Deep Science Publishing, 2026
https://doi.org/10.70593/978-93-7185-114-5_5

Chapter 5: Machine Learning for Process Optimization

5.1. Introduction

Machine learning has been successfully applied in various domains, where it is utilized to classify messages as spam or not, detect faces in images, or identify individuals from facial images. Despite the broad range of applications that can be tackled using machine learning, the most natural use case for machine learning is in data-rich environments, where massive amounts of data are available for the training and evaluation of models. The process industries have only recently started to benefit from data-driven techniques, as these industries have historically relied on the development of physics-based models for the prediction of process behavior.

Although some process systems are highly controlled and automated, enabling the online acquisition of large amounts of process data, these data often remain unused, with the primary function being archiving. In addition, the Internet of Things (IoT) is enabling process systems to become equipped with more sensors and data-acquisition tools than ever before, increasing the opportunity for the development of data-driven diagnostic, fault-detection, and fault-tolerant mechanisms. With the advancement of affordable computing and storage resources, data are being generated and stored at an unprecedented level, further contributing to the development of data-driven applications. Nevertheless, it remains a challenge for engineers to find the balance between leveraging data effectively while also maintaining a deep physics-based understanding of the underlying process.

5.1.1. Overview of the Study

Attention is being paid to the growing interest in the use of machine learning (ML) techniques to optimize processes in a range of engineering areas. Process optimization is a broad subject that can be understood as the identification of the most appropriate inputs for a process that reduces costs or enhances performance while satisfying constraints. Machine-learning methods are often employed for two different processes: predictive control, in which supervised learning predicts the system outputs from the inputs and controller operation minimizes an objective function; and sequential decision making, which considers a sequential decision-making problem that employs a model-free RL algorithm to optimize system inputs.

In the process industries, control systems for automatic operation have become to rely on feedback control systems that automatically seek optimal operating points for the process. Often, however, the objective function being minimized is unknown. In chemical processing, for example, the aim is often to ensure high quality at low cost while minimizing the environmental footprint. The adoption of IOT technologies on plant equipment and processes is now enabling continuous data acquisition from operational data, which can be employed to develop machine-learning models for predictive control that seek to improve performance by data-based operation or optimization of the control targets.

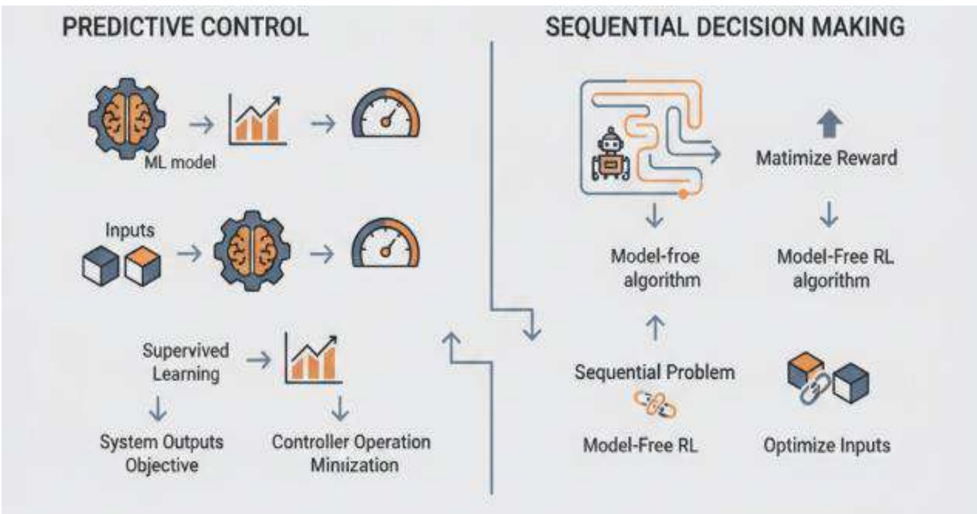

Fig 5.1: Data-Driven Synthesis: Integrating Predictive Control and Reinforcement Learning for Autonomous Industrial Process Optimization

5.2. Fundamentals of Process Optimization

Process optimization applies to systems modeled by either an objective function over a domain or a cost function that balances reward and penalty. An operation, experiment, or simulation is performed at every point in the domain or at every decision point in a sequence, and an outcome is recorded. Data decomposition, feature selection, and dimensionality reduction are machine learning methods that improve visualization, acquisition, and model performance.

