Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Designing in an Unpredictable World: Novel Methods for Uncertainty Characterization, Quantification, and Optimization in Process Engineering

Diederik Coppitters[a,*], Antoine Laterre[a], Mahdi Kchaou[a], Kevin Verleysen[b], Panagiotis Tsirikoglou[c], Jerome Stock[d], Matthias Weigold[d], Konstantinos Kyprianidis[e], Ward De Paepe[f], Francesco Contino[a]

[a] Institute of Mechanics, Materials and Civil Engineering (iMMC), Université catholique de Louvain (UCLouvain), Place du Levant, 2, 1348 Louvain-la-Neuve, Belgium
[b] Thermo and Fluid dynamics (FLOW), Vrije Universiteit Brussel (VUB), Pleinlaan 2, 1050 Brussels, Belgium
[c] Limmat Scientific AG, Weinbergstrasse 31, 8006 Zürich, Switzerland
[d] Institute for Production Management, Technology and Machine Tools (PTW), Technical University of Darmstadt, Otto-Berndt-Str. 2, 64287 Darmstadt, Germany
[e] Mälardalen University, Universitetsplan 1, Västerås, Sweden
[f] Thermal Engineering and Combustion Unit, University of Mons, Rue de l'Épargne 56, Mons, 7000, Belgium

* Corresponding Author: diederik.coppitters@uclouvain.be.

ABSTRACT

Computer-Aided Process Engineering (CAPE) has transformed how we analyze, design, and optimize energy processes. Yet, even advanced models rest on uncertain ground: their reliability depends on how well future operating environments are described—environments that are dynamic, complex, and deeply uncertain. In practice, uncertainty is often reduced to local parameter variations, driven by limited data, computational burden, and overconservative robust formulations. This narrow treatment creates a false sense of confidence: Designs that perform well in theory often fail in real-world operation. In a century marked by economic, climatic, and technological volatility, designing under uncertainty is no longer optional; it is essential.

We have developed approaches that place uncertainty at the core of energy process modeling and design. This paper provides an overview of these methods and how uncertainty can be explicitly represented, quantified, and embedded into the design process.

We present approaches to characterize uncertainties under data scarcity, including imprecise probabilities to distinguish between aleatory and epistemic uncertainty. To propagate uncertainty through computationally intensive models, we introduce surrogate-assisted techniques that exploit structural sparsity, enabling the analysis of systems with large numbers of uncertain inputs (100+) while mitigating the curse of dimensionality. These methods are integrated into optimization frameworks that target expected performance, robustness, and, for the first time, antifragility—systems that can benefit from variability rather than merely withstand it. We illustrate these approaches across applications ranging from detailed process models to system-level energy analyses, advancing a shift in CAPE toward designs suited for an unpredictable world.

Keywords: Uncertainty assessment, Robust design optimization, Process simulation, Antifragility.

INTRODUCTION

Computer-Aided Process Engineering (CAPE) tools, such as Aspen Plus, gPROMS, Thermoflex, and in-house high-fidelity thermodynamic simulators, have become central to process design and operation [1]. Engineers now evaluate complex heat integration networks, power-to-X routes, carbon capture systems, and flexible plants using nonlinear, tightly coupled models that embed thermodynamics, transport phenomena, and equipment constraints. These models can simulate performance with high accuracy, but they remain conditional on assumptions about their operating environment, such as variations in ambient temperature and humidity, cooling-water temperature, heat-transfer coefficients, fuel composition (carbon, hydrogen, lower-heating value), electricity prices, CAPEX and OPEX of emerging units, policy constraints, market regimes, and evolving sustainability targets. Despite this, uncertainty in many CAPE studies is still mostly handled through deterministic base-case optimization followed by local sensitivity checks or a small scenario sweep [2]. This approach persists for understandable reasons: (i) data about the operating environment is fundamentally incomplete which makes rigorous uncertainty characterization difficult [3], (ii) brute-force Monte Carlo becomes infeasible for expensive nonlinear flowsheets [4], and (iii) worst-case robust formulations often drive designs toward conservatism that is difficult to defend in engineering and economic terms [5]. This narrow treatment of uncertainty can create confidence in designs that are fragile when exposed to operational variability and imperfect foresight.

Over the past decade, we have developed methods that place uncertainty at the center of CAPE modeling and design. This paper documents a probabilistic uncertainty workflow that we use for nonlinear, technology-detailed CAPE models under limited information. We focus on uncertainty characterization under limited data, surrogate-assisted uncertainty propagation that remains tractable for expensive simulations, and robust design optimization that explicitly trades off expected performance and sensitivity to variability. We also present antifragility as a paradigm shift for optimization under uncertainty and ways to address structural uncertainty, where uncertainty is not only about parameter values but also about what the problem formulation should be, for example when stakeholder preferences or acceptability thresholds are uncertain. Finally, we conclude this paper with a discussion and perspectives.

