Optimization of Process Families for Deployment of Carbon Capture Processes using Machine Learning Surrogates

Georgia Stinchfield[a], Bashar L. Ammari[a], Joshua C. Morgan[c,d], John D. Siirola[b], Miguel Zamarripa[c,d], Carl D. Laird[a*]

[a]*Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA, 15232*
[b]*Sandia National Laboratories, P.O. Box 5800, Albuquerque, NM, 87185*
[c]*National Energy Technology Laboratory (NETL), 626 Cochran Mill Road, Pittsburgh, PA, 15236*
[d]*NETL Support Contractor, 626 Cochran Mill Road, Pittsburgh, PA, 15236*
**claird@andrew.cmu.edu*

Abstract

Traditional process design approaches focus on exploiting economies of scale but are inefficient when designing a large number of similar processes for decentralized applications. By optimizing each process individually, these approaches do not allow for manufacturing standardization or make use of economies of numbers. In this work, we design a family of processes (i.e., multiple processes) of a carbon capture facility with various performance requirements. We identify sub-components within the process and create a small set of sub-component designs that can be shared across all the processes in the process family. We formulate this optimization problem as a nonlinear generalized disjunctive program (GDP) and, in previous work, developed two approaches for reformulating and solving this problem: one based on full-discretization of the design space (Zhang et al., 2022) and one that used Machine Learning (ML) surrogates to replace the nonlinear process models (Stinchfield et al., 2022). Using ML surrogates to predict required system costs and performance allows us to reformulate the nonlinearities in the GDP to generate an efficient MILP formulation. In this work, we apply the ML surrogate approach to design a family of carbon capture systems to cover a set of different flue gas flow rates and inlet CO_2 concentrations, where we consider the absorber and stripper as sub-component types.

Keywords: Manufacturing, Optimization, Energy Systems, Machine Learning.

1. Introduction

Effectively mitigating climate change requires the broad deployment of green energy and industrial decarbonization processes. When designing *multiple* processes, a conventional design approach produces a set of unique unit designs for each process that can be expensive to manufacture. Modularity, a well-explored design approach, provides significant reductions in manufacturing costs and has proven beneficial in particular when designing multiple instances of a process (Baldea et al., 2017). However, modular designs are typically 'numbered-up' to achieve capacity requirements and often fail to fully exploit economies of scale. Rather than using traditional or modular design, we propose a hybrid approach that applies concepts from product family design to the design of process systems to gain the manufacturing benefits of modularity while still obtaining process customization and economies of scale.

In product family design (PFD), a product family is broadly defined as "a set of products that share one or more common 'element(s)' yet target a variety of different market segments" (Simpson et al., 2014). This approach is used extensively in automotive industries, for example, where a wide variety of products are available, but designed and assembled using a smaller, commonly designed platform of sub-components. As process vendors tackle the broad deployment of large numbers of processes for industrial decarbonization, they face similar design goals. We are extending the ideas of PFD into this area of process design. In our work, the "products" are the set of processes we must design. Each individual process is referred to as a process variant; each variant is associated with a set of conditions that the process design must meet. The "common 'element(s)'" are sub-components shared among the process variants. For our case study, we consider the design of a carbon capture system (CCS). Typically, a CCS is designed for a specific point source capture location (i.e., coal-fired power plants, natural gas combined cycle, cement, refineries, etc.) mainly based on the flue gas conditions and flowrate. PFD has the potential to accelerate the deployment of CCS for industrial processes through the simultaneous design of process families that share common sub-components. Our case study considers 12 process variants. Each process variant in the family is designed for a particular flue gas flow rate and CO_2 concentration. We wish to design each of these process variants using a common platform that includes shared sub-component types, in this case, the absorber and the stripper. We require one of each sub-component for each process variant (i.e., 12 absorbers and 12 strippers in total) but hope to meet all the process variant requirements with a small number of sub-component designs (e.g., here we consider a platform with only two absorber designs and two stripper designs from which to design all 12 process variants).

