FOPAM 2026
Proceedings of the 3rd Foundations of Process/Product Analytics and Machine Learning
Decatur, USA, July 26-30 (2026)

From Insight to Action: AI-Powered Decisions in the Chemical Industry

Ivan Castillo

Dow

ABSTRACT

Decision-making in the chemical industry is difficult because operators must choose actions under uncertainty: plant behavior is nonlinear and time-varying, measurements are noisy and incomplete, operating constraints and safety margins are strict, and economic objectives often conflict with each other. Data-driven AI methods can help turn high-frequency process data into predictions and recommendations, but in highstakes settings they are most useful when they leverage prior domain knowledge rather than treating the plant as a black box. In this work, prior knowledge is incorporated in two general ways. First, we use hybrid modeling, where deterministic structure from first principles (e.g., balances, thermodynamics, kinetics) is combined with data-driven learning to improve extrapolation, robustness, and interpretability. Second, we use preference-based learning, where expert judgment is captured qualitatively through pairwise comparisons or rankings of candidate options, enabling optimization even when a single numerical objective is hard to define. We illustrate these ideas with industrial decision-making examples that span physics-grounded monitoring of gradual degradation and preference-driven tuning of operational and control choices. In the hybrid modeling examples, mechanistic knowledge such as mass and energy balances and reaction kinetics is embedded alongside data-driven models to yield physically consistent predictions that remain reliable even when data are sparse or subject to drift. This work focuses on industrial systems that undergo gradual deterioration over time. Three representative case studies are considered: catalyst deactivation monitoring in continuous reactors, where an adiabatic reactor model is used to compute catalyst activity as a health indicator that is integrated into an empirical lifetime model; heat exchanger fouling forecasting in an ethylene oxide plant, where physics-informed features enable month-ahead prediction of fouling surrogates to support cleaning decisions; and remaining useful life prediction for pollutant scrubbers, where hybrid models capture slow degradation trends to inform maintenance planning. Together, these examples show that combining first-principles understanding with datadriven methods enables accurate characterization of degradation dynamics using limited data, improves extrapolation to unseen operating conditions, and increases confidence in maintenance-related decisions (Sansana et al., 2024; Venegas et al., 2024; Bui et al., 2022).