A FRAMEWORK FOR UNCERTAINTY-AWARE PROCESS DESIGN

In this section, we present a framework for uncertainty-aware process design (Figure 1). It covers methods for characterizing uncertainty, propagating uncertainty and robust design optimization.

Representing Uncertainty Under Limited Data

Uncertainty characterization translates knowledge about the parameter variation into a mathematical representation that will drive propagation and optimization. In CAPE, researchers repeatedly face the same obstacle [2]: Data about the operating environment is often sparse, inconsistent, or context-specific. For instance, electricity-price projections may exist only as a small set of scenarios, emerging unit costs may be based on few studies with incompatible assumptions, and ambient conditions may only be available at coarse data resolution. In this setting, analysts often default to common distributions such as uniform or normal [3]. In this way, they risk encoding strong assumptions, either by assigning zero probability to tail events or by extrapolating tails from limited central information. In nonlinear flowsheets, these assumptions can change optimized designs because feasibility, active constraints, and operating regimes can switch across the uncertain domain [6].

One way to represent uncertainty uses quantile-anchored parametric distributions by combining data with expert judgment. While frequently available, translating expert judgment into distributional parameters (e.g., mean, variance) is difficult: Even when experts can provide plausible ranges, they struggle to express how probability mass should be distributed within those ranges [7]. Expert sources are usually more consistent when asked for percentile statements [8]. Therefore, a pragmatic way is to elicit expert judgment on quantiles rather than attempt to obtain full distribution parameters directly, such as “a pessimistic value that is likely exceeded in only 5% of the scenarios.” Moreover, if the modeler commits to a distribution type (e.g., a lognormal distribution to represent the future electricity price), two distinct percentiles determine a unique distribution (in the common case of a two-parameter type), so the resulting distribution is fully anchored to explicit expert statements [9]. To illustrate this approach, we elicited two percentiles and distribution types from expert reports and characterized the corresponding distributions for grid electricity price, wind speeds and electricity demand in the assessment of a wind-powered community [8].

A second representation uses imprecise probabilities to treat distributional ambiguity explicitly. The same elicited quantiles may be consistent with a bounded distribution that has hard limits, or with a heavy-tailed distribution that assigns small but decision-relevant probability to extreme outcomes. When this ambiguity affects design choices, a single fitted curve can convey unwarranted certainty. Probability boxes (p-boxes) address this by representing uncertainty as a family of admissible distributions bounded by lower and upper cumulative distribution functions [10]. This representation is particularly useful when epistemic uncertainty dominates, because it treats uncertainty about the distribution itself as part of the model, not as an afterthought. It also supports a separation between aleatory variability, which reflects intrinsic randomness, and epistemic uncertainty, which reflects limited knowledge that could shrink with better information or clearer assumptions.

