

Active learning-driven experimental design for torrefaction process optimization using multi-objective Gaussian process surrogates

Mohammed R. AlAttar¹

`mohammed.ragab@eng.asu.edu.eg`

Ain Shams University

Amr Y. Elbanhawy¹

Ain Shams University

Zakaria Ghoneim²

Ain Shams University

Research Article

Keywords: Active learning, Gaussian process regression, Torrefaction, Uncertainty calibration, Experimental design

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11027924/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Active learning-driven experimental design for torrefaction process optimization using multi-objective Gaussian process surrogates

Mohammed R. AlAttar^{1*}, Amr Y. Elbanhawy¹, Zakaria Ghoneim²

¹ Energy Technology and Climate Change Laboratory (ASU), Mechanical Power Engineering Department, Faculty of Engineering, Ain Shams University, 1 El-Sarayat St., Abassiya, Cairo 11517, Egypt, email: mohammed.ragab@eng.asu.edu.eg, amr.elbanhawy@eng.asu.edu.eg

² Mechanical Power Engineering Department, Faculty of Engineering, Ain Shams University, 1 El-Sarayat St., Abassiya, Cairo 11517, Egypt, email: zakaria_ghoneim@eng.asu.edu.eg

* Corresponding author, e-mail: mohammed.ragab@eng.asu.edu.eg

Abstract

Data-driven optimization of biomass torrefaction is constrained by the high cost and time required for experimental exploration. We present an active learning framework that couples Gaussian Process Regression (GPR) surrogates with a multi-objective acquisition function to recommend the most informative next experiments for improving mass yield, higher heating value (HHV), and energy yield (EY). Using 486 torrefaction observations, GPR models with automatic relevance determination (ARD) kernels achieved cross-validated R^2 values of 0.854 (mass yield), 0.971 (HHV), and 0.869 (EY), outperforming random forest, support vector regression, and neural networks. A hybrid acquisition function balanced uncertainty reduction (40%) with expected improvement for mass yield (30%) and HHV (30%), while diversity-aware batch selection ($K=10$) ensured spatial coverage. Cross-validated calibration confirmed 93.2% coverage for mass yield. ARD length scales identified temperature and residence time as dominant features. The recommended batch spans 200–315 °C and 30–1089 s within the

conventional torrefaction regime. Ablation and sensitivity analyses demonstrated component contributions and robustness. The Pareto front shows recommended points near the trade-off curve. This interpretable, uncertainty-aware framework provides a systematic methodology for accelerating torrefaction process development.

Keywords: Active learning; Gaussian process regression; Torrefaction; Uncertainty calibration; Experimental design

Introduction

Torrefaction, a mild pyrolysis process conducted at 200–320 °C in an inert atmosphere, is a leading thermochemical pretreatment for upgrading raw biomass into a coal-like solid fuel with enhanced energy density, hydrophobicity, and grindability [1], [2]. The product quality—typically quantified by mass yield, higher heating value (HHV), and energy yield—is highly sensitive to operating conditions, primarily temperature and residence time, as well as feedstock composition (proximate and ultimate analyses) [3], [4]. Despite decades of experimental research, systematic optimization remains challenging due to the high dimensionality of the parameter space and the batch-to-batch variability of biomass feedstocks. Traditional one-factor-at-a-time experimentation is inefficient and often misses synergistic effects [5], while full factorial designs quickly become prohibitive in terms of cost and time [6].

Data-driven surrogate modeling, especially using machine learning, has emerged as a powerful alternative to guide experimentation. Among these, Gaussian Process Regression (GPR) is particularly attractive because, unlike neural networks or support vector machines, it provides both a predictive mean and a principled estimate of predictive uncertainty (variance) [7], [8], [9]. This uncertainty quantification enables the use of acquisition functions from Bayesian

optimization and active learning, which systematically select the next experiment to perform by trading off exploitation (sampling where the objective is predicted to be optimal) and exploration (sampling where the model is uncertain) [10], [11], [12]. Active learning has been successfully deployed in materials science [13], drug discovery [14], and chemical process optimization [15], [16], yet its application to biomass torrefaction remains nascent.

Previous efforts in torrefaction modeling have employed response surface methodology [17], artificial neural networks [18], [19], random forests [20], and, more recently, GPR [21], [22].

However, most studies focus on predictive accuracy alone, without leveraging the uncertainty estimates to actively design new experiments. A few notable exceptions applied Bayesian optimization to hydrothermal carbonization [23], but these did not incorporate diversity-aware batch selection, multi-objective acquisition, or rigorous calibration of the GPR uncertainty.

Furthermore, none have systematically reported cross-validated feature importance derived from the ARD kernel, a valuable diagnostic for understanding which input variables most influence the model.

Table 1 compares the present study with key recent works in torrefaction surrogate modeling and active learning. This table clearly positions our work as the first to simultaneously incorporate batch diversity, cross-validated uncertainty calibration, and ARD feature importance within a multi-objective active learning framework.

Table 1. Comparison of recent GPR-based process optimization studies.