An optimization problem may be stated in terms of either an objective function with constraints or a cost function. Process optimization typically involves multiple experiments under varying conditions to identify the set of values for the control variables that would provide the best outcome for the industry. The outcome of interest could be profit maximization or cost minimization. In practice, however, the objective function is usually not available, and it must be approximated using machine learning techniques. The approximation is seeded and validated using the process evaluation metric.

Machine learning can also be useful in reducing the data collection burden. Cost function or reinforcement-learning approaches are applicable when the optimization target is control sequence that maximizes cumulative reward or minimizes cumulative penalties during operation.

5.2.1. Objective Functions and Constraints

Process optimization combines data analytics with machine learning to achieve optimal performance. The application of these techniques is not limited to a single industry. The following sections discuss machine learning for process optimization in the context of control systems and industrial processes, describing the necessary assumptions and integration steps. Special emphasis is placed on the model development lifecycle and the underlying principles of chemical process and semiconductor manufacturing optimization.

Machine learning techniques for process optimization build upon the mathematical formulation of the optimization problem. Broadly speaking, optimization aims to find the input settings that minimize (or maximize) the values of the objective function while satisfying the constraints. In process systems, the input is typically a vector of control variables that define the operating conditions for the system or environment, such as temperature, pressure, and gas-phase composition. The objective function measures performance and is usually non-increasing in term of economic cost or energy consumption, but can also be an indirect measure such as environmental impact.

5.2.2. Evaluation Metrics

Machine learning methods learn behavior from data and consequently require sufficient historical data for learning. Consoli and Benassi (2021) reviewed research papers applying ML methods for process optimization in the process industries, which include research works from chemical, food, oil and gas, pulp and paper, utility, and semiconductor manufacturing. For a majority of the articles, process optimization is defined as the minimization or maximization of an objective function with respect to constraints. Research papers can be grouped according to the focus of optimization: predictive control, support for sequential decision making, or development of self-optimizing control schemes. A selection of articles representative of each group is summarized.

Process Control System

Machine learning methods can be integrated into process control systems in two major ways. First, common applications in industrial automation deal with predictive behavior under normal operating conditions. The process under control is represented mathematically, either using an explicit function or an implicit function inferred by statistical learning methods, and the trained model is used for predictive control. Classical optimization tools can then be exploited, namely, dynamic programming and non-linear model predictive control. Use of predictive control methods that use trained models is understood as supervised learning for optimization.

Second, processes in the chemical and semiconductor industries encounter system disturbances leading to deviations from normal operating conditions. Reinforcement learning is a driven approach, aiming to derive control actions from sensor measurements that optimize a long-term performance index. The performance of the controller can be quantified using performance metrics such as cumulative regret and cumulative cost.

5.3. Machine Learning Methods for Optimization

Supervised Learning for Predictive Control

Two types of control problem can be addressed by machine learning methods, specifically supervised learning and reinforcement learning. Supervised learning is applied, most commonly, to predictive control of chemical reactors. The purpose of the machine learning model is to predict an outcome or property as a function of several input variables, thus mimicking an input-output process model. The primary reason for developing such a predictive model is that direct control of the system is often very expensive or difficult to implement within a process control system. The true input-output model is not known, hence neural networks or other forms of supervised learning are trained on experimental or simulation data available from the true process model.

A simple example is a batch reactor in which the concentration of a minority species is controlled in the product stream. Direct control of the flow of this minority component is difficult because it travails through a limited volume of the reactor; hence, only small adjustments in the flow are possible during reaction. The concentration of this minor species is predicted from the major inputs by employing an artificial neural network (ANN). The ANN is trained on data acquired from the simulator of the real reactor process, which is a non-linear dynamic input-output continuous-time state space model. The impact of various input factors on the concentration of the minor species can be studied using the ANN, which is fast, simple to build, and easily retrainable when required.