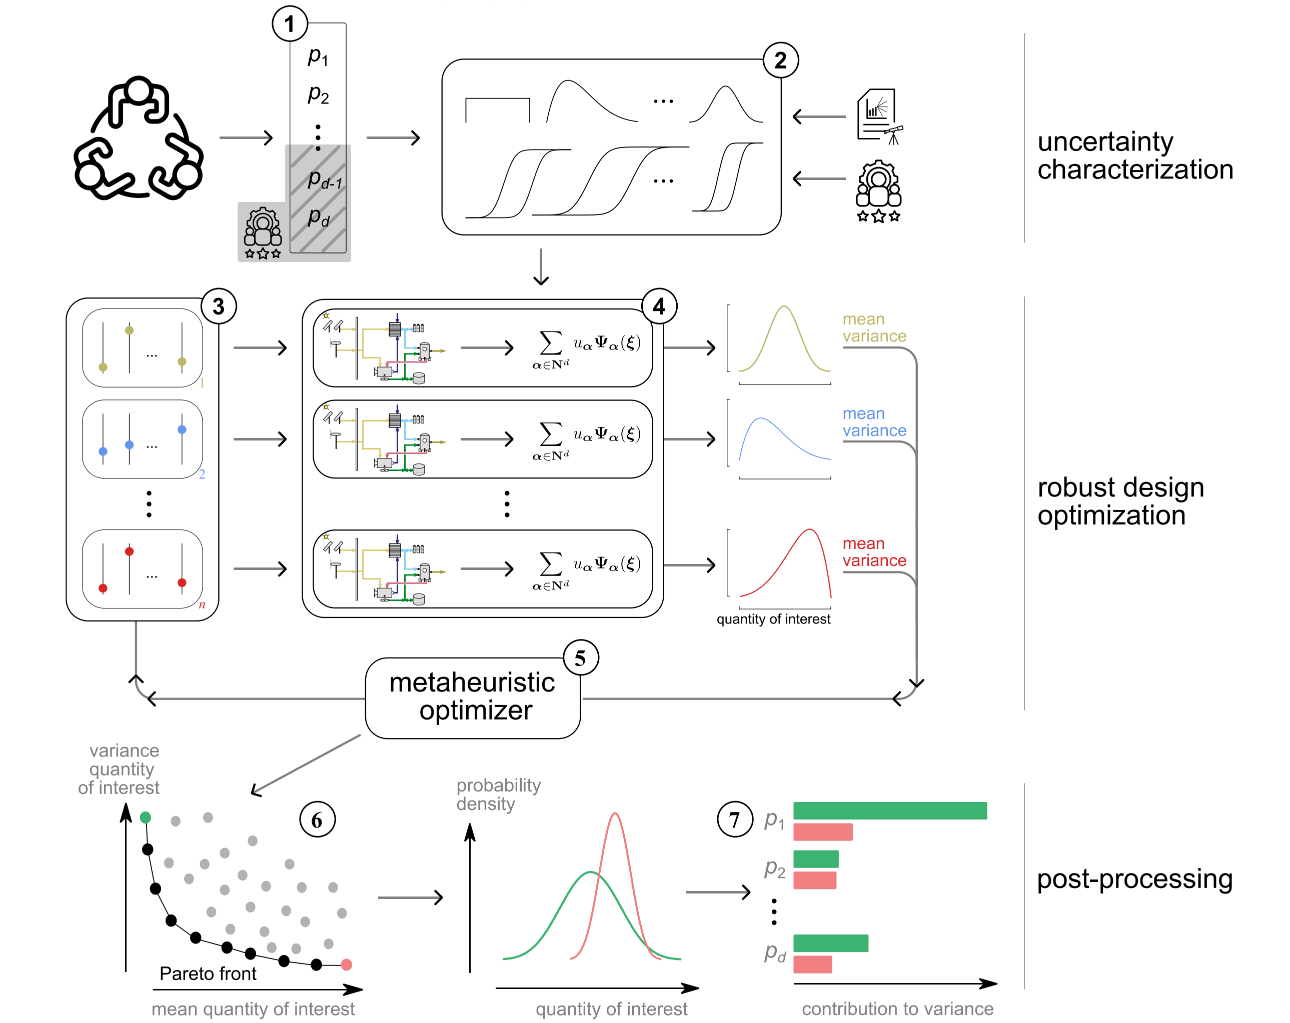

Figure 1: The schematic illustrates the proposed framework for uncertainty-aware process design. Experts first identify the parameters to be treated as uncertain (1). An optional expert-based screening may be applied to discard parameters expected to have negligible influence. Uncertainty is then characterized by combining data, projections, and expert judgment, using either precise probabilistic distributions or imprecise formulations when epistemic uncertainty is substantial (2). The optimization workflow then begins. An initial set of candidate designs is generated by exploring the design space, for example using Latin Hypercube Sampling (3). For each candidate design, the process model is evaluated under uncertainty by constructing a surrogate, typically a full or sparse Polynomial Chaos Expansion depending on model complexity and the number of retained uncertain inputs (4). Performance statistics such as the mean and variance, or bounds on these quantities in the imprecise case, are computed analytically from the surrogate and provided to a population-based metaheuristic optimizer as fitness values (5). The optimizer iteratively generates new design populations and updates the design set. After convergence, the resulting Pareto front quantifies the trade-off between expected performance and robustness (6). For selected Pareto-optimal designs, global sensitivity analysis is extracted analytically from the surrogate to identify the uncertain inputs that dominate performance variability and those with negligible influence (7).

We used probability boxes to represent the operating environment of a hybrid renewable energy system comprising a photovoltaic array, batteries, a heat pump, and thermal storage. For electricity and heat demand, aleatory uncertainty represented irreducible variability in future demand, while epistemic uncertainty reflected limited knowledge of occupant behavior and reliance on generic demand profiles. Energy prices were treated analogously, with aleatory uncertainty capturing plausible market evolution and epistemic uncertainty capturing ambiguity in tariffs and supplier choice. We propagated these probability boxes through the system model and used the aleatory component to drive robust design optimization of the levelized cost of exergy, because this variability cannot be reduced through additional information and can only be mitigated through design choices. In contrast, we used global sensitivity analysis on the epistemic component to identify which uncertain assumptions most affected predicted performance and robustness, thereby indicating where additional information would most effectively reduce decision-relevant uncertainty [11].