This approach captures the benefits of manufacturing standardization, akin to modularity, since a large number of process variants share the same common sub-component designs. At the same time, this retains economies of scale as the platform of manufactured sub-components is designed simultaneously with the set of process variants in the family. This approach has proven beneficial in other industries; companies such as Nissan, Toyota, and Boeing have reported substantial cost and time savings by using this approach (Simpson et al., 2014). However, despite documentation of significant success in other industrial manufacturing settings, PFD has largely not been applied to chemical process design. Furthermore, reported approaches have largely been heuristic, and there is a need for rigorous optimization approaches.

Our optimization problem simultaneously determines the individual *designs* of the sub-components in the platform (i.e., if we only offer two absorber designs, what should they be) and *which design* each process variant should use (i.e., for each of 12 process variants, which of the two absorber designs should each variant use). We first formulated this problem as a nonlinear GDP. We solved this problem by reformulating it as a MILP formulation based on full discretization of the design space (Zhang et al., 2022). We extended this approach, proposing another MILP reformulation that avoids discretization by utilizing ReLU activated Neural Network (NN) as piecewise linear surrogate models (Stinchfield et al., 2022, Ceccon et al., 2022). This significantly decreased pre-computation requirements for simulations and allowed us to search for an approximation of the continuous design space. In this paper, we further demonstrate this piecewise linear surrogate approach to optimize a process family of carbon capture systems, using both an ReLU activated NN and a linear model decision tree as surrogates. We design 12 carbon

capture systems with different flue gas flow rates and CO_2 concentrations with the absorber and stripper columns as the shared sub-components.

2. Optimization Formulation

The following formulation for optimization of a process family design was first proposed by Stinchfield, et al., (2022). A process variant is represented by v, and is characterized by the design requirements for that variant. The parameter set V represents all process variants. The boundary conditions of a process variant are parameterized in b_v. In our case study, $|V| = 12$ and the boundary conditions for each variant include the flue gas flow rate and CO_2 concentration. The total annualized cost of a process variant is represented by variable c_v. The objection function includes a weight for each process variant, w_v, that captures the expected sales or number of installations of each variant. The variable p_v captures the performance metrics of a particular process variant. In many cases, this will be an indicator function that determines whether a particular design is feasible for the variant. This ensures that an infeasible combination of sub-component designs is not selected for a process variant v. We use trained piecewise linear surrogate models, represented by the system of equations f^c and f^p, to predict the variable cost of the variant, c_v, and the performance indicator, p_v. Piecewise linear surrogates are used because they can be represented exactly as MILPs within the optimization formulation.

The set of common sub-component types considered in the platform is given by K. In our case study, $K = \{\text{absorber, stripper}\}$. The set J_k includes the sub-component designs for a particular sub-component of type k. The cardinality of J_k dictates the number of sub-component designs to consider in the platform for sub-component type k (i.e., if we want two absorber designs, then $J_k = \{1,2\}$ where $k = \text{absorber}$). The variable $\hat{d}_{k,j}$ represents the design variable values for design j of sub-component type k, while variable $d_{v,k}$ represents the design variable values of sub-component type k for process variant v. With this, if design j is chosen for sub-component type k in variant v, then $d_{v,k} = \hat{d}_{k,j}$.

$$min. \sum_{v \in V} w_v c_v \tag{1a}$$

$s.t.$

$$c_v = f^c\left(b_v, d_{v,1}, \dots, d_{v,k}\right) \qquad \forall v \in V \tag{1b}$$

$$p_v = f^p\left(b_v, d_{v,1}, \dots, d_{v,k}\right) \qquad \forall v \in V \tag{1c}$$

$$\bigvee_{j \in J_k} \begin{bmatrix} Y_{v,k,j} \\ d_{v,k} = \hat{d}_{k,j} \end{bmatrix} \qquad \forall v \in V, k \in K \tag{1d}$$

$$\hat{d}^L_{k,j} \le \hat{d}_{k,j} \le \hat{d}^U_{k,j} \qquad \forall k \in K, j \in J_k \tag{1e}$$

$$p^L_v \le p_v \le p^U_v \qquad \forall v \in V \tag{1f}$$

$$c_v \ge 0 \qquad \forall v \in V \tag{1g}$$

$$Y_{v,k,j} \in \{\text{True, False}\} \qquad \forall v \in V, k \in K, j \in J_k \tag{1h}$$

