MITSUBISHI ELECTRIC RESEARCH LABORATORIES
https://www.merl.com

Navigating the Trade-offs and Synergies of Economic and Environmental Sustainability Using Process Systems Engineering

Zhang, Qi; Avraamidou, Styliani; Paulson, Joel A.; Thakkar, Vyom; Wang, Zhenyu; Chiang, Leo; Braun, Birgit; Rathi, Tushar; Chakrabarty, Ankush; Sorouifar, Farshud; Tang, Wei-Ting; Guertin, France; Munoz, Paola; Sampat, Apoorva

TR2025-106 July 09, 2025

Abstract

This paper provides an overview of recent re- search efforts on the role of Process Systems Engineering (PSE) in advancing sustainability initiatives, particularly in achieving net-zero emissions and carbon neutrality. The paper is organized as a collection of four domains where PSE methodologies contribute to sustainability: (i) carbon monetization and low-carbon supply chains, where optimization and systems modeling help design cost-effective decarbonization strategies; (ii) circular economy and sustainable manufacturing, which leverage system-level optimization to minimize resource consumption and maximize economic viability; (iii) sustain- able land management and ecosystem services, where PSE approaches aid in quantifying trade-offs between land use, emissions, and economic feasibility; and (iv) advanced control technology, particularly in building energy management, where data-driven control strategies enhance energy optimization under significant uncertainty. In addition to reviewing relevant literature, this paper highlights common challenges across these domains and discusses future opportunities for integrating emerging technologies, such as generative AI and mixed-integer programming, into PSE-driven sustainability strategies.

American Control Conference (ACC) 2025

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

Navigating the Trade-offs and Synergies of Economic and Environmental Sustainability Using Process Systems Engineering

Qi Zhang[1], Styliani Avraamidou[2], Joel A. Paulson[3], Vyom Thakker[4], Zhenyu Wang[4], Leo Chiang[4], Birgit Braun[4], Tushar Rathi[1], Ankush Chakrabarty[5], Farshud Sorourifar[3], Wei-Ting Tang[3], Paola Munoz[2], Apoorva Sampat[4], and France Guertin[4]

Abstract— **This paper provides an overview of recent research efforts on the role of Process Systems Engineering (PSE) in advancing sustainability initiatives, particularly in achieving net-zero emissions and carbon neutrality. The paper is organized as a collection of four domains where PSE methodologies contribute to sustainability: (i) carbon monetization and low-carbon supply chains, where optimization and systems modeling help design cost-effective decarbonization strategies; (ii) circular economy and sustainable manufacturing, which leverage system-level optimization to minimize resource consumption and maximize economic viability; (iii) sustainable land management and ecosystem services, where PSE approaches aid in quantifying trade-offs between land use, emissions, and economic feasibility; and (iv) advanced control technology, particularly in building energy management, where data-driven control strategies enhance energy optimization under significant uncertainty. In addition to reviewing relevant literature, this paper highlights common challenges across these domains and discusses future opportunities for integrating emerging technologies, such as generative AI and mixed-integer programming, into PSE-driven sustainability strategies.**

I. Introduction

Sustainability initiatives, such as net-zero emissions and carbon neutrality, have gained significant momentum worldwide as governments, industries, and consumers seek to reduce environmental impact [1]. However, achieving these goals is not straightforward, as success requires overcoming a host of technological, economic, and operational challenges. Without effective strategies, sustainability efforts can lead to inefficient resource allocation, missed environmental targets, and economic disruptions that may discourage long-term commitments to sustainability. At the core of these challenges are two fundamental requirements for informed (data-driven) decision-making:

- **Accurate modeling of environmental and economic impacts:** Organizations must understand how their operations affect sustainability metrics while quantifying trade-offs in costs, energy use, emissions, and resource consumption. Without accurate models, decision-makers risk underestimating emissions, misallocating resources, or implementing ineffective policies. Conversely, well-calibrated models help identify strategic opportunities to reduce emissions while maintaining economic viability.
- **Optimal planning and control strategies for navigating uncertainty:** Sustainability is inherently dynamic; systems must adapt to fluctuations in energy demand, renewable energy availability, evolving regulatory policies, and shifting market conditions. Poorly designed decision-making frameworks can lead to inefficiencies, financial losses, or even unintended increases in emissions. In contrast, robust optimization and control strategies enable companies and policymakers to anticipate uncertainties and exploit system variability, leading to more resilient and cost-effective sustainability solutions.

A critical challenge in sustainability efforts is ensuring that environmental initiatives are financially viable. Many industrial sectors, including manufacturing, energy, and transportation, face pressure to reduce emissions while maintaining competitiveness. The scale of this challenge is reflected in estimates from McKinsey & Company, which project that reaching global net-zero targets will require \$3.5 trillion in annual investment [2]. If sustainability strategies are not designed with economic feasibility in mind, they may impose burdensome costs, slow adoption, or even set back global decarbonization efforts. However, as illustrated in some recent work [3]–[6], well-designed strategies can create *win-win* scenarios, where companies improve efficiency, reduce waste, and lower costs while also decreasing emissions.

Achieving sustainability at scale requires integrating *systematic, data-driven methodologies* that enable organizations to optimize resource allocation, make informed trade-offs, and develop long-term strategies that balance several factors such as economic growth and environmental responsibility. **Process Systems Engineering (PSE)** provides a structured approach to addressing these challenges by integrating mathematical modeling, optimization, control, and data analytics to guide decision-making in complex, uncertain environments [7]. PSE is not a single methodology but rather a unifying *framework* (or perspective) that combines diverse analytical tools to improve decision-making across various domains, from supply chain optimization to real-time process control. Traditionally, PSE has been applied in manufacturing, logistics, and energy systems to enhance efficiency,

[1]Department of Chemical Engineering and Materials Science, University of Minnesota, Minneapolis, MN 55455
[2]Department of Chemical and Biological Engineering, University of Wisconsin–Madison, Madison, MI 53707
[3]Department of Chemical and Biomolecular Engineering, Ohio State University, Columbus, OH 43210
[4]The Dow Chemical Company, Lake Jackson, TX 77566
[5]Mitsubishi Electric Research Laboratories, Cambridge, MA 02139

minimize waste, and optimize production processes. More recently, its principles have been extended to address sustainability challenges by helping organizations and policymakers navigate the trade-offs between environmental impact and economic considerations.

This paper explores four key domains where PSE methodologies can play a critical role in sustainability efforts:

1) **Carbon monetization and low-carbon supply chains (Section II)**: Companies are increasingly leveraging carbon monetization strategies, such as compliance carbon market, voluntary carbon market and internal carbon pricing, to create financial incentives for emission reductions. PSE methodologies help optimize low-carbon supply chain planning, ensuring that decarbonization strategies are both cost-effective and scalable.
2) **Circular economy and sustainable manufacturing (Section III):** The transition from a linear to a circular economy, where resources are reused, remanufactured, and recycled, requires system-level optimization to design closed-loop material flows and evaluate economic-environmental trade-offs in manufacturing systems. PSE provides tools for designing these strategies, minimizing resource consumption, and maximizing the economic viability of circular business models.
3) **Sustainable land management and ecosystem services (Section IV):** Land use decisions, such as reforestation, wetland restoration, and carbon sequestration, can play a complementary role to technological CO_2 reduction strategies. However, quantifying the trade-offs between land use, emissions, and economic feasibility requires systematic modeling and optimization approaches—areas where PSE methodologies are well-suited to provide insights.
4) **Advanced control technology for energy efficiency (Section V):** Many sustainability challenges involve real-time decision-making to minimize waste and improve energy efficiency. For example, Heating, Ventilation, and Air Conditioning (HVAC) systems account for nearly 50% of energy consumption in commercial buildings, with 30% of this energy often wasted due to suboptimal control. Advanced control strategies, particularly stochastic model predictive control (SMPC), offer a means of dynamically optimizing energy usage while accounting for uncertainty in demand, weather conditions, and renewable energy availability. Recent advancements in machine learning-enhanced control, such as generative time-series models, further improve the ability to forecast disturbances and optimize system performance adaptively.

By examining these four key domains, this paper aims to (1) review cutting-edge PSE methodologies that support sustainability efforts, (2) identify common challenges and methodological needs across these domains, and (3) discuss future opportunities for integrating PSE with emerging technologies (such as generative AI, Bayesian optimization, and mixed-integer programming) to further advance sustainable decision-making in both academic and industrial contexts.

II. Carbon Monetization and Its Impact on Low-Carbon Supply Chain Planning

Carbon monetization, as the name implies, is the strategic use of financial and economic incentives to encourage and ensure the reduction of carbon emissions. This allows industries involved to treat the process of reducing their carbon footprint, either directly or indirectly, as an opportunity to generate additional revenue or reduce costs, in addition to the more obvious reasons, such as the social benefit of decarbonizing the economy, adhering to government regulations, and attracting consumers and investors by being environmentally responsible.

One way governments monetize carbon reduction is through the carbon tax, also known as the emission tax, which imposes a direct cost on businesses and organizations based on the amount of CO_2 emissions they contribute. Although reducing greenhouse gas (GHG) emissions is necessary to mitigate irreversible climate change, it is unreasonable to expect all companies to develop and implement carbon capture or carbon-free technology immediately. This is primarily due to one or a combination of the following reasons: immature technology, issues with scale-up, or, in the case of many small to medium-sized companies, a lack of funds to invest in the research, development, and implementation of sustainable technologies. An increasingly common workaround is government-issued tradable emission allowances, also known as carbon permits. One allowance typically grants permission for the emission of one tonne of carbon dioxide equivalent (tCO_2e). These allowances are governed by national and international governments and must be mandatorily complied with by companies that fall under their jurisdiction; thus, these allowances are traded in what is called the compliance carbon market (CCM). For example, the EU Emissions Trading System (EU ETS) [8], established in 2005, uses a cap-and-trade system in which allowances are sold in auctions, and companies can later trade them among themselves to reduce costs or justify their emissions to avoid heavy penalties. The revenue generated from these allowances is used to fund the development of new low-carbon technologies as well as the expansion of existing ones. To further minimize emissions, the cap on available allowances across different sectors is reduced each year. Other examples of similar ETS include California's cap-and-trade program, which was launched in 2013 and regulates 80% of the state's GHG emissions [9], and China's ETS, which started operating in 2021 and is the world's largest in terms of emissions covered [10].

Another type of fast-emerging carbon market is the voluntary carbon market (VCM). VCM typically involves the trade of carbon offsets (offset credits), which are verifiable certificates issued to projects that fall into one of two categories: avoidance and removal projects. Avoidance projects generally substitute conventional GHG-emitting processes and involve the development and implementation of renewable technologies such as wind power, solar power, green H_2 production using electrolysis, and many others. Removal projects entail the direct removal of CO_2 from the atmosphere, such as direct air capture, afforestation, reforestation, and other sequestration projects. One carbon offset is generated when 1 tCO_2e of emissions is avoided or removed from the atmosphere. Companies with their own net-zero mission, usually driven by corporate social responsibility (CSR) to enhance their reputation and brand value, often invest in carbon offset-generating projects to compensate for their emissions and quickly attain their sustainability goals. Unlike CCM, participation in VCM is completely optional and not tied to government-regulated targets. Furthermore, in addition to fulfilling the CSR objectives of large corporations, VCM also provides an opportunity for small businesses and individuals to reduce their carbon footprint by purchasing carbon offsets, as they are traded in open markets. Lastly, note that the term "carbon credits" is often used loosely for both carbon allowances and offsets; however, one should pay attention to whether it is being referred to in the context of CCM or VCM.

In addition to the above-mentioned carbon markets, many companies are actively devising their own strategies to make optimal investment decisions for the future while hedging against potential financial risks posed by carbon taxes or caps imposed by government regulations. One such strategic tool gaining popularity is internal carbon pricing (ICP) [11], where a company assigns a monetary value to its GHG emissions, often done at regional levels. Effective utilization of ICP requires careful evaluation of the business's emissions profiles, including the section of the value chain that contributes the most to it, the objectives of whether the focus is to hedge against potential risks or drive investments in low-carbon technologies, or a mix of both, and, obviously, the uncertainty in how carbon prices might unfold in the future to meet the net-zero goal. The pricing mechanisms in ICP are generally based on a carbon fee, which involves a per-unit fee based on the amount of GHG emissions, or a shadow price, which involves a hypothetical payment for the purpose of analyzing and drafting future investment decisions.

