

Research Article

Thermodynamic Modelling with Machine Learning for Industrial-Scale Biogas Production with CHP Recovery

Shirin Naderi¹, Amirhossein Ahmadi^{2*}

¹Department of Structural, Geotechnical and Building Engineering, Politecnico di Torino, Turin, Italy

²Institute for Nanomaterials, Advanced Technologies and Innovation, Technical University of Liberec, Liberec, Czech Republic

*amirhossein.ahmadi@tul.cz

Abstract

Most process models of anaerobic-digestion (AD) biogas plants either resolve the biochemistry in detail while treating the combined heat and power (CHP) unit with fixed heating values, or resolve the thermodynamics while neglecting how operating conditions affect biological conversion. This study presents an open-source, digital-twin-oriented simulation framework that couples a multi-factor steady-state kinetic model of a co-digestion reactor (cattle manure, food waste and maize silage) with an exergy model of the biogas and CHP combustion, evaluated directly with the PYroMat property library. After calibration against 19 steady-state operating periods from two independent laboratory CSTR studies, the kinetic model reproduced the measured specific methane yield (SMY) with a cross-validated R^2 of 0.89, whereas an uncalibrated ADM1 implementation under-predicted the same data. The framework was applied to a 365-day industrial run of a 0.6 MW_{el} plant, a 4000-point machine-learning (ML) dataset and a 10,000-run Monte Carlo study. Gradient Boosting surrogates reproduced the mechanistic model with $R^2 = 0.990$ for SMY and 0.995 for CHP exergy efficiency, values that reflect surrogate fidelity rather than accuracy against plant data. Sobol' sensitivity analysis showed that digester temperature and organic loading rate govern SMY, whereas feedstock blend and organic loading rate govern exergy efficiency, a ranking confirmed by ML permutation importance. At the nominal point, the plant reaches 82.5% energy efficiency but only 40.7% exergy efficiency, mainly because recovered heat carries little exergy. All code, data and figures are released as a reproducible artefact.

Keywords: Biogas; Anaerobic Digestion; PYroMat; Exergy Analysis; Machine Learning; Monte Carlo Simulation; Sobol' Sensitivity Analysis; Combined Heat and Power

INTRODUCTION

Biogas generated from the anaerobic digestion (AD) of agricultural residues, livestock manure and the organic fraction of municipal and food waste is one of the few renewable-energy vectors that simultaneously supplies dispatchable electricity, recoverable heat, a vehicle-grade biomethane substitute and a nutrient-rich digestate, while treating an otherwise polluting waste stream [1]. Global biogas and biomethane output has been forecast to expand by roughly one third over the remainder of this decade as agricultural,

municipal and industrial actors pursue decarbonisation and circular-economy targets [1], and biogas already displaces a meaningful share of natural-gas demand in several European agricultural regions [2, 3]. Realising this potential at industrial scale, however, requires plant operators to reconcile two coupled but historically separately modelled sub-systems: the microbial conversion of heterogeneous feedstock inside the digester, and the downstream thermal conversion of the resulting biogas in a combined heat and power (CHP) engine-generator.

The biochemistry of AD is comparatively well understood. The IWA Anaerobic Digestion Model No. 1 (ADM1) formalised the coupled hydrolysis-acidogenesis-acetogenesis-methanogenesis pathway together with inhibition kinetics for pH, ammonia and long-chain fatty acids, and remains the reference mechanistic description of the process [4]. For practical, industrially oriented process design, however, reduced-order descriptions are usually preferred: batch biochemical methane potential (BMP) assays are still fitted with first-order or modified-Gompertz kinetics [5], the latter originally developed to describe bacterial growth curves [5] and subsequently adapted to cumulative methane production by [6] and a large number of co-digestion studies since. Reported BMP values depend strongly on substrate biochemistry: cattle manure typically yields on the order of 0.20-0.30 Nm³ CH₄ kg⁻¹ VS, food waste 0.35-0.50 Nm³ CH₄ kg⁻¹ VS, and maize silage 0.30-0.39 Nm³ CH₄ kg⁻¹ VS [7, 8], which motivates co-digestion strategies that blend a nitrogen-rich, low-yield substrate (manure) with carbon-rich, high-yield co-substrates to balance the carbon-to-nitrogen (C/N) ratio around the reported 20-30 optimum [7, 8]. Beyond substrate selection, the realisable fraction of the theoretical BMP in a continuously fed industrial digester is further reduced by deviations from the mesophilic (~35-38 °C) or thermophilic (~55 °C) temperature optima [9], by short hydraulic retention time (HRT) causing biomass washout, by volatile-fatty-acid (VFA) accumulation at excessive organic loading rate (OLR) [10], and by free-ammonia or pH-driven inhibition of methanogens [10, 11]; standard environmental-biotechnology treatments formalise these CSTR washout and inhibition phenomena [12].

By contrast, the thermodynamic performance of the CHP stage is very frequently treated with tabulated or assumed values: a fixed lower heating value (LHV) for methane, a fixed engine electrical efficiency, and, when exergy is considered at all, the standard chemical exergy of methane taken directly from reference tables such as those of Szargut, Morris and Steward [13]. This is defensible for methane alone, but the raw biogas leaving an industrial digester is a variable CH₄-CO₂-H₂S-N₂ mixture whose composition (typically 45-70% CH₄ by volume, with H₂S trace concentrations that can reach several thousand ppm) depends on feedstock and loading conditions [14], so that a fixed heating value silently discards a first-order source of variability that is precisely the one an operator can influence through feedstock blending. Rigorous, composition-resolved thermodynamic and exergy analysis of biogas CHP and biogas-fuelled cogeneration systems has nevertheless been demonstrated to be both feasible and informative. The present study addresses this gap directly by computing every enthalpy, entropy, adiabatic flame

temperature and physical-exergy value used in the CHP sub-model from the open-source PYroMat ideal-gas property library [15], which evaluates NASA-polynomial-based thermodynamic properties for the individual biogas constituents on the fly rather than from a single look-up constant, while retaining the well-established Szargut chemical-exergy reference values [13] for the chemical-exergy term.

A further, increasingly important, requirement for modern biogas-plant operation is fast, data-driven decision support: full mechanistic simulation (or a coupled kinetic-thermodynamic model of the type developed here) is well suited to design studies and sensitivity screening but is generally too slow, and too dependent on well-calibrated kinetic parameters, for real-time supervisory control. Machine-learning (ML) surrogates trained on either field data or simulated data have therefore been proposed extensively in the recent literature to predict biogas or methane yield from operating-condition inputs, using Random Forest, Gradient Boosting and artificial neural network (ANN) architectures with reported coefficients of determination typically in the $R^2 = 0.85\text{--}0.98$ range depending on dataset size and feature set [16–18]. A methodologically important, but frequently neglected, companion step to ML surrogate modelling is a rigorous accounting of parametric uncertainty and of which input variables actually drive the output variance; global, variance-based sensitivity analysis following Sobol' [19] and implemented through Monte-Carlo/Saltelli sampling schemes [19–21] provides exactly such an accounting, and - when applied to the same mechanistic model that generated the training data - offers a complementary, physically grounded cross-check of the feature-importance rankings extracted from the trained ML surrogate.