In practice, a precise probabilistic formulation is pragmatic, as it is simpler to implement, propagate, and communicate, but it necessarily collapses epistemic uncertainty into a single predictive model. This is often acceptable for parameters such as ambient conditions, standardized fuel properties, or short-term market fluctuations, where uncertainty is largely aleatory or where epistemic uncertainty is small relative to its effect on design decisions. We propose an imprecise representation when epistemic uncertainty is significant, meaning that uncertainty about the parameters or assumptions of the distribution itself is large enough to affect design conclusions. Typical examples include electricity or carbon prices over multi-decadal horizons, CAPEX of emerging technologies, or demand projections driven by poorly specified user behavior. In such cases, explicitly separating aleatory variability, which must be mitigated through design, from epistemic uncertainty, which can be reduced through additional information or specification, avoids masking decision-relevant uncertainty, despite the added conceptual and computational complexity.

Uncertainty Propagation and Quantification

Once input uncertainties have been characterized, whether through precise or imprecise probabilities, the next step is to propagate that uncertainty through the model to quantify its impact on the outputs of interest. The most straightforward approach is Monte Carlo Simulation (MCS), which samples uncertain inputs from their assigned distributions, evaluates the model for each sample, and estimates output statistics. Although this approach is conceptually simple and asymptotically convergent, it becomes impractical for many CAPE applications. MCS typically requires a large number (~10^{3-5}) of model evaluations to resolve output variability and tail behavior, while high-fidelity CAPE models are computationally expensive (minutes to hours) and often involve many uncertain inputs (>10). As a result, naive Monte Carlo propagation quickly becomes computationally intractable.

An effective alternative is to move to surrogate-assisted Uncertainty Quantification (UQ). The core idea is to replace numerous brute-force evaluations of an expensive simulator with a cheaper approximation that captures the dominant input–output relationships. Among available surrogate approaches, Polynomial Chaos Expansion (PCE) is particularly well suited to UQ because it represents model outputs as an expansion in orthogonal polynomials of the uncertain inputs [12]. Once a PCE surrogate is trained, it can be evaluated orders of magnitude faster (typically milliseconds) than the original model. An additional benefit of PCE is that uncertainty propagation naturally yields statistical moments on the output (mean, variance) and a global sensitivity analysis can be performed at negligible additional cost by exploiting the PCE coefficients. In practice, high Sobol' indices guide where further knowledge is necessary. For example, we applied PCE to the integration of Direct Air Capture (DAC) in a power-to-gas flowsheet to quantify uncertainty in energy efficiency, exergy efficiency and levelized cost under uncertain ambient conditions, sorbent thermodynamics, electricity price, and CAPEX/OPEX. The resulting Sobol' indices identified that electricity price and DAC CAPEX dominated the variance of levelized cost, while thermodynamic model parameters and ambient conditions had a secondary influence on the overall performance [13].

In CAPE applications, the uncertainty space typically includes many input parameters, which makes a full PCE impractical because the number of required expansion coefficients, and thus training simulations, grows combinatorially with both polynomial degree (linked to the nonlinearity of the input-output relation) and the number of uncertain parameters. Dimensionality reduction is therefore unavoidable.

Screening based on expert judgment or coarse sensitivity checks provides a pragmatic first reduction, but it is inherently subjective and can overlook higher-order interactions that only emerge in nonlinear regimes. Multi-fidelity PCE strategies alleviate this by using low-order surrogates to rank influential inputs before constructing higher-order models, but their effectiveness depends on the reliability of these preliminary approximations, which can be poor when responses are highly nonlinear or discontinuous. In a long-term energy-system planning model with many uncertain techno-economic inputs, we used a low-order PCE as a preliminary step to reduce an initial set of uncertain parameters to a smaller set of critical parameters before building the higher-fidelity surrogate and revealed that a small subset of fuel-price and demand-growth assumptions dominated the uncertainty on the system costs [14].