The objective, Eq. (1a), is to minimize the total weighted annualized cost of all process variants, $v \in V$. Piecewise linear surrogates predict the cost of each variant c_v, Eq. (1b), and the performance indicator p_v, Eq. (1c). Both surrogates are functions of the boundary

conditions at a particular variant, b_v, and the design of each sub-component type $k \in K$. The disjunctions in Eq. (1d) select which of the sub-component designs, $j \in J_k$, for a sub-component of type k are selected for process variant v. The number of disjunctions is equal to $|K| \times |V|$ because we make this decision for each sub-component k and each process variant v. The number of disjuncts corresponds to $|J_k|$ for a particular sub-component of type k, because we decide on which sub-component design, $j \in J_k$, to select. Eq. (1e) bounds the design variable of sub-component type k, which also corresponds to the ranges used to train the surrogate. Eq. (1f) captures the performance constraints for each process variant.

3. Numerical Case Study

We demonstrate this formulation by designing a process family of 12 aqueous monoethanolamine (MEA) solvent-based carbon capture variants. Each variant requires a design that captures at least 90% of CO_2 entering the facility. In this case study, each process variant v is described by one of three flue gas flow rates and one of four carbon dioxide concentrations. This led to $3 \times 4 = 12$ possible combinations, representing the boundary conditions for the 12 process variants. Considering the units in the capture facility flowsheet, we chose the common sub-component types to be the absorber and stripper, $K = \{\text{absorber, stripper}\}$, since they contribute heavily to the overall purchase cost and have large effects on the performance of the system, thereby directly affecting operating costs.

To gather surrogate training data, we discretized the design ranges for the absorber and stripper. For the absorber, we considered 10 diameters in the range of 0.3 m – 1.2 m. For the stripper, we considered 10 diameters in the range of 10 in.–55 in. We performed a simulation of the carbon capture system for all possible combinations of the process variant boundary conditions and sub-component sizes. This led to $12 \times 10 \times 10 = 1{,}200$ simulations. We used an Aspen Plus model of the aqueous MEA system for the simulations (Morgan et al., 2018). Each simulation was marked as *feasible* if the key performance requirements (e.g., CO_2 recovery) were met and *infeasible* otherwise. This set of data and corresponding *feasible, infeasible* labels were used to train a 15-node, single-layer ReLU activated NN to determine the performance indicator, p_v. The NN was imported directly in the formulation, as shown in Eq. (1c), using OMLT: Optimization and Machine Learning Toolkit (Ceccon et al., 2022). Purchase cost was estimated using the IDAES costing framework (Lee et al., 2021). Combined with estimates for annual operating cost, total annualized cost was used to train a linear decision tree (Ammari et al., 2022) and imported using OMLT, represented by Eq. (1b), to predict c_v of process variant v. Note that both surrogates are captured as piecewise linear functions, resulting in an MILP formulation.

For this case study, we considered two absorbers and two strippers for the sub-component designs (i.e., $J_k = \{1,2\}$ for each sub-component type). The problem was formulated using Pyomo (Bynum et al., 2021) and solved using Gurobi (Optimization, LLC Gurobi, et al., 2020). Additionally, we used OMLT and the Pyomo extension, Pyomo.GDP (Chen et al., 2022), to capture the disjunctions. The optimal designs for the two absorbers and two strippers are reported in Table (1), and Figure (1) shows which combinations of the two possible absorber designs and two possible stripper designs were selected for each of the 12 process variants. The optimization determined the sizes for the common sub-components. The larger absorber was used for the upper right corner of the variant process

space, while the larger stripper was needed only for the most demanding variant in the upper right-hand corner.