Although carbon markets are emerging as an effective tool for decarbonization, several challenges remain in their adoption and implementation. Some of the challenges are summarized below:

- Regulatory uncertainty: Since decarbonization demands global efforts, it requires governments of different countries to come together and draft consistent and effective regulatory policies, which is often a difficult task due to the complexities of international relations and varying levels of commitment to the decarbonization goal. This is especially a hurdle for the operations of multinational corporations and supply chains that transcend national boundaries, as they must navigate inconsistent regulations across different jurisdictions.
- Offsets credibility: Offsets from carbon reduction or removal projects are difficult to quantify due to their complexity and type. Although an increasing number of standards and registries, such as the Verified Carbon Standard [12] and the American Carbon Registry [13], are being designed to verify and certify these offsetting projects, ensuring that they actually reduce emissions, it remains a significant challenge because it directly impacts the price, perception, and thus the sale of these offsets.
- Carbon price volatility: Offsets sold in the market can have drastically different prices based on factors including: (1) Type of project – removal projects generate offsets that are more sought after and thus priced higher. (2) Leakage – a high-quality project must ensure that the emissions reduced are not shifted to another location or source. (3) Permanence risk – projects that reduce emissions in ways that are less likely to be reversed in the future generate higher-quality (and thus higher-priced) offsets. (4) Additionality – "additional" projects are those that would not have been financially feasible without the revenue from the offsets they generate. This allows companies to claim that their investment (from purchasing the offsets) directly enabled the emissions reduction. The ability to make such a claim often leads to a higher price for these offsets. (5) Co-benefits – projects that provide benefits beyond emissions reduction are perceived positively and generate higher-priced offsets. For example, reforestation supports wildlife in addition to carbon removal. (6) Vintage – offsets generated from recent projects are valued higher compared to similar projects launched in the distant past.
- Greenwashing: The misuse of sustainability claims due to the complexity of carbon credits or offset certifications can allow companies to exploit the system and present a false image of environmental responsibility without making genuine sustainability efforts.

To achieve decarbonization in a timely manner, long-term optimal capacity planning [14], [15] is crucial. In addition, as renewable technologies penetrate the chemical sector, considering low-carbon alternatives for heavily utilized products such as ammonia [16], [17] is essential. As such, drawing from the work by Rathi et al. [18], we discuss how a low-carbon ammonia supply chain can have significantly different investment and operational decisions, as well as costs, depending on the accountability and traceability method (more commonly known as the chain of custody model) chosen for tracking ammonia of different carbon intensities in the network. Chain of custody, as defined by ISO [19], "is a process by which inputs and outputs and associated information are transferred, monitored and controlled as they move through each step in the relevant supply chain." For a low-carbon product supply chain, the choice of chain of custody model governs the extent to which the product is categorized based on its characteristics (e.g., carbon intensity) at each stage of the supply chain. The four most commonly used chain of custody models are identity preservation, segregation, mass balance (MB), and book-and-claim (B&C); however, only MB and B&C are practically relevant to the chemical industry. In the MB model, products with different characteristics can mix together; however, the quantity of product with each characteristic is reconciled at each stage, ensuring that the consumer ultimately pays for the characteristic of the product it receives. In a low-carbon ammonia supply chain, this means that ammonia with different carbon intensities from various production sites is mixed during the distribution stage but the consumer pays for the final carbon intensity of the ammonia it physically receives. In contrast, B&C, also known as certificate or credit trading, differs from MB in that there is no bookkeeping at any stage of the supply chain. Instead, certificates are issued to producers of low-carbon ammonia, which can be purchased by consumers intending to claim low-carbon ammonia irrespective of the carbon intensity of the ammonia physically delivered to them. Note that this is similar to the concept of carbon offsets explained previously. Lastly, to ensure proper accounting in the B&C model, it is crucial to ensure that the low-carbon certificates generated and purchased correspond exactly to the amount of low-carbon product flowing in the supply chain.

We consider a case study of a low-carbon ammonia supply chain network spanning nine U.S. states, with 15 existing ammonia production sites, three potential production sites, and each state's geographical center acting as a consumer site (see Fig. 1). Each production site is modeled as a superstructure with the option of steam methane reforming (SMR) and electrolysis (ELC) for hydrogen production, a carbon capture unit (CCU), and air separation (ASU) and Haber-Bosch (HB) process units for ammonia production. Natural gas and electricity from renewable sources (solar and wind) are explicitly considered as resources. The expansion of the supply chain is analyzed over a 25-year planning horizon from 2025 to 2050, divided into five five-year time periods. The authors assume that ammonia demand increases by 15% in each period. The price of ammonia (assumed to be constant over the planning horizon) offered by each consumer as a function of carbon intensity is shown in Fig. 2. The carbon intensity ranges from 0 to 1, with lower values indicating "greener" ammonia. Specifically, a carbon intensity below 0.05 indicates green ammonia, from 0.05 to 0.4 (excluding) indicates blue ammonia, and 0.4 or above indicates gray ammonia. Consumers C, F, and I are assumed to be willing to pay a premium for low-carbon ammonia, meaning they are more interested in claiming the receipt of low-carbon ammonia. In the case of the B&C model, the higher premiums can also be interpreted as the additional cost of purchasing low-carbon certificates. The expansion planning problem for the two chain of custody models (MB and B&C) is formulated as a mixed-integer nonlinear programming model, with the objective of maximizing the net profit of operating the supply chain.

The ammonia distribution plots are shown in Fig. 3. The color of the edges represents the grade of ammonia being transported, the color of the producer nodes indicates the grade of ammonia they manufacture, and the color of the consumer nodes indicates the grade of ammonia they pay for (which is also the grade of ammonia delivered to them in the

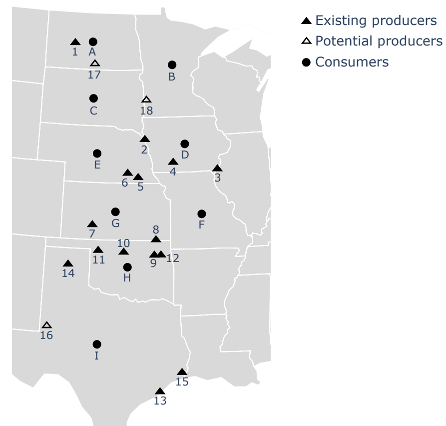

Fig. 1: Ammonia producers and consumers: producers are numbered 1 to 18, and consumers are labeled A to I [18].

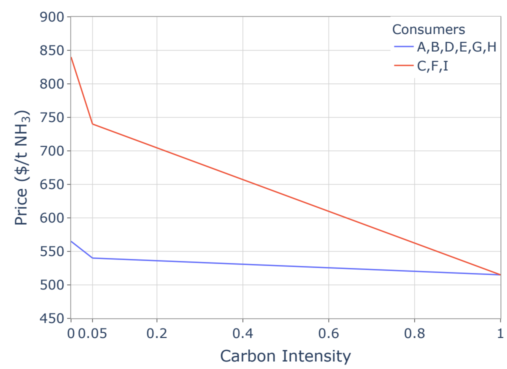

Fig. 2: Time-invariant consumer price profiles [18].

case of the MB model). The key observation here is that, in the MB model, each consumer pays for the ammonia they receive, whereas this is generally not true in the B&C model. For example, in the first time period, the high-premium-paying Consumer I pays for and receives blue ammonia in the MB model, whereas in the B&C model, it receives gray ammonia from Producers 13 and 15 but is able to claim green ammonia regardless. Many other similar examples can be seen in Fig. 3. Second, the ability to claim low-carbon ammonia via the purchase of certificates in the B&C model enables the avoidance of long-distance deliveries. For example, in the third time period, the high-premium-paying Consumers C and F receive part of their low-carbon ammonia from a relatively far-off Producer 7. In contrast, in the B&C case, they are able to claim green ammonia even when receiving gray ammonia from nearby producers. In the later time periods, this becomes more evident, with the MB case showing many more edges for long-distance deliveries unlike the B&C case. Lastly, as the initially installed SMR units retire and electrolysis becomes cost-competitive over time for hydrogen production, green ammonia trading increases

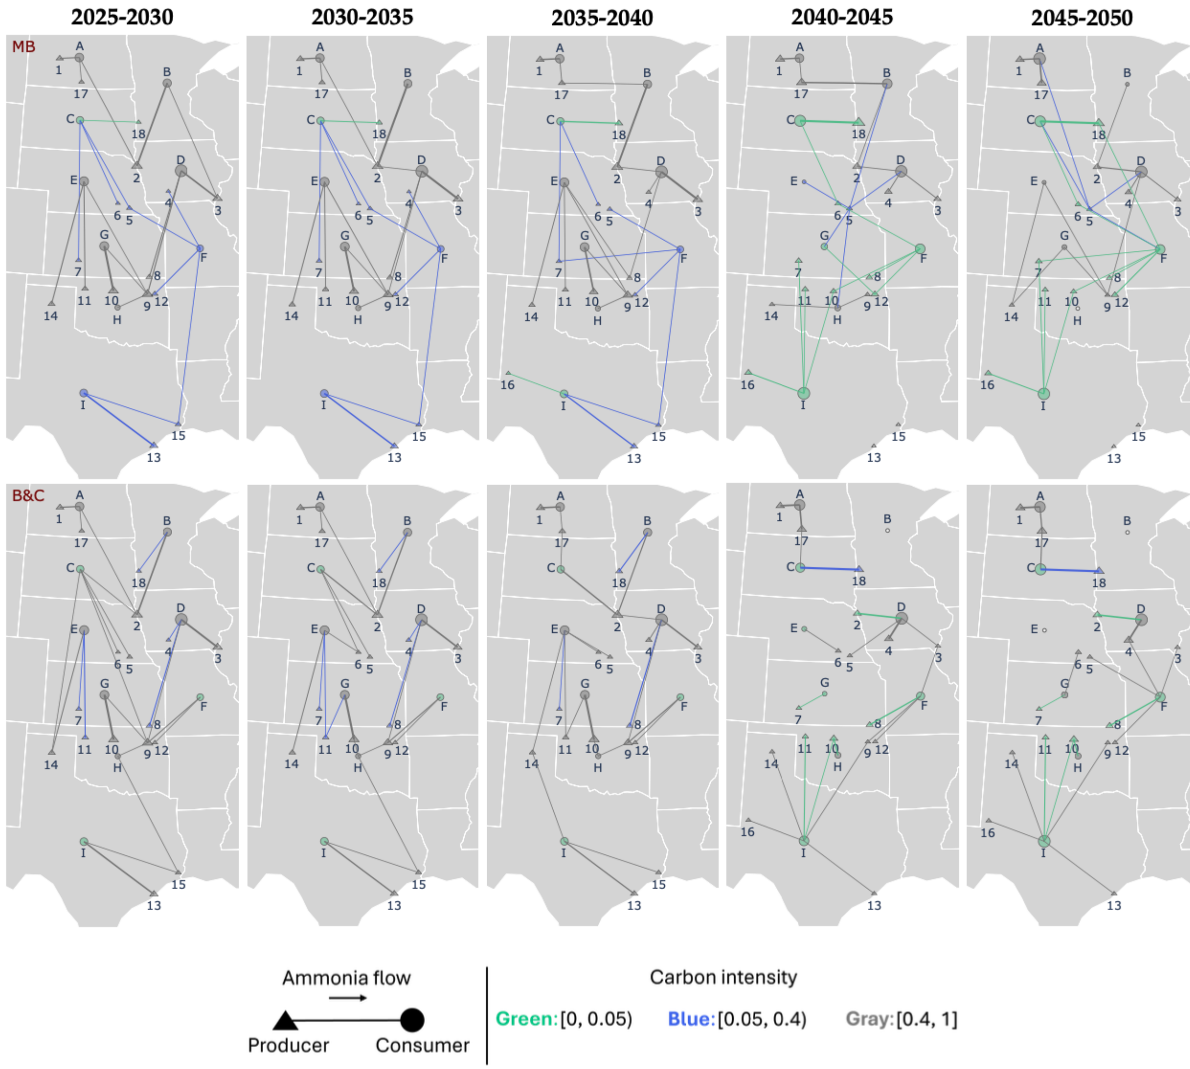

Fig. 3: Ammonia distribution network under the MB (top) and B&C (bottom) models. A missing node at any producer location indicates either no production units were set up, or all previously installed capacity has been retired [18].

significantly in the network during the final two time periods.