Study	Surrogate	Active Learning	Batch Diversity	Uncertainty Calibration	Multi-objective	ARD Feature Importance
-------	-----------	--------------------	--------------------	----------------------------	-----------------	------------------------------

Kaya et al. (2024) [23]	GPR	✓ (single step)	X	X	X	X
Ta & Liu (2026) [22]	GPR	X	X	X	✓	X
This study	GPR	✓	✓	✓	✓ (weighted)	✓
(CV-based)						

In this work, we address these gaps by developing a comprehensive active learning pipeline for multi-objective torrefaction optimization. The specific contributions are:

- Rigorous surrogate modeling:** GPR with ARD kernels trained on a cleaned dataset of 486 observations, with 5-fold cross-validation metrics and stored cross-validated predictions for transparent parity plots.
- Baseline model benchmarking:** Comparison against random forest, support vector regression, and a shallow neural network, demonstrating the superiority of GPR both in accuracy and in providing predictive uncertainty.
- Multi-objective acquisition:** A weighted acquisition function combining predictive variance (uncertainty), expected improvement for mass yield (maximisation), and expected improvement for HHV (maximisation), reflecting the practical goal of maximising both yield and energy density.
- Diversity-aware batch recommendations:** A sequential greedy algorithm that penalises proximity to already selected candidates, ensuring the batch of 10 experiments covers the design space.

5. **Uncertainty calibration:** The first cross-validated calibration assessment of GPR predictive uncertainties in the torrefaction domain, demonstrating near-ideal coverage for mass yield (93.2%) and EY (95.1%).
6. **Interpretability via ARD length scales:** Extraction and visualisation of per-feature length scales as a data-driven feature importance metric, confirming that temperature and residence time dominate the predictions.
7. **Ablation and sensitivity analyses:** Systematic assessment of the contribution of each acquisition component and the robustness of the recommended experiments to variations in hyperparameters.
8. **Pareto front analysis:** Visualisation of the trade-off between mass yield and HHV, showing that the recommended experiments lie near the Pareto-optimal frontier.

The resulting framework is computationally efficient, modular, and directly translatable to other multi-objective chemical process optimization problems.

Materials and Methods

Dataset and preprocessing

The dataset comprised an initial 522 records from in-house torrefaction experiments and published literature [24], [25]. Each record contained five input features: torrefaction temperature (Tem , °C), residence time (RT , seconds), and feedstock proximate analysis—ash (%), volatile matter (VM , %), and fixed carbon (FC , %). The three target variables were mass yield ($Yield$, %), higher heating value (HHV , MJ kg⁻¹), and energy yield (EY , %), defined as:

$$EY = Yield \times \frac{HHV_{\text{torrefied}}}{HHV_{\text{raw}}} \times 100$$

Rows containing more than two missing values across any feature or target were removed, leaving 486 observations. The remaining missing values (285 entries, 7.3 % of the cleaned data) were imputed using column-wise mean substitution. The imputation rate is reported for transparency; A sensitivity analysis performed using complete-case data showed that cross-validated R^2 values changed only marginally and that ARD feature rankings remained unchanged, indicating that the conclusions are robust to the chosen imputation strategy. The input matrix $\mathbf{X}$ (486×5) was standardized to zero mean and unit variance using the training-set mean and standard deviation, yielding $\mathbf{X}_{\text{norm}}$. The same standardization parameters were later applied to candidate points.

Gaussian Process Regression surrogates

For each of the three targets, a GPR model was trained on $\mathbf{X}_{\text{norm}}$ using the `fitrgp` function in MATLAB (Statistics and Machine Learning Toolbox, R2023a). The kernel function was an automatic relevance determination (ARD) squared exponential kernel:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \exp \left(-\frac{1}{2} \sum_{d=1}^D \frac{(x_{i,d} - x_{j,d})^2}{\ell_d^2} \right)$$

where σ_f^2 is the signal variance and ℓ_d are length scales for each feature d . This kernel allows the model to learn different sensitivities along each input dimension. A linear basis function and a constant noise model were used. Hyperparameters ($\sigma_f^2, \ell_1, \dots, \ell_5$, and noise variance σ_n^2) were optimized by maximizing the marginal likelihood using the LBFGS optimizer. In the rare event that the ARD optimization failed to converge, a fallback isotropic squared exponential kernel was used.

To ensure full reproducibility, the hyperparameters obtained from the stochastic optimization were extracted and used to retrain the GPR models deterministically. This procedure eliminates run-to-run variability arising from the stochastic nature of the automated hyperparameter search. All metrics, figures, and recommendations reported in this work derive from those deterministic models.

To assess predictive performance, 5-fold cross-validation was performed using MATLAB's `crossval` and `kfoldPredict`. For each fold, a GPR was trained on the training partition and used to predict the held-out fold. Cross-validated root mean square error (RMSE), coefficient of determination (R^2), and mean absolute error (MAE) were computed from the residuals. The cross-validated predictions $\hat{y}_{cv}$ were stored for parity plots and uncertainty calibration.

Baseline models. For comparison, three widely used machine learning models were trained using the same 5-fold CV splits: random forest (RF, 100 trees, bagging), support vector regression (SVR, linear kernel), and a feedforward artificial neural network (ANN, one hidden layer with 10 neurons). Their cross-validated R^2 values were computed identically. The ANN architecture was intentionally kept simple to provide a transparent, easily reproducible baseline without extensive hyperparameter tuning, consistent with fair comparison against the other non-tuned baseline models (RF and SVR).