The term "digital-twin-oriented" is used deliberately: the framework reproduces the physics and operating envelope of a representative industrial plant, but it is not synchronised with, or calibrated against, measured data from a specific physical installation. It is therefore intended as the modelling core onto which real-time data coupling can be added in future work, rather than as a digital twin in the strict sense of a continuously updated virtual counterpart of an operating plant. Building on these three strands of literature, the present work makes five contributions. First, it develops a coupled kinetic-thermodynamic model of an industrial co-digestion CHP plant in which every thermodynamic property is computed by direct evaluation of the PYroMat ideal-gas database [15] rather than by table look-up, so that composition-dependent effects (CH_4 content, trace H_2S , excess-air ratio) propagate consistently into the heating value, the adiabatic flame temperature and the exergy balance. Second, it exercises this model as a realistic, 365-day "industrial run" of a representative 0.6 MW_{el} plant subject to seasonal ambient-temperature and weekly feedstock-delivery variability. Third, it uses Latin-Hypercube sampling of the full operating envelope to generate a physically consistent, fully documented synthetic dataset on which Random Forest, Gradient Boosting and ANN surrogates are trained and evaluated as surrogates of the mechanistic model on held-out simulator outputs, and it separately propagates realistic operating-parameter uncertainty through the model via a 10,000-run Monte Carlo study. Fourth, it performs Sobol' global

sensitivity analysis of the mechanistic model and cross-checks the resulting parameter ranking against permutation feature importance obtained from the ML surrogates, providing two complementary lines of evidence, both derived from the same mechanistic model, for which operating variables matter most for methane yield and for CHP exergy performance. Fifth, it calibrates the two least-constrained efficiency factors of the kinetic sub-model against measured steady-state CSTR data and benchmarks the calibrated model quantitatively against those data and against an open-source ADM1 implementation, using RMSE, MAE, R^2 and normalised RMSE. The remainder of the paper is organised as follows: the Materials and Methods section describes the mechanistic model, the industrial-run design, the dataset-generation and Monte Carlo procedures, the ML architectures, the sensitivity-analysis methodology and the calibration and benchmarking procedure; the Results and Discussion section reports and discusses the results; and the Conclusion summarises the main findings, limitations and directions for future work.

MATERIALS AND METHODS

Reference Plant, Feedstock and Design Basis

The reference system is an agricultural co-digestion biogas plant with a 5500 m³ working-volume continuously stirred-tank digester, fed a blend of cattle manure, source-separated food waste and maize silage, and coupled to a reciprocating-engine CHP unit with heat recovery to a 90 °C hot-water loop. This configuration and scale are representative of mid-sized European agricultural biogas installations [7, 8]. Table 1 summarises the fixed design parameters and the nominal (design-point) operating condition used throughout as a reference case.

Table 1. Fixed Design Parameters and Nominal Operating Point of the Reference Industrial Plant.

Parameter	Symbol	Value	Unit
Digester working volume	V_{dig}	5500	m ³
CHP heat-recovery supply temperature	T_{hw}	90 (363.15)	°C (K)
Digester head-space pressure	p_{dig}	1.03	bar
Background inert N ₂ mole fraction in biogas	y_{N_2}	0.02	-
Nominal digester temperature	T_{dig}	38.0	°C
Nominal hydraulic retention time	HRT	28.0	days
Nominal organic loading rate	OLR	3.2	kg VS/m ³ /d
Nominal C/N ratio	C/N	27.0	-
Nominal digester pH	pH	7.1	-
Nominal feedstock mass fractions (manure/food/silage)	$f_m/f_i/f_s$	0.45/0.30/0.25	-
Nominal CHP excess-air ratio	λ	1.15	-
Reference (dead-state) temperature	T_0	300.0 (26.85)	K (°C)
Reference (dead-state) pressure	p_0	1.01325	bar

Table 2 depicts the substrate-specific biochemical methane potential (BMP) values adopted for the feedstock-blending sub-model, consistent with the ranges reported for these three substrate classes [7, 8], following the standardised batch BMP assay protocol of [22], building on established biodegradability-assessment methodology [23].

Table 2. Substrate-Specific Biochemical Methane Potential (BMP) Values Used in The Feedstock-Blending Sub-Model.

Substrate	BMP(Nm ³ CH ₄ /kgVS)	Literature range	Key reference(s)
Cattle manure	0.250	0.20-0.30	[7, 8]
Food waste (source-separated)	0.420	0.35-0.50	[7, 8]
Maize silage	0.355	0.30-0.39	[8]

Kinetic Sub-Model

Two complementary kinetic descriptions are used. For illustrative batch digestibility behaviour (Figure 1), the modified-Gompertz equation [5, 6] is used in its standard form:

$$B(t) = B_0 \cdot \exp\{-\exp[(R_m \cdot e/B_0) \cdot (\lambda - t) + 1]\} \quad (1)$$

where $B(t)$ is the cumulative specific methane yield at time t , B_0 the ultimate methane yield, R_m the maximum specific production rate and λ the lag-phase duration; e is Euler's number. For the continuously fed industrial digester, a steady-state, multi-factor efficiency formulation is used, in which the realisable specific methane yield (SMY) is the feedstock-weighted BMP multiplied by five bounded, (0,1)-valued efficiency factors representing, respectively, digestion temperature, HRT, OLR, C/N ratio and pH:

$$SMY = BMP_{mix} \times \eta_{T(T)} \times \eta_{HRT(HRT)} \times \eta_{OLR(OLR)} \times \eta_{CN(C/N)} \times \eta_{pH(pH)} \quad (2)$$

with BMP_{mix} the feedstock-mass-fraction-weighted average of the values in Table 2. η_T is a double-Gaussian function peaking at 37 °C (mesophilic) and 55 °C (thermophilic) [9]; $\eta_{HRT} = 1 - \exp(-HRT/12)$ captures biomass washout at short retention time [12]; $\eta_{OLR} = 1/(1+(OLR/5)^4)$ captures VFA-accumulation inhibition at high loading [10]; η_{CN} is a Gaussian centred on the C/N optimum of 27.5 [7, 8], and η_{pH} is a Gaussian centred on pH 7.0, reflecting reported ammonia/VFA pH-inhibition behaviour [10, 11]. This functional-factor approach is the standard reduced-order counterpart of the detailed inhibition kinetics embedded in mechanistic models such as ADM1 [4]. The optima of all factors and the temperature, HRT and OLR factors are taken from the cited literature, whereas the widths of the C/N and pH factors (60 and 2.46, respectively) were calibrated against measured steady-state CSTR data, as described in the calibration and benchmarking subsection below. The raw-biogas methane volume (mole) fraction x_{CH_4} and trace hydrogen-sulphide concentration are obtained from an empirical correlation in feedstock blend and OLR, calibrated to the composition ranges reported for full-scale co-digestion plants by [14] (CH₄ 45-70%, H₂S up to several thousand ppm).