Reinforcement Learning for Sequential Decision Making

Reinforcement learning possesses the ability to make sequential choices over time, resulting in policies intended to maximize cumulative reward or minimize cumulative cost. It operates in two distinct learning phases: exploration and exploitation. The first is related to the collection of data through a trial-and-error strategy, while the second is responsible for choosing the best decision based on the experience gained. It is applicable in situations where it is infeasible to rely solely on historical data or where the state space is challenging to model a transfer function. In reinforcement learning, a trial-and-error interactive control-learning mechanism guides an agent making sequential decisions. A reward is received after each action, indicating the immediate feedback of the actions taken. The agent's objective is to maximize the total rewards it receives. The evaluation is based on the reward structure, but there are no predefined optimal actions. Instead, the agent explores the possible actions, searching for the one that returns the maximum reward over time.

5.3.1. Supervised Learning for Predictive Control

Supervised machine-learning techniques have been successfully integrated into industrial optimization strategies by formulating optimization problems as predictive control tasks over an appropriately defined objective function. Predictive control methods maintain a process model built from data on the response surface that links an operating parameter or a set of parameters to the optimization target. The response surface is explored either in closed loop or in open loop using a small number of objective function evaluations.

The predictive control strategy can be applied to various problems in the process industries—minimizing emissions and maximizing yields for chemical plants, thermal efficiency for furnaces, or production rates for polymer processes. The optimization objective and the constraints are expressed mathematically, and the data required for training the predictive model are either available within the system or acquired through relatively inexpensive experiments.

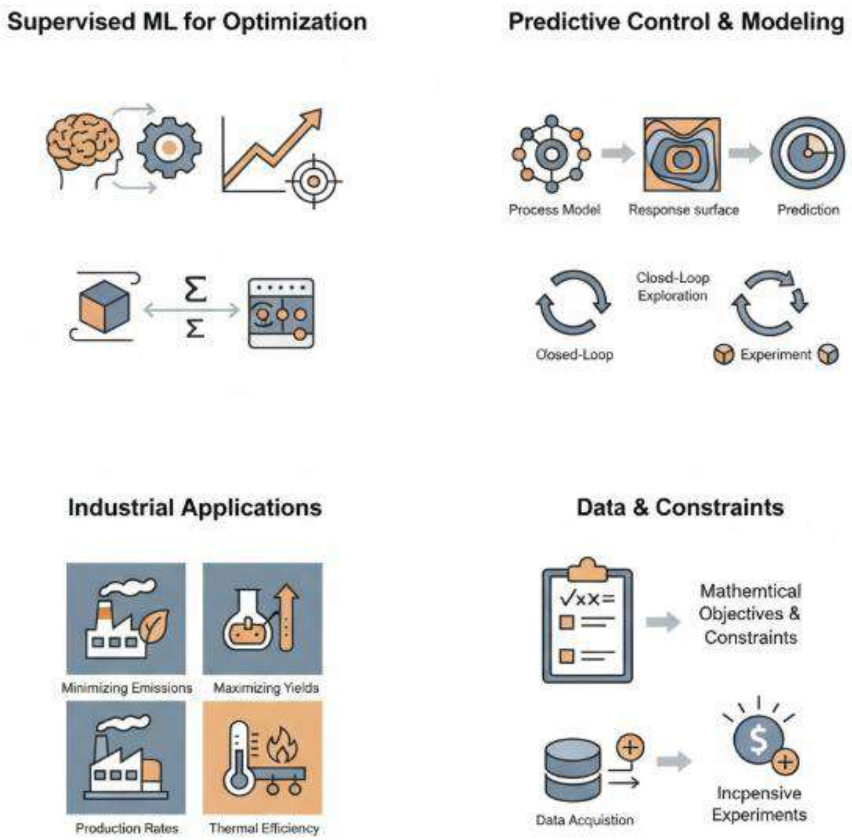

Fig 5.2: Data-Driven Predictive Control: Leveraging Supervised Machine Learning for Multiobjective Optimization in Process Industries

5.3.2. Reinforcement Learning for Sequential Decision Making

Reinforcement learning (RL) approaches sequential decision making as a problem of exploration and exploitation, modeled as a Markov Decision Process (MDP). An agent interacts with the environment by taking actions defined by a policy, a mapping from states to actions. Due to rewards being delayed, RL is often applied to complex tasks that require integrating knowledge over time, supporting skills that can yield long-term highest rewards. Experience sampling from the interaction data is vital to update the policy, choosing actions not only to maximize immediate reward but also to improve data selection towards efficient learning in this high-dimensional space.