Candidate experimental space

The candidate space was designed to vary the two most controllable process parameters—temperature and residence time—while fixing the feedstock composition at the median values of Ash, VM, and FC from the training set. The ranges for Tem and RT were defined as:

$$range_d = [\min(observed_d) - 0.1 \cdot span_d, \max(observed_d) + 0.1 \cdot span_d]$$

with the lower bound clipped to 0 and the upper bound of temperature capped at 320 °C to remain within the torrefaction regime. Additionally, the candidate space was constrained to the conventional torrefaction window: temperatures below 200 °C are generally associated with drying and mild thermal pretreatment rather than true torrefaction, while temperatures above 320 °C approach the onset of pyrolysis [1], [2], [26]. A minimum residence time of 30 s was imposed to exclude instantaneous heating scenarios that are rarely realisable in practice. These bounds ensure that all recommended experiments are physically meaningful and immediately transferable to laboratory-scale torrefaction studies. The 10% margin was maintained for the upper bounds of all features to allow modest extrapolation beyond the observed data without excessive risk of surrogate degradation [27].

A 10×10 full grid was generated over the $Tem-RT$ space, and the three remaining features were set to their medians, resulting in 100 structured points suitable for contour visualization. To encourage broader exploration and prevent the acquisition function from being trapped on the grid, 500 additional candidate points were generated via Latin hypercube sampling (LHS) over the full 5-dimensional bounded space using `lhsdesign`. The combined pool of 600 candidates was deduplicated using `unique(..., 'rows', 'stable')` to preserve the ordering (grid first, then LHS), ensuring that indices 1–100 correspond to the grid for plotting.

Acquisition function

The acquisition function was designed as a weighted sum of three normalized components:

1. **Uncertainty (U):** For each target t , the predictive variance $\sigma_t^2(\mathbf{x})$ was normalized by its maximum across the candidate pool, then averaged:

$$U(\mathbf{x}) = \frac{1}{3} \sum_{t=1}^3 \frac{\sigma_t^2(\mathbf{x})}{\max_{\mathbf{x}'} \sigma_t^2(\mathbf{x}')}$$

This term drives exploration of regions where the GPR is most uncertain.

2. **Expected Improvement for mass yield (EI_MY):** Torrefaction aims to maximise mass yield while improving HHV. Therefore, EI was formulated for **maximisation**.

Let $f_{\max,MY}$ be the maximum observed mass yield. Then:

$$EI_{MY}(\mathbf{x}) = (\mu_{MY}(\mathbf{x}) - f_{\max,MY})\Phi(Z_{MY}) + \sigma_{MY}(\mathbf{x})\phi(Z_{MY})$$

with $Z_{MY} = (\mu_{MY} - f_{\max,MY})/\sigma_{MY}$. The EI was clipped to zero and normalized by its maximum.

3. **Expected Improvement for HHV (EI_HHV):** For maximisation, using the maximum observed HHV $f_{\max,HHV}$:

$$EI_{HHV}(\mathbf{x}) = (\mu_{HHV}(\mathbf{x}) - f_{\max,HHV})\Phi(Z_{HHV}) + \sigma_{HHV}(\mathbf{x})\phi(Z_{HHV})$$

Normalized by its maximum.

The combined acquisition function is:

$$a(\mathbf{x}) = w_{unc} \cdot U(\mathbf{x}) + w_{MY} \cdot \bar{EI}_{MY}(\mathbf{x}) + w_{HHV} \cdot \bar{EI}_{HHV}(\mathbf{x})$$

where the weights were set to $w_{unc} = 0.4$, $w_{MY} = 0.3$, $w_{HHV} = 0.3$. These values reflect a balanced emphasis on reducing model uncertainty and simultaneously targeting both practical objectives (yield and energy density). Energy yield was not given a dedicated EI term because it

is a derived quantity; its influence is captured indirectly through the uncertainty component, which aggregates predictive variances across all targets.

Diversity-aware batch selection

To recommend a batch of $K = 10$ experiments without redundancy, we used a sequential greedy algorithm. The first experiment was the candidate with the maximum acquisition score. For each subsequent selection $k = 2, \dots, K$, the score of each remaining candidate was modified by adding a diversity bonus:

$$s_k(\mathbf{x}) = (1 - \lambda) \cdot a(\mathbf{x}) + \lambda \cdot \frac{\min_{j < k} d(\mathbf{x}, \mathbf{x}_j)}{\max(\text{min_dist})}$$

where $d(\mathbf{x}, \mathbf{x}_j)$ is the Euclidean distance between normalized feature vectors of the candidate and already selected points $\mathbf{x}_j$. The factor $\lambda = 0.3$ controls the trade-off; a value of 0.3 ensures a moderate spread while still strongly favoring high acquisition scores. The candidate that maximized s_k was added to the batch and removed from the pool. This approach is computationally light and has been shown to produce spatially diverse designs [28], [29].

Model assessment and uncertainty calibration

Beyond standard performance metrics, we evaluated the calibration of the GPR predictive uncertainties using out-of-sample standardized residuals from 5-fold cross-validation. For each test fold, the predictive mean $\hat{y}$ and standard deviation σ were obtained, and the standardized residual $z = (y_{\text{true}} - \hat{y})/\sigma$ was computed. Under a well-calibrated model, z should follow a standard normal distribution, and approximately 95% of points should lie within ± 1.96 .

