering the narrow operational ranges of most industrial process units. Unfortunately, this goal very quickly becomes problematic when using data-driven models that are not guaranteed to be accurate outside of the training range. In fact, highly flexible models such as neural networks tend to degrade in accuracy very quickly outside of the training range.

Incorporating information from first-principles that are true everywhere, not just within the training range, can greatly improve accuracy upon extrapolation, and at least ensure that the result does not violate those first-principles. This is even further improved by conjoining highly flexible models inside the training range with less flexible models (e.g., linear) outside of the training range.

4 **Interpretability**

Models that align with first-principles are generally easier to interpret and understand. This speaks to one of the primary challenges in process modeling in industry: is the model trusted by those who will employ its results? This is, of course, aided by enforcing first-principles in the model. A model that obeys well-known constraints such as mass and energy balances, positive flow rates, and composition fractions summing to one is much more likely to be trusted by someone who was trained to check for those constraints from the start of their engineering education. If any of these critical properties are violated, even by a small amount, the accuracy of the rest of the results is brought into question as well.

5 **Exact satisfaction of constraints**

Even with very high quality data, such as data generated from a process simulator, it is difficult to train a regression model that recovers key first-principles (such as a mass balance) exactly. When using noisy data from the plant, it is all but guaranteed that the data itself will deviate from known first-principles, since sensors in the plant tend to be imperfect. For example, it is very common to see negative flow rates in a process historian due to imperfect flow meter calibration.

Depending on the application that a machine learning model is to be used in, even small deviations from first-principles can cause convergence issues. For example, if a unit model is created that has even a very small mass balance violation, it is likely to perform quite poorly inside of a recycle loop in a process flowsheet simulator. Although most flowsheet simulators have simple checks such as normalizing composition fractions, this can cause error if done at the time of inference instead of at the time of training. Applying these constraints up-front at the time of training allows the user to have confidence in the model's performance when deployed in an application that requires them.

6 **Problem Statement Formulation**

The above items speak to the technical issues related to employing first-principles, however last, but certainly not least, is the human element. We still rely on engineers with a firm grounding in first-principles to formulate sound problem statements connected to business objectives and value, even if that problem is later solved by some AI method. Thus, it is beneficial to deliver easily usable AI tools directly to the engineers who are capable of intelligently guiding their application through alignment with first-principles.

https://doi.org/10.1021/acs.iecr.5c04353
Ind. Eng. Chem. Res. 2026, 65, 16960−16969

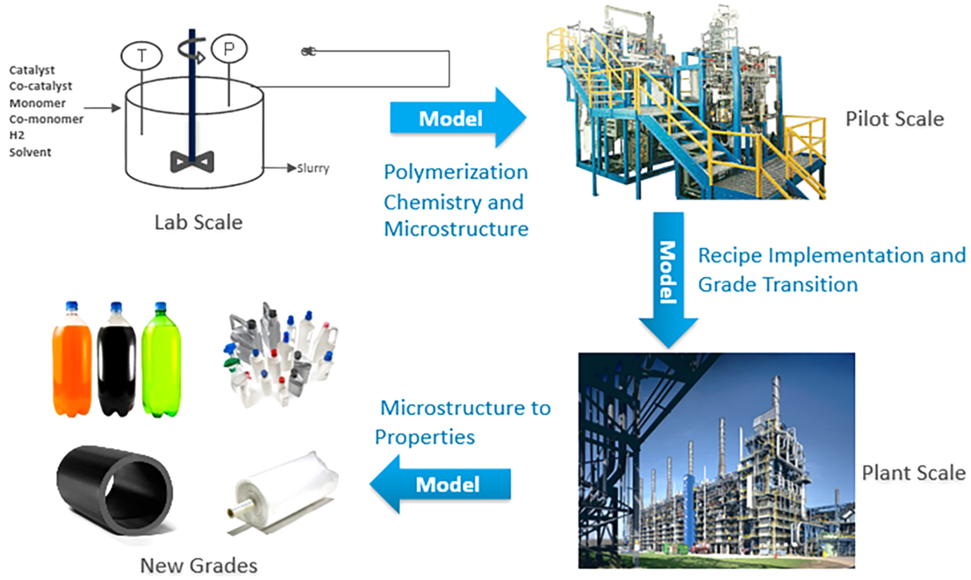

Figure 6. SABIC uses process modeling to enable value at scale.