RL has drawn attention for applications in process industries, especially when applied to problems forming an MDP. In these, the state is described by the system dynamics, such as in chemical reactor applications where RL has been used for active fault diagnosis or controller setup in batch operations. In semiconductor manufacturing, RL has been used to learn the control policy for cluster tools, modeling the tool as an MDP and learning the optimal cycle-time through more complex control policy architectures. In the latter, due to cycle-time-dependent operation of modules, the control policy switches between high-performing mode of operation and low-performing control, with little to none impact on scheduling control.

5.4. Integration with Process Industries

The connection with process industries is twofold. First, a distinct class of problems exists that directly involves the optimization of process systems. Second, the availability of physical equipment, along with advancements in the Internet of Things (IoT), provides an unprecedented opportunity for developing ML models for practical industrial applications. The present section focuses on these aspects of the connection with process industries.

Process industries encompass the plants in which chemical substances undergo chemical changes. These include the petroleum refining, petrochemical, chemical, and metal-producing industries. The production and manufacturing processes are handling-oriented; raw materials are moved from one place to another through equipment such as reactors, boilers, distillation columns, heat exchangers, evaporators, and so on. In these industries, digital automation systems are widely used, usually following the client-server architecture. Control computers gather real-time data on plant parameters from the field through controllers and sensors, and they send instructions to actuators to make the necessary adjustments, controlling the operation of plant units. Digital control systems provide crucial benefits in terms of productivity, safety, and environmental protection.

5.4.1. Control Systems and Industrial Automation

Manufacturing processes are frequently regulated by a distributed control system consisting of hundreds or thousands of linear feedback controllers and optimizer modules. Controller tuning may be performed manually or automatically but is not always optimal. In the absence of feedback loop interactions, the control structure may also be optimized in a coordinated fashion. However, qualitative modeling approaches such as dynamic matrix control and internal model control are more commonly used in control system design. Even with relatively complete first-principle models, employing derivatively active predictions based on these models often leads to sustained oscillations. In addition, control systems equipped with significant feedforward elements may require active damping strategies to improve process stability following disturbances.

Process plants also incorporate subsystem automation, such as supervisory control and data acquisition systems for sequential and batch control and alarm management systems and operator supervision for more complex operational modes. The incorporation of machine learning techniques into alarm systems, for example, has been explored. A data-based approach for generating operating procedures has also been proposed.

5.4.2. Internet of Things and Data Acquisition

The integration of process optimization and machine learning also requires the availability of historical datasets for model development. Quality data can be acquired in a cost-effective manner via the Internet of Things (IoT), which enables cloud-based connectivity for various types of embedded systems. In the context of process optimization, sensors embedded in the control systems gather process data and relay it via the Internet to data centers. These data repositories are used to store both structured and unstructured data originating from disparate sources.

Structured data encompasses the information stored in traditional relational databases. Unstructured sources, on the other hand, may consist of open-text documents presenting operational experiences or predictions about future developments in technologies and chemical production processes. Data mining techniques are gradually being developed to extract knowledge from these sources, which can be embedded in both the objective function and the constraints of process optimization models.

5.5. Model Development Lifecycle

A model development lifecycle provides a systematic approach for applying machine learning algorithms to process optimization. This lifecycle is general and applies to supervised learning for predictive control, reinforcement learning for sequential decision making, or any other branch of machine learning. With specialist expertise in the area of knowledge engineering, these stages can often be executed in sequence by a single team or group using available tools with relatively low costs. However, for areas with less mature combinations of engineering and data science knowledge, specialist teams may need to cooperate closely.

A model development lifecycle, adapted for process optimization, comprises two main stages. The first stage involves problem formulation and data acquisition, where the knowledge engineering community identifies an optimization objective and relevant factors affecting model performance; the second stage focuses on model training, validation, and testing, where the data science community applies machine learning to extract predictive or prescriptive patterns from the gathered data. A more detailed explanation of these stages is outlined below.

5.5.1. Problem Formulation and Data Acquisition

The first and most important step in any machine learning model development is problem formulation, which needs careful consideration. Two central questions are: (i) What problem is being solved? and (ii) Are there sufficient data available to solve it? In the context of process optimization, the problem is usually defined very formally and rigorously. The considerations for data acquisition are, however, process-dependent and they become crucial at this stage of development.