Thermodynamic and Exergy Sub-Model (PYroMat)

All thermodynamic properties of the raw biogas and of the CHP combustion reactants and products are computed directly from the open-source PYroMat ideal-gas property library [15], which implements NASA-polynomial-based property correlations for each pure species (CH_4 , CO_2 , H_2S , N_2 , O_2 , H_2O). The biogas is treated as an ideal-gas mixture of CH_4 , CO_2 , H_2S and N_2 ; mixture enthalpy and entropy are obtained by mole-fraction (Gibbs/Amagat) mixing of the pure-species PYroMat properties, which is exact for ideal gases. Three thermodynamic quantities central to the analysis are computed entirely from first principles rather than from tabulated constants: (i) the lower heating value of methane, obtained from a PYroMat energy balance of the reaction $\text{CH}_4 + 2 \text{O}_2 \rightarrow \text{CO}_2 + 2 \text{H}_2\text{O}(\text{g})$ evaluated at the dead state, which returned 50.02 MJ/kg - within 0.1% of the accepted textbook value of approximately 50.0 MJ/kg, validating the property-evaluation chain; (ii) the physical exergy of the biogas stream, $\text{ex}_{\text{ph}} = (h - h_0) - T_0(s - s_0)$, evaluated directly from PYroMat enthalpy and entropy at the digester temperature relative to the dead state ($T_0 = 300.0 \text{ K}$, $p_0 = 1.01325 \text{ bar}$; the dead-state temperature is set 1.85 K above the ISO-standard 298.15 K so that it lies within the valid temperature range of every ideal-gas species used, including H_2S , which is only defined by PYroMat from 300 K upward); and (iii) the adiabatic flame temperature of the biogas-air combustion reaction, obtained by solving the PYroMat-based reactant/product energy balance.

$$H_{\text{reactants}}(T_{\text{in}}) = H_{\text{products}}(T_{\text{ad}}) \quad (3)$$

for T_{ad} by bisection (`scipy.optimize.brentq`), where the reactant stream comprises the raw biogas plus combustion air supplied at excess-air ratio λ (fraction of stoichiometric O_2), and the product stream comprises CO_2 , H_2O , excess O_2 and N_2 . The standard chemical exergy of the biogas mixture is obtained from the pure-component standard chemical exergy values tabulated by Szargut, Morris and Steward [13] (methane: 831,650 kJ/kmol; carbon dioxide: 19,870 kJ/kmol; nitrogen: 720 kJ/kmol; hydrogen sulphide: 812,000 kJ/kmol), combined via ideal mixing plus the standard entropy-of-mixing correction, $RT_0 \sum(y_i \times \ln y_i)$. The CHP engine's electrical efficiency is represented by an empirical correlation, bounded to 26-44%, increasing with CH_4 content and mildly penalised by excess dilution air, consistent with reported ranges for biogas reciprocating engine-generators [7, 8]; the exergy of the electrical output is taken as 100% of its energy value, and the exergy of the recovered thermal output is discounted by the Carnot factor $(1 - T_0/T_{\text{hw}})$ evaluated at the 90 °C hot-water supply temperature. This reference formulation treats the heat as if it were delivered isothermally at the supply temperature and neglects the return-water temperature. When heat is delivered by hot water cooling from a supply temperature T_s to a return temperature T_r , its exergy is more accurately obtained by integrating over this temperature glide, which is equivalent to replacing T_{hw} by the thermodynamic mean temperature $T_{\text{lm}} = (T_s - T_r)/\ln(T_s/T_r) < T_s$. The effect of this refinement and of the supply temperature is quantified in a sensitivity analysis (Supplementary Figure S2 and Table S2). The Carnot factor evaluates the work-equivalent content of the delivered

heat and does not assume that the heat transfer from the engine exhaust and jacket to the hot-water loop is reversible; the exergy destroyed in that heat transfer is included, together with the combustion irreversibility, in the lumped exergy-destruction term. The overall exergy efficiency is then the ratio of exergy output (electricity plus discounted heat) to the total exergy input (chemical plus physical exergy of the biogas feed to the engine).

Industrial-Run, Dataset-Generation and Monte-Carlo Design

Three complementary numerical experiments were performed with the identical mechanistic model described in the kinetic and thermodynamic sub-model subsections above, implemented as a single, fully vectorised Python function so that the same physics underlies every experiment. (a) A 365-day, daily-resolution 'industrial run' propagates a stochastic seasonal ambient-temperature signal, a closed-loop digester-temperature controller with realistic control error, a weekday/weekend food-waste delivery pattern and small operational noise on OLR, HRT, C/N, pH and excess-air ratio, to obtain a full year of plant operation at daily resolution (365 records). (b) A machine-learning dataset was generated by Latin-Hypercube sampling (LHS) [24] of eight operating parameters - digester temperature, HRT, OLR, C/N ratio, pH, manure mass fraction, food-waste mass fraction and excess-air ratio - over the ranges given in Table 3 below, drawing 4000 design points; the resulting raw design matrix (before model evaluation) is released as the 'initial dataset', and the same design matrix propagated through the mechanistic model (adding SMY, biogas composition, CHP power, exergy and all intermediate quantities) is released as the 'full ML dataset'. (c) An independent Monte Carlo (MC) uncertainty-propagation study drew 10,000 realisations from approximately normal distributions centred on the nominal operating point of Table 1, with standard deviations chosen to represent realistic industrial measurement and control variability (e.g. ± 1.5 °C on digester temperature, ± 0.35 kg VS/m³/d on OLR, ± 0.15 pH units), propagated through the identical mechanistic model to obtain output probability distributions.

Table 3. Sampling Ranges of The Eight Operating Parameters Used for Latin-Hypercube ML-Dataset Generation (Experiment (B) of The Dataset-Generation Design) and Sobol' Global Sensitivity Analysis. The Nominal Design Point (Table 1) Lies Within Every Range

Parameter	Symbol	Lower bound	Upper bound	Unit
Digester temperature	T_{dig}	30.0	58.0	°C
Hydraulic retention time	HRT	15.0	45.0	days
Organic loading rate	OLR	1.5	5.5	kg VS/m ³ /d
C/N ratio	C/N	15.0	35.0	-
Digester pH	pH	6.6	7.8	-
Manure mass fraction	f_{manure}	0.15	0.70	-
Food-waste mass fraction	f_{food}	0.10	0.55	-
CHP excess-air ratio	λ	1.05	1.40	-

Maize silage mass fraction (f_{silage}) is not independently sampled; it is set as the remainder $1 - f_{\text{manure}} - f_{\text{food}}$, floored at 0.03, after which all three fractions are renormalised so that they sum exactly to unity (see the treatment of feedstock composition below).