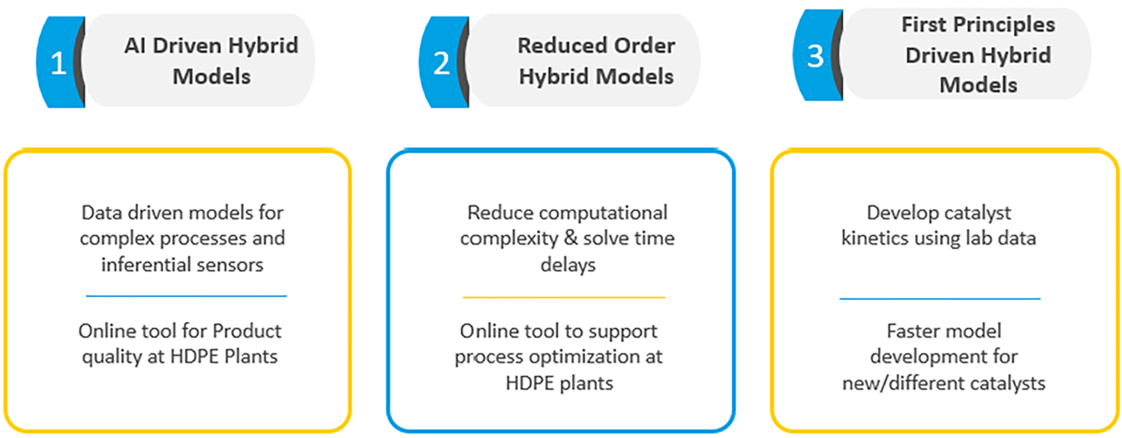

Figure 7. SABIC hybrid modeling applications.

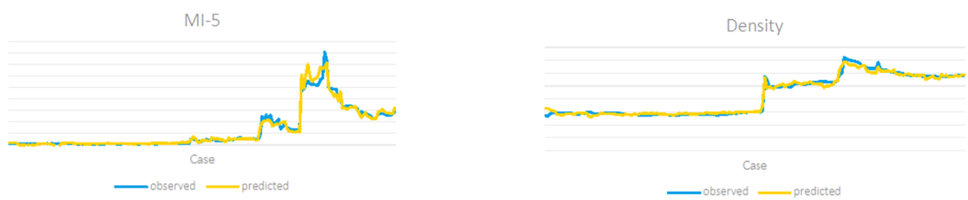

Figure 8. AI driven inferential sensor for properties predictions.

4. CASE STUDIES WITH ASPEN HYBRID MODELS

AspenTech has developed machine learning tools[11] that can easily be used in conjunction with existing AspenTech first-principles modeling tools. This technology, broadly referred to as Aspen Hybrid Models, is very general purpose and applies to all of the process modeling challenges mentioned in Section 2. In the remainder of this work, we show the application of Aspen Hybrid Models in collaboration with SABIC in the polymers space and with Saras S.p.A. in the refining space.

4.1. Case Study in Polymers with SABIC

Saudi Basic Industries Corporation (SABIC) is a global leader in the chemical process industries, with 65 plants globally, roughly 150 new products brought to market each year, $83 billion USD in assets, and the number two chemical brand value worldwide. Excellence in process modeling is a key component of SABIC's process development workflows, ranging from lab scale, to pilot scale, to plant scale, as shown in Figure 6. This supports the tenets of differentiated products, accelerated time-to-market, and superior asset reliability in order to drive business level success.

In this section, we will describe how hybrid modeling has impacted process modeling at SABIC through a multitude of projects. A summary of these projects is shown in Figure 7. Note that, in this context, "online" refers to a model deployed to an on-prem server and connected to a data historian in order to give real-time predictions, implemented via Aspen OnLine.

https://doi.org/10.1021/acs.iecr.5c04353
Ind. Eng. Chem. Res. 2026, 65, 16960−16969

Figure 9. Scope of reduced order model.

Option	Value	Benefit
Data Split (Training/Validation)	80/20	Monitor overfitting
# of Hidden Layer	1	Keep number of parameters low (N_{para} = 28 for 69 data points)
# of Neurons	4	
Activation Function	Leaky Relu	Avoid vanishing gradient
Algorithm	Adam	Good for non-convex problems
Learning Rate/Decay Rate	1E-4 / 0.99	Reduce oscillation in training loss
Batch Size	8	Reduce number of updates and training time required

Figure 10. Neural network parameters for first-principles driven hybrid model.

4.1.1. AI Driven Hybrid Model. In the first project, Aspen AI Model Builder (AIMB) was used to create an inferential sensor for polymer melt-index and density that can easily be deployed online for process monitoring. Predictions are then available on demand, as opposed to waiting for lab results. In order to create this model, 10,000 data points were collected spanning five years of operational data and reconciled with a first-principles Aspen Plus model (this is not a requirement for AIMB, but can be helpful depending on the available data quality). This data was then used in AIMB for LASSO[5] regression with nonlinear transformations, and the resulting model was deployed in Aspen Plus for usage as described in Figure 2. The resulting prediction accuracy can be seen in Figure 8. Exact values in the plots are not shown to maintain confidentiality, but "in-spec" for the melt index is considered to be 0.15−0.21 g/10min.

