

Kriging and Radial Basis Function Surrogate Models for a Fluidized-Bed Methane Pyrolysis Reactor

Orlu Ndamati Moses¹, Akuma Oji², Muwarure Peter³

¹PhD Candidate, Centre for Gas, Refining and Petrochemical Engineering, University of Port Harcourt
Email: [muwarure.peter\[at\]cgrpng.org](mailto:muwarure.peter[at]cgrpng.org)

²Professor, The Department of Chemical Engineering, University of Port Harcourt, Rivers State, Nigeria.
Email: [akuma.oji\[at\]uniport.edu.ng](mailto:akuma.oji[at]uniport.edu.ng)

³Lecturer, Centre for Gas, Refining and Petrochemical Engineering, University of Port Harcourt
Email: [Petermuwarure\[at\]gmail.com](mailto:Petermuwarure[at]gmail.com)

Abstract: Robust multi-objective optimization of mechanistic fluidized-bed reactor models is computationally prohibitive because each candidate design must be evaluated many times to estimate the mean and variance of its performance. This paper develops and validates computationally efficient surrogate models using Kriging (Gaussian process regression with a Matérn 5/2 kernel) and radial basis function (RBF) interpolation to approximate the high-fidelity fluidized-bed methane pyrolysis reactor model. A five-dimensional design space is sampled by Latin hypercube sampling with 300 training points and 50 independent test points. At each design point, 500 Monte Carlo samples are evaluated, giving 150,000 full model runs for training and 25,000 for testing. Six primary statistics (mean and standard deviation of methane conversion, H₂ mole fraction, and LCOH) are approximated. All six Kriging surrogates met the acceptance criteria ($R^2 \geq 0.9956$, $\text{NRMSE} \leq 1.41\%$, $\text{NMAE} \leq 0.69\%$) and outperformed the RBF benchmark for every output. The surrogates evaluate a design about 4×10^5 times faster than Monte Carlo with an adaptive ODE solver, reducing a robust optimization run from about 3.6 days to seconds. Validation against the full model was essential: retraining with 20 additional points reduced the maximum error on the robust front from 19% to below 4.5%.

Keywords: Surrogate modelling; Kriging; Gaussian process; fluidized bed; methane pyrolysis; robust multi-objective optimisation

1. Introduction

Hydrogen is predicted to be crucial to the low-carbon energy transition, however most production comes from fossil fuels (IEA, 2023). Methane pyrolysis, which produces a marketable carbon co-product and avoids direct CO₂ emissions, splits methane into hydrogen and solid carbon ($\text{CH}_4 \rightarrow \text{C(s)} + 2\text{H}_2$, $\Delta H_{298}^\circ = +74.85 \text{ kJ/mol}$) (Muradov & Veziroglu, 2005; Abbas & Wan Daud, 2010). Continuous catalyst addition and carbon withdrawal, superior temperature consistency, and industrial maturity make fluidized-bed reactors (FBRs) appealing (Kunii & Levenspiel, 1991; Ammendola et al., 2008; Dagle et al., 2017).

A systematic comparison of Kriging and RBF surrogates for six Monte Carlo statistics of a methane pyrolysis FBR, validation on independent test designs with five-fold cross-validation, quantification of computational savings compared to direct Monte Carlo, and proof that robust optimisation requires validation against the full model with adaptive retraining are the main contributions.

2. Literature Survey

The design of a methane pyrolysis FBR must balance high methane conversion, high hydrogen content in the output gas, and cheap hydrogen cost. Heating increases reaction rate and equilibrium conversion but deactivates catalysts and raises material costs. Increased catalyst supplies boost conversion but increase reactor size and expense. Higher pressures lower equilibrium conversion but reactor volume and compression

costs. Instead of a single response, multi-objective optimisation (MOO) finds a set of Pareto-optimal trade-off solutions (Deb, 2001; Deb et al., 2002).

However, conventional MOO assumes model parameters are known. Many key methane pyrolysis parameters are unknown. Kinetic constants like activation energy and pre-exponential factor differ between catalyst batches and labs. Operating history and carbon morphology affect deactivation. Empirical correlations with large scatter yield heat transfer coefficients and interphase mass transfer rates. The economics also depend on variable market pricing, particularly natural gas and carbon co-product sales (Dagle et al., 2017; Keipi et al., 2018). Once uncertainties are realised, a design that is ideal at nominal parameter values may perform badly or fail to satisfy product specifications. Robust optimisation explores designs that operate insensitively to uncertainty or meet constraints with high probability (Beyer & Sendhoff, 2007; Deb & Gupta, 2006).

3. Problem Definition

The fundamental impediment to effective mechanistic reactor model optimisation is computing expense. To assess performance mean and variance, each candidate design must be evaluated using 500 Monte Carlo samples. Most evolutionary algorithms like NSGA-II assess 15,000 candidate designs, therefore a robust optimisation requires millions of reactor simulations.