The carbon intensity of ammonia is directly dependent on the technology utilized for hydrogen production; thus, the expansion of these technologies, including carbon capture at each producer site, is highlighted in Fig. 4. The first (black) bar at each producer indicates any initially installed SMR capacity, and the next five bars, from left to right, correspond to the five time periods in chronological order, showcasing the expansion (or retirement) of SMR (gray bar), CCU (blue bar), and ELC (green bar). A missing bar for a time period indicates that either no production facility is set up at the location or all capacity has retired. It is observed that the MB model invests significantly more in ELC compared to the B&C model. This is primarily because the MB model can only benefit from the high premiums offered for carbon intensities <0.05 if it actually delivers ammonia of such low carbon intensity, which requires investing in hydrogen production technology capable of producing such low-carbon ammonia, i.e. ELC units. In contrast, recall that the B&C model relies on certificate trading and thus prefers investing in SMR and CCU because it can still take advantage of the high premiums by, for example, delivering higher-carbon-intensity ammonia to consumers who want to claim green ammonia (by purchasing certificates) and blue ammonia to consumers who pay for gray ammonia. The total investment in hydrogen and ammonia production technologies over the planning horizon under the two chain of custody models is summarized in Table I. In comparison to the MB model, the choice of the B&C model leads to approximately 5% more ammonia production, with 37.36% more investment in blue ammonia but 72.66% less investment in green ammonia production.

The percentage demand met in each time period under the two chain of custody models is summarized in Table II. For the given network of producers and consumers and the underlying price profiles, compared to the MB model, the B&C model results in investment decisions that lead to relatively more demand satisfaction over the last two time

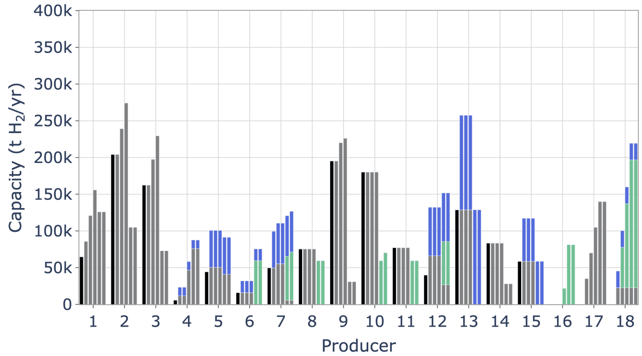

(a) MB.

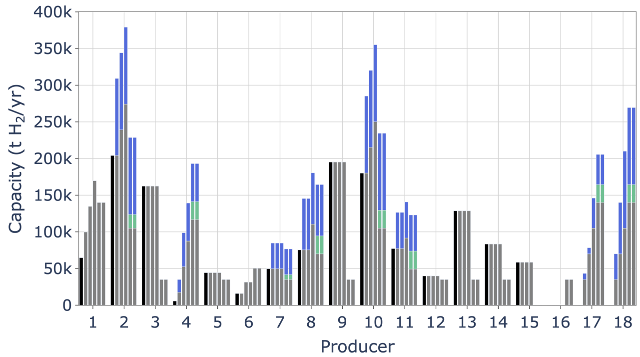

(b) B&C.

Fig. 4: Available capacities for hydrogen production technologies at all sites over the planning horizon. The first bar for each producer represents the initially installed SMR capacity (if any), while the remaining bars show the available capacities of each technology from time period 1 to 5 (left to right), color-coded as follows: SMR (gray), CCU (blue), and ELC (green). A missing bar indicates no available capacity, either because no units were installed or because previously installed units have been retired [18].

Technology	Unit	MB	B&C	% difference
SMR	kt H_2/yr	675	1,197	77.33
CCU	kt CO_2/yr	3,683	5,059	37.36
ELC	kt H_2/yr	629	172	-72.66
ASU+HB	kt NH_3/yr	7,341	7,709	5.01

TABLE I: Summary of total expansion for each technology over the planning horizon under the MB and B&C models [18].

periods, following the retirement of initial capacity. This can be primarily attributed to the greater flexibility offered by the B&C model in investment decisions, as it has a relatively less restricted certification system. The demand requested and met for each consumer over the planning horizon under both chain of custody models is shown in Fig. 5. Each consumer has five bars (one per time period, increasing from left to right), where the fill levels indicate the demand met, and the color represents the grade of ammonia claimed. Given the high premiums offered by Consumers C, F, and I, both the MB and B&C models fully satisfy their requested demand. However, unlike in the MB model, under the B&C model, these consumers are able to claim green ammonia in all time periods.

Model	Time period				
	2025-30	2030-35	2035-40	2040-45	2045-50
MB	100	100	100	56.29	49.59
B&C	100	100	100	59.88	52.07

TABLE II: Summary of the percentage demand met in each time period for the MB and B&C models [18].

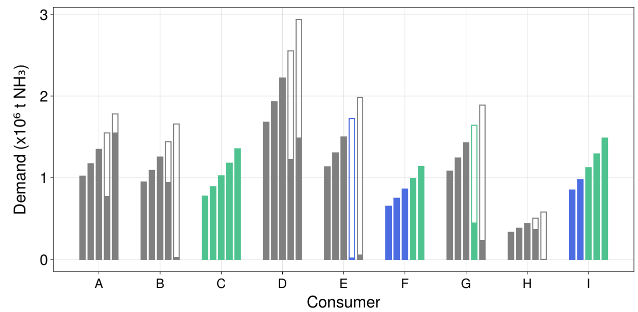

(a) MB.

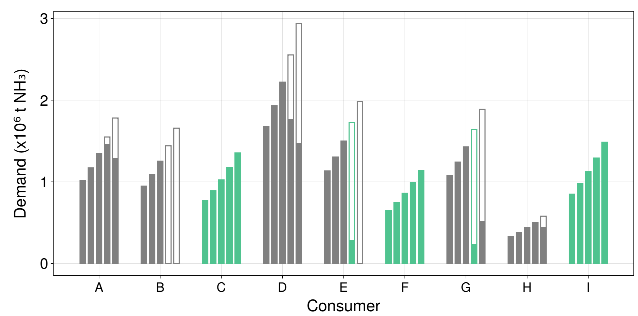

(b) B&C.

Fig. 5: Demand satisfied for each consumer over the planning horizon. For each consumer, the bars represent the requested demand for each time period (increasing left to right), with the fill level indicating the extent of demand met. Bar colors reflect the range in which the carbon intensity of the claimed ammonia falls [18].

An interesting observation in the case of the MB model is its tendency to overbuild CCU, increasing the likelihood of these units becoming redundant when SMR units are retired. For example, due to their proximity to high-premium-paying Consumer I, Producers 13 and 15 initially invest in CCU (see Fig. 4) to supply low-carbon ammonia (see Figs. 3 and 5), as electrolyzers are expensive during the initial time periods. However, without any additional investment in SMR, the installed CCU units become completely redundant after the SMR units are retired at the end of the third time period, as seen in Fig. 4. Similar observations can be made for Producers 5, 6, 7, and 12, whose usable CCU capacities reduce to 81.3%, 0%, 10%, and 40% of the net installed CCU capacities at these sites, respectively, in the last two time periods. In contrast, in the B&C model, any CCU capacity

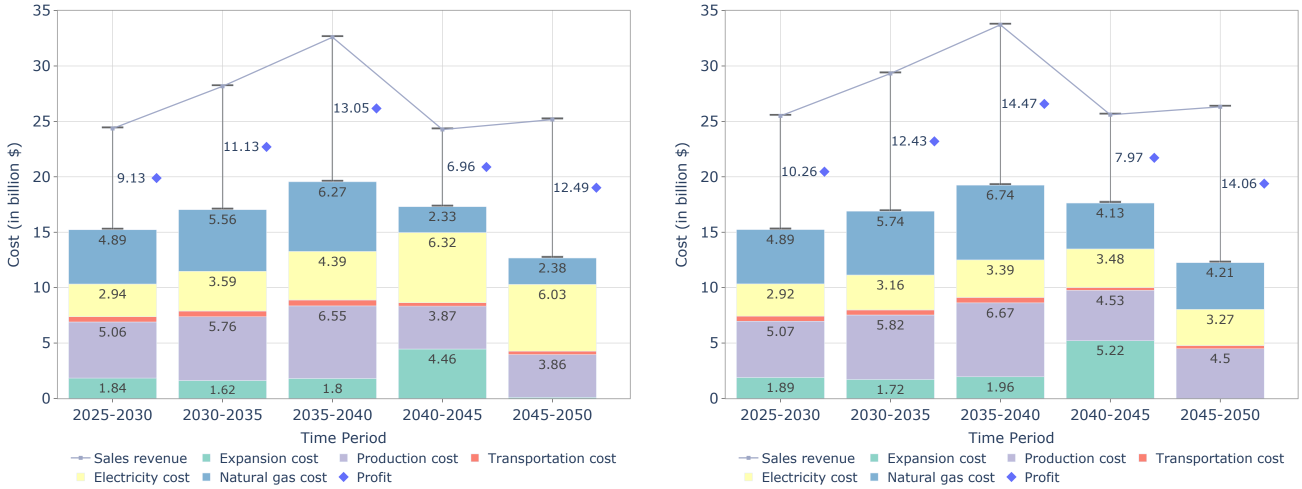

(a) MB.

(b) B&C.

Fig. 6: Breakdown of various costs and resulting profit for each time period in the MB and B&C models [18].

built during the planning horizon maintains 100% utilization. Hence, while the optimization model may theoretically favor overbuilding CCU, in practice, this could be inefficient due to wastage of resources.

Figure 6 depicts the breakdown of various costs under the two chain of custody models. The first major difference is that B&C model results in higher profits, which can be attributed to the flexibility offered by its underlying certification system, making it easier to claim low-carbon ammonia, which is sold at a much higher premium. Second, the MB model results in higher electricity costs, particularly in the final two time periods, due to greater investment in electrolyzers. Third, the higher natural gas costs in the B&C model during the last two time periods are primarily due to low-carbon ammonia production using SMR and CCU. Lastly, although transportation cost is rather small compared to other costs, B&C leads to significantly lower transportation cost per tonne of ammonia demand met, as shown in Table III.

Model	**Time period**				
	2025-30	**2030-35**	**2035-40**	**2040-45**	**2045-50**
MB	57.8	51.8	47.5	44.5	41.9
B&C	54.5	46.8	43.2	33.3	35.1

TABLE III: Summary of the transportation cost incurred per unit demand met ($/t NH_3) in each time period for the MB and B&C models [18].

Finally, Table IV summarizes the overall carbon intensity in each time period under the two chain of custody models. Although similar carbon intensities are observed across the two models, the B&C model results in the production of marginally higher-carbon-intensity ammonia during the last two time periods, primarily due to its preference for SMR+CCU to produce hydrogen for low-carbon ammonia, as opposed to the MB model, which prefers the lower-carbon option of ELC for hydrogen production.

Model	**Time period**				
	2025-30	**2030-35**	**2035-40**	**2040-45**	**2045-50**
MB	0.74	0.74	0.74	0.45	0.45
B&C	0.73	0.73	0.73	0.48	0.48

TABLE IV: Summary of the overall carbon intensity of ammonia produced in each time period for the MB and B&C models [18].

In this case study, PSE methodologies have been utilized to uncover the impact of various chain of custody models on investment and operational decisions. These insights are crucial for formulating optimal strategies and navigating the challenges associated with carbon monetization, ultimately aiding in the achievement of decarbonization goals.

III. Optimal Design for Circular Supply Chains

Despite growing awareness, environmental issues like resource depletion, water pollution, and greenhouse gas (GHG) emissions continue to threaten sustainability. These challenges are exacerbated by increasing population growth and industrialization, placing stress on natural resources and ecosystems. In response, supply chain objectives have evolved beyond merely meeting demand; they now emphasize the integration of sustainable practices, including the adoption circular economy principles to minimize environmental impacts.

Integrating sustainability aspects into circular supply chain design, redesign, modeling, or optimization requires systems engineering tools that can capture the complexity of new circular supply chains, combine optimization techniques and quantitative decision-making tools to find optimal alternatives. Promoting sustainability involves adopting circular economy practices—such as reusing, recycling, and re manufacturing—to close loops within the supply chains. Additionally, it entails adopting innovative processes and green resources that reduce resource consumption and minimize environmental impacts. However, while numerous options can be proposed, a comprehensive assessment is essential. Systems engineering frameworks are required to integrate the multiple environmental and economic assessment tools to identify optimal pathways and ensure that truly sustainable choices are made.

A. *Challenges in using systems engineering tools for circular economy supply chains*

Modeling the entire supply chain is essential because it consists of multiple interconnected stages, where decisions or events in one stage can significantly impact the others. A holistic perspective ensures accurate representations that optimize not only operational efficiency but also economic and environmental sustainability. For instance, a supplier might prefer a bio-based material for its natural origin, but if the production process or waste management involves energy- or water-intensive steps, the overall environmental impact could be worse than that of alternative materials, highlighting the importance of considering all stages of the supply chain. However, such models are inherently more complex, and incorporating circular processes further increases the challenge, as it introduces additional cycles and feedback loops into the system.