The resulting model is accurate, fast, lightweight, and easy to use in operations. The relevant use cases of this model include process monitoring, recipe development for new grades, recipe testing, and control. In this case, alignment with first-principles is retained by embedding the data driven model in the Aspen Plus flowsheet, where first-principles predictions become the inputs to the data driven model. Leveraging first-principles for the bulk of the predictions greatly reduces the overall amount of data required. Using this hybrid model, SABIC was able to predict the effect of the transition parameters to avoid off-spec production for an extended or unplanned period.

4.1.2. Reduced Order Hybrid Model. In the second project, a ROM was developed for an existing Aspen Plus model in order to deploy it for operational support, the scope of which is shown in Figure 9. The original model took around 30 min to solve, which is not sufficient for real-time decision making.

Data was generated for 3600 simulator cases via Aspen Multi-Case, which is a very useful tool for ROM creation that both defines the sampling strategy for a given operational range and manages multiple simulator runs in parallel. Generating this data set required approximately 18 h on a high-performance machine. This data was fed to AIMB for regression, as in the workflow shown in Figure 3. In this case, a LASSO method with nonlinear transformations and constraints was found to be sufficient, with an overall model fit of R^2 = 0.98 on a 20% testing set. Note that, through proprietary methods, constraints such as a mass balance are not only enforced during training, but also upon prediction. A built-in knowledge informed regression method was also used to enforce the known causal structure of the process, so that only physically logical independent variables may impact a given dependent variable. These first-principles guardrails ensure that the data-driven model gives trustworthy results that do not violate physical reality.

Again, the resulting model is accurate, lightweight, and convenient for use in operation. Computational time is reduced from roughly 30 min to a few seconds, thereby providing rapid responses in real time. This enabled live what-if analysis for grade transitions, which was proven to reduce off-spec production during trials.

The ROM and AI driven projects were accomplished in less than a month of effort total, which highlights the value of these tools to support agile responses to changing operational conditions.

4.1.3. First-Principles Driven Hybrid Model. In the final project, the AI training workflow within Aspen Plus was used to calibrate the kinetic model of a reactor. The result is also known as a first-principles driven hybrid model. The overall workflow is shown in Figure 4, with the model architecture shown in Figure 5.

This was done in order to adapt an existing Aspen Plus model to a new catalyst, for which limited commercial data is available. This is possible with relatively few data points since the mechanistic core of the model is retained, and the data driven portion of the model only needs to account for a marginal error gap. In this and most similar use cases, the necessary neural network is very small, with only a handful of neurons. The parameters of the neural network are shown in Figure 10.

https://doi.org/10.1021/acs.iecr.5c04353
Ind. Eng. Chem. Res. 2026, 65, 16960−16969

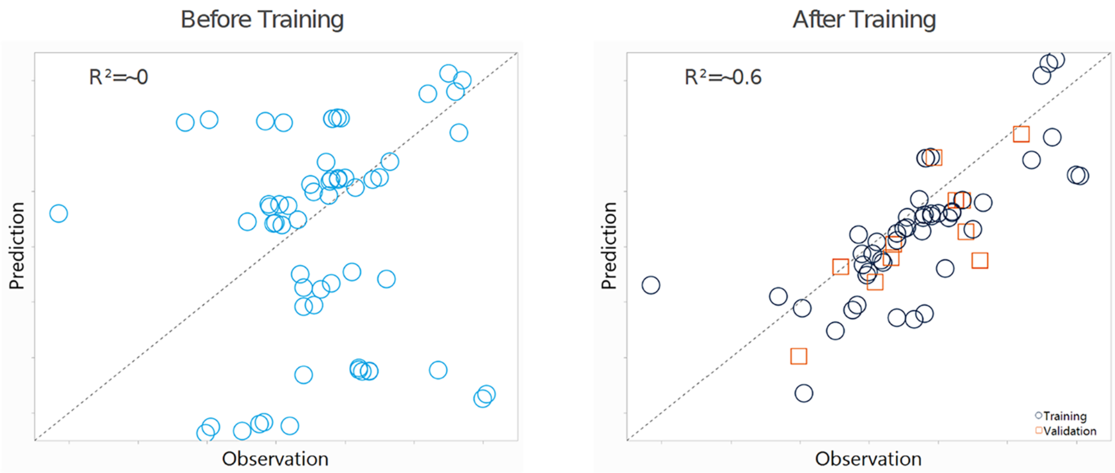

Figure 11. AI training in Aspen Plus gives significant accuracy improvements.