Sparse PCE addresses these limitations by integrating screening directly into surrogate construction, retaining only influential polynomial terms and interactions while discarding the rest, thereby controlling dimensionality without requiring explicit pre-screening or multiple surrogate stages and significantly reducing the number of required training simulations [15]. We used sparse PCE for a Thermoflex®-based retrofit study of a large thermal power plant where 243 uncertain inputs made conventional PCE infeasible, and the resulting Sobol' indices showed that fuel composition, lower heating value, and boiler efficiency parameters dominated variability in furnace temperature and efficiency, while most auxiliary parameters had negligible influence [16].

PCE can also be applied when input uncertainties are represented through probability boxes. In a parametric p-box setting, uncertainty in the statistical moments, such as the mean and variance, is represented by uniformly distributed variables to reflect bounded epistemic uncertainty. For any specific realization of these parameters, a corresponding formulation of the aleatory uncertainty is defined through the associated conditional distribution. Because this conditional formulation introduces dependence between epistemic and aleatory components, the aleatory variable is rewritten in standardized form, and the physical input is reconstructed by combining this standardized variable with the sampled mean and variance. This reformulation represents the uncertain input through a set of independent variables: one standardized variable capturing aleatory variability and uniform variables capturing epistemic uncertainty in the distribution parameters. A single PCE surrogate constructed on this augmented input space can then be sampled efficiently to compute bounds on the output distribution, yielding the probability box of the model response without resorting to a nested MCS [17]. In a PV–battery–heat-pump system with thermal storage, we represented mixed aleatory variability (demand and weather evolution) and epistemic ambiguity (occupant behavior and cost assumptions) with parametric p-boxes, then used the augmented PCE to propagate output bounds while using the aleatory part for design objectives and the epistemic part for sensitivity-guided information prioritization, showing that aleatory demand variability primarily affected the robustness, while epistemic uncertainty in electricity tariffs and user behavior dominated the epistemic uncertainty on the optimized costs [11].

Designing for Robustness

Design optimization under uncertainty is where all the pieces come together—and also where computational burden tends to peak. The goal is to identify design decisions (equipment sizes, operating conditions, etc.) that perform well across the range of uncertainties characterized earlier. A specific type of optimization under uncertainty is Robust Design Optimization (RDO), where the mean and variance of the quantities of interest obtained from UQ, such as costs, efficiencies and emissions, define the optimization objectives and constraints.

From a computational perspective, RDO inherits the computational cost of UQ. We therefore build on the surrogate-based propagation strategies introduced in the previous section and embed them within an optimization loop [19]. For a given candidate design, the design variables are held fixed, and the PCE is constructed by sampling the uncertain inputs around that design according to their prescribed probabilistic descriptions. Because the surrogate represents the full input–output map for that design, the statistical moments required for the objectives can be computed analytically from the PCE coefficients. Reliability-type constraints can be acquired in the same way by evaluating, for each design, the probability of violating constraints based on the propagated uncertainty [18]. In a high-fidelity RDO setting considered here, uncertainty propagation is thus performed separately for each candidate design evaluated by the optimizer. While this approach is computationally more demanding than global surrogates spanning both design and uncertainty spaces, it preserves accuracy in the presence of strong nonlinearities in the design space.

The optimization itself is performed using gradient-free, population-based algorithms, which are well suited to the black-box and often nonconvex nature of CAPE problems [20]. In our applications, multi-objective formulations are solved using NSGA-II, with objectives being the mean and variance of the quantity of interest, such as cost, efficiency and emissions. The resulting Pareto front quantifies the trade-off between the expected performance and robustness and provide a structured basis for optimized design selection.

The validity of this approach depends on the reliability of the surrogate models used during optimization. The fidelity is therefore enforced through an automated two-step procedure. First, the design space is explored using a space-filling sampling of the design variables, such as Latin Hypercube Sampling. For this initial set of designs, full PCEs are built for increasing polynomial orders, and the worst-case leave-one-out error over all sampled designs is evaluated. The smallest polynomial order that meets a predefined accuracy threshold is then selected and fixed for the remainder of the optimization. During the RDO iterations, sparse PCEs of this order are constructed for each candidate design, and their accuracy is monitored using leave-one-out error and stabilization of the estimated mean and standard deviation as training samples are added. If a candidate design leads to non-convergent or nonphysical model evaluations for one or more training samples, that design is classified as infeasible and assigned a penalized objective value.