The multi-scale, multi-faceted and interconnected nature of supply chains represent challenges into modeling and decision making. Given supply chains are complex interconnected networks that collaborate to efficiently produce and distribute goods or services, they involve multiple stakeholders which bring conflicting objectives into play. Additionally, considering supply chains with extended boundaries and multiple levels of complexity leads to large-scale problems that are difficult to solve. Expanding the boundaries introduces multiple scales, both in space and time, which can create challenges as each level has the potential to impact the others [20]. There are additional challenges specific for different types of supply chains. For instance, food and biomass supply chain face challenges like perishability that require precise coordination and introduce time sensitive constraints [21]. Also, chemical and pharmaceutical supply chains require many regulatory and safety constraints.

B. *Systems Engineering Tools for Circular Supply Chains*

1) Circular Economy Assessment Tools: To achieve effective decision-making, assessment tools and methods are essential to evaluate and compare different alternatives based on their economic and environmental impacts. To measure economic impacts, methods like techno-economic assessment (TEA) are used to evaluate the profitability of the processes [22]. For environmental assessment, one of the most common methods is life cycle assessment (LCA) which is a standardized tool that measures factors like emissions and health impacts related to the life cycle of products or services [23]. Nevertheless, to achieve a more holistic environmental evaluation, circularity metrics have been proposed focusing on different aspects of Circular Economy like material circularity and resource consumption. [24], [25].

An example of a holistic circularity assessment is the micro-level framework proposed by Baratsas et al [24], who proposed indicators and metrics based on categories derived from the CE goals. The categories defined are: waste, water and procurement, energy, emissions and spillages and durability. Originally applied to companies across various economic sectors, the framework tracks their circularity performance over multiple years. The circularity score ranges from 0 to 1, where a score of 1 represents perfect circularity and a score of 0 indicates complete linearity. The indicators, based on the GRI standards, enable the collection of information for the different companies. A circularity subindex is calculated for each category by assigning specific weights to the respective metrics. Finally, the overall circularity index is computed as a linear average of all the subindices.

Based on the previous work of Baratsas et al [24], the framework was adjusted to apply to multiple sectors, including plastic recycling [26], chemical production [27], energy carrier production and use [28], and food packaging waste management infrastructure [29]. Taking the food packaging waste management infrastructure as an example, to achieve a comprehensive environmental assessment, we expanded beyond an assessment of GHG emissions and included this circularity framework. The circularity assessment includes the production, waste management and transportation of different packaging an waste management options. The indicators and metrics applicable were selected and the durability category was replaced for a substitutability category with a metric based on the product's market value. The study focused on identifying the optimal packaging and waste management combination in terms of economic and environmental criteria [29].

To demonstrate how the circularity index is calculated, we will focus just on waste management, specifically the scenario of sending a polyethylene (PE) rigid container to mechanical recycling. First, data is collected on a per-container basis. For the energy category, the established indicators are total energy consumed and total renewable energy consumed. In this scenario, the total energy consumed is 0.14 MJ per container, and the total renewable energy consumed is 0.01384 MJ. The corresponding metrics are the percentage of renewable energy over total energy consumed and the total energy consumed, respectively. Based on the collected data, the first metric is calculated as 9.9%, while the second metric is 0.14 MJ. The first metric is already expressed as a percentage, but the second metric requires normalization. Using an upper bound defined as 1.5 times the average of all total energy consumption across all the waste management scenarios, the normalized value for this metric is 0.8170. Given that the total energy consumed is considered more significant than the percentage of renewable energy, the weights for the metrics are distributed as 60% and 40%, respectively. This results in an energy subindex score of 0.53. A similar process is applied to calculate the sub indices for the other categories: water (0.75), emissions (0.88), waste

(0.69), and substitutability (0.65). The overall circularity index is then computed as the linear average of these sub indices, yielding a score of 0.65. This value is subsequently compared with other waste management scenarios and stages to identify the most circular option.

2) Exploring trade-offs between environmental and economic impacts: Transitioning toward sustainability requires addressing environmental, economic, and social aspects simultaneously. To achieve this, it is essential to consider multiple objectives at once, which can be effectively accomplished through multi-objective optimization. This technique evaluates two or more conflicting objectives to generate a set of optimal solutions, enabling the exploration of relationships and trade-offs between them. Some of the methods commonly used to solve multi-objective optimization are the weighted sum method and the ϵ-constraint method [30]–[32]. The first method consists on assigning weights to the different objectives and then multiplying and summing over to obtain a single overall objective function. The ϵ-constraint method instead focuses on optimizing one of the objective functions while treating the other objective functions as constraints. This is achieved by solving the other objective functions separately, and defining a set of ”epsilon” limit. These ϵ limits act as bounds, ensuring that when you optimize one of the objective functions, the other objectives work as constraints and a set of feasible solutions for all the objectives can be obtained. The resulting Pareto plot will visually represents the trade-offs between the objectives, providing valuable insights to guide decision-making. Numerous studies have successfully applied multi-objective optimization to compare key aspects such as cost versus emissions or resource consumption [33]–[35].

3) Modeling Circular Supply Chains: Within the supply chain context there are different approaches for modeling and optimization. Graphs are used to model complex systems like supply chains given their ability to represent a networks of interconnected elements like nodes and edges. Being edges the connections between nodes. In the case of the supply chain, nodes can represent the different entities such as, suppliers, producers, retailers or if looking closely into a specific stage like waste management, the different processes involved like pretreatment and recycling. Superstructures are a graph based approach commonly used to model a representation of all the possible pathways within the system. Here, the edges between nodes represent the flows coming from one unit operation or entity (nodes) to another. After defining the representation with any of the different types of graphs, they are transformed into a mathematical programming model, followed by the application of a solving method to determine the optimal pathway.

4) Case study: The study mentioned earlier in the assessment section, which aims to identify the optimal packaging and waste management combination by considering economics, emissions, and circularity, serves as an example of a circular economy (CE) systems engineering framework. This framework integrates assessment and optimization tools. A superstructure approach is used to represent all the possible pathways between packaging, waste management technologies and products obtained. The study considers a variety of packaging options, including films and rigid containers made from materials such as polyethylene (PE) and glass. It also evaluates multiple waste management technologies, such as mechanical recycling, pyrolysis, solvent-targeted recovery and precipitation (STRAP), landfill, incineration, and a cleaning facility for returnable containers. The study combines life cycle assessment (LCA), techno-economic analysis (TEA), and the circularity assessment previously described. A mixed-integer linear programming (MILP) model is formulated with the objective functions of minimizing cost, minimizing emissions, and maximizing circularity.

In the single-objective optimization, the combination of multilayer film and the technology STRAP is identified as the most economically feasible alternative. The reusable glass container emerges as the least emitting alternative. However, if considering non-local distances (greater than 65 miles) the least emitting alternative becomes multilayer film sent to landfill. Given landfill has other environmental impacts that emissions do not include, the next step is to identify the most circular pathway which is selected as the reusable glass container. For the multi-objective optimization, the ϵ-constraint method is applied to solve the problem, generating a Pareto front that illustrates the trade-offs between the objectives. The following figure (Fig 7) presents an example of the Pareto plot for cost versus emissions, demonstrating that cost increases as emissions decrease, and vice versa. The trade-off obtained between circularity and emissions, as shown in figure 8 demonstrates that emissions and circularity do not always follow the same trend. For a local distance of 7 miles, the most circular option—the reusable glass container—is also the least emitting. However, for a non-local distance of 65 miles, the reusable glass container remains the most circular option, while the multilayer film becomes the least emitting alternative. Furthermore, the results indicate that when emissions increase, circularity also increases. This suggests that reducing emissions does not necessarily lead to improved circularity, highlighting the importance of considering both criteria to assess environmental impacts.

Another example of a CE systems engineering framework that considers implementing alternative processes and valorization pathways is proposed by Baratsas et al [34], and applied to a coffee supply chain case study. The framework encompasses the entire supply chain from coffee cherries to different end-products like whole beans, coffee beverages and soluble/instant coffee. Different alternatives for production are identified as well as valorization pathways for wastes and by-products. For instance, coffee waste and by-products, such as husks, pulp, and mucilage, can be converted into valuable resources like bio ethanol or biogas. All these stages are arranged into a RTN (Resource-Task-Network) representation, including conversion factors. An RTN is a graph-based approach that represent tasks and resources as nodes and their relationship as edges, forming a bipartite directed graph. Then, a mathematical MILP (mixed-integer linear programming) model of the supply chain is formulated

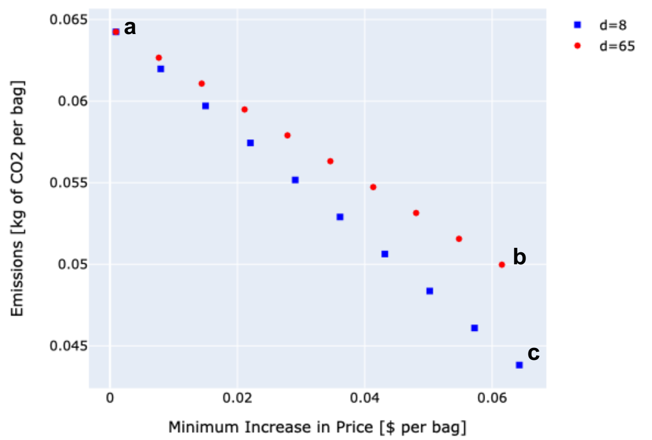

Fig. 7: Pareto front considering cost and emissions for a local and non local distance with the following optimal solutions. Most profitable pathway: a) Multilayer film & STRAP. Least emissions pathway: b) Multilayer film & Landfill c) Reusable glass container .

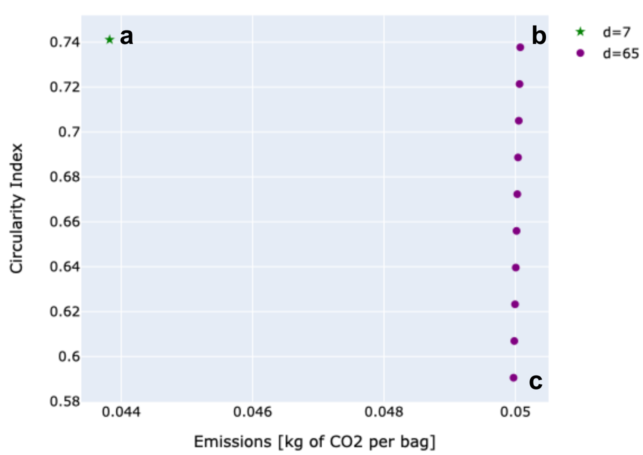

Fig. 8: Pareto front considering circularity and emissions for a local and non local distance with the following optimal solutions. a) Optimal pathway (emissions and circularity): Reusable glass container, b) Most circular pathway: Reusable glass container c) Least emissions pathway: Multilayer film & Landfill

based on this information. Objectives based on the CE goals are defined like maximizing the energy output of the supply chain, minimizing water consumption, coffee cherries consumption, waste generation and emissions. Single optimization and multi-objective optimization considering some of the objectives previously mentioned are explored for five coffee-product demand scenarios. A Pareto front was obtained from all the multi-objective optimization combinations, illustrating a set of feasible solutions and the trade-offs between the two objectives.

C. *Future perspective*

To tackle the multi-scale challenge inherent in supply chain modeling, future advances will focus on the use of surrogate models and the development of novel optimization algorithms, to efficiently solve these new larger scale and more interconnected models. Future developments also include systems engineering frameworks that integrate multi-level or multi-agent optimization, as the presence of multiple interconnected stakeholders introduces significant challenges in supply chain modeling. Accurately capturing real-world interactions and incorporating these considerations is essential for creating robust and realistic models [20]. Therefore, game theory-based approaches are needed, that can represent the multiple stakeholders as autonomous agents interacting in the same or at different hierarchical levels.

Regarding the assessment methods, more comprehensive environmental metrics are expected to be adopted, as relying solely on emissions has proven insufficient for capturing the full range of environmental impacts. Therefore, circularity metrics that encompass resource consumption, waste generation and product durability will be key for the assessments.