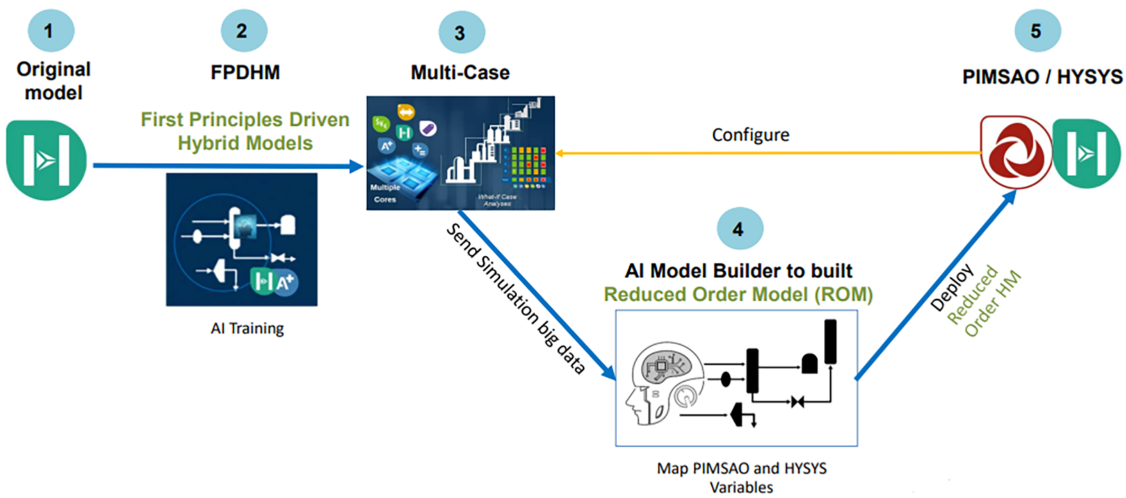

Figure 12. Hybrid modeling workflow for FCC simulation and planning model update.

Most results of this example cannot be shared, but the accuracy improvement is illustrated in Figure 11. This is a significant improvement given the limited data that was available, and this was accomplished before performing any full-scale trials. This was used to support recipe development and operational training, as well as serve as the basis for dynamic simulation of new catalyst transitions. First-principles provide a huge benefit in this example such that minimal data is needed to train an accurate model.

4.1.4. Summary. As shown in the preceding examples, hybrid models at SABIC have proven to reduce off-spec production, monitor hard-to-measure properties during production, capture dynamic transitions, and accelerate scale-up with new catalysts. From a combination of the previously mentioned projects, SABIC projects reduction of off-spec product by 0.45 kilo-tonnes per annum per trial at the time of scale up, as well as accelerated time to market for multiple new grades. Future work will focus on scaling these benefits to additional process units.

4.2. Case Study in Refining with Saras

Societá Anonima Raffinerie Sarde (Saras) is an Italian energy company with operations in refining and power generation. Saras operates the largest single-site refinery in the Mediterranean basin (300 kbbl/day), utilizing up to 40 different crudes from 150 cargo shipments from a wide range of countries per year. This wide variety in feedstocks presents a significant operational challenge in ensuring product quality and optimizing production.

The fluid catalytic cracker (FCC), a reactive process that converts heavy distillates from crude oil into gasoline and other petroleum products, is a key unit in the refinery that is highly sensitive to feed quality. Ensuring that an FCC model used in production planning optimization accurately captures what the FCC in the plant could produce from available feeds is critical to success. This particular model has a significant impact on operational margins, and its maintenance consumes much time and effort from process engineers, unit technologists, and production planners.

Recently, Saras has leveraged hybrid modeling to enable agile responses to changing operational conditions. The overall workflow employed is shown in Figure 12 and will be explained in the remainder of this section.