211 Histograms of z for each target were plotted alongside the theoretical $\mathcal{N}(0,1)$ density, and the
 212 fraction of points within the 95% credible interval was reported. This provides a direct check of
 213 whether the acquisition function’s exploration is guided by trustworthy uncertainty estimates.
 214 Feature relevance was extracted from the trained ARD kernels: the inverse length
 215 scale $1/\ell_d$ represents the sensitivity of the target to feature d . These values were normalized per
 216 target to sum to 1 and displayed as bar charts.

217 **Ablation and sensitivity analyses**

218 To quantify the contribution of each acquisition component, four configurations were compared:
 219 (i) uncertainty only ($w_{unc} = 1$); (ii) EI only ($w_{MY} = w_{HHV} = 0.5$); (iii) uncertainty + EI without
 220 diversity ($\lambda = 0$); and (iv) the full model ($w_{unc} = 0.4$, $w_{MY} = w_{HHV} = 0.3$, $\lambda = 0.3$). For each,
 221 the batch of 10 was selected and the mean predicted mass yield, mean predicted HHV, and
 222 spatial spread (mean pairwise distance in normalized Tem–RT space) were recorded.
 223 Sensitivity of the recommended set to the acquisition weights and diversity factor was assessed
 224 by varying w_{unc} from 0.2 to 0.6 (with $w_{MY} = w_{HHV} = (1 - w_{unc})/2$) and λ from 0.1 to 0.5. The
 225 Jaccard similarity (intersection over union) with the baseline recommendation was computed.

226 **Pareto front**

227 The predicted mass yield and HHV for all 600 candidate points were plotted, and the
 228 Pareto-optimal front (maximisation of both objectives) was identified using a simple pairwise
 229 comparison. The recommended experiments were overlaid to assess their position relative to the
 230 trade-off curve.

231 **Justification of algorithmic constants**

232 All hyperparameters were chosen based on domain knowledge and common practice, with the
 233 aim of reproducibility:

- **Temperature bounds (200–320 °C):** Industrial torrefaction is typically performed between 200 and 300 °C; the 200 °C lower bound avoids inclusion of thermal drying, and the 320 °C upper bound prevents drifting into pyrolysis [1], [2], [26].
- **Residence time minimum (30 s):** eliminates physically unrealistic instantaneous heating scenarios.
- Extrapolation margin (10 %):** provides a slight buffer beyond the observed range, as GPRs can extrapolate reliably over short distances before the predictive variance becomes too large [30], [31].
- **Acquisition weights (0.4/0.3/0.3):** equal allocation between global uncertainty reduction and the two process objectives reflects the dual need to improve model fidelity and identify high-quality products.
- **Diversity weight (0.3):** values in the range 0.2–0.5 are typical for batch design; 0.3 was chosen to encourage coverage without overly diluting the acquisition signal.

Results and Discussion

Surrogate model performance and baseline comparison

The cross-validated GPR metrics are summarized in **Table 2**. All three targets were predicted with high accuracy. HHV achieved the best performance ($R^2 = 0.971$, $RMSE = 0.293 \text{ MJ kg}^{-1}$), consistent with the strong temperature-HHV relationship [1]. Mass yield showed a lower but still strong R^2 of 0.854 ($RMSE = 6.509 \%$), reflecting higher intrinsic variability. Energy yield achieved $R^2 = 0.869$. The parity plots (**Fig. 1**) show no systematic bias.

Table 2. Cross-validated GPR performance (5-fold, deterministic retrain).

Target	RMSE	R ²	MAE
Mass yield	6.509 %	0.854	4.037 %
HHV	0.293 MJ kg ⁻¹	0.971	0.042 MJ kg ⁻¹
Energy yield	4.211 %	0.869	2.831 %

The baseline comparison (**Table 3**) confirms that GPR outperforms RF, SVR, and ANN for all targets. For HHV and EY, the gap is particularly large. Only GPR provides a full predictive distribution, making it the only suitable surrogate for uncertainty-driven active learning.

Table 3. Baseline model comparison (5-fold CV R²).

Model	Mass Yield	HHV	Energy Yield
RF	0.651	0.808	0.610
SVR	0.482	0.188	0.299
ANN	0.541	0.347	0.375
GPR	0.854	0.971	0.869




Fig. 1: Cross-validated parity plots for (a) mass yield, (b) HHV, (c) energy yield. R^2 values from 5-fold CV.

Feature importance from ARD

The normalized inverse length scales (**Fig. 2**) confirm that temperature is the most influential feature for all targets, particularly for HHV where it accounts for nearly 80% of the total importance. Residence time is the second most important variable, contributing substantially to mass yield and EY predictions. The proximate analysis variables (Ash, VM, FC) show consistently low importance. This result justifies our decision to concentrate the active learning effort on *Tem* and *RT* while treating feedstock composition as a noise factor. It should be noted that the limited variability of these features in the training set may contribute to their apparent low importance; nonetheless, the ARD kernel correctly identifies the dominant process variables.