A single optimisation run requires 7.5 million reactor simulations or 3.5 days of serial computation with 500 Monte Carlo samples per design. A less optimised implementation or more complex model could execute a model in 1 s, but it would take 87 days. Each optimisation run takes seconds with surrogate models.

4. Methodology/Approach

4.1 Reactor Model Summary

The high-fidelity model is a one-dimensional, steady-state, two-phase (Kunii–Levenspiel) model of a bubbling fluidised bed with a Geldart group B catalyst ($d_p = 200 \mu\text{m}$, $\rho_p = 2,500 \text{ kg/m}^3$; internal diameter = 0.10 m). The model includes gas–solid hydrodynamics, power-law kinetics, first-order catalyst deactivation, hydrogen inhibition, interphase mass transfer, and axial energy balance. Design is determined by five variables: reactor temperature (800–1,100 °C), pressure (1–5 bar), fluidisation ratio (2–5), catalyst mass (0.5–5 kg), and feed methane mole fraction (0.5–1.0). Model outputs methane conversion (X_{CH_4}), hydrogen mole fraction (Y_{H_2}), carbon generation rate, outlet temperature, heat duty, and hydrodynamic features.

4.2 Design of Experiments

Latin hypercube sampling was used to sample the five-dimensional design space with $N_{\text{train}} = 300$ training points and $N_{\text{test}} = 50$ test points. For each design point, a Monte Carlo approach was used with $N = 500$ samples, resulting in 150,000 full model runs for training and 25,000 for testing. Thus, 175,000 full model runs were needed for 300 training and 50 test configurations. The adaptive RK45 solution took 0.041 s each run, while the vectorised solver took 15.5 s, or 0.1 ms per run.

4.3 Surrogate Training

Following Architecture B (separate surrogates for each statistic), six primary surrogates were trained for μ_X , σ_X , μ_Y , σ_Y , μ_{LCOH} and σ_{LCOH} as functions of the five variables (scaled to [0, 1]).

- 1) Primary surrogate (Kriging): Gaussian process regression with a constant $\times$ Matérn 5/2 kernel with automatic relevance determination (one length-scale per input) plus a small white-noise (nugget) term, implemented with scikit-learn's Gaussian Process Regressor (Pedregosa et al., 2011), which is equivalent to MATLAB fitrgp. Outputs were standardised. Hyperparameters were estimated by maximising the log marginal likelihood with L-BFGS-B and four random restarts.
- 2) Benchmark surrogate (RBF): thin-plate spline radial basis function interpolation with a linear polynomial tail (SciPy RBF Interpolator; Virtanen et al., 2020).

4.4 Validation Criteria

Surrogates were accepted when, on the 50 independent test points, they met all three criteria $R^2 > 0.92$, normalised RMSE $< 5\%$ and normalised MAE $< 3\%$. NRMSE and NMAE were normalised by the range of the test response:

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$$

$$\text{NRMSE} = \frac{\sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}}{\gamma_{\max} - \gamma_{\min}} \times 100\%$$

$$\text{NMAE} = \frac{\frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|}{\gamma_{\max} - \gamma_{\min}} \times 100\%$$

Where n is the number of independent test designs; y_i and $\hat{y}_i$ are the full-model and surrogate values for design i , respectively; $\bar{y} = (1/n) \sum_i y_i$; $\gamma_{\max} = \max_i(y_i)$; $\gamma_{\min} = \min_i(y_i)$; and $\gamma_{\max} - \gamma_{\min}$ is the range of the test response. R^2 is dimensionless; NRMSE and NMAE are reported in percent.

Parity plots and five-fold cross-validation (with fixed hyperparameters) were also examined. If any primary surrogate failed, 50 infill points were to be added where the Kriging predictive standard deviation was largest (the exploration component of the expected-improvement criterion; Jones et al., 1998), followed by retraining.

4.5 Software and Hardware

MATLAB R2024a was utilised for reactor model (ode45) reference implementation and base-case computations. The vectorised model, RK45 verification, RBF, Kriging, LHS support and Sobol analysis, NSGA-II, MOEA/D, hypervolume and IGD, and graphics (Matplotlib) were done in Python 3. The complete study ran on an 8-GB RAM two-core cloud workstation. The 300×500 DOE was solved in under 20 seconds using the vectorised solution. Each surrogate-based NSGA-II run took 4 s.

5. Results

5.1 Data Generation

Results from 300 training designs show a wide performance range: $\mu_X = 0.114$ to 0.978, $\sigma_X = 0.011$ to 0.134, $\mu_Y = 0.167$ to 0.971, $\mu_{\text{LCOH}} = 1.625$ to 9.664 USD/kg, $\sigma_{\text{LCOH}} = 0.772$ to 3.932 USD/kg, and $\mu_{\text{CO}_2} = 3.531$ to 138.992 kg/kg. The vectorized solution ran 175,000 full-model runs in 15.5s, or 0.1ms each run, while the adaptive RK45 solver took 0.041 s. This 400-fold data generation speedup is essential for DOE feasibility.