IV. Sustainable Land Management

Land-use transformations have the potential to supplement the technological systems like CO_2 capture to help meet corporate sustainability goals and greenhouse gas emission targets. In addition, they support several water and nature conservation initiatives. These land-use transformations include reforestation, restoration, remediation, wetland, pond construction, etc., each capable of providing natural-capital and environmental benefits through ecosystem services [36]. Since corporate entities own numerous land properties and there are several possible land-use transformations, it is impossible to manually analyze each solution. It is also a challenge to obtain the appropriate ecological data from various sources to define a robust corporate decision-making strategy backed by scientific methods under internal degrees of freedom and constraints [37]. Systems engineering, mathematical optimization and data science approaches can effectively explore all the feasible solutions [38], prioritize and recommend cost-favorable land management strategies.

In this section, we will present a digital PSE toolkit for sustainable land management that can help in making informed decisions, which exemplifies how these approached can help an organization progress towards their environmental goals and corporate social responsibilities. We have developed a unique methodology to leverage geographic information system (GIS) capabilities in a mathematical optimization setting using ArcGIS Pro (for spatial analytics) and AIMMS Pro (for programming and application deployment). The developed tool that has the capabilities to i) analyze environmental risk at global locations and identify priority sites for implementing nature-based solutions; and to ii) optimize land-use transformation strategies at a selected site for maximizing value from ecosystem services. Examples of environmental risk metrics considered include water stress indicators from Aqueduct 4.1 [39], biodiversity hotspots from integrated Biodiversity Assessment Tool (iBAT), air quality measurements, and 2040 temperature projections. On a site level we leverage county level land parcel data and national land cover database (NLCD) [40] to characterize the land types and identify the best land transformation strategies using a mathematical optimization formulation. These objectives are accomplished through a multi-stage procedure which is depicted in figure 9.

With the help of this digital toolkit, one can prioritize the selection of properties for climate action based on environmental risk hotspots, such as water stress, bio-diversity risk, air/water pollution, etc. This also involves bucketing the sites under similar risk categories for multiple risk indicators through multi-variate statistical measures such as hierarchical clustering on principal components (HCPC) [41], as shown using an example in figure 10. This risk assessment methodology is deployed through interactive views on the global map to visualize the sites, their rankings for environmental risk and categories.

The Land-use change optimization toolkit (second level in figure 9 requires a more elaborate procedure to first estimate the baseline Ecosystem service 'value' at currently owned the land parcels (areas owned by a single entity) at selected sites. This step entails collecting the land ownership data from internal sources and county tax records, overlaying them onto National Land Cover databases to identify the fractions of different land categories such as Shrubs, Grasslands, Deciduous or Evergreen forests, Developments, Open Water, etc. in each land parcel. The area of each land cove category in a parcel is then used to estimate the ecosystem service value it delivers through a surrogate linear model fitted on the ecological assessments for individual land parcels in the Ecosystem Services Identification & Inventory (ESII) Tool [42]. The ESII tool is developed by The Nature Conservancy in partnership with Dow and Eco Metrix solutions group [43] and has been used to effectively generate information on the ecosystem service performance of a specified landscape requiring field data collections. Once the baseline land cover information for company-owned land parcels is obtained, the tool identifies improvement opportunities to optimize environmental and economic benefit of the property (greenbelt, surplus, operating, and remediation) under budgetary constraints that can be targeted towards: a. Increasing carbon sequestration b. Improving water quality and quantity c. Enhancing air quality (kgNOx or Particulate matter) d. Promoting biodiversity growth through green cover This is achieved through a constrained Knapsack optimization formulation that solves for optimal transitions in land parcels (wherever possible, user defined) to general maximum value from these ecosystem services (multi-objective) at lowest cost. The deployed web-application allows for selection of weights to different ecosystem service values and also generate Pareto-optimal solutions for user-selected objectives. Figure 11 demonstrates the workflow of the land-use change optimization toolkit. Finally, to analyze remediation options at the level of individual land parcels, the toolkit also provides ways to perform hypothesis testing using time-series datasets on land cover images. For example, in a sparsely forested scrubland, a hypothesis about higher natural reforestation where green corridors exists, is studied to help remediation teams determine which land parcels to focus on for low-cost remediation. Several such analyses on soil type, regional water availability risk forecasts, etc. can be conducted on the developed toolkit for the land parcel of interest.

The proposed digital solution demonstrates the proactiveness to align with Taskforce on Nature-related Financial Disclosure's (TNFD) Locate, Evaluate, Assess and Prepare (LEAP) approach [44] that is meant to guide organizations assess nature-related issues, from dependencies and impacts to risks and opportunities. The mathematical frameworks and PSE approaches adopted in this digital toolkit have the potential to provide model-guided strategies and actions to meet corporate goals and targets in the Water and Nature space more effectively.

V. The Role of Advanced Control Systems in Meeting Sustainability Goals

Real-time decision-making and control are fundamental to achieving sustainability objectives across diverse sectors, including industrial manufacturing, transportation, and energy systems. These systems operate in highly dynamic environments, where external factors (such as weather conditions, fluctuating demands, and evolving market constraints) must be effectively managed. Traditional setpoint tracking and rule-based control strategies are often inadequate in such contexts, as they lack the adaptability required to respond to rapid variations and leverage opportunities for improving efficiency. In contrast, *model-based* and *optimization-driven* advanced control frameworks empower decision-makers to proactively adjust operations by systematically incorporating forecasts of external variables, enforcing operational constraints, and balancing conflicting objectives such as cost, energy efficiency, and system performance.

Among the many applications of advanced control, *building energy management* serves as an illustrative example of both the challenges and opportunities in real-time decision-making for sustainability. Buildings account for approximately 40% of global energy consumption and 35% of carbon emissions [45], largely due to their reliance on energy-intensive heating, ventilation, and air conditioning (HVAC) systems. As the power sector transitions toward cleaner energy sources, buildings will play an increasingly critical role in global decarbonization efforts. However, uncontrolled or inefficiently managed building operations contribute to significant energy waste. Studies suggest that up to 30% of HVAC energy is wasted due to suboptimal control policies, ultimately hindering progress toward net-zero emissions goals [46], [47].

A key challenge in modern building energy management is achieving a balance between occupant comfort (often quantified through temperature, humidity, and air quality constraints) and energy efficiency. This trade-off is complex due to the influence of external factors such as ambient

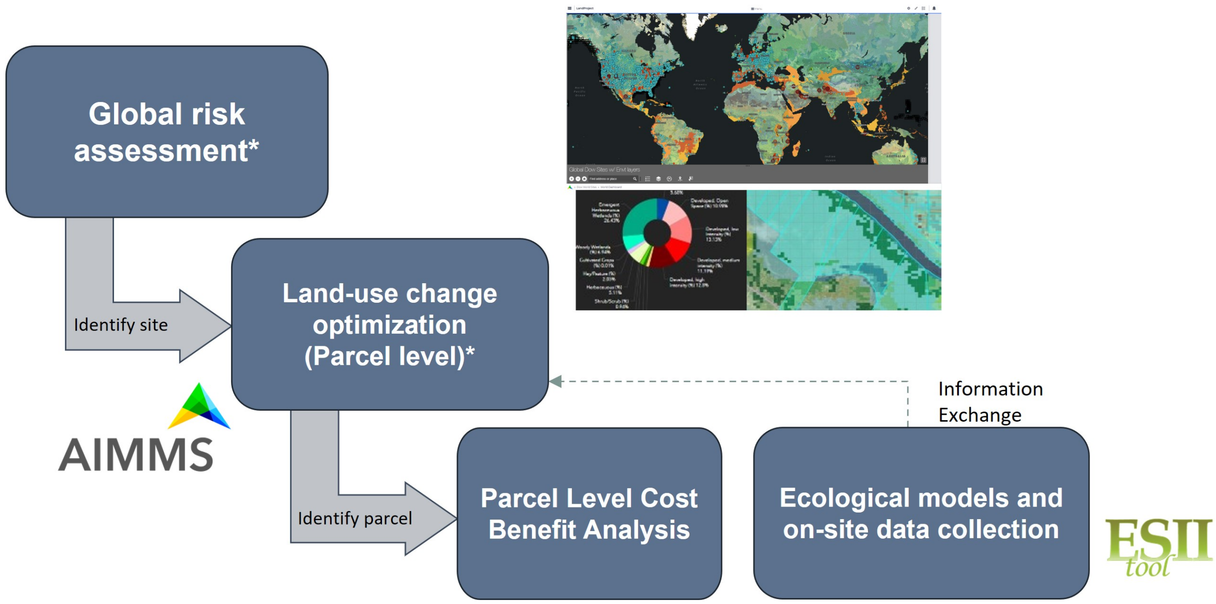

Fig. 9: Multi-stage procedure for land-asset management, deployed as web application on company servers with dashboards to visualize and interpret environmental risk, and identify optimal land-use transformation strategies.

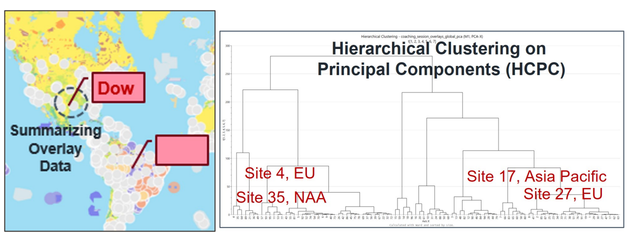

Fig. 10: Prioritization of sites for remediation action is a key step for capital allocation. This approach demonstrates how open-source environmental risk datasets and statistical approaches can be leveraged to aid in decision making.

weather, dynamic occupant behavior, and time-varying energy prices. Model predictive control (MPC) has emerged as a powerful tool for addressing this challenge, given its ability to optimize multivariable systems while enforcing constraints [48]. MPC leverages a predictive model of the building to determine optimal control actions over a receding time horizon, allowing energy consumption to be minimized without violating comfort constraints.

Despite its promise, deploying MPC (or its variants) in real-world building environments poses two key challenges. First, predictive models must account for disturbances such as occupant-driven heat loads, internal appliance usage, and weather variations. These disturbances are inherently stochastic, and failing to capture them accurately can lead to constraint violations and degraded control performance. Second, standard MPC approaches often rely on a limited set of predefined disturbance scenarios, which may not capture the full variability of real-world conditions. With the increasing availability of sensor data from modern buildings, recent research has explored the use of *deep generative models* to synthesize realistic time-series disturbances, thereby improving the robustness of control strategies.

This section highlights recent advancements in integrating *stochastic* MPC (SMPC) with data-driven scenario generation for building energy management. Specifically, we summarize our recent contributions in the following areas:

1) **Generative Time-Series Modeling for Building Disturbances**: We have developed probabilistic deep generative models, trained on real building datasets, to capture stochastic disturbances such as occupant-driven heat loads and environmental effects [49].
2) **Incorporating Generative Models into SMPC Architectures**: We have proposed a novel approach to integrate generative time-series models into SMPC, ensuring improved control performance while maintaining computational efficiency [50].
3) **Sampling Rare but Influential Scenarios via Active Learning**: We have developed a sample-efficient active learning strategy to identify *limiting scenarios*, i.e., extreme conditions under which SMPC may approach performance boundaries [51]–[53].

By examining the case of building energy control, we demonstrate how advanced control systems – enabled by high-fidelity modeling, probabilistic forecasting, and real-time optimization – can contribute to achieving sustainability goals. Similar methodologies can be extended to other domains requiring high adaptivity under uncertainty, reinforcing the broader role of advanced control in sustainable systems engineering. For a broader discussion of integrating machine learning into MPC, see [54].