4.2.1. FCC Model Calibration. Traditionally, FCC models are calibrated using snapshots of data. For each available data set, kinetic parameters are calculated, typically from a square parameter estimation problem. Then, the parameters are averaged across each data set, and these average values are then used for prediction. This limits the predictive capabilities of the model across a broad operational range.

https://doi.org/10.1021/acs.iecr.5c04353
Ind. Eng. Chem. Res. 2026, 65, 16960−16969

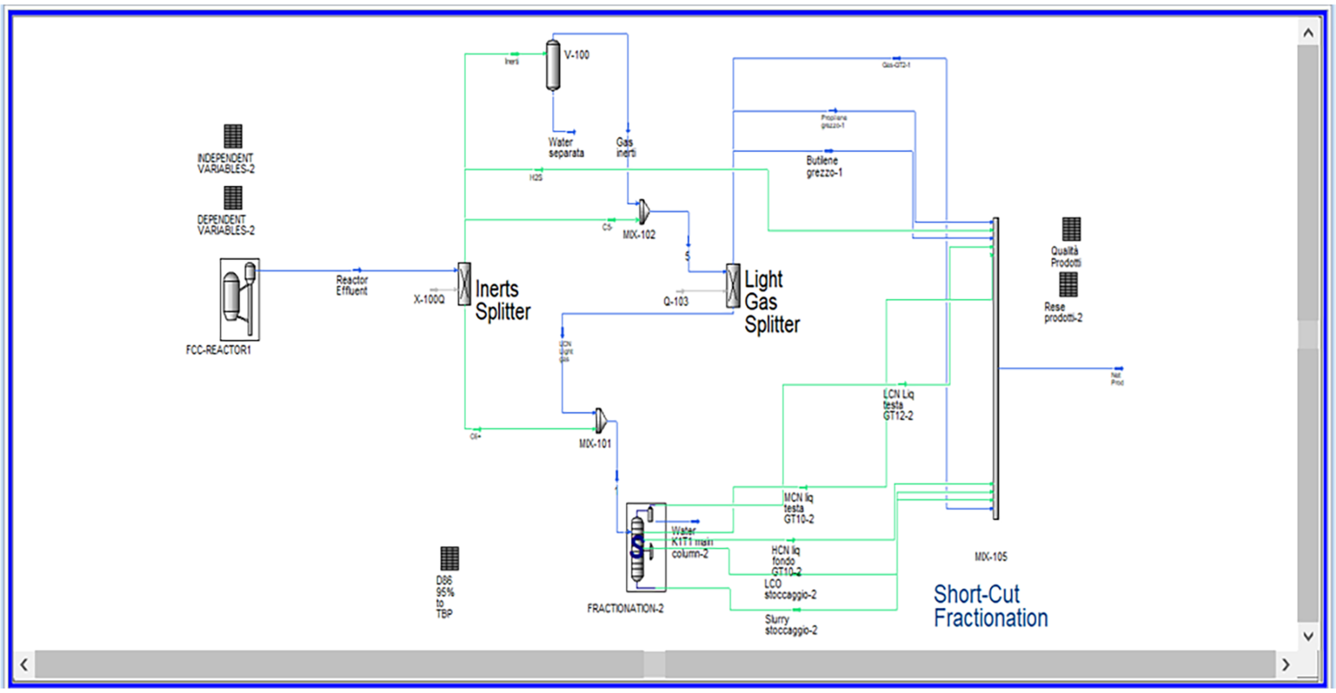

Figure 13. Original Aspen HYSYS flowsheet for FCC and fractionation.

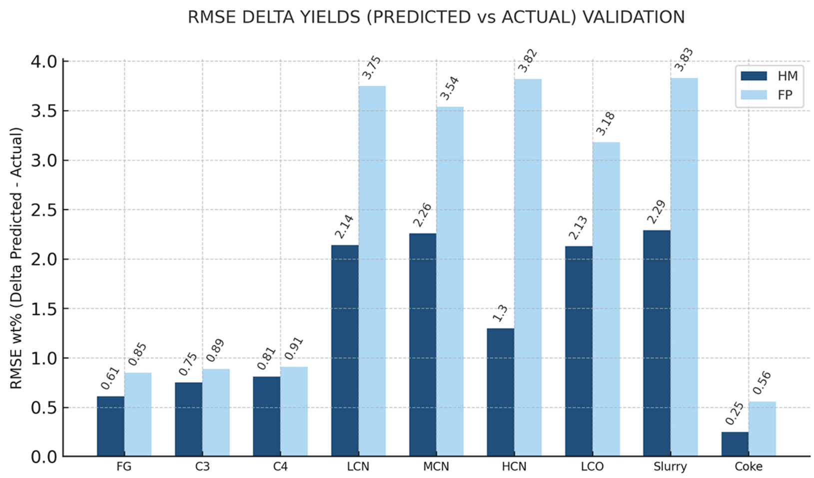

Figure 14. Hybrid modeling improves yield predictions.