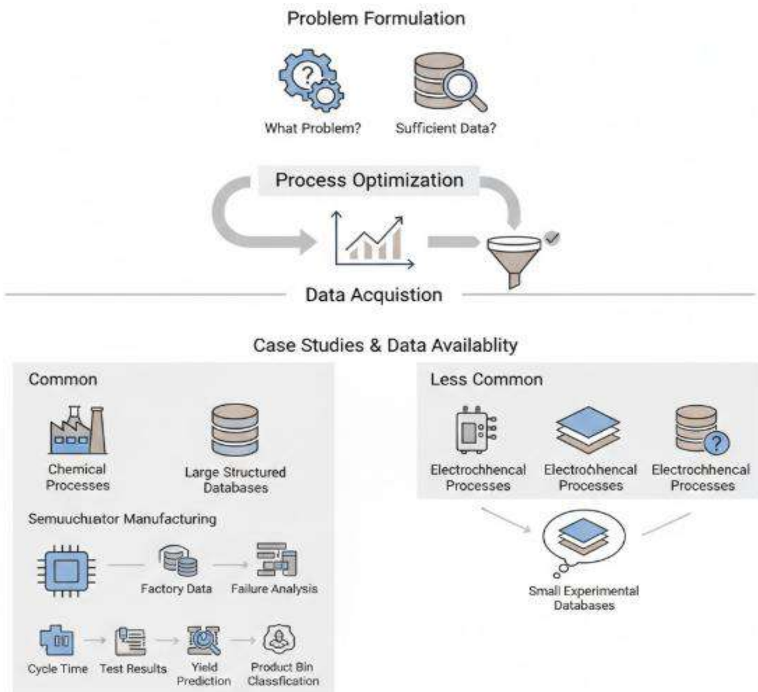

Fig 5.3: Formulating Data-Driven Success: A Comparative Analysis of Problem Definition and Data Acquisition Across Process Industries

Among process industries, chemical processes appear to be the most popular. Chemical process data is routinely stored in large, structured databases. Semiconductor device manufacturing is another area where industrial automation systems have been widely deployed, and massive production and test data accumulated. Successful semiconductor industry case studies show that the combination of data from factory control systems, product test results and product failure analysis can be exploited simultaneously using machine learning techniques for different applications like cycle time prediction, yield prediction, product bin classification etc. Less common are the applications of data-driven models in electrochemical processes such as polymeric membrane fuel cells, and galvanoplastics (the deposition of metals and other materials) on deep skin for micro-electromechanical systems. These models are based on small experimental databases, and a literature survey shows that data-driven models for process optimization are still scarce.

5.5.2. Model Training, Validation, and Testing

Training a machine learning model consists of estimating its unknown parameters using a training data set to minimize generalization error. Generalization error is the difference between a model's performance on training data and the performance on new, unseen data. Therefore, a good strategy for training is to use as much data for training as possible, while also reserving separate data different from the training data set for validating the model at step 4. The simplest method for doing this is to split the entire data set randomly into a training and validation set with an 80:20 ratio. For some applications, stratified sampling should be performed for categorical response variables. For complicated or expensive processes where data acquisition is challenging or cost prohibitive, this strategy may not be appropriate. In such cases, k-fold cross validation can be performed, where the data set is split into k subsets. For each subset, the model is trained on the other (k–1) subsets and validated on the subset kept aside. The final performance of the model on the validation set is then the average of the k separate validation performances. Finally, if desired, a model can be trained on the entire data set and, if sufficient data exists, tested on a different data set that has not been used for training.

During training, hyperparameters often need to be tuned to achieve good performance. Hyperparameters are parameters not learned directly by the model during training, and their values need to be determined a priori and are often domain specific. Two methods often used for tuning hyperparameters are grid search and random search. Both strategies can be computationally expensive because they require the training of multiple models. Grid search is an exhaustive approach where each hyperparameter in a grid of hyperparameters is assigned all possible values. The combination that produces the best model on the validation set is retained. Random search samples the value of each hyperparameter from a set specified by the user; it does not guarantee the best solution but is often more efficient than grid search. The. performance of the model is validated with a validation set different from the training set.

5.6. Case Studies

This section reviews recent and significant examples of machine learning applied to process optimization in process industries, namely, semiconductors and chemical production.