The adaptive RK45 solver required 0.041s per full-model run. The vectorised solver completed the entire 175,000-run design of experiments in 15.5s wall-clock time, which is equivalent to $15.5/175,000 = 8.9 \times 10^{-5} \text{ s}$ per run, i.e. approximately 0.1ms per run. The per-run speed-up is therefore $0.041/(8.9 \times 10^{-5}) \approx 460 \times$, or about $400 \times$ using the rounded 0.1ms value. The 15.5s figure is a batch total for 175,000 runs, not a single-run time, and should not be compared directly with the 0.041s single-run adaptive-solver time.

5.2 Accuracy of the Kriging and RBF Surrogates

All six Kriging surrogates met the criteria: $R^2 > 0.9956$, NRMSE $\leq 1.41\%$, NMAE $\leq 0.69\%$. The mean responses (μ_X and μ_Y) were almost equal ($R^2 = 0.9999$). The models

accurately predicted standard deviations ($R^2 = 0.997-0.999$), a noisier design function. The RBF benchmark had lower

accuracy for all outputs, particularly μ_{LCOH} ($R^2 = 0.958$, NRMSE = 4.27%).

Table 4.8 Test-Set Accuracy of the Primary Surrogates (50 Independent Designs)

Output	Kriging R^2	Kriging NRMSE (%)	Kriging NMAE (%)	Kriging 5-fold CV R^2	RBF R^2	RBF NRMSE (%)	Criteria met
$\mu_{X_{CH_4}}$ (mean methane conversion)	0.9999	0.20	0.12	0.9999	0.9966	1.53	Yes
$\sigma_{X_{CH_4}}$ (SD methane conversion)	0.9973	1.41	0.69	0.9977	0.9806	3.81	Yes
$\mu_{Y_{H_2}}$ (mean H_2 mole fraction)	0.9999	0.24	0.14	0.9999	0.9975	1.20	Yes
$\sigma_{Y_{H_2}}$ (SD H_2 mole fraction)	0.9994	0.67	0.36	0.9991	0.9836	3.58	Yes
μ_{LCOH} (mean LCOH, USD/kg)	0.9956	1.38	0.64	0.9936	0.9581	4.27	Yes
σ_{LCOH} (SD LCOH, USD/kg)	0.9990	0.59	0.40	0.9973	0.9644	3.58	Yes

Kriging with a Matérn 5/2 kernel provided higher predictive accuracy than RBF interpolation. Kriging outperforms RBF for smooth, low-dimensional responses with limited training sets, according to Sacks et al. (1989) and Rasmussen and Williams (2006). No infill points were needed at surrogate training, showing the 300-point LHS DOE fit the original fit (Jones et al., 1998; Forrester, 2008). The catalyst mass and temperature have the shortest optimised Matérn length-scales, indicating the highest non-linear responses. Temperature and catalyst mass affected conversion and cost most, as in the parametric analysis. Sample the design space more thoroughly in these dimensions for accuracy. Automatic relevance determination (ARD) by Rasmussen and Williams (2006) finds relevant inputs using length-scales.

5.3 Computational Savings

Surrogates outperform adaptive Monte Carlo by 4×10^5 times every robust evaluation, reducing optimization time from 3.6 days to seconds. Training all surrogates (RBF and CV) took 38 s, evaluating six at 15,000 designs took 0.73 s, and each surrogate design took 49 μ s, compared to adaptive Monte Carlo's 20.5 s per design. This speedup allows improving formulations, algorithm benchmarks, carbon, and deactivation scenarios. The robust optimization investigation would have been computationally impossible without surrogates. Surrogate-assisted optimization is needed for computationally intensive reactor model design. Forrester et al. (2008) and Bhosekar and Ierapetritou (2018) found that surrogates reduce optimisation computational costs by orders of magnitude. The above-average speed-up of 4×10^5 confirms the accuracy of Kriging surrogates and the effectiveness of vectorized implementation.