Because the three feedstock fractions are constrained to sum to unity, they cannot all be sampled independently. The manure and food-waste fractions were therefore sampled over the ranges of Table 3, and the silage fraction was obtained as the remainder; where the sampled manure and food-waste fractions summed to more than 0.97, the remainder was floored at 0.03 and all three fractions were renormalised. This procedure guarantees physically valid compositions (non-negative fractions summing exactly to one, to machine precision), but it is active in 15.9% of the LHS samples and 15.6% of the Sobol' samples (0.01% of the Monte Carlo samples), in which the manure and food-waste fractions are rescaled by 0.05 on average (at most 0.15) from their sampled values. It also concentrates these samples in a narrow band of low silage fraction (0.024–0.030, the lower bound falling slightly below the floor because renormalisation follows the floor) along the line $f_{\text{manure}} + f_{\text{food}} \approx 0.97$, and it makes the final fractions mutually correlated ($r = -0.74$ for manure-silage and -0.58 for food-silage; Supplementary Figure S1). The consequences of this treatment for the sensitivity and feature-importance results were assessed with the robustness checks described in the following two subsections.

Machine-Learning Surrogate Models

Three regression architectures were trained on the 4000-point LHS dataset to predict two headline outputs - the specific methane yield (SMY) and the overall CHP exergy efficiency - from the nine feedstock/operating features (the eight LHS variables plus the derived silage mass fraction): (i) Random Forest regression [16, 25], an ensemble of decorrelated regression trees; (ii) Gradient Boosting regression [17, 25], a sequential, gradient-descent-in-function-space ensemble of shallow regression trees; and (iii) a shallow, fully connected artificial neural network (32-16 hidden units, ReLU activation, Adam optimiser, early stopping), trained on standardised features. All models were implemented with scikit-learn [25]. For each target, an 80/20 train-test split was used for the headline metrics (coefficient of determination R^2 , root-mean-square error RMSE, mean absolute error MAE and mean absolute percentage error MAPE), complemented by 3-fold cross-validation across the full dataset to check for overfitting. Because both the training and the test data are generated by the mechanistic model itself, these metrics measure how faithfully each surrogate reproduces the simulator (surrogate fidelity); they do not measure how accurately either the simulator or the surrogate predicts the behaviour of a real plant, which would require benchmarking against measured data. Permutation feature importance (15 repeats, scored on held-out R^2) was extracted from the best-performing model for each target as a model-agnostic ranking of input relevance, to be cross-checked against the Sobol' sensitivity ranking described in the following subsection. Because the three feedstock fractions are compositionally dependent, the importance of the feedstock composition as a whole was additionally evaluated by permuting the three fractions jointly (whole rows), which keeps every permuted composition physically valid.

Global Sensitivity Analysis (Sobol')

Global, variance-based sensitivity analysis was performed directly on the mechanistic model (not on the ML surrogate) using the Sobol' method [19,20] as implemented in the open-source SALib library [21], with Saltelli quasi-random sampling of the eight-dimensional parameter space (Table 3) at a base sample size of $n = 1024$, giving $1024 \times (8+2) = 10,240$ model evaluations per output (second-order indices were not computed, consistent with standard practice when only main-effect screening is required). First-order (S1) and total-order (ST) sensitivity indices were computed for three headline outputs: specific methane yield, CHP exergy efficiency and electrical power output. S1 quantifies the fraction of output variance attributable to a parameter acting alone; ST additionally captures interaction effects with all other parameters, so that $ST > S1$ indicates non-negligible interactions. The Sobol' decomposition requires statistically independent inputs; the sampled manure and food-waste fractions satisfy this requirement, but after renormalisation their indices describe the sampled coordinates rather than the exact final fractions of the clipped samples. As a robustness check, the analysis was therefore repeated with an independent, clip-free parameterisation of the blend, in which the manure fraction (0.15-0.70) and the food-waste share of the non-manure feed, s (0.20-0.90), are sampled independently and $f_{\text{silage}} = (1 - s)(1 - f_{\text{manure}}) \geq 0.03$ holds by construction.

Model Calibration and Benchmarking

The kinetic sub-model was calibrated and benchmarked against 21 steady-state operating periods of laboratory CSTRs co-digesting cattle manure and maize silage, taken from two independent studies: study in [26] (5 L working volume, 37 °C, manure alone and a 4:1 manure/maize-silage blend, OLR 1.1-5.6 kg VS m⁻³ d⁻¹, HRT 15-30 d; 14 periods) and [27] (4 L working volume, 35 °C, maize shares of 0-67% of feed VS, OLR 2-3 kg VS m⁻³ d⁻¹, HRT 23-30 d; 7 periods). For each period, the reported temperature, HRT, OLR, feed composition (VS basis) and digester pH were used as model inputs, and C/N was used where reported in [26]; the measured SMY was taken directly or computed as the reported specific biogas yield multiplied by the measured CH₄ fraction. The authors in [26] have seen that blend ratio was converted to a VS-basis maize-silage fraction of 0.22 by assuming a dry-matter basis, consistent with the reported organic content of the mixture. Two periods described by their authors as unstable or overloaded (high VFA) were excluded from calibration, because the steady-state model does not represent VFA accumulation dynamics, and are reported separately. Only the widths of the C/N and pH factors were fitted, by non-linear least squares on the remaining 19 periods; all other parameters, the substrate BMPs and the CHP/exergy sub-model were left unchanged. Predictive accuracy was quantified by leave-one-out cross-validation and by leave-one-source-out validation (fitting on one study and predicting the other). As a mechanistic reference, the ADM1 as implemented in the open-source PyADM1ODE package [28, 29] was run to steady state with its library characterisations of cattle manure and maize silage (i) at 45 operating points spanning 35-42 °C, OLR 1.5-5.5 kg VS m⁻³ d⁻¹ and manure shares of 30-70% of feed VS at the reference digester volume, and (ii) at the 21 experimental operating points, with the

manure diluted to reproduce the experimental HRT. Agreement was quantified by the root-mean-square error (RMSE), the mean absolute error (MAE), the coefficient of determination (R^2) and the normalised RMSE (NRMSE = RMSE divided by the mean of the reference values).

Software and Reproducibility

The complete pipeline - mechanistic model, industrial run, dataset generation, Monte Carlo propagation, machine-learning training and Sobol' analysis - is implemented in an openly released Python script, complemented by separate scripts for calibration, ADM1 benchmarking, composition diagnostics and heat-exergy sensitivity, built on PYroMat [15] for thermodynamics, SALib [21] for sensitivity analysis, scikit-learn [25] for machine learning, and NumPy, SciPy and pandas [30–32] for numerical and data-handling infrastructure. A fixed random seed was used throughout so that every reported number, table and figure in this manuscript is exactly reproducible from the accompanying code and data files.




Figure 1. Modified-Gompertz Batch Cumulative Methane Yield Curves Illustrating the Kinetic Behaviour of AD at Psychrophilic-Mesophilic Transition, Mesophilic Optimum and Thermophilic Digestion Temperatures (Eq. 1).

Table 4 shows a representative excerpt of the raw 4000-point Latin-Hypercube 'initial dataset' described in the dataset-generation subsection, prior to propagation through the mechanistic model; the complete design matrix and the corresponding full ML dataset (with all simulated outputs appended) accompany this manuscript as CSV files.