Processes for manufacturing semiconductors have undergone constant advances in both miniaturization and manufacturing efficiency. These developments require not only appropriate production procedures but also an effective production management system and effective device control. Many complex parameters—such as material properties of glass masks, etching conditions, and residual layer thickness—affect the production yield at different stages, and the optimal values differ from step to step. The relationship between these parameters and the yield function is extremely complicated and difficult; however, it is essential to grasp the viewpoint of yield at every step. Therefore, machine learning, which can handle complicated functions without establishing a concrete mathematical model, is applied to the problem of guiding step-by-step production to yield improvement in semiconductor manufacturing by considering optimally matching parameters of the entire sequence of processes, thus avoiding consideration of intermediate near-optimum states.

A chemical plant operates an organic compound separation process. The optimum rotation speed of the distillation column reboiler has significant effects on the production yield and energy efficiency. A senior operator usually makes a decision by experiencing a series of production stations, real-time observing the separation qualities and communicating with other senior operators. A decision support system is proposed to realize the senior operators' heuristics by developing and applying reinforcement learning, automating the production quality inspection and establishing a multi-agent communication platform for the rotation speed adjustment. The service request is opened and handled by multi-agent systems instead of a manual operator, and different rotation speeds are proposed to production station supervisors for checking. By receiving the supervisor feedback, the reward function is refined and independent agents for different textbooks are trained. The results indicate that the decision support system effectively simulates the heuristics of senior operators.

5.6.1. Chemical Process Optimization

Chemical processes involve the conversion of raw materials into refined products or energy carriers. Chemical transformation can be used to produce fuels, solvents, and other chemicals, but without optimization, chemical production generally incurs high energy, raw material, and waste treatment costs, which leads to lower profit margins. On the equipment level, processes such as oil refining can have more than 1,000 units of operation, and the optimizer considers the physics of these interactions when concurrently adjusting many decision variables, both continuous and discrete. The functionalities of chemical processes are further complicated by environmental and economic issues caused by energy consumption. Reducing energy consumption, especially CO2 emissions, maintenance losses, explosion, and waste treatment, while ensuring product quality and system reliability, has become even more critical in chemical process plants.

With the growing availability of real-time data streams from the process internet of things (PIoT), smart chemical process optimization solutions can be developed through the synergy of advanced optimization and machine learning (ML). Unlike traditional optimization frameworks that make some baseline assumptions, effective methodologies with abundant sensor data are sought. Recently, reinforcement learning (RL) has been employed in chemical process optimization problems, considering a new dynamic decision-making framework where the process continuously works and optimization decisions are made in real-time. Instead of merely following a trajectory created by a process model, the RL agent actively interacts with the process real-time, providing a fundamentally different way to optimize state-of-the-art chemical processes.

5.6.2. Semiconductor Manufacturing

In the semiconductor industry, mask designs must be optimized to continue producing chips at ever-smaller scales. Optical proximity correction (OPC) combines geometric edits, such as adding serifs to line-ends, with exposure- and illumination-specific corrections. The corrections must be carefully managed to ensure the design can be accurately manufactured and does not degrade chip performance. Due to a lack of explicit gradient information and the need for complex handling, OPC designers often use guides that rely on intuition. For four-layer designs at standard nodes, machine learning has been applied to learn correct edits from characterizations of chip-quality-critical polygons.

Deep neural networks successfully learned to classify patterns into widely-used categories for which known geometric changes are reliable. Using the labeled samples, a reinforcement learning approach with a simple specialized oracle was adopted. The resulting system offered a more consistent and faster alternative for optimally OPCing patterns. Similar approaches have also shown promise for solving other semiconductor manufacturing problems, including control design for the chemical-mechanical polishing of substrate surfaces and predictive process monitoring for wafer faults.

