Journal of Applied Computational Intelligence

eISSN: 3136-6937

Vol. 4, No. 1, January, 2026

Gaussian Process Surrogates and Uncertainty Aware NSGA II for Optimizing Wind Turbine Blade Pyrolysis Parameters

Odysseas Dellas[1,*], Sotiris Kafatos[1], Thanasis Petrou[2], Achilleas Kallis[1]

[1] *Department of Computer Science, University of Crete, Rethymno, 74100, Greece*

[2] *Department of Computer Science and Engineering, Democritus University of Thrace, Xanthi, 67100, Greece*

Received 12 November 2025; Revised 28 January 2026; Accepted 10 Feburary 2026; Published 13 Feburary 2026

Abstract

Recovery of valuable products from used wind turbine blades requires operating conditions that balance condensable yield, energy demand and the preservation of reinforcing fibres. This study develops an uncertainty-aware optimization framework combining Gaussian process surrogates with the nondominated sorting genetic algorithm II. The case study considers temperature, heating rate, holding time and nitrogen flow within a bounded operating domain. Response definitions and uncertainty conventions are applied consistently throughout the analysis. Separate surrogate models predict liquid yield, specific energy demand and retained fibre strength, with conservative prediction bounds used for screening. The evaluation protocol specifies 120 development cases, 40 calibration cases, and 40 independent assessment cases; the displayed diagnostic subsets provide a numerical assessment of the decision rules. A typical compromise has a liquid yield of 27.4 wt%, an energy requirement of 6.20 MJ per kg of dry feed, and a retained strength of 76.0%. Compared with the designated reference setting, these values are an increase of 3.4 percentage points in yield and a reduction of 15.1% in energy. A conservative lower strength bound of 72.0% exceeds the acceptance threshold of 70.0%. Fifteen data panels show the variations in response, calibration requirements, trade-offs and operating sensitivity of the operating alternatives. The framework provides a reproducible specification linking process performance and surrogate uncertainty for subsequent engineering verification.

Keywords: Wind Turbine Blade Pyrolysis, Gaussian Process Regression, Uncertainty Aware Optimization, NSGA II, Fibre Strength Retention

1. Introduction

The fibre-reinforced polymer architecture of end-of-life wind turbine blades creates substantial challenges for material separation and recycling [1]. Research on blade recycling identifies a source of feedstock heterogeneity and treatment complexity in the combination of crosslinked matrices, reinforcing fibres, adhesive joints and protective layers [2]. Thermochemical recovery studies investigate the redistribution of organic matter in condensable products, non-condensable gases and deposits on the fibre-rich residue under pyrolysis conditions [3]. Research on recovered glass fibres has studied the effects of heat and surface conditions on mechanical properties that affect their reuse [4]. Therefore, the selection of operating parameters is related to product recovery, utility demand and reinforcement quality. An increase in the severity of treatment will increase the energy demand and thermal exposure for fibre matrix removal. The engineering objective therefore combines product recovery, energy demand and reinforcement quality. An effective operating policy should have a uniform definition of liquid yield, process energy and retained strength to allow for

* Corresponding author: dellas.dy@cs.uoc.gr

http://dx.doi.org/10.68406/jaci.2026.vol4iss1no6:68-82

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

comparison of competing settings on the same basis. Feed composition and the chosen process boundary will also affect this comparison to some extent, given different considerations for upstream preparation and downstream fibre treatment. To be suitable for actual application, the recommended settings of the optimization method should still connect to the operating instructions, product cost accounting and quality requirements for the recovered reinforcing fibres.

Operating space comprises interacting thermal and transport controls. Heating rate influences exposure during decomposition, holding time determines exposure after the target temperature is reached, and carrier gas flow affects transport through the reactor and condensation train. Pyrolysis-atmosphere experiments have shown that the conditions of the gas phase affect both matrix decomposition and recovered reinforcement [5]. Evaluating several responses across multiple operating settings requires a substantial experimental budget. Sequential experimentation can reduce this problem when the selected model represents both mean response and predictive uncertainty. Gaussian process regression is a probabilistic model that has a predictive spread based on assumed covariance and observed data [6]. Independent checks are needed to verify narrow prediction intervals because they are based on model assumptions and design coverage. Multi-objective evolutionary optimization provides another type of solution by retaining multiple operating choices in the form of a set of competing objective values. Nondominated sorting genetic algorithm II is used to rank the alternatives and maintain diversity along the trade-off front [7]. Surrogate-assisted optimization uses inexpensive response predictions to guide search, which can preferentially select overoptimistic predictions in poorly supported regions. Explicit uncertainty handling and targeted evaluation are therefore needed to assess promising candidates. Work on uncertainty-guided surrogate search has motivated the explicit treatment of this interaction [8]. Therefore, the practical problem of blade pyrolysis is how to handle the uncertainty in candidate ranking and selection under constraints on the experimental budget for liquid recovery, energy demand and fibre preservation.

Coupling probabilistic prediction with constrained search yields a specific research question: How should predictive uncertainty be accounted for in the ranking of blade pyrolysis settings under the conflict of material recovery and energy demand, given a fibre quality requirement? A decision system has been developed for this problem in a transparent engineering case study. The decision variables are temperature, heating rate, holding time and nitrogen flow. Liquid yield and specific energy consumption form the objective pair, while retained fibre strength defines the acceptance requirement. Separate data partitions are used for model development, interval calibration and the final test; thus, fitting, uncertainty adjustment and independent assessment retain distinct roles. Conservative response bounds connect the surrogate outputs to candidate ranking and acceptance. An exploratory queue is used for information collection, and an acceptance archive saves candidates that meet the quality screen. Response arrays and diagnostic coordinates are used in the analysis for reproducibility. The contribution is a process-specific integration of established surrogate and evolutionary components, with explicit accounting boundaries, uncertainty conventions, and separate exploration and acceptance rules. The framework specifies response definitions, uncertainty conventions and candidate selection procedures before experimental evaluation. It also identifies the observations needed to compare the method with mean-only surrogate search and directly evaluated multi-objective search under matched budgets.

The analysis is based on the problem of material recovery and includes models and operating decisions to address this problem. Section 2 introduces thermal blade recycling, the construction of a probabilistic surrogate model and evolutionary optimization, and provides the necessary support for connecting process response prediction with quality-constrained search. Section 3 sets up the process boundary and operating domain, establishes the response accounting conventions, and develops the Gaussian process representation. Separation of latent uncertainty and observation noise is employed to address the problem of predicting performance and future acceptance values. Section 4 introduces a conservative objective ranking, the retained strength screen, evolutionary variation and the selection of a representative trade-off. The roles of exploratory evaluations and approved operating candidates are separated, and the recommendation can be traced back to the original data. Section 5 examines response patterns, prediction diagnostics, candidate trade-offs and sensitivity. Each figure panel is related to a number of observations and the corresponding implications for operating selection and related uncertainty assumptions. Section 6 presents the methodological conclusions, identifies current limitations and outlines experimental and modelling work for assessing the framework with measured blade pyrolysis data. The organization connects the optimization decisions to the response definitions and the uncertainty assumptions of these decisions.

2. Related Work

2.1. Thermal recovery and fibre quality

Research on the thermal recycling of composite blades separates several tasks that are often grouped as material recovery. Matrix decomposition reduces the proportion of organic matter, condensation gathers a liquid phase, and further processing may be required to eliminate residues from the recovered fibres. Research on fibre recovery by blade pyrolysis has provided a foundation for the above studies [9]. It can be used to optimize the process by adjusting the energy requirements and changing the quality of different product streams at different treatment stages. To compare the liquid yields of composites and resins, a common denominator is needed. Similarly, the strength-retaining effect is also

influenced by the reference fibre, conditioning method, gauge length and broken filament treatment. Consistent definitions and test conditions are required to attribute differences to improved operating performance. Therefore, the current formula is based on a dry-feed basis for yield and an explicit pair for strength, and secondary cleaning is excluded from the main process.

Previous studies also examine how feed preparation and reactor configuration affect treatment performance. Xu et al. [10] investigated swelling-assisted low-temperature pyrolysis for glass fibre recovery, showing how pretreatment can modify the conditions required for thermal recycling. Transfer to the untreated blade piece needs to consider the extra material and energy. Resin chemistry, filler content and thermal gradients can alter apparent response trends. These factors become particularly important when a model is fitted to a small experimental sample and applied beyond the sampled conditions. Thus, most optimization problems require a local domain of applicability and explicitly bounded operating conditions. The research gap addressed here is the decision procedure in that area: how to keep competing product and energy outcomes while avoiding the treatment of uncertain fibre quality predictions as established acceptance evidence.