A. Generative time-series disturbance modeling

Effective control policies require accurate models of system disturbances. However, disturbances in building energy

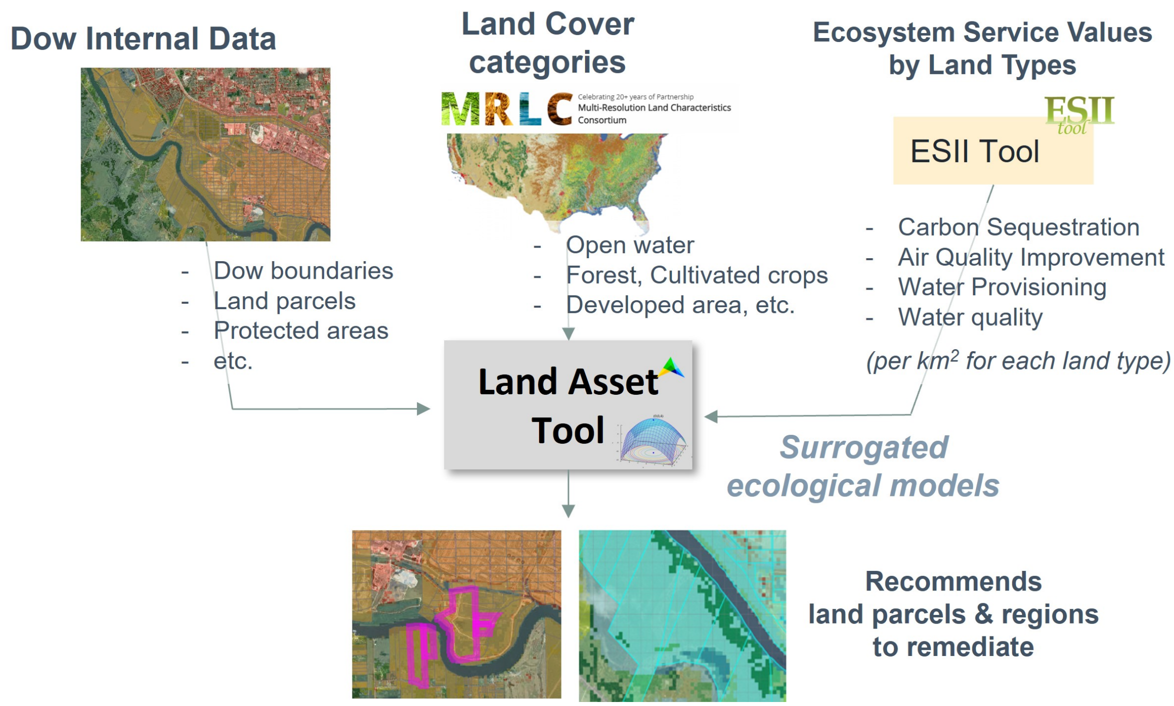

Fig. 11: Workflow of the land-use change optimization toolkit.

systems (e.g., internal heat loads, appliance usage, and ventilation demands) are inherently stochastic and difficult to model using first-principles approaches, as they depend on human occupancy patterns that are highly uncertain. These disturbances exhibit complex temporal dependencies, making traditional deterministic models insufficient for capturing real-world variability.

Recent advances in *probabilistic deep generative models* offer a promising solution by enabling the synthesis of realistic time-series disturbances from historical building data. One such approach is the Regularized Adversarial Fine-Tuned Variational Autoencoder-GAN (RAFT-VG) architecture [49], which has been trained on the real-world SUSTIE building dataset [55]. As shown in Figure 12, RAFT-VG combines the structural advantages of Variational Autoencoders (VAEs) with the high-fidelity generative capabilities of Generative Adversarial Networks (GANs). The model training consists of two primary stages: first, a VAE learns a latent representation of the input time-series data, capturing its statistical structure. In the second stage, adversarial fine-tuning refines the latent space by employing a discriminator network that enhances the realism of generated samples. Additionally, RAFT-VG employs regularization techniques to ensure consistency between the encoder and decoder, preserving the integrity of synthesized data distributions.

The ability to generate realistic disturbance scenarios has several advantages: (i) it enables robust *scenario-based optimization* within SMPC; (ii) it provides synthetic data for training and testing controllers under diverse conditions; and (iii) it enhances privacy preservation by allowing synthetic data to replace sensitive building operation records.

B. Incorporating generative models into SMPC

A fundamental advantage of SMPC is its ability to optimize control policies while explicitly accounting for uncertainty in disturbances [56], [57]. However, the effectiveness of SMPC depends on the quality of the disturbance models used in scenario generation. Conventional approaches often rely on static or nominal disturbance distributions, which fail to capture real-world variations, leading to overly conservative or suboptimal control policies.

By integrating time-series disturbance forecasts made by generative AI methods, such as RAFT-VG, into the SMPC framework, we can more effectively balance risk and reward – ensuring that control actions are neither overly aggressive (which may lead to downstream constraint violations) nor overly conservative (which may waste energy and reduce performance). In our recent work [50], we propose a computationally tractable SMPC method that utilizes a conditional VAE for disturbance generation (another type of generative AI method). The key idea is to generate a large number of realistic scenarios that are incorporated into a so-called *scenario tree*, a structured representation of uncertainty commonly used in stochastic programming [58]. Instead of constructing an exhaustive scenario tree (where disturbances can take on many finite realizations at each time step, leading to exponential growth in problem size), we strategically restrict branching to the first stage. This approach captures the most relevant uncertainty while keeping the optimization computationally feasible. Figure 13 illustrates the impact of generative model-informed SMPC compared to classical SMPC approaches. The adaptive strategy achieves performance comparable to a perfect (but unachievable) forecast

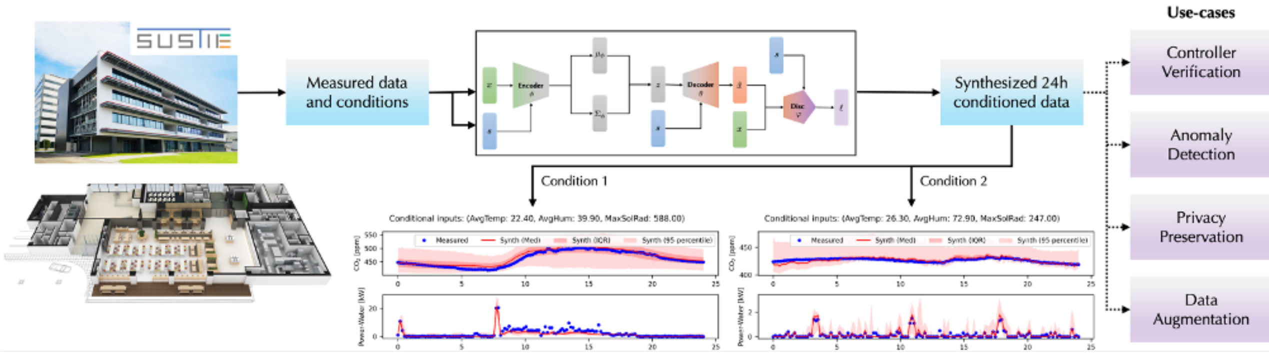

Fig. 12: Overview of the RAFT-VG generative model, trained on the SUSTIE dataset to generate realistic building disturbance scenarios. The synthesized data supports multiple applications, including controller verification, anomaly detection, privacy preservation, and data augmentation.

while significantly outperforming non-adaptive methods.

The rationale behind restricting scenario branching at the first stage can be understood through the lens of dynamic programming (DP) [59], [60]. MPC can be interpreted as performing a single step of Newton's method on the Bellman equation, using an approximate value function whose accuracy depends on the horizon length. A key observation from this perspective is that the first-stage decision must be computed as accurately as possible, as it determines the initial Newton step, while subsequent steps can be approximated with lower fidelity since they primarily serve to refine the initial decision. This principle suggests that, while exact disturbance modeling is critical in the near term, approximate representations can suffice in later stages, allowing for computational savings without significant loss in control performance.

This observation also aligns with the *exponential decay of sensitivity (EDS)* property [61], [62], which states that the impact of parametric perturbations in a dynamic system decays exponentially as one moves backward in time. From an SMPC perspective, this implies that uncertainties far in the future have a negligible impact on present decisions, reinforcing the validity of using high-fidelity disturbance models only in early-stage scenario generation. Interestingly, a similar approximation strategy has been shown to improve decision-making in competitive settings such as chess [63], where a reduced branching factor in early moves (analogous to first-stage scenario tree refinement) leads to significantly stronger gameplay without incurring excessive computational overhead. By adopting this structured approach to scenario tree design, generative model-informed SMPC can maintain a practical computational footprint while still capturing the essential variability needed for robust real-time control.

C. Efficiently sampling rare but influential scenarios

While optimizing control performance under typical conditions is essential, it is equally critical to evaluate controllers against rare but high-impact scenarios that may expose vulnerabilities in the system. These edge cases, often overlooked in standard scenario-based optimization, provide valuable insights into the robustness of control strategies and help identify potential failure modes. However, systematically discovering such extreme conditions is computationally challenging due to the combinatorial nature of scenario exploration. Since closed-loop simulations (especially those using high-fidelity physics-based models) can be computationally expensive, sample-efficient methods are needed to uncover limiting scenarios within a practical computational budget.

Sequential active learning is a powerful framework for systematically identifying new simulation conditions that maximize information gain. Active learning is closely related to optimal experiment design (OED) [64], where the objective is to maximize knowledge about an underlying system using the fewest possible experiments. In the context of control robustness assessment, this translates to selecting the most informative disturbance scenarios to probe system behavior under extreme but realistic conditions.

One approach for efficiently identifying such scenarios is InfoBAX, an information-theoretic Bayesian algorithm execution framework [65]. InfoBAX is a general-purpose methodology applicable to a variety of decision-making tasks, including optimization, model validation, and robustness assessment (effectively generalizes Bayesian optimization beyond global optimization). In recent work, we demonstrated how InfoBAX can efficiently discover limiting scenarios for SMPC architectures [51]. A key advantage of InfoBAX is its use of mutual information criteria to select scenarios that provide maximal insight into system vulnerabilities. Rather than relying on brute-force Monte Carlo sampling, which often requires an infeasibly large number of simulations, InfoBAX refines its scenario selection over successive iterations, focusing on cases that are most informative for assessing worst-case performance. By prioritizing extreme conditions, such as rare weather anomalies, unexpected occupancy surges, or HVAC system faults, InfoBAX enables efficient controller verification while significantly reducing computational costs.

Figure 14 illustrates the effectiveness of InfoBAX in extracting high-impact scenarios. Compared to naive exhaustive search strategies, InfoBAX achieves the same level of insight into system performance while using 35x fewer simulations

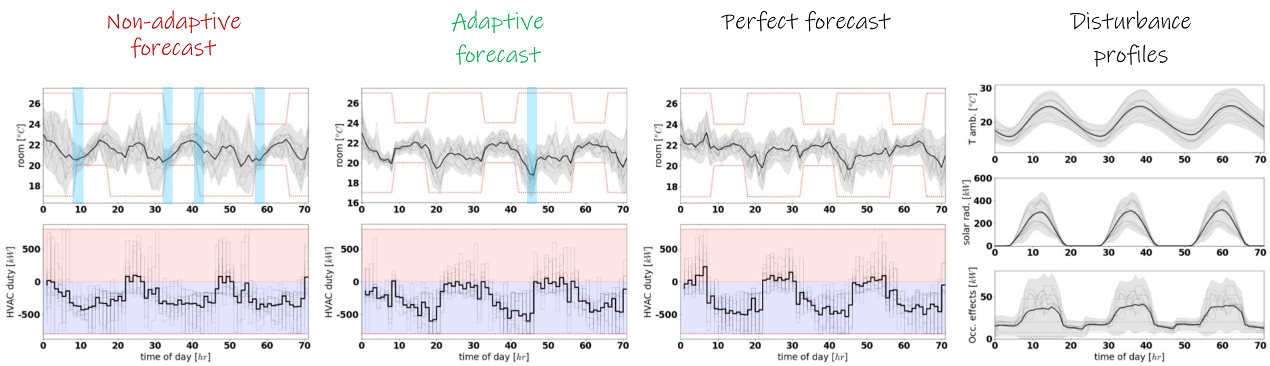

Fig. 13: Comparison of classical SMPC, generative model-informed SMPC, and MPC with a perfect forecast. Blue bands represent periods of constraint violation. The proposed generative model-enhanced approach significantly reduces energy consumption while maintaining comfort constraints.

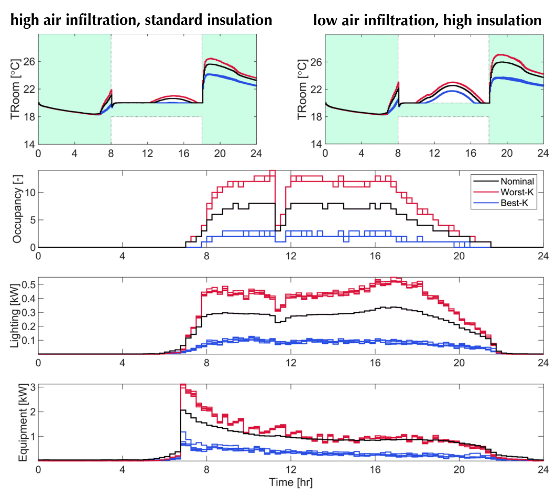

Fig. 14: Evaluation of InfoBAX-based scenario selection for testing robustness of a tuned SMPC controller. The method identifies extreme but realistic cases that lead to the best and worst closed-loop performance outcomes, providing an efficient strategy for controller verification.

than brute-force methods. The experiments were conducted using Modelica [66], a high-fidelity modeling framework widely employed for simulating building energy systems. The ability to efficiently identify extreme scenarios allows for targeted controller refinements, improving resilience to unexpected disturbances without incurring prohibitive computational costs.