Table 4. Representative Excerpt (First 6 of 4000 Records) of The Latin-Hypercube-Sampled Machine-Learning Design Matrix ('Initial Dataset')

ID	T _{dig} (°C)	HRT (d)	OLR (kgVS/m ³ /d)	C/N	pH	f _{manure}	f _{food}	λ	f _{silage}
1	41.03	27.04	5.14	30.20	7.17	0.287	0.525	1.355	0.188
2	31.71	21.53	1.61	23.20	7.25	0.278	0.356	1.384	0.366
3	30.63	15.19	1.82	18.78	7.58	0.435	0.176	1.235	0.389
4	50.33	27.02	5.08	30.94	6.91	0.479	0.368	1.262	0.152
5	55.98	30.05	2.29	29.26	7.35	0.623	0.145	1.103	0.232
6	36.44	38.41	2.49	32.70	6.70	0.569	0.404	1.294	0.027

RESULTS AND DISCUSSION

Nominal Operating Point: Kinetic, Thermodynamic and Exergy Performance

In the nominal operating point of Table 1, the feedstock-weighted theoretical BMP is 0.327 Nm³ CH₄/kg VS, of which the multi-factor efficiency model (Eq. 2) realises 76.0% (individual factors: $\eta_T = 0.984$, $\eta_{HRT} = 0.903$, $\eta_{OLR} = 0.856$, $\eta_{CN} = 1.00$, $\eta_{pH} = 0.998$), giving a specific methane yield of 0.249 Nm³ CH₄/kg VS. With a volatile-solids load of 17,600 kg VS/day, this corresponds to a raw biogas flow of 8205 Nm³/day at a CH₄ content of 53.4% and a trace H₂S concentration of 1320 ppm, both squarely within the ranges reported for full-scale co-digestion plants [14]. The PYroMat-evaluated lower heating value of pure methane, 50.02 MJ/kg (802,557 kJ/kmol), matched the accepted textbook value to within 0.1%, confirming the correctness of the property-evaluation chain described in the thermodynamic and exergy sub-model subsection. Combusted at the nominal excess-air ratio of 1.15, the biogas-air mixture reaches a PYroMat-computed adiabatic flame temperature of 1959 K (Figure 2) - consistent with typical reported flame temperatures for lean-mixture biogas combustion, which are appreciably lower than for natural gas owing to the CO₂ dilution of the fuel.

On the energy side, the nominal-point fuel energy input (1814 kW, LHV basis) is converted at an electrical efficiency of 34.6% and a recoverable-thermal efficiency of 47.9%, giving 627.5 kW_{el} and 869.6 kW_{th} - an overall first-law (energy) efficiency of 82.5%, in line with typical CHP data sheets for reciprocating biogas engines. On the exergy side, the picture is markedly less favourable: of the 1913.8 kW of exergy entering the engine (chemical exergy 451,666 kJ/kmol biogas dominating over a physical-exergy contribution of only 48.3 kJ/kmol at the digester temperature), only 778.7 kW (40.7%) is recovered as useful exergy (electricity plus Carnot-discounted heat), while 1135.1 kW (59.3%) is destroyed in the combustion and heat-transfer processes or lost with the exhaust and to the surroundings, which the lumped CHP model does not separate. The gap between the two efficiencies is caused almost entirely by the recovered heat. Electricity is pure exergy and contributes nearly equally on both bases (34.6% of the energy input and 32.8% of the

exergy input, the small difference arising because the chemical exergy of the biogas exceeds its LHV by 5.5%). The recovered heat, by contrast, accounts for 47.9% of the energy input but only 7.9% of the exergy input, because at a delivery temperature of 90 °C only 17.4% of it (the Carnot factor $1 - T_0/T_{hw}$) is work-equivalent. First-law accounting therefore treats low-temperature heat as equivalent to electricity, whereas second-law accounting discounts it according to its temperature - a distinction that is only visible once the analysis is built on actual thermodynamic property evaluation rather than on heating-value bookkeeping alone. Because the exergy efficiency depends on how the heat is valued, it is also sensitive to the assumed thermal boundary conditions. Evaluating the heat exergy over the temperature glide of a typical 90/70 °C hot-water loop, rather than isothermally at the 90 °C supply temperature, lowers the exergy efficiency from 40.7% to 39.6%; a 70/50 °C low-temperature network lowers it to 37.3%, and a 110/80 °C loop raises it to 41.2%. Over supply temperatures of 70-110 °C and supply-return differences of 0-40 K, the exergy efficiency ranges from 36.0% to 42.7%. The reference value of 40.7% is therefore an upper estimate for the stated 90 °C supply temperature, while the qualitative conclusion - that less than half of the input exergy is recovered and that the heat contributes little exergy - holds across all realistic conditions.

PYroMat-computed adiabatic flame temperature
vs. excess-air ratio and biogas CH₄ content




Figure 2. Pyromat-Computed Adiabatic Flame Temperature as a Function Of CHP Excess-Air Ratio (Λ) And Biogas CH₄ Mole Fraction, Obtained By Solving The Reactant/Product Energy Balance (Eq. 3) At Each Grid Point.

One-Year Industrial-Run Performance

The 365-day simulated industrial run (Figure 3) confirms that the nominal-point figures above are representative of a full annual operating envelope rather than an isolated

design point: over the simulated year the plant delivered 5374 MWh of electricity and 7475 MWh of recoverable heat, corresponding to an average electrical output of 614 kW - i.e. close to full nominal capacity, reflecting the baseload nature of biogas CHP operation. The seasonal ambient-temperature cycle (approximately -2 to $+22$ °C) produces a compensating digester-temperature control trajectory (36.6 – 39.4 °C under closed-loop control) and a mild winter-loading adjustment; the weekday/weekend food-waste delivery pattern is visible as short-period ripple superimposed on the seasonal trend in feedstock composition, biogas flow and CHP power output. No process instability or runaway inhibition was observed at any point in the simulated year, consistent with the plant being operated inside the favourable region of the temperature-OLR efficiency landscape characterised in the response-surface analysis below. Figure 4 presents the same annual run as a three-dimensional trajectory of ambient temperature and electrical power output over the year, making the coupling between the seasonal ambient cycle and the plant's power output visually explicit.




Figure 3. One-Year (365-Day), Daily-Resolution Industrial-Scale Simulation of The Reference Plant: (Top) Ambient and Digester Temperature; (Middle) Raw Biogas Flow Rate; (Bottom) CHP Electrical And Recoverable Thermal Power Output.

3-D annual operating trajectory of the industrial CHP plant




Figure 4. Three-Dimensional Annual Operating Trajectory of The Industrial CHP Plant (Day of Year, Ambient Temperature, Electrical Power Output), Coloured By Day Of Year, Illustrating The Coupled Seasonal-Diurnal Structure of the Simulated Industrial Run.

Response-Surface Behaviour of the Mechanistic Model

Figures 5-7 present three-dimensional response surfaces of the mechanistic model over pairs of operating parameters, holding all other variables at their nominal values (Table 1). Figure 5 shows the specific methane yield as a function of digester temperature and HRT: the surface exhibits the expected double-peaked temperature dependence (mesophilic optimum near 37-38 °C and a secondary thermophilic optimum near 55 °C, separated by a trough around 45-48 °C where neither microbial community is well adapted [9]) combined with a monotonically increasing, saturating dependence on HRT as washout risk recedes [12].