2.2. Probabilistic surrogate modelling

Gaussian process models can be used to interpolate costly observations and add uncertainty to the predictions. The covariance structure expresses assumptions about smoothness, characteristic variation scales, and the relationship among adjacent operating conditions. The foundational probabilistic formulation of Gaussian processes distinguishes latent function uncertainty from observation noise [11]. This distinction determines how repeated measurements contribute to uncertainty estimation. Repetition can estimate the variability of the observation process, and new settings can improve the coverage of the response surface. Use of an uncertainty band in an operating constraint requires an explicit definition of the variability it represents. Independent models are convenient when the output units and noise levels of the outputs are different; cross-response dependencies need to be explicitly modelled. A model that predicts attractive yield and low energy separately may assign these outcomes to combinations that are jointly less likely than the marginal intervals suggest.

Reliable interpretation of fitted variance requires assessment of model adequacy and predictive calibration. Cabral et al. [12] developed diagnostic statistics to assess model assumptions and examined how deviations from those assumptions affect statistical inference. This assessment is important because small point prediction errors can coexist with poorly calibrated uncertainty, particularly near the boundaries of the design domain. Hyperparameter fitting, output standardization, and the selection of a validation split consequently influence both predictive performance and the conservatism of the acceptance policy. Methodological work on data-driven interval calibration provides tools to adjust uncertainty statements when exchangeability and a suitable calibration sample can be justified [13]. Batch effects, chemistry changes, and input-dependent errors require separate assessment. The remaining engineering challenge is to incorporate calibrated prediction intervals into evolutionary search while respecting the evidential limits of a finite assessment sample.

2.3. Evolutionary optimization under limited evidence

Multiobjective optimization is used to address operating problems that require improving several functions simultaneously. Within an evaluated candidate set, a nondominated solution has no competing candidate that is at least as good in every objective and strictly better in at least one. The original NSGA-II study provides the historical basis for elitist nondominated sorting and diversity preservation [14]. Here it supports the ranking mechanism, while the recent studies in this section address surrogate modelling and uncertainty-aware application. The structure is suitable for process engineering, and if energy prices, product requirements or equipment limitations change, a different operating point can be selected. The credibility of the candidate front depends on the reliability of the underlying function evaluations. A surrogate model can be used to reduce the number of costly observations during optimization. The reliability of a densely sampled candidate front is related to the physical evidence supporting the predicted response.

Surrogate-assisted evolutionary methods address this problem by means of selective re-evaluation and model management. Research on expensive multi-objective optimization is related to selecting new evaluations for the search progress and uncertainty [15]. Exploring an uncertain region to improve the model and recommending an operating condition serve distinct purposes. Exploration may prioritize candidates with high uncertainty, while operating recommendations generally require acceptable conservative bounds. Combining the above roles may lead to a recommendation based on the convenience of obtaining sample information. A new structure for the exploration queue and the acceptance archive has been established to organize the comparison process. It is a way to specify the probabilistic results, a preserved strength threshold, a conservative objective ranking, and traceable operating alternatives. The numerical analysis provides the necessary data for validation of the proposed combination with existing optimization methods.

3. Process Definition and Gaussian Process Surrogates

3.1. Boundary and case study design

The start of the process boundary is dried and size-reduced blade composite entering an inert batch reactor, and the end is separated condensate, noncondensable gas and solid residue. The numerical energy balance excludes nitrogen production, upstream blade cutting, downstream fibre resizing and secondary oxidation. The reported energy reduction therefore applies to this reactor and separation boundary. Extending the comparison to a recycling facility requires adding the excluded utility and treatment loads using a consistent functional unit. Liquid response is the condensate mass divided by total dry composite feed mass, including the reinforcement fraction. Feedstocks with different resin contents therefore require composition-matched comparisons or an additional resin-normalized response. The energy input is the external supply of heat and electricity within the specified area per kilogram of dry feed. Retained strength is the recovered filament strength expressed as a percentage of a matched untreated reference. A measured comparison requires the same conditioning, gauge length, diameter measurement method, and treatment of filament failures for both groups. The analysis is for a glass fibre composite feed with a fixed composition; extension to carbon fibre spars or mixed blade waste requires a separate assessment.

Four continuous controls are specified: temperature from 430 to 570 degrees Celsius, heating rate from 5 to 25 degrees Celsius per minute, holding time from 15 to 75 minutes, and nitrogen flow from 0.4 to 1.2 litres per minute. These limits define the modelling domain. The retained strength threshold is set at 70%. A nominal design allocates 120 cases to development, 40 to calibration, and 40 to final assessment. The allocation separates model development, interval calibration, and independent assessment. The distinction prevents an attractive plot from being mistaken for a completed model validation. Replicates, batch identifiers, and instrument records support traceability and assessment of variability.

Figure 1 presents the process and information boundary, linking feed preparation, controlled heating, product separation, response recording, and the three data partitions. The flow distinguishes material transfer from information transfer so that a later implementation can identify where uncertainty enters the workflow. Reactor inputs are recorded before modelling, and measured outputs remain traceable to a case identifier. The surrogate depends on the product balance and mechanical test as its response evidence. Response definitions and measurement procedures are established before surrogate fitting. This ordering matters because a numerical model can fit inconsistently defined targets with apparently small errors while preserving an underlying accounting mistake.

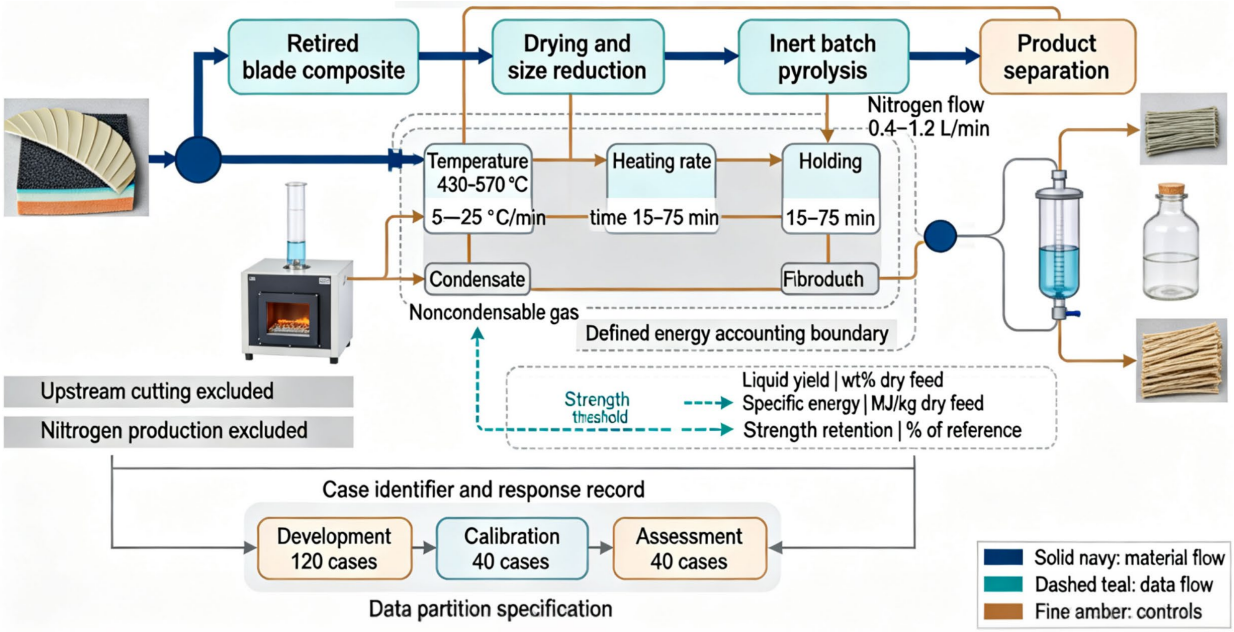


Figure 1. Process boundary and data flow from blade preparation to partitioned surrogate records

Each input is linearly scaled to the unit interval using the declared lower and upper limits. The response equation below separates the latent process function from observation noise. The subscript identifies liquid yield, specific energy demand, or retained strength, and the covariance describes uncertainty in the corresponding latent response. Independent Gaussian process models are specified for the three standardized outputs, linking the process records to the probabilistic representation used in the search.

$$f_j(\boldsymbol{x}) \sim \mathcal{GP}\left(0, k_j(\boldsymbol{x}, \boldsymbol{x}')\right), y_j = f_j(\boldsymbol{x}) + \varepsilon_j \tag{1}$$