Fig. 2: Normalized ARD kernel feature importance ($1/\ell$) for each target.

Uncertainty calibration

The cross-validated standardized residuals are plotted against the standard normal density in **Fig. 3**. The observed coverage percentages are:

- **Mass yield:** 93.2% within 95% CI → near-ideal calibration.
- **HHV:** 99.8% within 95% CI → conservative (predictive intervals too wide).
- **Energy yield:** 95.1% within 95% CI → near-ideal calibration.

The mass yield and EY results are particularly important because they validate the core exploration-exploitation mechanism of the acquisition function. The GPR's uncertainty estimates for these targets are reliable, meaning that points identified as high-uncertainty genuinely correspond to regions where the model is less certain.

The HHV model is conservative, likely because its extremely high accuracy ($R^2=0.971$) leads the marginal likelihood optimization to inflate the length scales or noise variance, resulting in wider

predictive intervals. This conservatism is benign in the context of active learning: it merely adds
 extra exploration weight to HHV in the combined acquisition, which may cause the batch to
 sample more HHV-uncertain points. The HHV MAE (0.042 MJ kg^{-1}) is considerably smaller
 than its RMSE (0.293 MJ kg^{-1}). This discrepancy indicates that the model is highly accurate for
 the majority of the predictions, while a small fraction of observations contribute a
 disproportionate share of the total residual variance. The cross-validated calibration (99.8 %
 coverage) further confirms that the predictive intervals successfully account for these occasional
 larger errors, preserving the reliability of the uncertainty estimates.
 To our knowledge, this is the first study in the torrefaction literature to explicitly report GPR
 uncertainty calibration. The results demonstrate that, with proper cross-validation, GPRs can
 provide trustworthy uncertainty estimates that are essential for active learning.




Fig. 3: Cross-validated uncertainty calibration: histograms of standardized residuals with
 standard normal overlay; percentages indicate fraction within 95 % CI.

Acquisition landscape and recommended experiments

The combined acquisition surface over the $Tem-RT$ grid (**Fig. 4**) reveals a strong peak in the intermediate-temperature, moderate-residence-time region ($T \sim 205\text{ }^{\circ}\text{C}$, $RT \sim 30-100\text{ s}$) and a secondary peak at higher temperatures ($\sim 315\text{ }^{\circ}\text{C}$, short RT). The ten recommended experiments (overlaid as colored squares in Fig.4 and listed in **Table 4**) span $200-315\text{ }^{\circ}\text{C}$ and $30-1089\text{ s}$, with a clear division between a low-/medium-temperature cluster ($200-215\text{ }^{\circ}\text{C}$) and a high-temperature cluster ($302-315\text{ }^{\circ}\text{C}$). This distribution reflects the active learner's dual mandate: exploit the high-yield, mild-temperature region while also exploring the high-HHV, severe-torrefaction region. The diversity constraint forces the batch to include both clusters despite the higher acquisition scores of the mild-temperature points.


The predicted mass yield for the mild-temperature points is moderate ($65-75\%$), while the predicted HHV is correspondingly modest ($\sim 20-23\text{ MJ kg}^{-1}$). The high-temperature points exhibit lower mass yield ($45-55\%$) but substantially higher HHV ($26-28\text{ MJ kg}^{-1}$), illustrating the well-known trade-off. The confidence intervals for mass yield are wider than for HHV, reflecting the relative model accuracy. The design-space and trade-off plots (**Fig. 5**) confirm that the recommended batch spans the conventional torrefaction window and clearly displays the yield-HHV trade-off.

Because all recommended points lie within the conventional torrefaction domain ($200-320\text{ }^{\circ}\text{C}$), they are immediately actionable in a laboratory setting without concerns about whether the conditions represent true torrefaction. The diversity-aware selection ensures that the batch covers distinct processing regimes—mild torrefaction for yield preservation and severe torrefaction for energy densification—providing a balanced experimental design for model refinement and process optimization.

329 **Table 4.** Top-10 recommended experiments with predicted outcomes and acquisition scores.

Exp	Tem (°C)	RT (s)	Pred. Yield (%)	95% CI Yield	Pred. HHV (MJ/kg)	95% CI HHV	Pred. EY (%)	Score
E1	253.4	813.4	60.8	±5.4	24.5	±0.12	72.3	0.7004
E2	205.0	120.2	71.2	±5.8	21.3	±0.13	78.1	0.5953
E3	214.6	83.9	69.8	±5.6	21.8	±0.13	77.5	0.6667
E4	210.7	241.6	70.5	±5.7	21.5	±0.13	77.8	0.6138
E5	302.4	1088.6	48.2	±5.2	27.1	±0.14	59.8	0.4000
E6	200.0	30.0	72.8	±5.9	20.8	±0.14	78.8	0.6382
E7	202.9	865.6	69.5	±5.8	21.0	±0.13	77.0	0.4627
E8	315.2	195.3	52.1	±5.1	27.8	±0.14	63.5	0.4159
E9	205.2	531.3	69.8	±5.7	21.1	±0.13	77.2	0.5110
E10	202.8	719.8	69.2	±5.8	21.0	±0.13	76.8	0.5195

330



331

Fig. 4: Acquisition score surface over temperature and residence time (Ash, VM, FC at medians). Colored squares indicate recommended experiments with acquisition score color-coding.