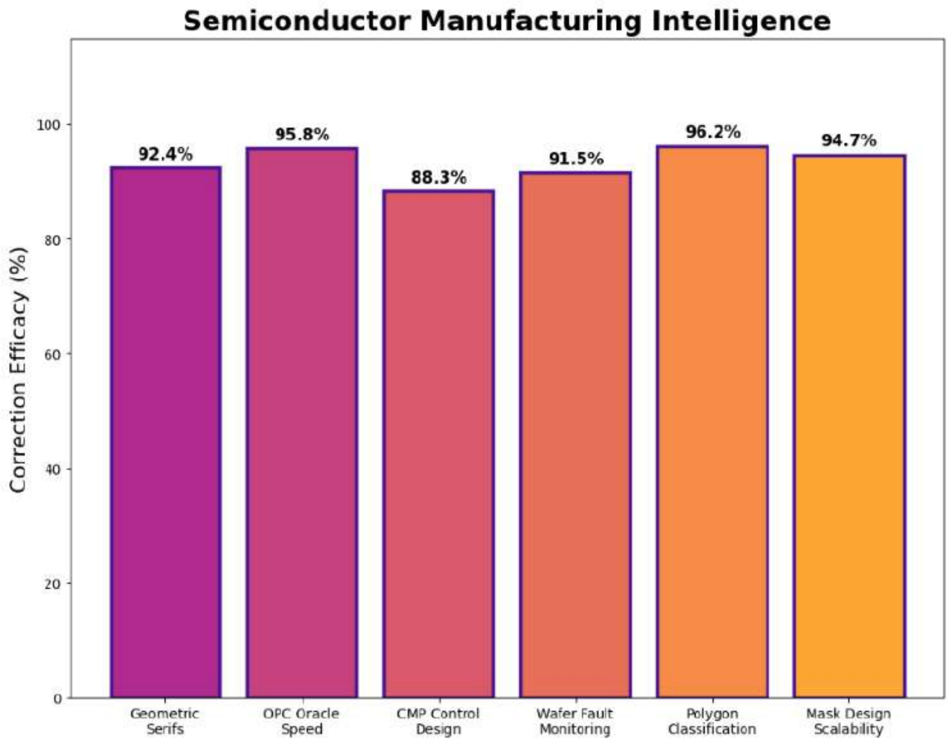

Fig 5.4: Semiconductor Manufacturing Intelligence

Real-time graphical performance-adaptive editing is also being integrated into game engines via an option trade-off between performance and quality, all while maintaining a rich user experience. Developing high-performance gaming-application CMU computation systems based on high-efficiency focus-deformable triangulation mesh technology and applying flexible user-gnome Sketch-Up models for performance-adaptive editing remain topics for further exploration.

5.7. Conclusion

Machine learning has the potential to make a major impact in process industries through optimization-oriented applications. Advances in control systems and industrial automation make it feasible to deploy models learned from data in real-time applications. The Internet of Things facilitates acquisition of the vast amounts of data typically required to train machine learning models. While academic research in the area is still relatively nascent, numerous pertinent applications have already been reported.

Two broad categories of machine learning applications focused on process optimization have emerged in the literature. The first encompasses supervised learning applications that assist or automate the operation of a process during its dynamic response to a change in set points. Reinforcement learning applications, on the other hand, seek to determine sequential control policies that deliver optimal performance over a specified operating horizon. A pragmatic understanding of how machine learning can contribute to process optimization focuses on specific aspects of the model development lifecycle: problem formulation and data acquisition, model training, validation, and testing, and integration into existing systems. Careful attention to these considerations ensures that models are not only accurate but also generalizable, interpretable, and able to synthesize sample-efficient control policies.

5.7.1. Final Thoughts and Future Directions

Despite its infancy, small applications have been demonstrated successfully in a variety of process sectors. A large proportion of research is devoted to chemical processes, especially predictive control methods based on supervised learning. Other areas of application include semiconductor and wood product manufacturing. The fundamental structure of these applications is largely similar, taking advantage of an integrated control system, an Internet of Things structure for data collection, and an often substantial body of historical data. Interest in these methods is therefore expected to continue growing.

However, it seems likely that the economic case for wider adoption will limit the range of applications. The techniques are embedded in specific novel products that can be applied to an entire range of similar process situations, with the resulting appeal that bespoke model development need not be undertaken for every promise. Nonetheless, their approach to truly sustainable production ultimately aims to build methods that can handle any number of requirements in an effective manner. It remains to be seen how close the result is to generalizing prediction-cum-control systems beyond the specfic area of interest to encompass all manufacturing process industries.

References

Radanliev, P., De Roure, D., Walton, R., & Van Kleek, M. (2021). AI systems safety and cybersecurity. Computers & Security, 102, 102192.

Guntupalli, R. (2025, August). Cloud-Native AI: Challenges and Opportunities in Infrastructure Security. In 2025 International Conference on Artificial Intelligence and Machine Vision (AIMV) (pp. 1-4). IEEE.