3.2. Covariance structure and response scaling

Input scaling gives the covariance length parameters a comparable interpretation across temperature, time, rate, and flow. A squared exponential covariance with separate length scales specifies spatial correlation in the kernel equation. This smoothness assumption makes the influence of each normalized input inspectable and requires checking against observed response variation. Alternative covariance families should be evaluated if held out residuals indicate abrupt changes or systematic local error. Hyperparameter selection is restricted to the development data, with multiple initializations and a small numerical diagonal term to stabilize covariance factorization. The diagonal stabilization term supports numerical conditioning and must be recorded separately from observation noise. Reproduction of a fitted model requires the amplitude, four length scales, noise estimate, stabilization magnitude, and optimization initialization for each response; numerical fitted values are outside the present model specification.

$$k_j(\mathbf{x}, \mathbf{x}') = s_j^2 \exp\left[-\frac{1}{2}\sum_{r=1}^{4} \frac{(x_r - x_r')^2}{\ell_{jr}^2}\right] \tag{2}$$

The covariance amplitude determines the scale of latent variation, while the individual length parameters determine how quickly responses can vary along each normalized coordinate. Output standardization is reversed before errors or engineering constraints are evaluated. Predictions for energy must therefore return to the energy unit used in the process boundary, and strength predictions must return to percentage points. Constraints are applied after restoring engineering units. The development covariance matrix and candidate covariance vector provide the quantities needed for conditional prediction. Their values depend on the fitted amplitude, length scales, and observation noise. Consistent scaling across these quantities preserves the interpretation of uncertainty when a candidate is evaluated in engineering units.

3.3. Conditional prediction and calibration

Conditioning on the development observations yields the predictive equation for the mean and latent variance. Here, the training covariance matrix includes the observation noise contribution on its diagonal, while the candidate covariance vector links the evaluation point to the development cases. The target vector contains standardized responses for the output being modelled.

$$\begin{aligned} \mu_j(\boldsymbol{x}) &= \boldsymbol{k}_*^\top \boldsymbol{K}_j^1 \boldsymbol{y}_j, \\ \sigma_j^2(\boldsymbol{x}) &= k_j(\boldsymbol{x}, \boldsymbol{x}) - \boldsymbol{k}_*^\top \boldsymbol{K}_j^{-1} \boldsymbol{k}_* \end{aligned} \tag{3}$$

The implementation should evaluate the covariance system through Cholesky factorization. The inverse notation specifies the mathematical operation implemented through this factorization. A predictive interval for a future observation includes the observation variance in addition to the latent variance. The framework uses latent uncertainty when comparing mean process responses and observation uncertainty when screening whether a future measured strength could fall below the threshold. The objective penalties use the latent standard deviation, whereas strength screening uses the square root of latent variance plus observation variance. Observation variance enters this screening spread once, and each uncertainty statement retains its own decision interpretation. If repeated tests show that variability depends on temperature or batch, a constant noise term would need revision.

Calibration data are reserved for examining interval coverage and, where justified, selecting an inflation factor that widens the fitted prediction spread. The final assessment sample remains reserved for independent evaluation. Diagnostic checks include paired predictions and targets, residual distributions, interval coverage, and the relation between error and distance from the development design. A model with low average error can still be unsuitable for constrained optimization if it underestimates uncertainty close to the strength threshold. These diagnostics support assessment of local prediction reliability and interval performance. Application requires traceable links between plotted points and independent measurements, disclosure of replicate counts, and explicit estimation of coverage uncertainty.

4. Uncertainty Aware Multiobjective Search

4.1. Conservative objectives and acceptance

The search seeks high liquid yield and low energy demand while screening retained strength. To use a common minimization convention, the yield objective changes sign. Uncertainty enters through conservative bounds: predicted yield is reduced by a multiple of its spread, energy is increased, and retained strength is reduced before comparison with the acceptance limit. The multiplier is an explicit decision parameter. Its probabilistic interpretation depends on the surrogate assumptions and calibration evidence. A multiplier of two is used as a decision setting for examining

conservatism. A nominal Gaussian tail interpretation requires an adequate predictive distribution and independent calibration; the multiplier alone establishes no verified acceptance probability. The resulting objective equation places a penalty on predictive spread for both responses.

$$\boldsymbol{F}(\boldsymbol{x}) = [-\mu_Y(\boldsymbol{x}) + \beta\sigma_Y(\boldsymbol{x}), \mu_E(\boldsymbol{x}) + \beta\sigma_E(\boldsymbol{x})] \quad (4)$$

Objective values are scaled using fixed engineering ranges before distance-based decisions are made. These ranges are fixed before method comparison to preserve a consistent definition of improvement as the archive evolves. Strength uses an observation scale spread when the intended decision concerns acceptance of a future test result. A candidate whose mean strength exceeds the threshold can therefore remain unacceptable when its uncertainty is large. Conversely, a less productive setting can remain attractive when it has a larger margin to the acceptance boundary. Feasibility is assessed separately from objective ranking through the lower-bound equation, with the acceptance threshold applied in the original strength units.

$$L_S(\boldsymbol{x}) = \mu_S(\boldsymbol{x}) - \beta\sigma_{S,obs}(\boldsymbol{x}) \geq S_{min} \quad (5)$$

A lower bound above the threshold establishes model-based acceptance under the specified screening convention. Operational approval and fibre reuse require separate process and material qualification. Other restrictions, including temperature limits, gas delivery capacity, and residence time, are handled as hard bounds on the decision variables. Adding several quality conditions requires assessment of their dependence before marginal bounds can support a joint probability statement. The acceptance archive stores the complete predicted response vector, uncertainty components, decision variables, and model version for each retained candidate.

4.2. Evolutionary ranking and selective evaluation

The proposed implementation begins with a space filling population inside the admissible input box. A population of 80 candidates and 60 generations defines the proposed configuration. A repeatable implementation must additionally record crossover and mutation probabilities, their distribution indices, the seed sequence, and the allocation of costly evaluations. Binary tournament selection compares feasibility first, then nondomination rank, then crowding distance. Simulated binary crossover and polynomial mutation generate offspring, with bounds repaired before surrogate evaluation. Feasible candidates dominate infeasible ones; among infeasible candidates, the smaller normalized violation is preferred. Parent and offspring populations are combined before truncation, preserving elitism while retaining spread along the conservative front. Search settings, random seeds, stopping conditions, and the number of expensive evaluations must be disclosed when numerical performance is eventually reported.

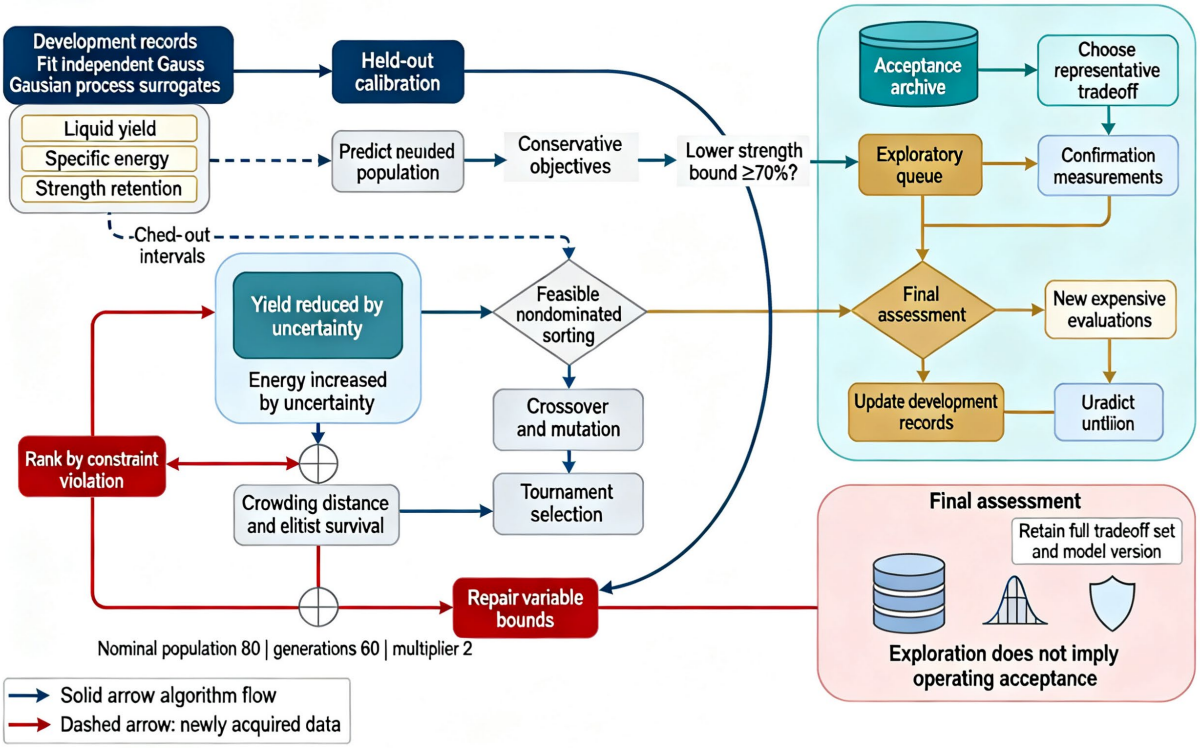


Figure 2. Uncertainty aware evolutionary search with separate exploratory and acceptance archives

Figure 2 describes the optimization loop from model fitting through conservative screening, ranking, variation, and selective new evaluation. An exploratory candidate queue is maintained separately from the recommendation archive.