We implemented the full workflow in an open-source framework called RHEIA [21] and applied it to several CAPE applications with distinct process structures and decision variables. In a grid-connected photovoltaic system with battery and hydrogen storage, RDO shifted the preferred designs from photovoltaic-only systems toward battery and hydrogen storage when robustness was valued. For a community-scale system, the robust PV–battery–hydrogen design increased the mean levelized cost of electricity from 269 €/MWh to 363 €/MWh (≈+35%), while reducing its standard deviation from 55.6 €/MWh to 41.9 €/MWh (≈ −25%), explicitly trading expected cost for reduced sensitivity to electricity-price uncertainty [19]. In a wind-powered power-to-ammonia plant, RDO revealed a clear trade-off between productivity and robustness: the design maximizing mean annual ammonia production achieved 3.2 times higher average output, but its standard deviation of annual production was 2.6 times larger than that of the robust design, indicating substantially higher sensitivity to wind-speed and reactor-temperature uncertainties [22]. Finally, for a renewable-powered hydrogen refueling station and bus fleet, converting 54% of the diesel buses to hydrogen-powered buses reduced the standard deviation of the levelized cost by 36%, reduced the mean carbon intensity by 46%, and reduced its standard deviation by 51%, at the expense of an 11% increase in mean cost [23].

Limitations

The workflow has three main limitations related to computational cost, input dependency, and accuracy.

First, the computational cost can still be high because the RDO loop requires building, validating, and sampling a surrogate for each candidate design. This nested structure can have computational overhead because of expensive models, large populations, many generations, and higher polynomial orders.

Second, standard PCE construction typically assumes independent uncertain inputs. Many CAPE problems feature correlated drivers (e.g., demand–price coupling, weather covariation, or jointly evolving techno-economic assumptions). Ignoring dependence can misrepresent joint extremes, constraint-violation risk, and sensitivity indices. Accounting for dependence is possible (e.g., via transformations or copula-based modeling), but it introduces additional modeling choices.

Third, PCE accuracy degrades when outputs are non-smooth or discontinuous in the uncertain inputs, which can arise from regime switching, active-constraint changes, phase changes, and feasibility boundaries. In these cases, moment estimates, tail metrics, and Sobol' indices may become unreliable unless more complex surrogate strategies are used (e.g., piecewise surrogates or classification-plus-regression), which reduces the simplicity of the approach.

EMERGING DIRECTIONS IN UNCERTAINTY-AWARE DESIGN

While RDO provides a principled way to control sensitivity to uncertainty, it implicitly treats uncertainty as a source of risk. As a result, optimization strategies that focus exclusively on robustness tend to penalize all variability symmetrically, which can suppress configurations that would perform better under favorable operating conditions. In addition, RDO primarily addresses uncertainty in model parameters and implicitly assumes that the chosen model structure and objective formulation capture all decision-relevant aspects. In practice, however, factors such as social acceptance and stakeholder preferences may be difficult, if not impossible, to encode, yet they can strongly influence whether an optimized design can be implemented. The following sections therefore examine two complementary extensions: antifragility, which explicitly accounts for asymmetric responses to variability, and approaches that address structural and preference uncertainty by exploring sets of near-optimal alternatives rather than a single solution.

Antifragility: A Paradigm Shift in Designing Under Uncertainty

Optimization under uncertainty in CAPE has traditionally focused on robustness, that is, identifying designs whose performance is least sensitive to uncertain operating conditions. While effective when uncertainty is well characterized, this paradigm becomes limiting when the magnitude and structure of future techno-economic uncertainty are themselves poorly known. Moreover, robustness treats variability symmetrically: it controls downside risk, but it also suppresses upside opportunity.

An emerging alternative is antifragile design, which shifts the objective from minimizing harm due to uncertainty to exploiting favorable variability—an antifragile system does not merely withstand uncertainty; it benefits from it [24]. We have explored antifragility as an extension of RDO by introducing objectives that reward beneficial variability rather than penalizing variability indiscriminately [8]. One such metric is the Upside Potential Ratio (UPR), which favors designs that achieve substantially better-than-average outcomes while still limiting unacceptable losses. For common performance indicators, encouraging upside variability corresponds to the possibility of higher efficiency, lower cost, or lower emissions than expected under nominal conditions. We also propose using the skewness of the performance distribution as an objective to quantify antifragility. Positive skewness indicates limited downside exposure combined with rare but significant gains, and it captures how a design responds when the actual operating environment is more uncertain than anticipated.

Embedding these objectives within a multi-objective optimization framework yields designs with qualitatively different behavior from conventional robust solutions. In the design of a wind turbine for a community, antifragile optimization led to larger installed capacity than the robust design, deliberately accepting worse median performance and higher variability, but exhibited strongly asymmetric outcomes: when the actual uncertainty exceeded initial assumptions, it delivered substantially lower costs than both the deterministic and robust designs [8]. In a residential heat-pump dimensioning case for a single-family dwelling, antifragile objectives similarly shifted optimal designs away from those minimizing median cost. Designs with higher antifragility featured larger thermal capacity and storage, which increased cost but limited downside exposure and enabled gains under volatile electricity prices and demand [25].