Figure 6 shows the CHP exergy efficiency as a function of OLR and food-waste mass fraction: exergy efficiency falls with increasing OLR (through the OLR-dependence of the CH₄ fraction and hence of the engine's electrical efficiency correlation) and rises with food-waste fraction (which increases CH₄ content and therefore fuel quality), in agreement with the Sobol' and permutation-importance rankings presented below, both of which identify feedstock blend and OLR as the dominant drivers of exergy performance.

Figure 7 shows the raw biogas flow rate as a function of OLR and C/N ratio: biogas output scales essentially linearly with organic loading (more feedstock, more gas), whereas the surface is practically flat along the C/N axis, because the C/N factor calibrated against measured CSTR data exhibits no material effect over the sampled range.

3-D response surface: specific methane yield
vs. digester temperature and HRT




Figure 5. Three-Dimensional Response Surface of The Specific Methane Yield (SMY) as a Function of Digester Temperature and Hydraulic Retention Time (HRT), all Other Parameters Held at Their Nominal Values.

3-D response surface: CHP exergy efficiency
vs. organic loading rate and feedstock mix




Figure 6. Three-Dimensional Response Surface of The Overall CHP Exergy Efficiency As A Function of Organic Loading Rate (OLR) And Food-Waste Mass Fraction, all Other Parameters Held at Their Nominal Values.




Figure 7. Three-Dimensional Response Surface of The Raw Biogas Flow Rate as a Function of Organic Loading Rate (OLR) And C/N Ratio, All Other Parameters Held at Their Nominal Values.

Monte Carlo Uncertainty Propagation

Propagating realistic operating-parameter uncertainty through the mechanistic model ($n = 10,000$ realisations) yields the output distributions summarised in Table 5 and Figures 8 and 9. The specific methane yield has a mean of $0.237 \text{ Nm}^3 \text{ CH}_4/\text{kg VS}$ with a 5th-95th percentile interval of (0.198, 0.271), i.e. a relative spread of approximately $\pm 15\%$ around the mean at the two-sided 90% level, driven mainly by the temperature and OLR sensitivity documented in the Sobol' analysis. By contrast, the CHP exergy efficiency is tightly distributed (mean 0.406, 5th-95th percentile interval (0.399, 0.413)), because it depends on the operating conditions only through the biogas CH_4 content, which is governed by OLR and feedstock blend, and the assumed operational (as opposed to design-space) uncertainty on these variables is comparatively modest.

The adiabatic flame temperature distribution (mean 1960 K, 5th-95th percentile interval (1880, 2045) K) mirrors the CH_4 -content variability propagated from the feedstock-blend and OLR uncertainty. These results show that uncertainty in upstream biological performance (SMY) does not translate one-to-one into uncertainty in downstream thermodynamic performance (exergy efficiency): the two quantities share only one strong driver (OLR), and temperature uncertainty, which dominates the SMY spread, has no effect on exergy efficiency - a finding that is only made visible by propagating the same distribution through a single, internally consistent model of both sub-systems.




Figure 8. Three-Dimensional Monte Carlo Uncertainty-Propagation Point Cloud (n = 10,000) In Digester-Temperature/HRT/Specific-Methane-Yield Space, Coloured by CHP Exergy Efficiency.




Figure 9. Monte Carlo (n = 10,000) Output Probability Distributions for Specific Methane Yield, Electrical Power Output, CHP Exergy Efficiency and Adiabatic Flame Temperature, With Mean And 5th/95th Percentile Markers

Table 5. Monte Carlo (n = 10,000) Output Distribution Summary Statistics

Output	Mean	Std. dev.	5th pct.	Median	95th pct.
SMY (Nm ³ CH ₄ /kg VS)	0.2374	0.0223	0.1981	0.2393	0.2706
Biogas flow (Nm ³ /d)	7768.5	704.2	6480.8	7872.9	8722.3
Electrical power (kW)	592.38	50.53	497.81	599.71	660.72
Recoverable heat (kW)	823.17	71.56	690.97	833.68	919.04
Exergy efficiency (-)	0.4062	0.0044	0.3989	0.4061	0.4134
Adiabatic flame T (K)	1960.1	48.8	1880.0	1959.1	2044.9

Global Sensitivity Analysis: Sobol' Indices vs. Machine-Learning Feature Importance

Table 6 reports the Sobol' total-order sensitivity indices (ST) for the three headline outputs, obtained from 10,240 mechanistic-model evaluations per output; Figures 10 and 11 visualise the same results as, respectively, a three-dimensional and a two-dimensional bar chart. For the specific methane yield, digester temperature (ST = 0.653) and OLR (ST = 0.323) dominate, with HRT a minor contributor (ST = 0.038) and feedstock blend, pH, C/N ratio and excess-air ratio essentially negligible (ST ≤ 0.014). The small pH and C/N indices are a direct consequence of the calibration against measured CSTR data, which showed no detectable C/N effect between C/N 13.5 and 27.5 and only a weak pH effect over pH 6.75–7.5, so that these two factors remain close to unity over the sampled ranges. The large temperature index must be interpreted with care: it arises mainly from the trough between the mesophilic and thermophilic optima of the literature-based temperature factor (around 45–48 °C), which lies inside the 30–58 °C sampling range, and this factor could not be benchmarked because all available steady-state data were obtained at 35–37 °C. For the CHP exergy efficiency, OLR (ST = 0.433) and the feedstock mass fractions (manure ST = 0.310, food waste ST = 0.209) dominate, with excess-air ratio a minor contributor (ST = 0.047) and digester temperature, HRT, C/N and pH all statistically indistinguishable from zero, because, in the present model structure, exergy efficiency depends on the operating conditions only indirectly, through their effect on biogas CH₄ content and hence on engine electrical efficiency, and OLR and feedstock blend are precisely the variables that determine CH₄ content in the composition correlation. OLR is thus the only variable that strongly affects both outputs.

This Sobol' ranking is consistently reproduced by the permutation feature importance extracted from the best-performing (Gradient Boosting) ML surrogate for each target (Figure 12): for SMY, the ML model ranks digester temperature (importance 1.205) and OLR (0.677) as dominant, followed by HRT (0.069), in the same order as the Sobol' ST indices; for exergy efficiency, it ranks OLR (0.926), manure fraction (0.580), food-waste fraction (0.372) and excess-air ratio (0.087), again matching the Sobol' ranking exactly. Because the two methods are both derived from the same mechanistic model, their agreement is an internal-consistency check rather than independent evidence: it confirms that the ML surrogate has learned the input-output structure of the process model rather

than an artefact of the sampling design, while the accuracy of that structure against measured data is assessed separately in the calibration and benchmarking section.