Candidates in the queue may be uncertain and temporarily infeasible because their purpose is to reduce ambiguity in a useful region. Their archive status records an information gathering role. New evaluations should be chosen using a declared mixture of front coverage, uncertainty, and distance from existing cases. The mixture weights, batch size, tie-breaking rule, and update schedule form part of the implementation record; the current description specifies their roles rather than a uniquely reproducible acquisition policy. After each evaluation batch, the development model is updated while the final assessment set remains untouched. This separation supports evaluation of how newly acquired information changes the quality of operating decisions.

4.3. Compromise choice and comparison protocol

Selection of a representative operating point follows construction of the conservative archive. Each objective is mapped to the fixed normalized range and compared with the corresponding ideal coordinate. The analysis uses equal weights for this reporting step, while retaining the full front for other preferences. Equal weighting is a reporting preference whose effect depends on the fixed normalization endpoints. Reporting those endpoints and a weight-sensitivity check is required for a reproducible distance ranking; candidate C is examined here as the designated compromise, while economic selection requires application-specific valuations. Among admissible candidates, the representative point minimizes distance to the ideal coordinates. This choice is expressed by the weighted-distance equation, whose weights encode the stated reporting preferences.

$$D(\boldsymbol{x}) = \sqrt{\sum_{j=1}^{2} w_j \left[\tilde{F}_j(\boldsymbol{x}) - \tilde{F}_{j,\,ideal}\right]^2} \tag{6}$$

An eventual experimental comparison should include direct search with measured evaluations, mean only surrogate search, and uncertainty aware search under the same expensive evaluation budget. Random seeds should be paired across methods, and performance should be measured after reevaluating recommended candidates using the same response protocol. Archive hypervolume, feasible recommendation rate, and evaluation count answer different questions and should be reported separately. Measured reevaluation determines whether a larger surrogate front corresponds to improved process performance. Individual strength failures should be reported alongside average performance. The proposed comparison separates objective performance from constraint satisfaction. The displayed policy coordinates establish numerical trade-offs under stated assumptions; evidence of algorithmic superiority requires the matched-budget evaluations and repeated trials described here.

Stopping criteria should combine the expensive evaluation budget with a check on recent improvement and remaining uncertainty. Stopping because two successive fronts look similar can be misleading if both rely on the same poorly informed surrogate region. A practical implementation would retain a final batch for confirmatory runs at the selected compromise and at nearby operating settings. Those runs would assess the operating margin around the recommended setting and its sensitivity to small control deviations. The final assessment would report failures as well as successes and preserve all evaluated settings in the data record. This protocol supports comparison of the proposed method with alternatives using consistent evaluation criteria.

5. Experimental Data and Engineering Analysis

5.1. Process responses and predictive accuracy

Thermochemical recovery studies define liquid yield relative to a specified feed basis to support comparisons between operating conditions [16]. Figure 3(a) shows yield of 19.2, 23.8, 27.4, 28.1, and 27.6 wt% at 450, 470, 490, 510, and 530 degrees Celsius, respectively. The gain falls from 4.6 percentage points over the first temperature increment to 0.7 percentage points between 490 and 510 degrees Celsius, followed by a 0.5 percentage point decrease. This response identifies diminishing recovery gains near the upper part of the temperature range. The respective interval half widths are 1.3, 1.1, 0.9, 1.0, and 1.4 percentage points. The intervals at 490 and 510 degrees Celsius overlap. Establishing a preferred temperature from their yield difference requires paired uncertainty assessment and confirmation measurements; the central maximum alone provides insufficient evidence for that selection. Fibre recovery research treats retained mechanical properties as a separate quality response [17]. Figure 3(b) records strength retention decreasing from 86 to 61% across the same sweep. At 490 degrees Celsius, retention is 76%, providing a six-percentage point central margin above the 70% requirement. The margin becomes negative at 510 degrees Celsius, where retention reaches 69%. Temperature selection therefore depends on the joint recovery and quality response, with the fibre threshold restricting the usefulness of the small yield gain near the peak.

Figure 3(c) gives specific energy demands of 5.10, 5.55, 6.20, 6.95, and 7.80 MJ per kilogram across the temperature sweep. Successive increments rise from 0.45 to 0.65, 0.75, and 0.85 MJ per kilogram, indicating an increasing energy penalty for each additional temperature step within the defined accounting boundary. Between 490 and 510 degrees

Celsius, a 12.1% increase in energy accompanies a 2.6% relative increase in central liquid yield. The different rates of change sharpen the tradeoff near the productive region. Studies of composite pyrolysis product distributions motivate examination of nonmonotonic response patterns and their dependence on thermal exposure [18]. Figure 3(d) resolves this interaction using five temperatures and five holding times. At 490 degrees Celsius, the map gives yields of 26.2, 27.1, 27.4, 27.1, and 26.2 wt% at 21, 33, 45, 57, and 69 minutes. The central holding time exceeds either extreme by 1.2 percentage points, while the neighbouring times differ by 0.3 percentage points. This broad local plateau supports confirmation of an operating region around 45 minutes. The symmetric map describes the displayed interaction. Attribution to decomposition kinetics, vapour transfer, or secondary reactions requires time-resolved gas release, vapour residence information, and condensate composition measured across the same temperature–time region.

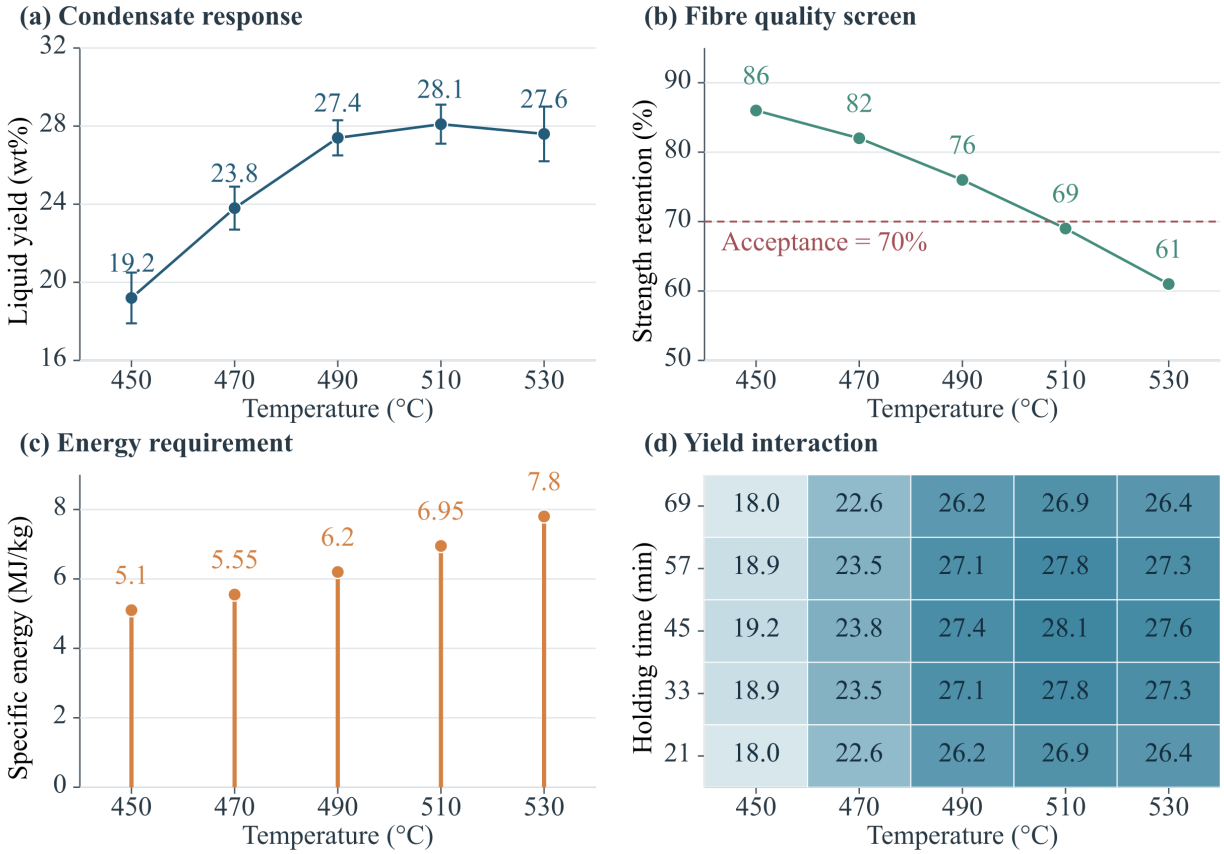


Figure 3. Process responses: (a) liquid yield with response intervals; (b) retained strength and acceptance threshold; (c) specific energy demand; (d) temperature and holding time yield map