The model calibration workflow described in Figure 4 significantly improves this status quo. This allows the creation of a first-principles driven hybrid model that accurately represents the unit across the entire operational range. For this purpose, Saras utilized 151 FCC test run data points collected from June 2019 to November 2023. This highlights a key difficulty in the industry. Quality data points in refining are typically only available from dedicated test runs with lab analysis, which contrasts significantly with the messaging of "big data" that is often repeated. By leveraging all of the first-principles built into the existing Aspen HYSYS FCC model, it is possible to use surprisingly small amounts of data to account for error gaps and create a highly accurate model. As such, Saras was able to create a hybrid model of the architecture shown in Figure 5, using a very small scale neural network, starting from an existing HYSYS model shown in Figure 13. Independent variables include operational conditions such as feed conditions, riser temperature, and catalyst cooler duty. Dependent variables include product yields and properties. Variables predicted by the neural network to be fed to the first-principles model include kinetic parameters. Performance improvements from AI training were significant and will be summarized in the next section. Note that 25 of the above-mentioned data points were held out as a test set to be used for quantifying accuracy on unseen data. Similar to Section 4.1.3, the minimal amount of plant data required is a huge benefit gained from leveraging the first-principles of HYSYS's built in FCC model.

4.2.2. Planning Model Update. Subsequent to creating the calibrated HYSYS model, a ROM is created for utilization in a broader scope production planning optimization problem in PIMS-AO. This was done via the workflow in Figure 3. Aspen Multi-Case was used to generate simulation data from the calibrated model for 1000 sample points. This data is then used in AIMB to fit the dependent variables directly against independent variables across a broad operational range. In this

https://doi.org/10.1021/acs.iecr.5c04353
Ind. Eng. Chem. Res. 2026, 65, 16960−16969

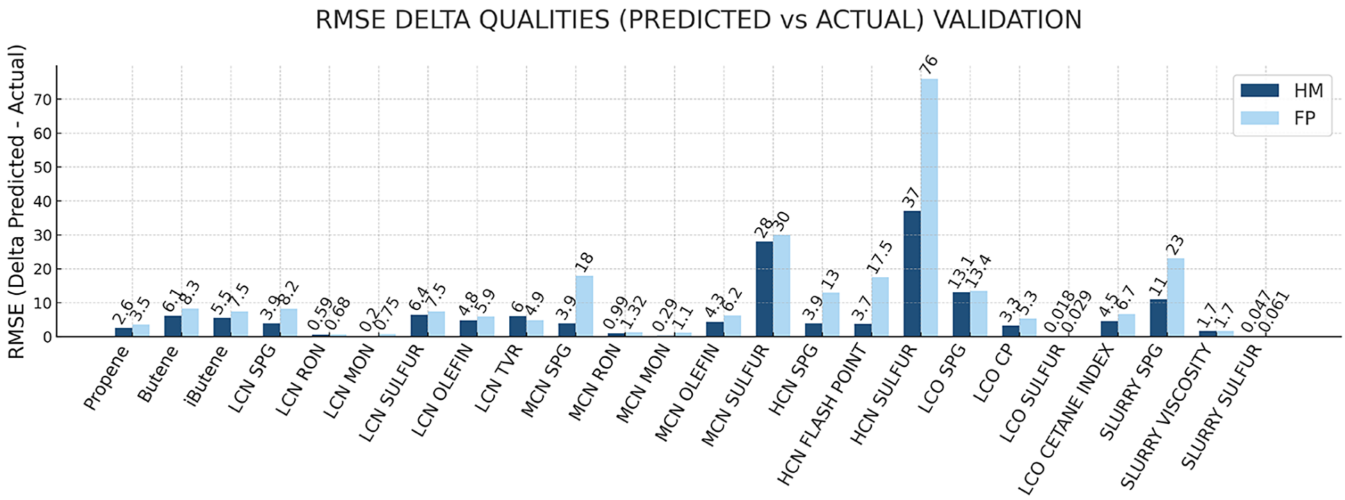

Figure 15. Hybrid modeling improves quality predictions.