The robustness checks show that the treatment of feedstock composition does not affect the conclusions for specific methane yield or electrical power: with the independent, clip-free blend parameterisation, the SMY indices of digester temperature ($ST = 0.646$) and OLR (0.320) are practically unchanged, and the blend remains negligible. For exergy efficiency, the independent parameterisation assigns a larger combined index to the blend (manure fraction $ST = 0.546$, food share $ST = 0.124$) and a smaller one to OLR (0.308) than the original analysis, partly because it spans a wider food-waste range (0.06 - 0.76). Consistently, jointly permuting the three fractions in the ML surrogate gives a composition importance of 1.03 , above that of OLR (0.93). Feedstock composition, taken as a whole, is therefore the largest driver of exergy efficiency, followed by OLR, while the split of its influence between individual fractions depends on how the dependent fractions are parameterised. The near-zero individual importance of the silage fraction reflects its redundancy (it is fully determined by the other two fractions) rather than a physical effect: removing it as an ML feature leaves the test R^2 unchanged (0.989 and 0.995). Finally, the surrogates are as accurate on the clipped boundary samples as on the remaining samples (test $R^2 = 0.989$ vs. 0.990 for SMY and 0.995 vs. 0.996 for exergy efficiency), so the local concentration of samples does not degrade the ML results.

3-D Sobol' total-order sensitivity indices across parameters and model outputs




Figure 10. Three-Dimensional Bar Chart of Sobol' Total-Order Sensitivity Indices (ST) for Specific Methane Yield, CHP Exergy Efficiency and Electrical Power Output, Across all Eight Sampled Operating Parameters




Figure 11. Sobol' First-Order (S1) and Total-Order (ST) Sensitivity Indices for Specific Methane Yield, CHP Exergy Efficiency and Electrical Power Output (Two-Dimensional Companion View of Figure 10).

Table 6. Sobol' Total-Order (ST) Global Sensitivity Indices for the Three Headline Model Outputs (Base Sample $n = 1024$; 10,240 Model Evaluations per Output).

Parameter	ST (SMY)	ST (Exergy eff.)	ST (P _{el})
Digester temperature	0.653	<0.001	0.782
HRT	0.038	<0.001	0.046
OLR	0.323	0.433	0.128
C/N ratio	0.001	<0.001	0.001
pH	0.005	<0.001	0.006
Manure mass fraction	0.014	0.310	0.036
Food-waste mass fraction	0.004	0.209	0.015
Excess-air ratio	<0.001	0.047	<0.001




Figure 12. Permutation Feature Importance (15 Repeats, Held-Out R^2 Scoring) of the Best-Performing Gradient Boosting Surrogate for Specific Methane Yield and CHP Exergy Efficiency, Consistent with the Sobol' Ranking of Table 6

Machine-Learning Surrogate Model Performance

Table 7 reports the performance of the three regression architectures for both targets, evaluated as surrogates of the mechanistic model on held-out simulator outputs (20% test split).

Table 7. Surrogate-Fidelity Metrics of The Three Regression Architectures for Both Machine-Learning Targets, Evaluated on Held-Out Mechanistic-Model Outputs (20% Test Split) and by 3-Fold Cross-Validation

Target	Model	R ² (test)	RMSE (test)	MAE (test)	CV R ² (mean +/- std)
SMY	Random Forest	0.980	0.0095	0.0074	0.978 +/- 0.001
SMY	Gradient Boosting	0.990	0.0070	0.0055	0.989 +/- 0.001
SMY	MLP (ANN)	0.978	0.0101	0.0078	0.951 +/- 0.010
Exergy eff.	Random Forest	0.988	0.0015	0.0012	0.984 +/- 0.001
Exergy eff.	Gradient Boosting	0.995	0.0009	0.0007	0.996 +/- 0.0004
Exergy eff.	MLP (ANN)	0.980	0.0020	0.0015	0.916 +/- 0.008

Gradient Boosting achieved the best performance on both targets (SMY: $R^2 = 0.990$, RMSE = 0.0070 Nm³ CH₄/kg VS; exergy efficiency: $R^2 = 0.995$, RMSE = 0.00093), narrowly outperforming Random Forest ($R^2 = 0.980$ and 0.988, respectively) and the shallow ANN ($R^2 = 0.978$ and 0.980, respectively). Three-fold cross-validation scores were closely consistent with the held-out test scores for the tree-based models (Gradient Boosting cross-validated $R^2 = 0.989$ and 0.996), while the ANN showed a larger cross-validation spread (0.951 +/- 0.010 for SMY and 0.916 +/- 0.008 for exergy efficiency), consistent with the higher variance of small, shallow neural-network fits relative to ensembles of trees on tabular data of this size. Both targets are learned with $R^2 \geq 0.98$ because each depends smoothly on a small effective number of inputs (digester temperature and OLR for SMY; OLR and feedstock blend for exergy efficiency); the slightly lower SMY accuracy reflects its bimodal mesophilic/thermophilic temperature dependence, which is harder to approximate from a finite sample than the near-monotonic exergy response. These R^2 values should therefore be interpreted as surrogate-to-simulator fidelity, not as evidence that the SMY or exergy efficiency of a real plant could be predicted to this accuracy from operating data; the accuracy of the underlying model against measured data is quantified separately in Table 8.




Figure 13. Parity (Predicted vs. Mechanistic-Model True Value) Plots for The Best-Performing (Gradient Boosting) Surrogate Models on the Held-Out Test Set: (a)-Specific Methane Yield; (b)-Overall CHP Exergy Efficiency.




Figure 14. Side-By-Side Three-Dimensional Comparison of (a) The Mechanistic (Pyromat + Kinetic) Model and (b) The Gradient Boosting Machine-Learning Surrogate, for the Specific-Methane-Yield Response Surface Over Digester Temperature and pH, All Other Variables At Their Nominal Values (The Response is Nearly Flat Along The pH Axis After Calibration of the pH Factor).

Calibration and Quantitative Benchmarking

Table 8 and Figure 15 summarise the calibration and benchmarking results. With the original, literature-based factor widths, the model under-predicted the measured SMY severely ($R^2 = -5.18$, $\text{NRMSE} = 69\%$), almost entirely because the C/N and pH factors penalised conditions under which the laboratory reactors operated normally: at the measured manure C/N of 13.5 the original C/N factor reduced the predicted yield by about 80%, whereas the reactor produced $0.22 \text{ Nm}^3 \text{ CH}_4 \text{ kg}^{-1} \text{ VS}$. After calibration, the model reproduced the 19 stable periods with $R^2 = 0.90$ in-sample and $R^2 = 0.89$ under leave-one-out cross-validation ($\text{RMSE} = 0.018 \text{ Nm}^3 \text{ CH}_4 \text{ kg}^{-1} \text{ VS}$, $\text{NRMSE} = 9\%$), including the pronounced decline of SMY with increasing OLR in the Varol and Ugurlu data. The calibrated C/N width reached the upper bound of the search range, which means that the data show no detectable C/N effect between C/N 13.5 and 27.5, and the calibrated pH factor is close to unity over the observed pH range of 6.75–7.5. Leave-one-source-out validation succeeded when predicting in [27] from the [26] data ($R^2 = 0.69$, $\text{NRMSE} = 6\%$) but failed in the opposite direction ($R^2 = -5.15$), because only the [26] study reports C/N and the C/N factor is therefore identifiable from that dataset alone. Over all 21 periods, including the two unstable ones, the calibrated model achieved $R^2 = 0.78$ ($\text{NRMSE} = 13\%$); the largest single error occurred for the overloaded period with 67% maize, for which the model predicted 0.233 against a measured $0.153 \text{ Nm}^3 \text{ CH}_4 \text{ kg}^{-1} \text{ VS}$, since it does not represent VFA accumulation. The empirical biogas-composition correlation was less accurate (CH_4 fraction $\text{RMSE} = 0.085$, $R^2 = -1.54$), under-predicting the measured CH_4 content by about 6 percentage points on average. The ADM1 reference, run with the default PyADM1ODE substrate characterisations and without calibration, under-predicted the experimental SMY ($R^2 = -2.70$, $\text{NRMSE} = 52\%$), chiefly because its library cattle manure is considerably

less degradable than the manures used in both studies; accordingly, the disagreement between the proposed model and ADM1 over the 45-point design ($R^2 = -1.18$, NRMSE = 35%) reflects differences in substrate characterisation and in the treatment of overload as much as model structure, and ADM1 in this configuration should not be regarded as a more accurate reference than the measured data.