D. *Perspectives for future research*

The integration of generative AI with physics-based modeling and advanced control strategies holds significant potential for driving sustainability in building energy management and beyond. However, several key challenges and open research directions remain, particularly in ensuring that these advanced methods transition from theoretical development to real-world impact. One of the most critical challenges is the deployment and large-scale validation of these approaches in real building infrastructure. While generative AI and SMPC methods have demonstrated impressive results in simulation, their adoption in industry requires rigorous testing under real operating conditions. This necessitates collaborations between researchers, industry stakeholders, and policymakers to establish standardized benchmarks, develop open-access datasets, and conduct pilot studies that assess long-term performance, robustness, and scalability.

A major hurdle in real-world deployment is handling *non-stationarity*, where the statistical properties of building disturbances (such as occupant behavior, energy demand, and environmental conditions) evolve over time. Current generative models often rely on training data that may become outdated as buildings undergo operational changes. Continual learning strategies, potentially inspired by reinforcement learning (RL) and other online adaptation techniques, will be crucial for ensuring that generative models remain relevant over extended periods. One possible avenue is the development of self-updating generative models that integrate streaming data in real time while preserving past knowledge through techniques like experience replay or meta-learning.

The evolution of generative AI architectures also presents a promising direction for future research. Transformer-based time series models [67]–[69] and large-scale foundation models [70], [71] trained on diverse energy and environmental datasets could provide a more generalizable approach to disturbance modeling. These architectures have shown remarkable success in capturing long-range dependencies in sequential data and could enhance the ability to model complex interactions between weather patterns, energy consumption, and occupant-driven loads. However, the question remains: what types of disturbances are best handled by different generative AI architectures? A systematic investigation into the strengths and weaknesses of various models in different operational settings will be essential for guiding their adoption in real-world control applications.

Beyond generative modeling, advancements in SMPC formulation and computational efficiency are also needed to facilitate real-time implementation. The challenge of balancing computational tractability with control performance is particularly critical when deploying SMPC on embedded or cloud-based control systems. Since disturbance forecasting and SMPC, respectively, involve real-time inference and optimization, algorithmic efficiency must be carefully considered. Leveraging hardware-aware optimization techniques, such as pruning and quantization for neural network-based forecasting models, can help ensure that these approaches operate within the constraints of real-world control hardware. Additionally, automated code generation methods for embedded optimization have become increasingly popular, enabling high-performance control execution with minimal manual tuning. A recent example of this is the work in [72], which co-designs (S)MPC policies alongside its hardware implementation to ensure both software and hardware constraints are simultaneously considered. Such co-design strategies will likely become more important in the future, particularly as energy-efficient computing becomes a greater concern in sustainability-driven applications.

Another valuable research direction is the integration of *distributed control* architectures for multi-building or grid-interactive systems, e.g., [73]–[75]. While most advanced control strategies focus on optimizing individual buildings, a more holistic approach would involve coordinating multiple buildings in a shared energy ecosystem. This could include strategies for demand response, real-time grid integration, and cooperative energy storage management, where generative models provide uncertainty-aware forecasts for distributed control policies. The combination of federated learning (e.g., [76], [77]) with SMPC could allow for decentralized training of generative models without compromising privacy, making it possible for buildings to share insights on energy demand patterns without directly exchanging sensitive operational data.

Finally, an important challenge that extends beyond technical considerations is the broader adoption of these methods in policy and regulatory frameworks. Many building control systems are subject to strict regulations, which may pose barriers to the deployment of AI-driven decision-making. Future research should explore ways to incorporate explainability and interpretability [78]–[81] into advanced control strategies, ensuring that generative model-based forecasting and SMPC decisions can be audited, understood, and trusted by human operators. Bridging the gap between cutting-edge research and practical implementation will require interdisciplinary efforts that involve not only control theorists and AI researchers but also experts in economics, policy, and human-computer interaction.

VI. Conclusion

Process System Engineering methodologies are critical to achieve any of the sustainability objectives companies and societies are pursuing while maintaining and facilitating economic growth. The complexity of the systems and their interactions require mathematical models to enable informed decision making.

In this paper, we have reviewed and illustrated how PSE contributes to four critical domains: carbon monetization and low-carbon supply chains, circular economy and sustainable manufacturing, sustainable land management, and advanced control technology. Through these examples, we highlight that each domain presents unique challenges and opportunities for implementing PSE strategies but the insights related to resource allocation optimization, environmental impact minimization, and economic viability assurance are critical for our sustainability journey.

In low-carbon supply chains, PSE plays a pivotal role in designing strategies for carbon monetization, allowing companies to navigate regulatory landscapes and leverage financial incentives for emission reductions. The case study on low-carbon ammonia supply chains illustrates how different chain of custody models, such as Mass Balance and Book-and-Claim, impact investment decisions and operational efficiency. These insights underscore the importance of selecting appropriate certification systems to balance cost, carbon intensity, and demand satisfaction.

For circular supply chains, the integration of circular economy principles with PSE methodologies facilitates the optimization of resource use and waste reduction. The case study on food packaging waste management highlights the use of systems engineering frameworks to evaluate trade-offs between economic, environmental, and circularity metrics. The results emphasize the necessity of multi-objective optimization and comprehensive assessment tools to identify sustainable pathways.

In the domain of sustainable land management, PSE approaches are essential for quantifying the trade-offs between land use, emissions, and economic feasibility. By employing system-level modeling and optimization, decision-makers can develop strategies that align with ecological and economic goals, such as carbon sequestration and biodiversity preservation.

Finally, advanced control systems, particularly in building energy management, demonstrate the potential of PSE methodologies to enhance energy efficiency and reduce emissions. The integration of generative AI with stochastic model predictive control (SMPC) provides robust solutions for managing uncertainties in real-time operations. These advancements highlight the importance of combining high-fidelity modeling with data-driven forecasting to achieve optimal control performance.

The advancements in computational capabilities along with generative AI and machine learning will continue to compliment systems thinking and modeling, and enable the optimization of more complex interactions which is needed for better decision-making to minimize the environmental impact of human production. Future research should focus on addressing challenges related to computational efficiency, non-stationarity, and regulatory frameworks to facilitate the real-world deployment of these advanced methodologies. By fostering interdisciplinary collaboration, PSE can serve as a catalyst for sustainable transformation across industries, contributing to global efforts to combat climate change and promote environmental stewardship.

REFERENCES

[1] Schwarz, Alan, "Forbes net-zero leaders 2024." [Online]. Available: https://www.forbes.com/lists/net-zero-leaders/

[2] McKinsey & Company, "What is net zero?" [Online]. Available: https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-net-zero

[3] M. Beckmann, S. Hielscher, and I. Pies, "Commitment strategies for sustainability: How business firms can transform trade-offs into win–win outcomes," *Business Strategy and the Environment*, vol. 23, no. 1, pp. 18–37, 2014.

[4] C. A. Van der Byl and N. Slawinski, "Embracing tensions in corporate sustainability: A review of research from win-wins and trade-offs to paradoxes and beyond," *Organization & Environment*, vol. 28, no. 1, pp. 54–79, 2015.

[5] U. Shah and B. R. Bakshi, "Accounting for nature's intermittency and growth while mitigating NO_2 emissions by technoecological synergistic design—application to a chloralkali process," *Journal of Advanced Manufacturing and Processing*, vol. 1, no. 1-2, p. e10013, 2019.

[6] U. Shah, J. A. Paulson, and B. R. Bakshi, "Real-time synergies between homeostatic technological and homeorhetic ecological systems by multiscale MPC and Bayesian optimization," *Industrial & Engineering Chemistry Research*, vol. 63, no. 49, pp. 21 389–21 403, 2024.

[7] E. N. Pistikopoulos, A. Barbosa-Povoa, J. H. Lee, R. Misener, A. Mitsos, G. V. Reklaitis, V. Venkatasubramanian, F. You, and R. Gani, "Process systems engineering – the generation next?" *Computers & Chemical Engineering*, vol. 147, p. 107252, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0098135421000302

[8] European Commission, "EU emissions trading system (EU ETS)," 2025, accessed: 2025-02-16. [Online]. Available: https://climate.ec.europa.eu/eu-action/eu-emissions-trading-system-eu-ets/about-eu-ets_en

[9] California Air Resources Board, "Cap-and-trade program," 2025, accessed: 2025-02-16. [Online]. Available: https://ww2.arb.ca.gov/our-work/programs/cap-and-trade-program

[10] International Carbon Action Partnership (ICAP), "China national emissions trading system (ETS)," 2025, accessed: 2025-02-16. [Online]. Available: https://icapcarbonaction.com/en/ets/china-national-ets

[11] S&P Global, "Internal carbon pricing: Stress testing for climate risk," 2025, accessed: 2025-02-16. [Online]. Available: https://www.spglobal.com/spdji/en/documents/education/talking-points-internal-carbon-pricing-stress-testing-for-climate-risk.pdf

[12] Verra, "Verified Carbon Standard (VCS)," 2025, accessed: 2025-02-16. [Online]. Available: https://verra.org/programs/verified-carbon-standard/

[13] ACRC Carbon, "ACRC Carbon," 2025, accessed: 2025-02-16. [Online]. Available: https://acrcarbon.org

[14] T. Rathi and Q. Zhang, "Capacity planning with uncertain endogenous technology learning," *Computers & Chemical Engineering*, vol. 164, p. 107868, 2022.

[15] M. A. McDonald and C. T. Maravelias, "Sequential decision-making under uncertainty for long-term energy transition planning," *iScience*, vol. 27, no. 12, 2024.

[16] H. Wang, M. Palys, P. Daoutidis, and Q. Zhang, "Optimal design of sustainable ammonia-based food–energy–water systems with nitrogen management," *ACS Sustainable Chemistry & Engineering*, vol. 9, no. 7, pp. 2816–2834, 2021.

[17] H. Wang, P. Daoutidis, and Q. Zhang, "Ammonia-based green corridors for sustainable maritime transportation," *Digital Chemical Engineering*, vol. 6, p. 100082, 2023.

[18] T. Rathi, C. Brzakala, H. Wang, and Q. Zhang, "Assessing the impact of chain of custody models on the long-term planning of low-carbon ammonia supply chains," *Industrial & Engineering Chemistry Research*, vol. 64, no. 3, pp. 1680–1699, 2025.

[19] ISO, "ISO 22095:2020 Chain of Custody – General Terminology and Models," Standard, International Organization for Standardization, Geneva, Switzerland, 2020. [Online]. Available: https://www.iso.org/standard/72532.html

[20] S. Avraamidou, S. G. Baratsas, Y. Tian, and E. N. Pistikopoulos, "Circular economy - a challenge and an opportunity for process systems engineering," *Computers & Chemical Engineering*, vol. 133, p. 106629, Feb 2020.

[21] R. Accorsi, A. Gallo, and R. Manzini, "A climate driven decision-support model for the distribution of perishable products," *Journal of Cleaner Production*, vol. 165, pp. 917–929, 2017.

[22] R. Zhou, M. Jin, Z. Li, Y. Xiao, D. McCollum, and A. Li, "Techno-economic analysis and network design for co2 conversion to jet fuels in the united states," *Renewable and Sustainable Energy Reviews*, vol. 210, p. 115191, 2025.

[23] M. S. Akhtar, R. Dickson, and J. J. Liu, "Life cycle assessment of inland green hydrogen supply chain networks with current challenges and future prospects," *ACS Sustainable Chemistry & Engineering*, vol. 9, no. 50, pp. 17 152–17 163, 2021.

[24] S. G. Baratsas, E. N. Pistikopoulos, and S. Avraamidou, "A quantitative and holistic circular economy assessment framework at the micro level," *Computers and Chemical Engineering*, vol. 160, 4 2022.

[25] E. Pauer, B. Wohner, V. Heinrich, and M. Tacker, "Assessing the environmental sustainability of food packaging: An extended life cycle assessment including packaging-related food losses and waste and circularity assessment," *Sustainability (Switzerland)*, vol. 11, 2 2019.

[26] E. Chialdikas, A. d. C. Munguıa-López, H. Aguirre-Villegas, and S. Avraamidou, "A framework for the evaluation of the circularity of plastic waste management systems: A case study on mechanical recycling of hdpe," 2023.

[27] S. G. Baratsas, N. Masoud, V. A. Pappa, E. N. Pistikopoulos, and S. Avraamidou, "Towards a circular economy calculator for measuring the "circularity" of companies," in *Computer Aided Chemical Engineering*. Elsevier, 2021, vol. 50, pp. 1547–1552.