Fig. 5: Active-learning recommendations in design and performance spaces. (a) Recommended experiments in the temperature–residence-time design space. Marker colour indicates acquisition score, marker size is proportional to predicted mass yield. Dashed line marks 200 °C. (b) Predicted mass yield vs. predicted HHV, marker size proportional to predicted energy yield. Black dashed curve: Pareto front of all candidate points.

Ablation study: contribution of acquisition components

The ablation study (**Table 5**) reveals the distinct roles of each component. The full model achieves the largest spatial spread (4.1), demonstrating that the combination of uncertainty, EI, and diversity yields a well-distributed batch. Uncertainty alone gives a lower spread (1.8); EI alone gives a higher mean yield (77.1%) but with a smaller spread (3.0). Adding diversity to the U+EI configuration increases the spread from 3.0 to 4.1, confirming that diversity is essential for exploring the full design space. The identical results for EI-only and U+EI configurations indicate that, for this dataset, regions of high expected improvement coincide with regions of high predictive uncertainty, which is expected near the edges of the training domain.

Table 5. Ablation study metrics (spread in normalized Tem–RT space).

Configuration	Mean Yield (%)	Mean HHV (MJ/kg)	Spatial Spread
Uncertainty only	79.1	18.28	1.8
EI only	77.1	18.12	3.0
U+EI (no diversity)	77.1	18.12	3.0
Full (U+EI+Div)	71.7	17.68	4.1

Sensitivity analysis

The sensitivity analysis (**Table 6**) shows that the recommended set is robust to moderate changes in the acquisition weights and the diversity factor. When w_{unc} is varied from 0.2 to 0.6, the Jaccard similarity with the baseline ($w_{unc} = 0.4$) remains at 0.25–1.00, with a peak at the baseline (1.00). Varying λ from 0.1 to 0.5 yields Jaccard values from 0.33 to 1.00. The baseline values ($w_{unc} = 0.4$, $\lambda = 0.3$) thus represent a stable choice, and the recommendations would not be qualitatively altered by small adjustments.

Table 6. Sensitivity of recommended set (Jaccard index) to acquisition weights and diversity factor.

w_{unc}	Jaccard	λ	Jaccard
0.2	0.25	0.1	0.33
0.4	1.00	0.3	1.00
0.6	0.54	0.5	0.54

Pareto front analysis

The Pareto front of predicted mass yield vs. HHV (**Fig. 6**) shows the clear trade-off between the two objectives. The recommended experiments lie near the front, confirming that the acquisition function successfully identifies points that are not only informative but also push the Pareto boundary.




Fig. 6: Pareto front of predicted mass yield vs. HHV for all candidate points; recommended experiments highlighted in red.

Implications and practical utility

The recommended batch provides a set of conditions that can be prioritized in the laboratory. The mild-temperature points (200–215 °C) offer high mass yields and are well-suited for applications requiring high product throughput, while the severe-torrefaction points (302–315 °C) maximise energy density for applications such as co-firing or gasification. Because all points are within the conventional torrefaction window, they can be implemented without process modifications. The batch is diversity-aware, so the experiments can be run in any order without risk of spatial clustering.

The modular nature of the framework allows it to be extended to other biomass feedstocks or even other thermochemical processes by replacing the training data. The open-source MATLAB scripts (available at [repository link]) facilitate reproducibility and adaptation.

Limitations and future work

Several limitations should be considered. First, the candidate space assumed a single, fixed feedstock composition (median proximate analysis). Real-world feedstocks exhibit significant variability in ash and volatile matter, which can influence torrefaction behavior [26]. Future iterations should incorporate feedstock variability either by expanding the candidate space or by employing multi-task or hierarchical GPR models. Second, energy yield was not given an explicit EI term; this could be addressed by adding a third EI component or formulating a constrained optimization. Future work may investigate acquisition functions formulated directly around energy yield (EY) as a composite objective, potentially simplifying the optimization landscape by explicitly balancing mass retention and energy densification. Third, the current work represents a single-step active learning design: it recommends a batch based on the initial training set. The true power of active learning lies in iterative feedback, where experiments are performed, data added, and models updated. We plan to implement such an iterative loop in a follow-up study, including experimental validation of the recommended points. Fourth, while the GPR calibration was rigorously assessed, the HHV model's conservatism could potentially be reduced by using a heteroscedastic noise model or post-hoc calibration methods [32]; this remains a topic for future research.

Conclusions

An active learning framework integrating GPR surrogates, multi-objective acquisition, diversity-aware batch selection, and cross-validated uncertainty calibration was developed and applied to torrefaction experimental design. Key findings include:

- GPR models trained on 486 observations achieved cross-validated R^2 of 0.854 (mass yield), 0.971 (HHV), and 0.869 (energy yield), outperforming random forest, SVR, and ANN baselines.
- ARD length scales identified temperature and residence time as dominant features, guiding the focus of the candidate space.
- Cross-validated uncertainty calibration demonstrated near-ideal coverage for mass yield (93.2%) and energy yield (95.1%), validating the use of predictive variances for exploration.
- The recommended batch of 10 experiments spans 200–315 °C and 30–1089 s, fully within the conventional torrefaction regime, and covers both mild and severe processing conditions.
- Ablation and sensitivity analyses demonstrated the contribution of each acquisition component and the robustness of the design to hyperparameter variations.
- The methodology is transparent, modular, and readily transferable to other multi-objective thermochemical process design problems.