Probabilistic regression assessment separates the accuracy of central predictions from the quality of their uncertainty estimates [19]. Figure 4(a) pairs eight target yields of 19.2, 21.0, 23.8, 25.2, 27.4, 28.1, 27.6, and 24.7 wt% with predictions of 19.8, 20.5, 24.1, 24.5, 27.0, 28.8, 26.9, and 25.3 wt%. Prediction errors extend from negative 0.7 to positive 0.7 percentage points, and the calculated root mean square error is 0.58 percentage points. Four predictions exceed their targets and four fall below them; the mean signed error is approximately negative 0.01 percentage points. This near balance indicates little aggregate offset among the displayed pairs. It provides limited information about local bias because opposing errors can cancel, and the eight displayed pairs represent a diagnostic subset rather than the complete reserved assessment record. The largest absolute errors occur at targets of 25.2, 28.1, and 27.6 wt%, spanning the more productive portion of the displayed range. Local assessment in that region is consequential because the optimizer preferentially searches settings with favourable predicted yield. The eight pairs characterize the displayed diagnostic sample; assessment of the reserved case population requires its complete prediction record and corresponding targets.

Figure 4(b) compares coverage values of 52.5, 65.0, 77.5, 87.5, and 92.5% with nominal levels of 50, 60, 70, 80, and 90%. The corresponding excesses are 2.5, 5.0, 7.5, 7.5, and 2.5 percentage points. Prediction interval evaluation considers coverage together with interval width because excessively broad intervals can satisfy coverage requirements while providing weak discrimination between candidates [20]. Here, the largest discrepancy occurs at the intermediate nominal levels, making a single high coverage checkpoint insufficient to characterize the whole calibration pattern. Under a 40-case assessment design, each case changes empirical coverage by 2.5 percentage points. For that assumed denominator, the displayed excesses are arithmetically equivalent to one, two, three, three, and one additional covered cases relative to the nominal counts. Verification of empirical coverage requires the case-level inclusion records and

uncertainty on the resulting proportion. That resolution makes the calibration sample size important when interpreting apparently precise percentages. A subsequent assessment should retain the interval endpoints and coverage outcome for every case, permitting direct examination of whether the discrepancy concentrates near the operating boundary or in the productive region. Interval width should be examined over the same subsets to determine whether additional coverage comes from informative uncertainty estimates or broadly expanded bounds.

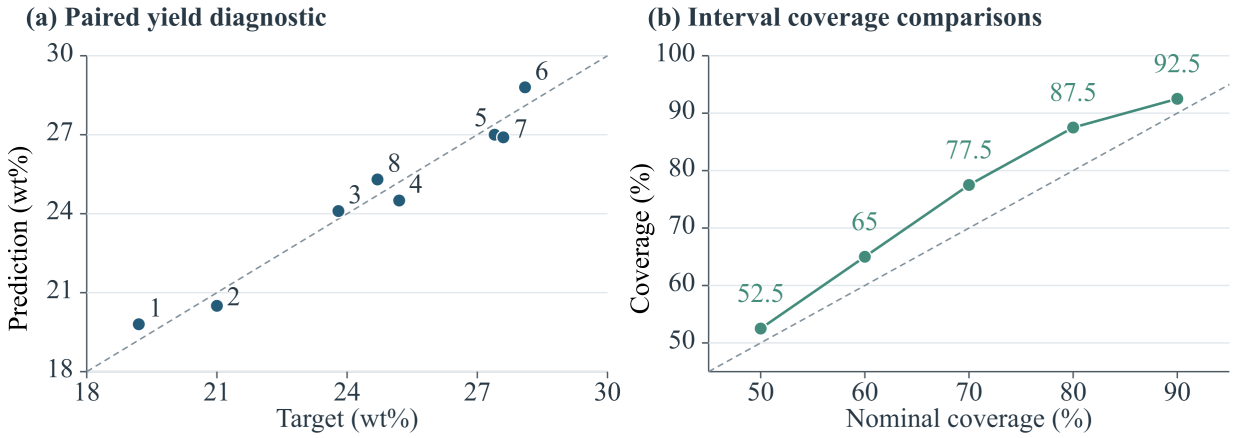


Figure 4. Prediction diagnostics: (a) target and prediction pairs; (b) interval coverage across nominal levels

5.2. Uncertainty diagnostics and candidate selection

Figure 5(a) displays twelve standardized residuals at negative 1.5, negative 1.2, negative 0.8, negative 0.6, negative 0.3, negative 0.1, positive 0.1, positive 0.3, positive 0.6, positive 0.8, positive 1.2, and positive 1.5. The series is symmetric around zero, with six values on either side and all values within two standardized units. Its symmetry is consistent with balanced directional error in this diagnostic series. The six histogram bins contain one, two, three, two, two, and two residuals, using edges of −1.8, −1.2, −0.6, 0, 0.6, 1.2, and 1.8. Bins are left-closed and right-open, with the final upper edge included; horizontal positions represent residual intervals rather than individual observations. The unequal central bin count reflects the left-closed boundary convention: values exactly on the edge of the bin enter the right-side bin. Interpretation therefore benefits from the listed residual coordinates as well as the histogram. Surrogate validation methods examine residual structure across operating regions to identify local model mismatch [21]. The same overall histogram could conceal a negative offset at high temperatures and a positive offset at low temperatures, so a complete validation should retain temperature, holding time, and batch identifiers for each residual. That information would help distinguish broad uncertainty miscalibration from a localized response feature requiring a different covariance structure.

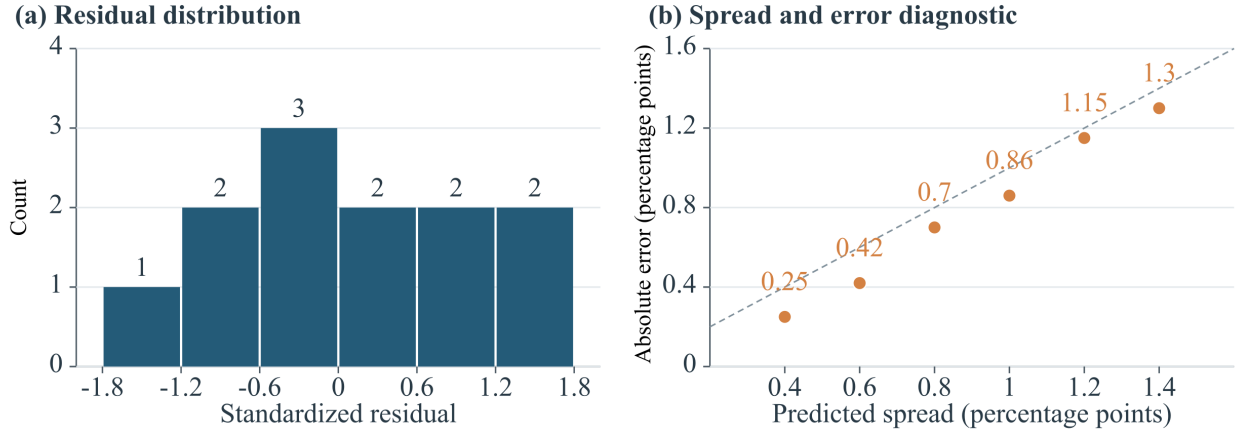


Figure 5. Uncertainty diagnostics: (a) standardized residual histogram; (b) predicted spread against absolute prediction error

Figure 5(b) relates predicted spreads of 0.4, 0.6, 0.8, 1.0, 1.2, and 1.4 percentage points to absolute errors of 0.25, 0.42, 0.70, 0.86, 1.15, and 1.30 percentage points. Absolute error increases with spread, indicating that the displayed uncertainty ordering identifies progressively less accurate predictions. Error to spread ratios range from approximately 0.63 to 0.96, with every point below the equality diagonal. The gap between spread and error ranges from 0.05 to 0.18 percentage points, so the comparison concerns relative ordering as well as absolute scale. An uncertainty estimate that differentiates candidates can guide exploratory evaluations towards settings where further information is valuable. Evaluation under data shift requires separate attention when feed chemistry or process configuration changes [22]. The

observed ordering within this diagnostic series therefore motivates checks on an independent feed batch and on cases near the edge of the input domain. Repeated response measurements at matched settings would help separate observation variability from uncertainty caused by limited development coverage. That separation affects whether a larger spread should prompt additional replication or exploration of a new operating condition.