[28] K. J. Martinez, P. A. Munoz-Briones, and S. Avraamidou, "Circular economy assessment of viable for fuels for mobility," in *Computer Aided Chemical Engineering*. Elsevier, 2024, vol. 53, pp. 2335–2340.

[29] P. A. Munoz-Briones, A. del C. Munguia-Lopez, K. L. Sanchez-Rivera, V. M. Zavala, G. W. Huber, and S. Avraamidou, "Optimal design of food packaging considering waste management technologies to achieve circular economy," in *Foundations of Computer Aided Process Design*, 7 2024, pp. 820–828.

[30] L. Ng, N. Chemmangattuvalappil, V. Dev, and M. Eden, "Chapter 1 - mathematical principles of chemical product design and strategies," in *Tools For Chemical Product Design*, ser. Computer Aided Chemical Engineering, M. Martín, M. R. Eden, and N. G. Chemmangattuvalappil, Eds. Elsevier, 2016, vol. 39, pp. 3–43. [Online]. Available: https://www.sciencedirect.com/science/article/pii/B9780444636836000010

[31] M. Ehrgott, "A discussion of scalarization techniques for multiple objective integer programming," *Annals of Operations Research*, vol. 147, pp. 343–360, 10 2006.

[32] G. Mavrotas, "Effective implementation of the ϵ-constraint method in multi-objective mathematical programming problems," *Applied Mathematics and Computation*, vol. 213, pp. 455–465, 7 2009.

[33] S. Validi, A. Bhattacharya, and P. J. Byrne, "A case analysis of a sustainable food supply chain distribution system - a multi-objective approach," *International Journal of Production Economics*, vol. 152, pp. 71–87, 2014.

[34] S. G. Baratsas, E. N. Pistikopoulos, and S. Avraamidou, "A systems engineering framework for the optimization of food supply chains under circular economy considerations," *Science of the Total Environment*, vol. 794, 11 2021.

[35] S. Karayılan, Özge Yılmaz, Çınar Uysal, and S. Naneci, "Prospective evaluation of circular economy practices within plastic packaging value chain through optimization of life cycle impacts and circularity," *Resources, Conservation and Recycling*, vol. 173, 10 2021.

[36] S. Searchinger, T. Wirsenius, T. Beringer, and P. Dumas, "Assessing the efficiency of changes in land use for mitigating climate change," *Nature*, vol. 564, no. 7735, pp. 249–253, 2018.

[37] D. Müller and D. K. Munroe, "Current and future challenges in land-use science," *Journal of Land Use Science*, vol. 9, no. 2, pp. 133–142, 2014.

[38] B. R. Bakshi and J. Fiksel, "The quest for sustainability: Challenges for process systems engineering," *AIChE Journal*, vol. 49, no. 6, pp. 1350–1358, 2003.

[39] P. Reig, T. Shiao, and F. Gassert, "Aqueduct water risk framework," World Resources Institute, Tech. Rep., 2013.

[40] C. G. Homer, J. A. Fry, and C. A. Barnes, "The national land cover database," US Geological Survey, Tech. Rep., 2012.

[41] M. Argüelles, C. Benavides, and I. Fernández, "A new approach to the identification of regional clusters: hierarchical clustering on principal

components," *Applied Economics*, vol. 46, no. 21, pp. 2511–2519, 2014.

[42] ESII Tool, "Dow case study: Ecosystem services identification & inventory (esii) tool," https://www.esiitool.com/case-studies/, accessed: 2024-10-17.

[43] The Dow Chemical Company, "Dow and the nature conservancy celebrate 10 years of groundbreaking collaboration to value nature in business strategy," 2021. [Online]. Available: https://investors.dow.com/en/news/news-details/2021/Dow-and-The-Nature-Conservancy-celebrate-10-years-of-groundbreaking-collaboration-to-value-nature-in-business-strategy

[44] Sustainability in Business, "The leap framework for nature-related financial disclosures," https://www.sustainabilityinbusiness.org/blogs/the-leap-framework-for-nature-related-financial-disclosures, 2023.

[45] NREL, "NREL Researchers Reveal How Buildings Across the United States Can–And Could–Use Energy," 2023, accessed: 2025-02-21.

[46] K. H. Lee and J. E. Braun, "Model-based demand-limiting control of building thermal mass," *Building and Environment*, vol. 46, no. 3, pp. 428–436, 2011.

[47] International Energy Agency, "Buildings – Breakthrough Agenda Report 2023," 2023, accessed: 2025-02-21. [Online]. Available: https://www.iea.org/reports/breakthrough-agenda-report-2023/buildings

[48] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model predictive control: Theory, computation, and design*. Nob Hill Publishing Madison, WI, 2017, vol. 2.

[49] A. Salatiello, Y. Wang, G. Wichern, T. Koike-Akino, Y. Ohta, Y. Kaneko, C. Laughman, and A. Chakrabarty, "Synthesizing building operation data with generative models: Vaes, gans, or something in between?" in *Companion Proceedings of the 14th ACM International Conference on Future Energy Systems*, 2023, pp. 125–133.

[50] F. Sorouifar, J. A. Paulson, Y. Wang, R. Quirynen, C. R. Laughman, and A. Chakrabarty, "Bayesian forecasting with deep generative disturbance models in stochastic MPC for building energy systems," in *IEEE Conference on Control Technology and Applications*. IEEE, 2024, pp. 414–419.

[51] A. Chakrabarty, L. Vanfretti, W.-T. Tang, J. A. Paulson, S. Zhan, S. A. Bortoff, V. Deshpande, Y. Wang, and C. R. Laughman, "Assessing building control performance using physics-based simulation models and deep generative networks," in *IEEE Conference on Control Technology and Applications*. IEEE, 2024, pp. 547–554.

[52] W.-T. Tang, A. Chakrabarty, and J. A. Paulson, "BEACON: A Bayesian optimization strategy for novelty search in expensive black-box systems," *arXiv preprint arXiv:2406.03616*, 2024.

[53] W.-T. Tang, A. Chakrabarty, and J. Paulson, "TR-BEACON: Shedding light on efficient behavior discovery in high-dimensional spaces with Bayesian novelty search over trust regions," in *NeurIPS 2024 Workshop on Bayesian Decision-making and Uncertainty*.

[54] A. Mesbah, K. P. Wabersich, A. P. Schoellig, M. N. Zeilinger, S. Lucia, T. A. Badgwell, and J. A. Paulson, "Fusion of machine learning and mpc under uncertainty: What advances are on the horizon?" in *In Proceedings of the American Control Conference*. IEEE, 2022, pp. 342–357.

[55] M. Electric, "SUSTIE," accessed: 2025-02-21. [Online]. Available: https://www.mitsubishielectric.com/en/about/rd/sustie/index.html

[56] T. A. N. Heirung, J. A. Paulson, J. O'Leary, and A. Mesbah, "Stochastic model predictive control—how does it work?" *Computers & Chemical Engineering*, vol. 114, pp. 158–170, 2018.

[57] J. A. Paulson, E. A. Buehler, R. D. Braatz, and A. Mesbah, "Stochastic model predictive control with joint chance constraints," *International Journal of Control*, vol. 93, no. 1, pp. 126–139, 2020.

[58] D. M. de la Penad, A. Bemporad, and T. Alamo, "Stochastic programming applied to model predictive control," in *Proceedings of the IEEE Conference on Decision and Control*. IEEE, 2005, pp. 1361–1366.

[59] D. Bertsekas, "Newton's method for reinforcement learning and model predictive control," *Results in Control and Optimization*, vol. 7, p. 100121, 2022.

[60] D. P. Bertsekas, "Model predictive control and reinforcement learning: A unified framework based on dynamic programming," *IFAC-PapersOnLine*, vol. 58, no. 18, pp. 363–383, 2024.

[61] S. Shin and V. M. Zavala, "Diffusing-horizon model predictive control," *IEEE Transactions on Automatic Control*, vol. 68, no. 1, pp. 188–201, 2021.

[62] S. Shin, M. Anitescu, and V. M. Zavala, "Exponential decay of sensitivity in graph-structured nonlinear programs," *SIAM Journal on Optimization*, vol. 32, no. 2, pp. 1156–1183, 2022.

[63] A. Gundawar, Y. Li, and D. Bertsekas, "Superior computer chess with model predictive control, reinforcement learning, and rollout," *arXiv preprint arXiv:2409.06477*, 2024.

[64] D. V. Lindley, "On a measure of the information provided by an experiment," *The Annals of Mathematical Statistics*, vol. 27, no. 4, pp. 986–1005, 1956.

[65] W. Neiswanger, K. A. Wang, and S. Ermon, "Bayesian algorithm execution: Estimating computable properties of black-box functions using mutual information," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8005–8015.

[66] S. E. Mattsson, H. Elmqvist, and M. Otter, "Physical system modeling with modelica," *Control Engineering Practice*, vol. 6, no. 4, pp. 501–510, 1998.

[67] B. Lim, S. Ö. Arık, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *International Journal of Forecasting*, vol. 37, no. 4, pp. 1748–1764, 2021.

[68] Q. Wen, T. Zhou, C. Zhang, W. Chen, Z. Ma, J. Yan, and L. Sun, "Transformers in time series: A survey," *arXiv preprint arXiv:2202.07125*, 2022.

[69] S. Ahmed, I. E. Nielsen, A. Tripathi, S. Siddiqui, R. P. Ramachandran, and G. Rasool, "Transformers in time-series analysis: A tutorial," *Circuits, Systems, and Signal Processing*, vol. 42, no. 12, pp. 7433–7466, 2023.

[70] Y. Liang, H. Wen, Y. Nie, Y. Jiang, M. Jin, D. Song, S. Pan, and Q. Wen, "Foundation models for time series analysis: A tutorial and survey," in *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, 2024, pp. 6555–6565.

[71] J. A. Miller, M. Aldosari, F. Saeed, N. H. Barna, S. Rana, I. B. Arpinar, and N. Liu, "A survey of deep learning and foundation models for time series forecasting," *arXiv preprint arXiv:2401.13912*, 2024.

[72] K. J. Chan, J. A. Paulson, and A. Mesbah, "A practical multiobjective learning framework for optimal hardware–software co-design of control-on-a-chip systems," *IEEE Transactions on Control Systems Technology*, 2024.

[73] N. Lefebure, M. Khosravi, M. H. de Badyn, F. Bünning, J. Lygeros, C. Jones, and R. S. Smith, "Distributed model predictive control of buildings and energy hubs," *Energy and Buildings*, vol. 259, p. 111806, 2022.

[74] C. J. Bay, R. Chintala, V. Chinde, and J. King, "Distributed model predictive control for coordinated, grid-interactive buildings," *Applied Energy*, vol. 312, p. 118612, 2022.

[75] E. Rezaei, H. Dagdougui, and M. Rezaei, "Distributed stochastic model predictive control for peak load limiting in networked microgrids with building thermal dynamics," *IEEE Transactions on Smart Grid*, vol. 13, no. 3, pp. 2038–2049, 2022.

[76] K. Bonawitz, P. Kairouz, B. McMahan, and D. Ramage, "Federated learning and privacy: Building privacy-preserving systems for machine learning and data science on decentralized data," *Queue*, vol. 19, no. 5, pp. 87–114, 2021.

[77] R. A. Sater and A. B. Hamza, "A federated learning approach to anomaly detection in smart buildings," *ACM Transactions on Internet of Things*, vol. 2, no. 4, pp. 1–23, 2021.

[78] L. H. Gilpin, D. Bau, B. Z. Yuan, A. Bajwa, M. Specter, and L. Kagal, "Explaining explanations: An overview of interpretability of machine learning," in *IEEE International Conference on Data Science and Advanced Analytics*. IEEE, 2018, pp. 80–89.

[79] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, "Explainable AI: A review of machine learning interpretability methods," *Entropy*, vol. 23, no. 1, p. 18, 2020.

[80] G. Vilone and L. Longo, "Notions of explainability and evaluation approaches for explainable artificial intelligence," *Information Fusion*, vol. 76, pp. 89–106, 2021.

[81] L. Sharkey, B. Chughtai, J. Batson, J. Lindsey, J. Wu, L. Bushnaq, N. Goldowsky-Dill, S. Heimersheim, A. Ortega, J. Bloom, S. Biderman, A. Garriga-Alonso, A. Conmy, N. Nanda, J. Rumbelow, M. Wattenberg, N. Schoots, J. Miller, E. J. Michaud, S. Casper, M. Tegmark, W. Saunders, D. Bau, E. Todd, A. Geiger, M. Geva, J. Hoogland, D. Murfet, and T. McGrath, "Open problems in mechanistic interpretability," *arXiv preprint arXiv:2501.16496*, 2025.