Addressing Structural and Preference Uncertainty

Decision-makers are rarely concerned only with the parameters and objectives explicitly modeled. Factors such as social acceptance, siting constraints, permitting risk, or political feasibility are difficult, if not impossible, to encode quantitatively, yet they can strongly influence whether an optimized design can be implemented. As a result, presenting only a single optimized solution, or even a narrow Pareto front, can be misleading. It embeds one particular problem formulation and worldview, while masking structural uncertainty, that is, uncertainty about the appropriate system layout or decision framing.

In the energy modeling literature, Modeling to Generate Alternatives (MGA) is an established response to this problem. MGA has been used primarily in Linear Programming settings to produce sets of near-optimal solutions that are structurally different, allowing decision-makers to choose among alternatives when unmodeled preferences become important [26].

Our work extended MGA concepts to nonlinear multi-objective optimization, generating portfolios of near-optimal designs under preference uncertainty and using these portfolios to reveal which design features are "must-haves" versus genuine degrees of freedom [27]. In practice, this was done by first computing the Pareto-optimal solutions and then solving additional optimization problems that were constrained to remain within a small, predefined degradation of the objective values while forcing variation in configuration choices and key design variables. The resulting set of designs was then analyzed by comparing configurations and parameter ranges across the ensemble. Applied to a Carnot-battery system, this procedure showed that some design choices, such as recuperated subcritical Heat Pump–Organic Rankine Cycle layouts for high-efficiency operation, appeared systematically across all near-optimal solutions and can be considered must-haves. Other choices, including temperature lift, superheating level, and specific operating points, varied widely with little impact on performance, indicating real choices. However, further convergence analysis and benchmarking are needed to assess the effectiveness of the method, and additional development of the post-processing is required to extract clearer and more compelling messages from the large volume of results.

CONCLUSION

Designing energy and process systems in an increasingly volatile operating environment requires uncertainty to be treated as a core modeling and optimization dimension rather than as a secondary sensitivity check. As operating environments become more volatile and assumptions harder to defend, design methods must explicitly acknowledge limited information, structural ambiguity, and asymmetric responses to variability. In this paper, we presented a coherent workflow for uncertainty-aware CAPE that spans uncertainty characterization under limited information, surrogate-assisted uncertainty propagation for computationally expensive models, and optimization frameworks that explicitly account for variability. Together, these methods advance a shift in CAPE from deterministic optimization toward design strategies that remain credible, informative, and actionable under uncertainty.

ACKNOWLEDGEMENTS

D.C. acknowledges the support of the Fonds de la Recherche Scientifique - FNRS [CR 40016260]. M.K. acknowledges the support of Wallonie-Bruxelles International – WBI.

AUTHOR IDENTIFIERS

Author ORCIDs:
Diederik Coppitters: 0000-0001-9480-2781
Antoine Laterre: 0000-0003-4226-1594
Mahdi Kchaou: 0009-0000-5840-409X
Kevin Verleysen: 0000-0002-7763-7208
Jerome Stock: 0009-0006-7908-069X
Matthias Weigold: 0000-0002-7820-8544
Konstantinos Kyprianidis: 0000-0002-8466-356X
Ward De Paepe: 0000-0001-5008-2946
Francesco Contino: 0000-0002-8341-4350

REFERENCES

1. Gani R et al. Process systems engineering: modeling and simulation. Ullmanns Encycl Ind Chem 1–54 (2012) https://doi.org/10.1002/14356007.o22_o06
2. van der Spek M et al. Uncertainty analysis in the techno-economic assessment of CO_2 capture and storage technologies: Critical review and guidelines for use. Int J Greenh Gas Control 100:103113 (2020) https://doi.org/10.1016/j.ijggc.2020.103113
3. Mavromatidis G et al. A review of uncertainty characterisation approaches for the optimal design of distributed energy systems. Renew Sustain Energy Rev 88:258–277 (2018) https://doi.org/10.1016/j.rser.2018.02.020
4. Zhang J. Modern Monte Carlo methods for efficient uncertainty quantification and propagation: A