Multiobjective decision analysis separates the generation of competing alternatives from the preference used to select a representative point [23]. Figure 6(a) presents candidates A to F with energy demands of 5.50, 5.80, 6.20, 6.60, 7.00, and 7.40 MJ per kilogram and yields of 23.5, 25.2, 27.4, 28.0, 28.4, and 28.6 wt%. Moving from B to C raises yield by 2.2 percentage points for an additional 0.40 MJ per kilogram. The same energy increment from C to D produces a gain of 0.6 percentage points, followed by gains of 0.4 and 0.2 percentage points over the next two increments. This decreasing incremental return makes the region around C relevant to compromise selection. Consistent objective definitions and reference settings are required for meaningful multiobjective comparisons [24]. The designated reference gives 24.0 wt% yield at 7.30 MJ per kilogram; candidate C therefore provides a 3.4 percentage point yield increase and a 15.1% energy reduction relative to those reference values. Candidate C corresponds to 490 degrees Celsius, 15 degrees Celsius per minute, 45 minutes, and 0.8 litres per minute. Table 1 lists the numerical responses and acceptance outcomes for all six candidates.

Table 1. Numerical comparison of candidate operating alternatives and conservative acceptance outcomes

Candidate	Liquid yield (wt%)	Energy (MJ/kg)	Mean strength (%)	Lower strength (%)	Accepted ≥70%
A	23.5	5.50	84	81	Yes
B	25.2	5.80	80	77	Yes
C	27.4	6.20	76	72	Yes
D	28.0	6.60	73	69	No
E	28.4	7.00	71	67	No
F	28.6	7.40	68	64	No

Figure 6(b) gives lower strength bounds of 81, 77, 72, 69, 67, and 64% for candidates A to F. Relative to the 70% requirement, the respective margins are positive 11, positive 7, positive 2, negative 1, negative 3, and negative 6 percentage points. A, B, and C remain acceptable under the conservative screen. D and E meet the central strength requirement at 73 and 71%, respectively, while their lower bounds fall below the threshold. The difference between central and conservative screening thus changes the classification of two of the six alternatives. Table 1 should be read using the lower-bound criterion: a central strength above the threshold does not by itself establish acceptance, and C's remaining margin is a model-based screening margin. Constraint handling methods can use violation magnitude to rank infeasible candidates during evolutionary selection [25]. D has the smallest violation and may consequently remain informative for further evaluation, although its current acceptance status is negative. For C, the 76% mean and two percentage point observation spread yield a 72% lower bound at the specified multiplier. Its positive margin supports acceptance under that convention, while the smaller margin relative to B indicates greater sensitivity to predictive bias. These differences provide a quantitative basis for choosing between higher recovery and a wider quality margin under the purchaser's strength specification.

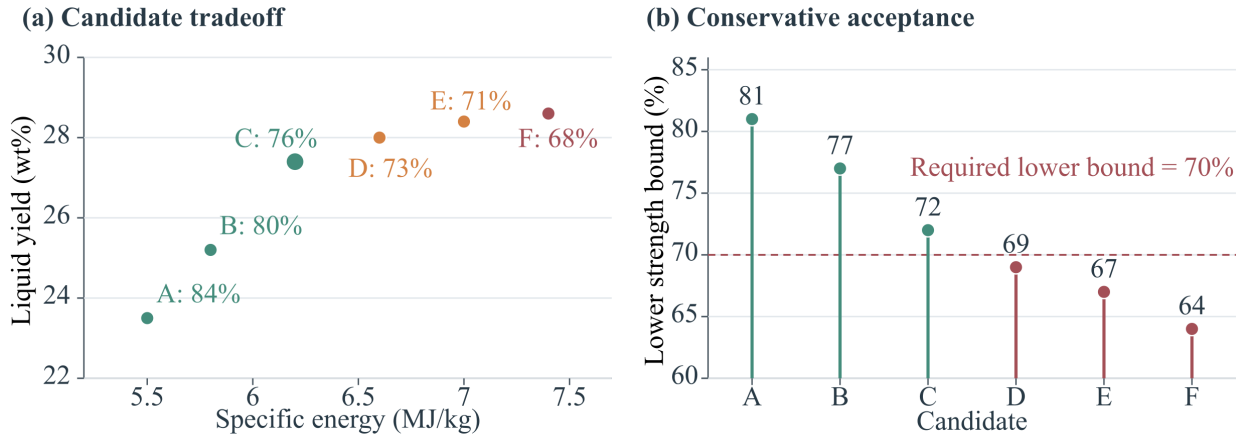


Figure 6. Operating alternatives: (a) energy and yield tradeoff with strength encoded by colour; (b) conservative strength bounds and acceptance threshold

5.3. Policy sensitivity and operating margins

Robust process design distinguishes uncertainty in operating conditions from uncertainty in predicted responses [26]. Figure 7(a) compares Reference, Mean Only, Fixed Margin, and Adaptive Margin policies with yields of 24.0, 28.3, 26.6, and 27.4 wt%, respectively. Adaptive Margin gains 3.4 percentage points over Reference and 0.8 percentage points over Fixed Margin, while its central yield is 0.9 percentage points below Mean Only. This descriptive ranking quantifies the central yield difference associated with the stated policy alternatives. Establishing a policy performance difference requires a common trial protocol and uncertainty estimates derived on the same basis. Figure 7(b) gives corresponding energy demands of 7.30, 6.00, 6.45, and 6.20 MJ per kilogram. Adaptive Margin reduces energy by 1.10 MJ per kilogram relative to Reference and by 0.25 MJ per kilogram relative to Fixed Margin. Its energy exceeds Mean Only by 0.20 MJ per kilogram. The displayed half widths are 0.30, 0.50, 0.28, and 0.20 MJ per kilogram, giving intervals of 7.00–7.60, 5.50–6.50, 6.17–6.73, and 6.00–6.40 MJ per kilogram. Overlap between the latter three intervals makes their central ranking a decision input that should be assessed alongside acceptance performance and interval interpretation.

Figure 7(c) gives acceptable recommendation frequencies of 60, 70, 85, and 90% for the four policies. Adaptive Margin exceeds Mean Only by 20 percentage points and Fixed Margin by five percentage points. Read together with the central yield differences, this comparison makes the reliability cost of pursuing the highest recovery explicit. The frequency denominator, assessment conditions, and definition of an accepted recommendation are required before these percentages can be interpreted as estimated success probabilities. The policy frequencies and the following learning-curve ranges are descriptive coordinates; trial-level outcomes and refitting records are required to establish their sampling interpretation. Learning curve analysis relates surrogate prediction error to development design size [27]. Figure 7(d) gives median yield errors of 1.80, 1.35, 1.02, 0.79, and 0.62 percentage points for 40, 60, 80, 100, and 120 development cases. The reduction across the full range is 65.6%, with successive absolute improvements of 0.45, 0.33, 0.23, and 0.17 percentage points. The declining increments identify diminishing gains from each additional batch. At 120 cases, the box spans approximately 0.53–0.71 percentage points and the whiskers span 0.43–0.81, following the stated relative spreads. These descriptive ranges help assess the remaining error scale against differences between neighbouring candidate yields, while evaluation budgets determine whether further sampling is justified.