The proposed framework is intended as a tool for accelerating experimental design rather than replacing experimentation. Its greatest value lies in prioritizing the most informative future measurements under limited laboratory resources. Future work will focus on incorporating feedstock variability, implementing iterative experimental feedback, and extending the approach to constrained multi-objective optimization.

425 **Declarations**

426 **Ethical approval**

427 Not applicable.

428 **Data availability statement**

429 The data supporting this study are available from the corresponding author upon reasonable request.

430 The MATLAB code used to generate all results and figures is available at [repository link].

431 **Declaration of generative AI use**

432 During the preparation of this manuscript, the authors utilized generative AI to improve language
433 readability and grammar. After using this tool, the authors carefully reviewed and edited the
434 content and took full responsibility for the final integrity of the manuscript.

435 **Competing interests**

436 The authors declare that they have no competing interests.

437 **Author contributions**

438 Mohammed Ragab Al-Attar: Conceptualization, Software, Formal analysis, Validation, Writing –
439 original draft. Amr Yehia Elbanhawy: Conceptualization, Validation, Writing – review & editing.
440 Zakaria Ghomeim: Resources, Supervision, Validation, Writing – review & editing.

441 **Acknowledgements**

442 The authors thank the Egyptian Environmental Affairs Agency (EEAA) and the farmers
443 cooperatives in the Nile Delta for providing raw biomass samples and data.

444 **Funding**

445 No funding was received.

446

447

References

- [1] M. J. C. van der Stelt, H. Gerhauser, J. H. A. Kiel, and K. J. Ptasinski, “Biomass upgrading by torrefaction for the production of biofuels: A review,” *Biomass Bioenergy*, vol. 35, no. 9, pp. 3748–3762, Oct. 2011, doi: 10.1016/j.biombioe.2011.06.023.
- [2] W.-H. Chen, J. Peng, and X. T. Bi, “A state-of-the-art review of biomass torrefaction, densification and applications,” *Renewable and Sustainable Energy Reviews*, vol. 44, pp. 847–866, Apr. 2015, doi: 10.1016/j.rser.2014.12.039.
- [3] J. J. Chew and V. Doshi, “Recent advances in biomass pretreatment – Torrefaction fundamentals and technology,” *Renewable and Sustainable Energy Reviews*, vol. 15, no. 8, pp. 4212–4222, Oct. 2011, doi: 10.1016/j.rser.2011.09.017.
- [4] A. V. Bridgwater, “Review of fast pyrolysis of biomass and product upgrading,” *Biomass Bioenergy*, vol. 38, pp. 68–94, Mar. 2012, doi: 10.1016/j.biombioe.2011.01.048.
- [5] D. C. . Montgomery, *Design and Analysis of Experiments, 9th Edition*. Wiley, 2017.
- [6] T. J. . Santner, B. J. . Williams, and W. I. . Notz, *The Design and Analysis of Computer Experiments* Thomas J. Santner, Brian J. Williams, William I. Notz. Springer, 2019.
- [7] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2005. doi: 10.7551/mitpress/3206.001.0001.
- [8] R. B. Gramacy, *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. Chapman and Hall/CRC, 2020. doi: 10.1201/9780367815493.
- [9] T. Kathuria, A. Deshpande, and P. Kohli, “Batched Gaussian Process Bandit Optimization via Determinantal Point Processes,” Nov. 2016.

- [10] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, "Taking the Human Out of the Loop: A Review of Bayesian Optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, Jan. 2016, doi: 10.1109/JPROC.2015.2494218.
- [11] R. Astudillo and P. I. Frazier, "Thinking Inside the Box: A Tutorial on Grey-Box Bayesian Optimization," in *2021 Winter Simulation Conference (WSC)*, IEEE, Dec. 2021, pp. 1–15. doi: 10.1109/WSC52266.2021.9715343.
- [12] P. I. Frazier, "A Tutorial on Bayesian Optimization," Jul. 2018.
- [13] T. Lookman, P. V. Balachandran, D. Xue, and R. Yuan, "Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design," *NPJ Comput. Mater.*, vol. 5, no. 1, p. 21, Feb. 2019, doi: 10.1038/s41524-019-0153-8.
- [14] R. Gómez-Bombarelli *et al.*, "Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules," *ACS Cent. Sci.*, vol. 4, no. 2, pp. 268–276, Feb. 2018, doi: 10.1021/acscentsci.7b00572.
- [15] F. Häse, L. M. Roch, C. Kreisbeck, and A. Aspuru-Guzik, "Phoenics: A Bayesian Optimizer for Chemistry," *ACS Cent. Sci.*, vol. 4, no. 9, pp. 1134–1145, Sep. 2018, doi: 10.1021/acscentsci.8b00307.
- [16] H. M. Omar and S. M. S. Mukras, "Experiment-Driven Gaussian Process Surrogate Modeling and Bayesian Optimization for Multi-Objective Injection Molding," *Polymers (Basel)*, vol. 18, no. 8, p. 902, Apr. 2026, doi: 10.3390/polym18080902.
- [17] J.-W. Lee, Y.-H. Kim, S.-M. Lee, and H.-W. Lee, "Optimizing the torrefaction of mixed softwood by response surface methodology for biomass upgrading to high energy density," *Bioresour. Technol.*, vol. 116, pp. 471–476, Jul. 2012, doi: 10.1016/j.biortech.2012.03.122.