5.4 Validation on the Robust Front and Retraining

After surrogate-based robust multi-objective optimization, 20 representative designs on the robust Pareto front were evaluated with the full Monte Carlo reactor model using 500 samples per design and compared with the surrogate predictions. The initial maximum error on the robust front was 19%; the largest initial errors were 16.1% for σ_X , 19.0% for μ_{LCOH} , and 17.9% for σ_{LCOH} . The largest discrepancy occurred at the high-conversion corner ($T = 1,100^\circ\text{C}$, $P = 5\text{bar}$, $W_{\text{cat}} = 5\text{kg}$), a sparsely sampled region of the 300-point DOE. To reduce this error, 20 additional training designs were selected by uncertainty sampling; among candidate designs on and near the robust front, the designs with the largest Kriging predictive standard deviation were chosen, consistent with the exploration component of the expected-improvement

criterion (Jones et al., 1998). Each additional design was evaluated with the full model using the same 500-sample Monte Carlo procedure, and the six surrogates were retrained with 320 training designs. After retraining, the same 20 robust-front designs were re-evaluated with the full model and Monte Carlo ($N = 500$); the maximum error fell to 4.50% (approximately 4.5%), and the checked errors ranged from 0.13% to 4.550%. the 50 independent test designs were not used for retraining; therefore, the test-set metrics in Section 5.2 remain independent measures of predictive accuracy, whereas the errors reported here are validation errors on the optimized robust fronts.

6. Conclusion

This study developed Kriging surrogates for Monte Carlo statistics of a methane pyrolysis fluidized-bed reactor and benchmarked them against RBF interpolation. On the 50 independent test designs, all six Kriging surrogates met the acceptance criteria ($R^2 > 0.9956$, maximum NRMSE = 1.41%, maximum NMAE = 0.69%) and outperformed the RBF benchmark for every output. These values quantify predictive accuracy on independent test designs, not accuracy on the optimized robust front. The surrogates reduced a robust optimization run from about 3.6 days to seconds. Validation against the full model on the robust front was essential: the initial maximum error was 19%, and after adding 20 uncertainty-selected training designs and retraining, the maximum error on the robust front fell to approximately 4.5%. The framework is applicable to uncertain gas-solid reactors relevant to the Nigerian gas sector.

7. Future Scope

Future work should re-fit E_a , k_o , and k_d to time-resolved fluidized-bed measurements on the iron catalyst before using absolute performance values for design. Additionally, the rate expression should be extended with hydrogen inhibition and contact-efficiency correction for poorly fluidizing materials (Geldart C). Architecture B and the validation-retraining loop should be applied to uncertain gas-solid reactors relevant to the Nigerian gas sector, such as gas-to-liquids, dry reforming, and chemical looping.

References

- [1] International Energy Agency (IEA). Global hydrogen review 2023. Paris: IEA; 2023.

- [2] Muradov NZ, Veziroğlu TN. From hydrocarbon to hydrogen-carbon to hydrogen economy. *Int J Hydrogen Energy*. 2005;30(3):225-37.
- [3] Abbas HF, Wan Daud WMAW. Hydrogen production by methane decomposition: a review. *Int J Hydrogen Energy*. 2010;35(3):1160-90.
- [4] Kunii D, Levenspiel O. *Fluidization engineering*. 2nd ed. Boston: Butterworth-Heinemann; 1991.
- [5] Ammendola P, Chirone R, Ruoppolo G, Russo G, Solimene R. Some issues in modelling methane catalytic decomposition in fluidized bed reactors. *Int J Hydrogen Energy*. 2008;33(11):2679-94.
- [6] Dagle RA, Dagle V, Bearden MD, Holladay JD, Krause TR, Ahmed S. An overview of natural gas conversion technologies for co-production of hydrogen and value-added solid carbon products. Report No.: ANL-17/11; PNNL-26726. Argonne, IL: Argonne National Laboratory; Richland, WA: Pacific Northwest National Laboratory; 2017.
- [7] Deb K. *Multi-objective optimization using evolutionary algorithms*. Chichester: Wiley; 2001.
- [8] Deb K, Pratap A, Agarwal S, Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans Evol Comput*. 2002;6(2):182-97.
- [9] Keipi T, Tolvanen H, Konttinen J. Economic analysis of hydrogen production by methane thermal decomposition: comparison to competing technologies. *Energy Convers Manag*. 2018; 159: 264-73.
- [10] Beyer HG, Sendhoff B. Robust optimization – a comprehensive survey. *Comput Methods Appl Mech Eng*. 2007;196(33-34):3190-218.
- [11] Deb K, Gupta H. Introducing robustness in multi-objective optimization. *Evol Comput*. 2006;14(4):463-94.
- [12] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: machine learning in Python. *J Mach Learn Res*. 2011; 12: 2825-30.
- [13] Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods*. 2020; 17: 261-72.
- [14] Jones DR, Schonlau M, Welch WJ. Efficient global optimization of expensive black-box functions. *J Glob Optim*. 1998;13(4):455-92.
- [15] Forrester AIJ, Sobester A, Keane AJ. *Engineering design via surrogate modelling: a practical guide*. Chichester: Wiley; 2008.
- [16] Bhosekar A, Ierapetritou M. Advances in surrogate based modeling, feasibility analysis, and optimization: a review. *Comput Chem Eng*. 2018; 108: 250-67.