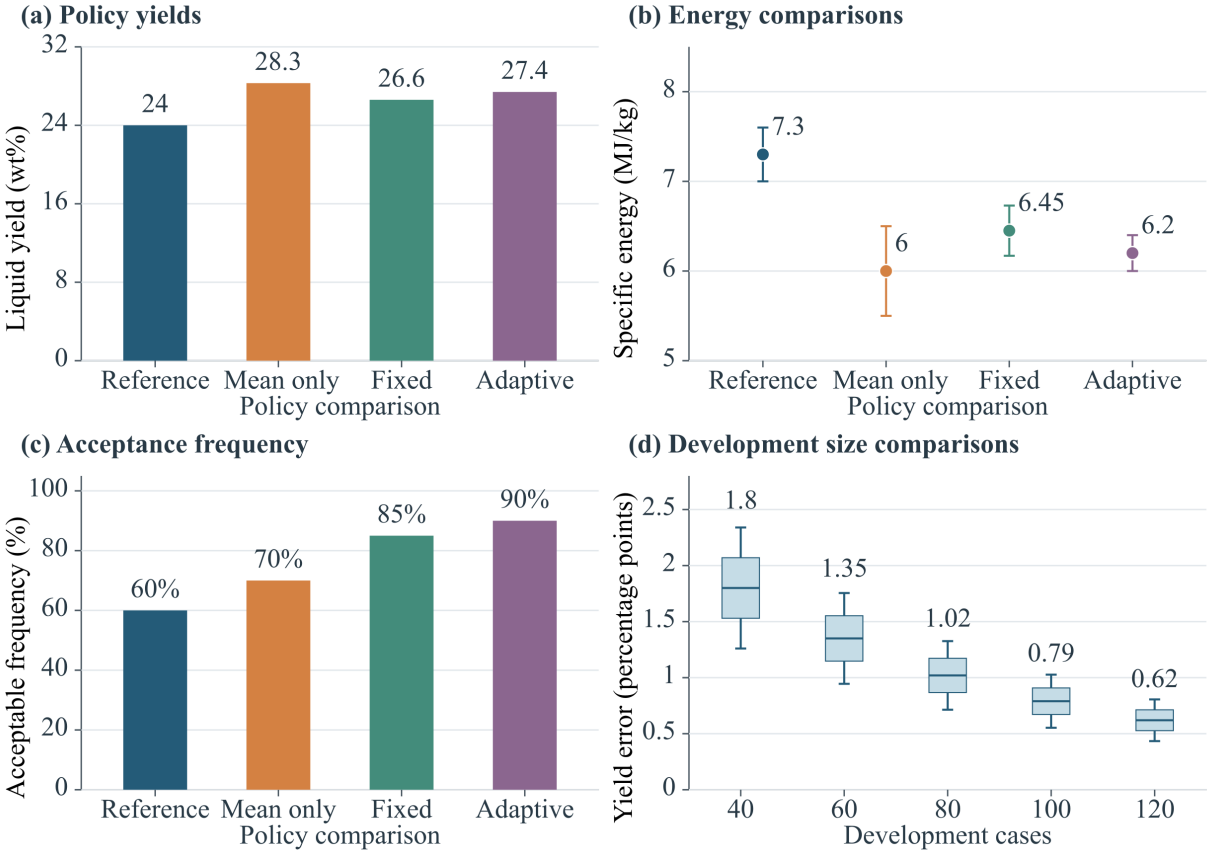


Figure 7. Policy and data size comparisons: (a) liquid yield; (b) specific energy and descriptive intervals; (c) acceptable recommendation frequency; (d) yield error across development sizes

Figure 8(a) examines candidate C with uncertainty multipliers of 0, 0.5, 1.0, 1.5, and 2.0, resulting in conservative strength margins of six, five, four, three, and two percentage points, respectively. Each increase of 0.5 consumes one percentage point of margin because the reported strength spread is two percentage points. Robust optimization methods examine this dependence to identify how uncertainty penalties change the admissibility of a fixed decision [28]. At the

largest displayed multiplier, the lower strength bound remains 72%, leaving two percentage points above the requirement. A downward shift of two percentage points in the predicted mean would exhaust that remaining margin. Alternatively, with the mean fixed at 76%, increasing the spread from two to three percentage points would place the lower bound exactly at the threshold. The two perturbations identify distinct validation priorities: estimation of local bias and estimation of observation variability. The operating setting remains fixed throughout the sweep, so the linear margin change follows algebraically from the screening convention. The sweep tests sensitivity to the uncertainty penalty; evidence of tolerance to temperature, holding-time, or flow deviations requires responses at perturbed operating settings. Its interpretation concerns sensitivity of acceptance to uncertainty assumptions, while control deviations require evaluation of neighboring process settings.

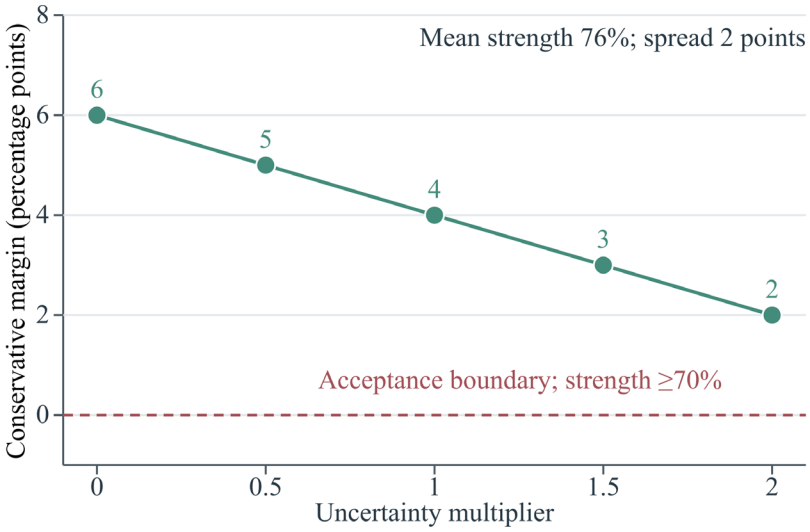


Figure 8. Robustness sweep

Experimental design for expensive response systems motivates confirmation measurements around a selected operating region [29]. For candidate C, nearby temperatures and holding times would test whether the two-percentage point conservative margin persists under achievable control tolerances. Repeated central runs and paired reference tests across feed batches would help distinguish systematic prediction offset from filament variability. The mass and energy boundaries would remain consistent across these runs, with raw force, filament diameter, collected condensate mass, and utility readings retained for assessment. A confirmation campaign should allocate sufficient observations to estimate uncertainty on the same scale as the remaining acceptance margin. If local mean bias approached two percentage points, the candidate would require additional evidence or a revised margin before recommendation. If strength spread approached three percentage points, the present central estimate would reach the conservative acceptance boundary at the selected multiplier. Under those conditions, candidate B's seven percentage point lower-bound margin would become relevant to the operating decision despite its lower yield. The comparison would then depend on the measured quality tolerance and control capability required to sustain each setting.

6. Conclusion

This study develops a process-specific framework coupling Gaussian process surrogates with uncertainty-aware evolutionary optimization of wind turbine blade pyrolysis. The formulation treats liquid recovery and energy demand as separate objectives and uses retained strength to constrain the acceptable operating range. Conservative objective values guide candidate ranking, and a lower strength bound determines admission to the recommendation archive. An independent exploratory queue can be used to obtain information about unselected candidate processes. A case study shows how these differences can be conveyed by traceable response specifications, diagnostic modules and operating modes. The framework provides a transparent decision specification and traceable numerical criteria for subsequent validation. Its present contribution is the formulation and interpretation of uncertainty-aware operating choices; qualified operation depends on the measurement and implementation evidence identified in the preceding sections.

The current system has a limited operating range and a fixed feed specification. Physical response trend, uncertainty calibration, comparative algorithm performance and operating window confirmation all require measurement validation. The fixed feed composition and process boundary limit the range of plant variability represented in the analysis. Separate response models require further development to capture cross-response dependence, while abrupt changes associated with incomplete decomposition or heterogeneous samples may require alternative covariance structures. The acceptance rule is based on both the correctness of the uncertainty estimate and the applicability of the selected fibre quality indicator. Confirmation measurements are needed to ensure the reliability of the recommended front.

Future work will expand reactor assessment and independent product characterization for other feed compositions. Multiple tests of the feed batch can be conducted to distinguish between measurement errors, variations in materials, and sampling-related uncertainties. Further development will explore alternative covariance structures, dependent response models, and candidate selection policies with the same experimental budget. Confirmation experiments in the chosen trade-off region will determine the actual margin for the operator. Extension of the boundary to the downstream fibre treatment and utility supply would then enable an assessment of whether the proposed conditions are still feasible given the requirements for recovery quality and the entire process.

7. Declarations

7.1 Author Contributions

Odysseas Dellas contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Sotiris Kafatos contributes to validation, analysis, investigation, data collection, draft preparation. Thanasis Petrou contributes to conceptualization, validation, analysis. Achilleas Kallis contributes to validation, analysis, investigation, data collection. All authors have read and agreed with the manuscript before its submission and publication.

7.2 Funding and Acknowledgments

This research received no specific financial support from any funding agency.

7.3 Institutional Review Board Statement

Not applicable.

7.4 Informed Consent Statement

Not applicable.

7.5 Date Availability

Not applicable.