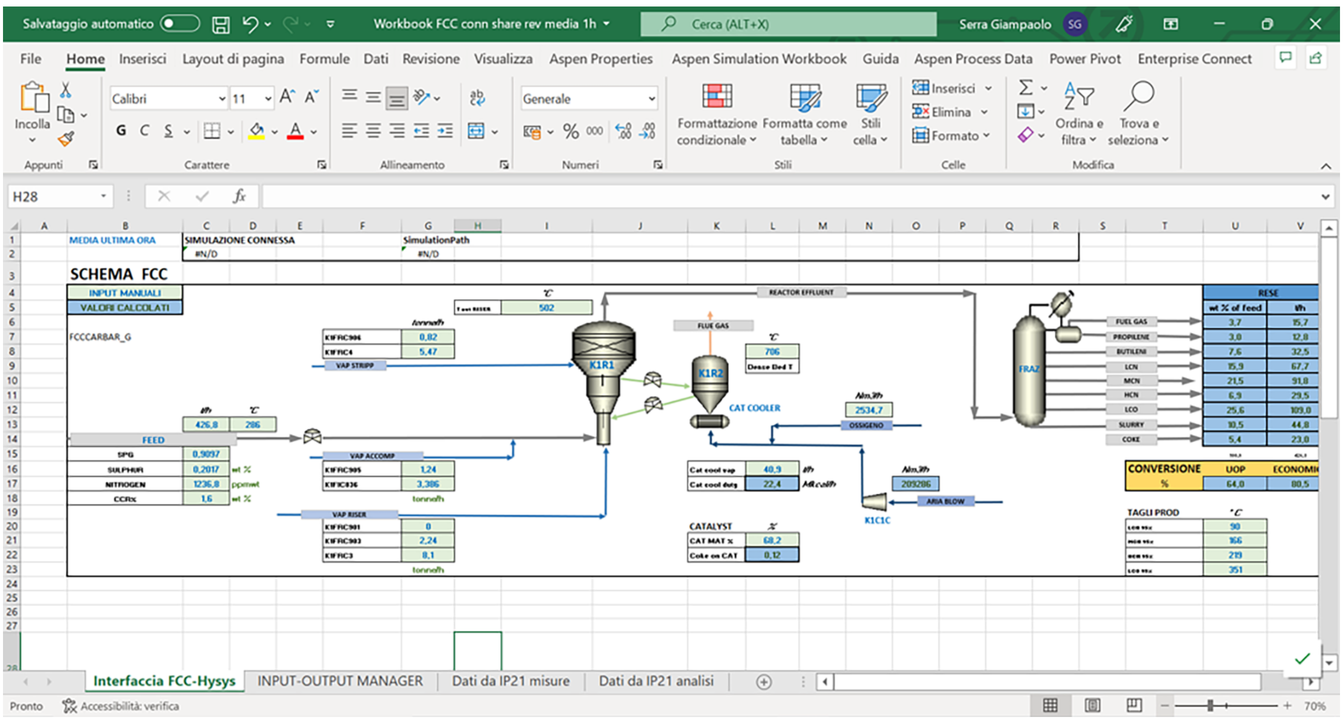

Figure 16. ROM deployed via Aspen simulation workbook.

case, LASSO with nonlinear transformations and constraints was found to be highly effective. The resulting model can be deployed in both HYSYS and PIMS-AO with minimal computational requirements. Accuracy results on the previously mentioned test are shown for yields in Figure 14 and for properties are shown in Figure 15. The plots compare the resulting reduced-order hybrid model directly against the original first-principles model. As can be seen, accuracy improvements are significant. In this case, incorporating first-principles via constraints (such as the mass balance and non-negative flow rates) is critical for adoption by production planners. These guardrails ensure that the resulting optimization solution is consistent with physical reality.

Another goal of this project was to create a process model that can easily be used in operations by HYSYS nonexperts, such as plant operators. This was accomplished with Aspen Simulation Workbook, which provides a convenient interface to the ROM, as shown in Figure 16. Since the ROM is a very robust and computationally simplified model, it can be reliably used for this purpose without convergence issues or other expert-level complications, thereby allowing any plant personnel to easily make use of predictions for operations support.

4.2.3. Summary. Now that the workflow is established, Saras expects a total time of 4 days, end-to-end, to update the planning model when necessary. Although the long-term economic benefits of more accurate planning models have not yet been studied at Saras, the accuracy improvements shown here are very promising. The goals, achievements, and next steps of this work are summarized in Figure 17.

https://doi.org/10.1021/acs.iecr.5c04353
Ind. Eng. Chem. Res. 2026, 65, 16960−16969

Goal	Achievement	Next Steps
Better alignment of HYSYS simulation with actual data.	Yes, accuracy of FCC model has been increased for properties and yields.	Apply the technology to other use cases in the Refinery Digital Twin. Monitor performance over time to define updating frequency.
Embed HM in HYSYS to close the workflow.	Yes, the model was successfully embedded in HYSYS and also in the PIMS environment (workflow completed entirely).	Evaluate performance through parallel testing with the current PIMS submodel. Compare to actual plant performance.
Develop "no simulation" user interface.	Yes, original model, FPDHM & ROM HM embedded in Aspen Simulation Workbook.	Evaluate alternative tools beyond Excel Workbook – Digital Twin.

Figure 17. Summary of hybrid modeling success at Saras.