Table 1. Optimal Designs for Common Units

Sub-component k	Design 1 $j_1 \in J_k$	Design 2 $j_2 \in J_k$
Absorber	0.3 m. diameter	0.61 m. diameter
Stripper	0.254 m. diameter	1.075 m. diameter

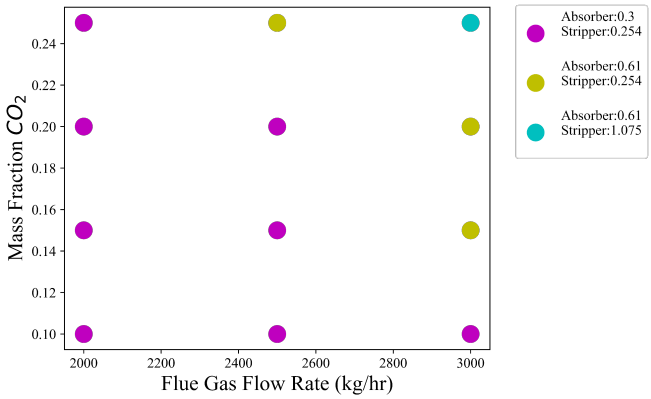

Figure 1. Process Family Design

Since the optimization is performed with surrogate models, we verified this result using the original model to simulate each of the 12 process variant boundary conditions with the corresponding set of assigned absorber and stripper designs. The total annualized verified cost was $\$1.29 \times 10^6$, which is an overall surrogate error of approximately 2%.

Table 2. Optimization Problem Specs

Num. of Binary Variables	191
Num. of Continuous Variables	273
Num. of Constraints	720
Total Annualized Cost (Obj.)	$\$1.26 \times 10^6$

The combination of training time for both the ReLU NN surrogate and linear decision tree was approximately 300 seconds, while the optimization formulation itself took approximately 1 second along with the problem structure as described in Table (2).

4. Conclusions and Future Work

We demonstrated a surrogate-based approach for optimally designing a process family of 12 carbon capture variants. We considered a platform of common sub-component designs that considered two absorbers and two strippers, whereas a conventional design approach would have generated 12 unique absorber designs and 12 unique stripper designs. This approach uses surrogates to predict required system costs and performance. This removes the nonlinearities in the GDP and allows the transformation to an MILP. The machine learning surrogates were easily included in the optimization with OMLT, and the overall formulation solved in under a second, demonstrating tractability in this approach.

For future work, we plan on expanding the sub-component design ranges and incorporating more sub-component types. Furthermore, we plan on incorporating an estimation of savings due to economies of numbers within the optimization formulation, exploring uncertainty in weights, and quantifying the degree of flexibility our designs offer across the design space.

Acknowledgments: This effort is part of the U.S. Department of Energy's (DOE) Institute for the Design of Advanced Energy Systems (IDAES) supported by the DOE Office of Fossil Energy and Carbon Management's Simulation-Based Engineering/Crosscutting Research Program.

References

Ammari, B., et al., 2022, Machine Learning Surrogates with OMLT and ideas for Improved Design and Analysis of Energy Systems, INFORMs Annual Meeting 2022, Indianapolis, IL.

Baldea, M., et al., 2017, Modular manufacturing processes: Status, challenges, and opportunities. *AIChE journal 63(10)*, 4262-4272

Bynum, M. L., et al., 2021, Pyomo- Optimization Modeling in Python. *Springer.*

Chen, Q., et al., 2022, Pyomo gdp: an ecosystem for logic based modeling and optimization development. *Optimization and Engineering 23(1)*, 607-642-

Ceccon, F., et al., 2022, Omlt: Optimization & machine learning toolkit, submitted, preprint on arXiv.

Lee, A., et al., 2021, The idaes process modeling framework and model library- flexibility for process simulation and optimization. *Journal of Advanced Manufacturing and Processing 3(3)*

Morgan, J.C., et al., 2018. Development of a rigorous modeling framework for solvent-based CO_2 capture. part 2: steady-state validation and uncertainty quantification with pilot plant data. *Industry & Engineering Chemistry Research 57*, 10464-10481.

Optimization, LLC Gurobi, et al., 2020, Gurobi optimizer reference manual, *URL http://www.gurobi.com12.*

Stinchfield, G., et al., 2022, Optimization-based Approaches for Design of Chemical Process Families Using ReLU Surrogates, submitted to Foundations of Computer Aided Process Operations, 2022.

Wilberforce, T., et al., 2019, Outlook of carbon capture technology and challenges, *Science of the total environment*, 56-72

Zhang, C., et al., 2022, Optimization-based design of product families with common components. *Computer Aided Chemical Engineering, Proceedings of 14th International Symposium on Process Systems Engineering (PSE 2021+)*. Volume 49, pp. 91-96. Elsevier.