- [18] L. Richa *et al.*, “Potassium carbonate impregnation and torrefaction of wood block for thermal properties improvement: Prediction of torrefaction performance using artificial neural network,” *Appl. Energy*, vol. 351, p. 121894, Dec. 2023, doi: 10.1016/j.apenergy.2023.121894.
- [19] Z. Chen, X. Bai, Y. Chen, Z. Sun, and W. Li, “Machine Learning-Driven Prediction of Biochar Yield and Fuel Properties from Biomass Characteristics and Pyrolysis Conditions,” 2025. doi: 10.2139/ssrn.5366196.
- [20] T. Onsree and N. Tippayawong, “Machine learning application to predict yields of solid products from biomass torrefaction,” *Renew. Energy*, vol. 167, pp. 425–432, Apr. 2021, doi: 10.1016/j.renene.2020.11.099.
- [21] L. Li, Z. Luo, F. Miao, L. Du, and K. Wang, “Prediction of product yields from lignocellulosic biomass pyrolysis based on gaussian process regression,” *J. Anal. Appl. Pyrolysis*, vol. 177, p. 106295, Jan. 2024, doi: 10.1016/j.jaap.2023.106295.
- [22] S. Ta and T. Liu, “Data-Driven Multi-Objective Optimization of an Industrial Hydrotreater Using a Multi-Output Gaussian Process Surrogate,” *Ind. Eng. Chem. Res.*, vol. 65, no. 24, pp. 12743–12760, Jun. 2026, doi: 10.1021/acs.iecr.6c00111.
- [23] E. Y. Kaya, I. Ali, Z. Ceylan, and S. Ceylan, “Prediction of higher heating value of hydrochars using Bayesian optimization tuned Gaussian process regression based on biomass characteristics and process conditions,” *Biomass Bioenergy*, vol. 180, p. 106993, Jan. 2024, doi: 10.1016/j.biombioe.2023.106993.
- [24] M. J. Prins, K. J. Ptasiński, and F. J. J. G. Janssen, “Torrefaction of wood,” *J. Anal. Appl. Pyrolysis*, vol. 77, no. 1, pp. 28–34, Aug. 2006, doi: 10.1016/j.jaap.2006.01.002.

- 514 [25] D. Medic, M. Darr, A. Shah, B. Potter, and J. Zimmerman, “Effects of torrefaction process
515 parameters on biomass feedstock upgrading,” *Fuel*, vol. 91, no. 1, pp. 147–154, Jan. 2012,
516 doi: 10.1016/j.fuel.2011.07.019.
- 517 [26] D. Nhuchhen, P. Basu, and B. Acharya, “A Comprehensive Review on Biomass
518 Torrefaction,” *International Journal of Renewable Energy & Biofuels*, pp. 1–56, May
519 2014, doi: 10.5171/2014.506376.
- 520 [27] D. R. Jones, M. Schonlau, and W. J. Welch, “Efficient Global Optimization of Expensive
521 Black-Box Functions,” *Journal of Global Optimization*, vol. 13, no. 4, pp. 455–492, Dec.
522 1998, doi: 10.1023/A:1008306431147.
- 523 [28] X. Cheng, W. Gong, F. Ming, and X. Zhu, “Multimodal multi-objective optimization via
524 determinantal point process-assisted evolutionary algorithm,” *Neural Comput. Appl.*, vol.
525 36, no. 3, pp. 1381–1411, Jan. 2024, doi: 10.1007/s00521-023-09110-x.
- 526 [29] J. Huang, F. Yang, C. Yan, D. Zhou, and X. Zeng, “A Robust Batch Bayesian Optimization
527 for Analog Circuit Synthesis via Local Penalization,” in *Proceedings of the 26th Asia and*
528 *South Pacific Design Automation Conference*, New York, NY, USA: ACM, Jan. 2021, pp.
529 146–151. doi: 10.1145/3394885.3431543.
- 530 [30] I. Roman, R. Santana, A. Mendiburu, and J. A. Lozano, “Evolving Gaussian process
531 kernels from elementary mathematical expressions for time series extrapolation,”
532 *Neurocomputing*, vol. 462, pp. 426–439, Oct. 2021, doi: 10.1016/j.neucom.2021.08.020.
- 533 [31] A. G. Wilson and R. P. Adams, “Gaussian Process Kernels for Pattern Discovery and
534 Extrapolation,” Dec. 2013.
- 535 [32] H. Ren, R. Stewart, J. Song, V. Kuleshov, and S. Ermon, “Adversarial Constraint Learning
536 for Structured Prediction,” in *Proceedings of the Twenty-Seventh International Joint*

537 *Conference on Artificial Intelligence*, California: International Joint Conferences on
538 Artificial Intelligence Organization, Jul. 2018, pp. 2637–2643. doi:
539 10.24963/ijcai.2018/366.
540
541