5. CONCLUSION

In this work, we have shown the reasoning and justification of combining first-principles and data-driven techniques, as well as examples of real-world value with our industrial collaborators. These are not just technological explorations, but rather solutions directly targeted to address core business challenges. This approach applies to many use cases across the industry, including a variety of examples in chemicals, energy, and other fields. Even more use cases will be enabled by continued innovation that has the potential to transform our industry. Most importantly, a collaborative effort across academia and industry is necessary to deliver the education, technology, and industrial toolsets that will enable this transformation. It is clear that the successful AI systems of the future will be those that understand and implement rigorous first-principles in conjunction with human domain experts as direct collaborators.

■ AUTHOR INFORMATION

Corresponding Author

Devin Griffith − *Director of Product Management, Emerson AspenTech,, Houston, Texas 77042, United States;* orcid.org/0009-0008-8619-0971; Email: devin.griffith@emerson.com

Authors

Saud Abdulaziz Al-Ghwainem − *Research Scientist, SABIC Corporate Technology and Innovation, Riyadh 11422, Saudi Arabia*

Giampaolo Serra − *Digital Applications Developer, Saras S.p.A., Sarroch 09018, Italy*

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.iecr.5c04353

Notes

The authors declare no competing financial interest.

■ REFERENCES

(1) Schweidtmann, A. M.; Zhang, D.; von Stosch, M. A review and perspective on hybrid modeling methodologies. *Digital Chem. Eng.* **2024**, *10*, No. 100136, DOI: 10.1016/j.dche.2023.100136.

(2) Sharma, N.; Liu, Y. A. A hybrid science-guided machine learning approach for modeling chemical processes: A review. *AIChE J.* **2022**, *68*, No. e17609, DOI: 10.1002/aic.17609.

(3) Bradley, W.; Kim, J.; Kilwein, Z.; Blakely, L.; Eydenberg, M.; Jalving, J.; Laird, C.; Boukouvala, F. Perspectives on the integration between first-principles and data-driven modeling. *Comput. Chem. Eng.* **2022**, *166*, No. 107898, DOI: 10.1016/j.compchemeng.2022.107898.

(4) Mukherjee, A.; Zavala, V. M. Physics-constrained machine learning for chemical engineering. *Curr. Opin. Chem. Eng.* **2026**, *51*, No. 101228, DOI: 10.1016/j.coche.2026.101228.

(5) Tibshirani, R. Regression Shrinkage and Selection via the Lasso. *J. R. Stat. Soc.: Ser. B* **1996**, *58*, 267−288.

(6) Kaleem, J.; Al Ali, E.; Allanaboyina, V.; Srinivas, R.; Sheikh, S.; Gebesy, E. Utilize a process simulation digital twin to optimize condensate yield. Hydrocarbon Processing https://www.hydrocarbonprocessing.com/magazine/2024/october-2024/digital-technologies/utilize-a-process-simulation-digital-twin-to-optimize-condensate-yield/. Accessed: 05−23−2026.

(7) Emerson and Aramco Deploy AI Solution Aimed at Higher Refinery Yield Volume and Efficiencies. https://www.emerson.com/en/corporate/news/2026/emerson-aramco-deploy-ai-solution-higher-refinery-yield-volume. Accessed: 05−23−2026.

(8) Al Jaberi, F. M.; Al Musharfy, M.; Saha, B.; Taher, A.; Alsuwaidi, A. M.; Al Blooshi, M.; Mantripragada, S.; Thamizha, V. S.; Polonia, R. D.; Periyasamy, D.; Jain, P.; Sheikh, S. AI/ML-based hybrid modeling for enhanced resid fluid catalytic cracking optimization. Hydrocarbon Processing https://www.hydrocarbonprocessing.com/magazine/2025/september-2025/special-focus-refining-technologies/ai-ml-based-hybrid-modeling-for-enhanced-resid-fluid-catalytic-cracking-optimization. Accessed: 05−23−2026.

(9) Gallardo, G.; Pérez-Uresti, S. I.; Varvarezos, D. K. On the use of surrogate models for the production planning of biofuels via hydroprocessed esters and fatty acids (HEFA) process. *Comput. Chem. Eng.* **2026**, *210*, No. 109649, DOI: 10.1016/j.compchemeng.2026.109649.

(10) Hornik, K.; Stinchcombe, M.; White, H. Multilayer feedforward networks are universal approximators. *Neural Networks* **1989**, *2*, 359−366.

(11) Chan, W. K. C.; Fischer, B.; Chen, H.; Bhakta, A. R.; Mobed, P. Combining Machine Learning With Domain Knowledge And First Principles For Modeling In The Process Industries.WO Patent WO2020/227383 A1, 2020.