Kief, M. G., & Bick, G. (2021). Digital transformation in financial services. Springer.

AI Powered Fraud Detection Systems: Enhancing Risk Assessment in the Insurance Sector. (2023). American Journal of Analytics and Artificial Intelligence (ajaai) With ISSN 3067-283X, 1(1). https://ajaai.com/index.php/ajaai/article/view/14

Eling, M., Nuessle, D., & Staubli, J. (2022). Artificial intelligence along the insurance value chain. Journal of Risk and Insurance, 89(2), 1–38.

Gottimukkala, V. R. R. (2023). Privacy-Preserving Machine Learning Models for Transaction Monitoring in Global Banking Networks. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 633-652.
Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers. International Journal of Forecasting, 37(4), 1748–1764.
Varri, D. B. S. (2023). Advanced Threat Intelligence Modeling for Proactive Cyber Defense Systems. Available at SSRN 5774926.
Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. Nature, 518, 529–533.
Garapati, R. S. (2022). Web-Centric Cloud Framework for Real-Time Monitoring and Risk Prediction in Clinical Trials Using Machine Learning. Current Research in Public Health, 2, 1346.
Polamarasetti, S., Kakarala, M. R. K., kumar Prajapati, S., Butani, J. B., & Rongali, S. K. (2025, May). Exploring Advanced API Strategies with MuleSoft for Seamless Salesforce Integration in Multi-Cloud Environments. In 2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC) (pp. 1-9). IEEE.
Hafner, D., Lillicrap, T., Norouzi, M., & Ba, J. (2021). Mastering atari with discrete world models. Advances in Neural Information Processing Systems, 34, 1–14.
Nagubandi, A. R. (2024). Breakthrough Real-Time AI-Driven Regulatory Intelligence for Multi-Counterparty Derivatives and Collateral Platforms: Autonomous Compliance for IFRS, EMIR, NAIC, SOX & Emerging Regulations. Journal of Information Systems Engineering and Management, 9.
Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. International Journal of Robotics Research, 32(11), 1238–1274.
Davuluri, P. N. (2020). Improving Data Quality and Lineage in Regulated Financial Data Platforms. Finance and Economics, 1(1), 1-14.
Bostrom, N. (2014). Superintelligence: Paths, dangers, strategies. Oxford University Press.
Emerging Role of Agentic AI in Designing Autonomous Data Products for Retirement and Group Insurance Platforms. (2025). MSW Management Journal, 34(2), 1464-1474.
Amodei, D., Olah, C., Steinhardt, J., et al. (2016). Concrete problems in AI safety. arXiv.
Floridi, L., Cowls, J., Beltrametti, M., et al. (2018). AI4People—An ethical framework for a good AI society. Minds and Machines, 28(4), 689–707.
Bommasani, R., Hudson, D. A., Adeli, E., et al. (2022). On the opportunities and risks of foundation models. Stanford Institute for Human-Centered Artificial Intelligence.
Kushvanth Chowdary Nagabhyru. (2023). Accelerating Digital Transformation with AI Driven Data Engineering: Industry Case Studies from Cloud and IoT Domains. Educational Administration: Theory and Practice, 29(4), 5898–5910. https://doi.org/10.53555/kuey.v29i4.10932
LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521, 436–444.
Yandamuri, U. S. (2021). A Comparative Study of Traditional Reporting Systems versus Real-Time Analytics Dashboards in Enterprise Operations. Universal Journal of Business and Management.
Chen, Y., & Zhang, L. (2022). Data engineering practices for real-time analytics. IEEE Transactions on Services Computing, 15(4), 2288–2302.
Ortega, P. A., & Stocker, A. A. (2016). Human decision-making under uncertainty. Current Opinion in Behavioral Sciences, 5, 100–107.
Amistapuram, K. (2024). Generative AI in Insurance: Automating Claims Documentation and Customer Communication. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 461–475. https://doi.org/10.61841/turcomat.v15i3.15474
Radanliev, P., De Roure, D., Walton, R., & Van Kleek, M. (2021). AI systems safety and cybersecurity: A systematic mapping study. Computers & Security, 102, 102192.

Gama, J., Žliobaitė, I., Bifet, A., et al. (2020). Concept drift adaptation. ACM Computing Surveys, 46(4), 44.