7.6 Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

7.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

8. References

[1] Meng, F., Hawkin, J. L., Gast, L., Job, S., Ó Brádaigh, C. M., Cullen, J. M., Leeke, G. A., McKechnie, J., Poulsen, A. K., Rothman, R., Ryan, A. J., Todd, I., Walton, A., Harper, G. D. J., Yang, Z., & Zhang, L. (2025). End-of-life management for wind turbines. Nature Reviews Clean Technology, 1(10), 677–698. https://doi.org/10.1038/s44359-025-00097-3

[2] Hu, Y., Zhang, Y., Li, Y., Wang, Y., Li, G., & Liu, X. (2024). Wind turbine blade recycling: A review of the recovery and high-value utilization of decommissioned wind turbine blades. Resources, Conservation and Recycling, 210, Article 107813. https://doi.org/10.1016/j.resconrec.2024.107813

[3] Royuela, D., Martínez, J. D., López, J. M., Callén, M. S., García, T., Verdejo, R., Murillo, R., & Veses, A. (2024). Pursuing the circularity of wind turbine blades: Thermochemical recycling by pyrolysis and recovery of valuable resources. Journal of Analytical and Applied Pyrolysis, 181, Article 106657. https://doi.org/10.1016/j.jaap.2024.106657

[4] Lu, Q., Ji, H., Lu, Y., Yang, J., Chen, W., Li, J., Li, W., & Xu, M. (2025). Transition-metal-assisted pyrolysis to recover glass fibers from end-of-life wind turbine blades. Journal of Analytical and Applied Pyrolysis, 189, Article 107081. https://doi.org/10.1016/j.jaap.2025.107081

[5] Zhang, B., Zhang, S., Yang, Z., Liu, W., Wu, B., Huang, M., & Liu, B. (2024). Pyrolysis process and products characteristics of glass fiber reinforced epoxy resin from waste wind turbine blades. Composites Part B: Engineering, 287, Article 111803.

https://doi.org/10.1016/j.compositesb.2024.111803

[6] Manfredi, P. (2024). Probabilistic uncertainty propagation using Gaussian process surrogates. International Journal for Uncertainty Quantification, 14(6), 71–104. https://doi.org/10.1615/int.j.uncertaintyquantification.2024052162

[7] Zheng, W., & Doerr, B. (2025). Approximation guarantees for the nondominated sorting genetic algorithm II (NSGA-II). IEEE Transactions on Evolutionary Computation, 29(4), 891–905. https://doi.org/10.1109/tevc.2024.3402996

[8] Kim, H., & Sanz-Alonso, D. (2025). Enhancing Gaussian process surrogates for optimization and posterior approximation via random exploration. SIAM/ASA Journal on Uncertainty Quantification, 13(3), 1054–1084. https://doi.org/10.1137/24m1677009

[9] Xu, M.-X., Ji, H.-W., Wu, Y.-C., Meng, X.-X., Di, J.-Y., Yang, J., & Lu, Q. (2024). Recovering glass fibers from waste wind turbine blades: Recycling methods, fiber properties, and potential utilization. Renewable and Sustainable Energy Reviews, 202, Article 114690. https://doi.org/10.1016/j.rser.2024.114690

[10] Xu, M., Yang, J., Ji, H., Wu, Y., Li, J., Di, J., Meng, X., Jiang, H., & Lu, Q. (2024). Recovering high-quality glass fibers from end-of-life wind turbine blades through swelling-assisted low-temperature pyrolysis. Waste Management, 187, 179–187. https://doi.org/10.1016/j.wasman.2024.07.020

[11] Liang, T., Wang, F., Wang, S., Li, K., Ma, X., & Mo, X. (2024). Towards unified aleatory and epistemic uncertainty quantification for machinery health prognostic through sequential heteroscedastic Gaussian process regression. Advanced Engineering Informatics, 62, Article 102719. https://doi.org/10.1016/j.aei.2024.102719

[12] Cabral, R., Bolin, D., & Rue, H. (2025). Robustness, model checking, and hierarchical models. Journal of the Royal Statistical Society Series B: Statistical Methodology, 87(3), 632–652. https://doi.org/10.1093/jrsssb/qkae107

[13] Papadopoulos, H. (2024). Guaranteed coverage prediction intervals with Gaussian process regression. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(12), 9072–9083. https://doi.org/10.1109/tpami.2024.3418214

[14] Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation, 6(2), 182–197. https://doi.org/10.1109/4235.996017

[15] Aghaei pour, P., Hakanen, J., & Miettinen, K. (2024). A surrogate-assisted a priori multiobjective evolutionary algorithm for constrained multiobjective optimization problems. Journal of Global Optimization, 90(2), 459–485. https://doi.org/10.1007/s10898-024-01387-z

[16] Yousef, S., Eimontas, J., Zakarauskas, K., Stasiulaitiene, I., Striūgas, N., & Tuckute, S. (2025). Catalytic pyrolysis of wind turbine blades waste for plasticizers recovery and its life cycle assessment. Journal of Environmental Management, 395, Article 127690. https://doi.org/10.1016/j.jenvman.2025.127690

[17] Ren, K., Liu, S., Kong, S., Ma, S., & Feng, R. (2026). Thermal recycling of wind turbine blade composites: Optimizing pyrolysis and oxidation for fiber recovery. Composites Part A: Applied Science and Manufacturing, 202, Article 109460. https://doi.org/10.1016/j.compositesa.2025.109460

[18] Zhang, Y., Zhang, Z., Ma, W., Chen, P., Bai, B., Li, L., Lai, Y., & Wang, X. (2025). Pyrolysis and oxidation characteristics and energy self-sustaining process design of retired wind turbine blades. Journal of Analytical and Applied Pyrolysis, 185, Article 106872. https://doi.org/10.1016/j.jaap.2024.106872

[19] Wang, Y., & Politis, D. N. (2026). Model-free bootstrap and conformal prediction in regression: Conditionality, conjecture testing, and pertinent prediction intervals. Journal of Nonparametric Statistics, 38(1), 311–345. https://doi.org/10.1080/10485252.2025.2603398

[20] Hajibabaee, P., Pourkamali-Anaraki, F., & Hariri-Ardebili, M. A. (2024). Adaptive conformal prediction intervals using data-dependent weights with application to seismic response prediction. IEEE Access, 12, 53579–53597. https://doi.org/10.1109/access.2024.3387858

[21] Ginsbourger, D., & Schärer, C. (2025). Fast calculation of Gaussian process multiple-fold cross-validation residuals and their covariances. Journal of Computational and Graphical Statistics, 34(1), 1–14. https://doi.org/10.1080/10618600.2024.2353633

[22] Trottner, L. (2024). Covariate shift in nonparametric regression with Markovian design. Information and Inference: A Journal of the IMA, 13(2), Article iaae011. https://doi.org/10.1093/imaiai/iaae011

[23] Pajasmaa, J., Miettinen, K., & Silvennoinen, J. (2025). Group decision making in multiobjective optimization: A systematic literature review. Group Decision and Negotiation, 34(2), 329–371. https://doi.org/10.1007/s10726-024-09915-8

[24] Zhang, H., Yi, W., Li, H., & Wang, S. (2025). Performance metrics evaluation of multi-objective optimization for transportation. Electronic Research Archive, 33(6), 3968–3988. https://doi.org/10.3934/era.2025176

[25] Cheng, X., Yang, Y., Zhu, C., Kong, X., & Liu, J. (2026). A constraint violation-based adaptive switching surrogate-assisted evolutionary algorithm for expensive constrained multiobjective optimization. Swarm and Evolutionary Computation, 104,

Article 102372. https://doi.org/10.1016/j.swevo.2026.102372

[26] Feng, Z., Wang, J., & Huang, W. (2026). An active learning multi-objective robust optimization framework via Gaussian process models under uncertainty. Engineering Optimization, 58(4), 976–999. https://doi.org/10.1080/0305215x.2025.2495800

[27] Dayimu, A., Simidjievski, N., Demiris, N., & Abraham, J. (2024). Sample size determination for prediction models via learning-type curves. Statistics in Medicine, 43(16), 3062–3072. https://doi.org/10.1002/sim.10121

[28] Li, Y., Yorke-Smith, N., & Keviczky, T. (2024). Machine learning enabled uncertainty set for data-driven robust optimization. Journal of Process Control, 144, Article 103339. https://doi.org/10.1016/j.jprocont.2024.103339

[29] Malluhi, B., Fezai, R., Kravaris, C., Nounou, H., Al-Rawashdeh, M., & Nounou, M. (2024). Guided experimental design for static nonparametric modeling. Chemical Engineering Science, 298, Article 120327. https://doi.org/10.1016/j.ces.2024.120327