survey. WIREs Comput Stat 13:e1539 (2021) https://doi.org/10.1002/wics.1539
Bertsimas D, Sim M. The price of robustness. Oper Res 52:35–53 (2004) https://doi.org/10.1287/opre.1030.0065
5. Sahinidis NV. Optimization under uncertainty: State-of-the-art and opportunities. Comput Chem Eng 28:971–983 (2004) https://doi.org/10.1016/j.compchemeng.2003.09.017
6. Grigore B et al. Methods to elicit probability distributions from experts: A systematic review of reported practice in health technology assessment. PharmacoEconomics 31:991–1003 (2013) https://doi.org/10.1007/s40273-013-0092-z
7. Coppitters D, Contino F. Optimizing upside variability and antifragility in renewable energy system design. Sci Rep 13:9138 (2023) https://doi.org/10.1038/s41598-023-36379-8
8. Cook JD. Determining distribution parameters from quantiles. UT MD Anderson Cancer Center Dept Biostatistics (2010)
9. Ferson S, Ginzburg LR. Different methods are needed to propagate ignorance and variability. Reliab Eng Syst Saf 54:133–144 (1996) https://doi.org/10.1016/S0951-8320(96)00071-3
10. Coppitters D et al. Robust design optimization of a photovoltaic-battery-heat pump system with thermal storage under aleatory and epistemic uncertainty. Energy 229:120692 (2021) https://doi.org/10.1016/j.energy.2021.120692
11. Sudret B. Global sensitivity analysis using polynomial chaos expansions. Reliab Eng Syst Saf 93:964–979 (2008) https://doi.org/10.1016/j.ress.2007.04.002
12. Coppitters D et al. Energy, exergy, economic and environmental (4E) analysis of integrated direct air capture and CO_2 methanation under uncertainty. Fuel 344:127969 (2023) https://doi.org/10.1016/j.fuel.2023.127969
13. Rixhon X et al. The role of electrofuels under uncertainties for the Belgian energy transition. Energies 14:4027 (2021) https://doi.org/10.3390/en14134027
14. Abraham S et al. A robust and efficient stepwise regression method for building sparse polynomial chaos expansions. J Comput Phys 332:461–474 (2017) https://doi.org/10.1016/j.jcp.2016.12.019
15. De Meulenaere R et al. Uncertainty quantification for thermodynamic simulations with high-dimensional input spaces using sparse polynomial chaos expansion: Retrofit of a large thermal power plant. Appl Sci 13:10751 (2023) https://doi.org/10.3390/app131910751
16. Schöbi R, Sudret B. Global sensitivity analysis in the context of imprecise probabilities (p-boxes) using sparse polynomial chaos expansions. Reliab Eng Syst Saf 187:129–141 (2019) https://doi.org/10.1016/j.ress.2019.02.017
17. De Paepe W et al. Robust operational optimization of a typical micro gas turbine. Energy Procedia 158:5795–5803 (2019) https://doi.org/10.1016/j.egypro.2019.01.549
18. Coppitters D et al. Robust design optimization and stochastic performance analysis of a grid-connected photovoltaic system with battery storage and hydrogen storage. Energy 213:118798 (2020) https://doi.org/10.1016/j.energy.2020.118798
19. Tsirikoglou P et al. Comparison of metaheuristic algorithms on robust design optimization of a plain-fin-tube heat exchanger. AIAA Proc 2017:2017-3827 (2017) https://doi.org/10.2514/6.2017-3827
20. Coppitters D et al. RHEIA: Robust design optimization of renewable hydrogen and derived energy carrier systems. J Open Source Softw 7:4370 (2022) https://doi.org/10.21105/joss.04370
21. Verleysen K et al. How can power-to-ammonia be robust? Optimization of an ammonia synthesis plant powered by a wind turbine considering operational uncertainties. Fuel 266:117049 (2020) https://doi.org/10.1016/j.fuel.2020.117049
22. Coppitters D et al. How can renewable hydrogen compete with diesel in public transport? Robust design optimization of a hydrogen refueling station under techno-economic and environmental uncertainty. Appl Energy 312:118694 (2022) https://doi.org/10.1016/j.apenergy.2022.118694
23. Taleb N. 'Antifragility' as a mathematical idea. Nature 494:430 (2013) https://doi.org/10.1038/494430e
24. Frank M et al. A software framework for optimization under stochastic uncertainty to achieve antifragility. Procedia CIRP
25. Lombardi F et al. Near-optimal energy planning strategies with modeling to generate alternatives to flexibly explore practically desirable options. Joule (2025) https://doi.org/10.1016/j.joule.2025.102144
26. Laterre A et al. Designing small-scale Rankine Carnot batteries that suit your preferences: A near-optimal approach. J Energy Storage 141:118650 (2026) https://doi.org/10.1016/j.est.2025.118650