Table 8. Calibration and Benchmarking Metrics for Specific Methane Yield ($\text{Nm}^3 \text{CH}_4 \text{ kg}^{-1} \text{VS}$) Unless Otherwise Stated

Comparison	n	RMSE	MAE	R^2	NRMSE (%)
Proposed model, original factors vs. experiment (stable)	19	0.138	0.129	-5.18	69.0
Proposed model, calibrated, in-sample (stable)	19	0.018	0.014	0.90	8.9
Proposed model, calibrated, leave-one-out CV (stable)	19	0.018	0.015	0.89	9.2
Leave-source-out: fitted on [26], predicting [27]	5	0.014	0.012	0.69	6.1
Leave-source-out: fitted on [27], predicting [26]	14	0.153	0.146	-5.15	79.1
Proposed model, calibrated vs. experiment (all periods)	21	0.025	0.018	0.78	12.8
ADM1 (PyADM1ODE, default substrates) vs. experiment	21	0.104	0.094	-2.70	52.4
Proposed model (calibrated) vs. ADM1, 45-point design	45	0.064	0.054	-1.18	35.4
CH_4 mole fraction: proposed model vs. experiment	21	0.085	0.078	-1.54	15.7




Figure 15. Parity Plots of Specific Methane Yield: (a) Calibrated Proposed Model Versus ADM1 (PyADM1ODE) at 45 Steady-State Operating Points, with pH at the Set-Point and pH Taken From ADM1; (b) Calibrated Proposed Model and ADM1 Versus 21 Measured Steady-State CSTR Operating Periods [26, 27]

Comparison with the Literature and Practical Implications

The nominal-point electrical efficiency (34.6%) and overall energy efficiency (82.5%) obtained here fall within the ranges commonly reported for biogas reciprocating-engine CHP installations, and the simulated annual capacity factor (electrical output close to nameplate for most of the year) is consistent with the baseload duty cycle typical of agricultural biogas plants [7, 8]. The $R^2 = 0.980$ -0.995 range achieved by the tree-based ML surrogates lies at the upper end of, but is not directly comparable with, recent reports of

Random Forest and Gradient Boosting models for anaerobic-digestion methane-yield prediction [18], which is expected given that the present training data are drawn from a smooth, physically consistent simulator rather than from noisy field measurements - a distinction that should be kept in mind when transferring the reported accuracy figures to a field-deployment context, where sensor noise, unmeasured confounders and process drift will degrade achievable accuracy relative to the values reported here. From a practical standpoint, the analysis suggests that: (i) plant operators seeking to maximise methane yield should prioritise holding the digester close to its temperature optimum, avoiding the mesophilic-thermophilic transition zone, and keeping OLR below the onset of overload, whereas within the ranges tested C/N and pH had no measurable effect on yield once the process was stable; (ii) operators seeking to maximise CHP exergy performance should prioritise feedstock blend and OLR, essentially independently of digester temperature, HRT, C/N or pH; and (iii) OLR is the one lever that acts on both objectives, in the same direction: increasing it lowers both the specific methane yield and the exergy efficiency while raising total throughput, so the practical trade-off is between plant output and per-unit conversion quality, which a multi-objective operating strategy informed jointly by the mechanistic model and by its ML surrogate (evaluated against held-out mechanistic-model outputs), which can be evaluated in microseconds for real-time supervisory use, can help resolve.

The framework has several limitations that bound the interpretation of the above findings. The multi-factor efficiency formulation of Eq. (2) is a reduced-order surrogate for the full biochemical reaction network resolved by mechanistic models such as ADM1 [4] and does not resolve transient inhibition dynamics (e.g. VFA accumulation followed by recovery) explicitly; accordingly, it over-predicted the yield of the one overloaded experimental period (0.233 vs. 0.153 Nm³ CH₄ kg⁻¹ VS). Its benchmark against measured data covers manure/maize-silage blends at 35-37 °C and OLR 1.1-5.6 kg VS m⁻³ d⁻¹ only. The food-waste BMP and the temperature factor, including the mesophilic-thermophilic transition that produces the large temperature Sobol' index, therefore remain literature-based and untested. The C/N factor is constrained by a single dataset, and its calibrated width reached the upper bound of the search range, so the result should be read as the absence of a detectable C/N effect over C/N 13.5-27.5 rather than as an identified optimum. The empirical biogas-composition correlation under-predicts the measured CH₄ content of the experimental feeds by about 6 percentage points on average, and the cattle-manure BMP of Table 2 (0.250 Nm³ CH₄ kg⁻¹ VS) lies above the 0.212 Nm³ CH₄ kg⁻¹ VS measured by Seppälä et al. [27]. The CHP electrical-efficiency correlation is likewise an empirical fit to typical manufacturer data rather than a full thermal/mechanical engine-cycle simulation; the recovered-heat fraction is also independent of the hot-water temperatures, all recovered heat is assumed to be used, and the exergy destruction is not resolved into its combustion and heat-transfer components; and the synthetic ML/Monte-Carlo datasets, while physically consistent by construction, cannot capture unmodelled real-world phenomena such as sensor drift, feedstock heterogeneity within a nominal category, or foaming/scum-related process upsets. Moreover, the framework is not yet a digital twin in

the strict sense: it has no bidirectional data link to a physical plant, and the 365-day run is driven by synthetic ambient-temperature and feedstock-delivery profiles rather than by measured plant records. Extending the thermodynamic core developed here to a full ADM1-resolved biochemistry module, to a detailed engine-cycle (e.g. Otto-cycle) CHP model, to benchmarking against food-waste co-digestion and thermophilic-range data, and to online coupling with, and periodic recalibration against, field-metered plant data (e.g. SCADA records of biogas flow, CH₄ content, digester temperature and CHP power output), which would turn the present framework into a full plant digital twin, are natural directions for future work.

SUMMARY AND CONCLUSION

This study presented an integrated, open-source, digital-twin-oriented simulation framework that couples a multi-factor kinetic model of industrial co-digestion with a thermodynamically rigorous, PYroMat-based property and exergy model of the downstream CHP unit, and exercised it as a 365-day industrial run, a 4000-point Latin-Hypercube machine-learning dataset, a 10,000-run Monte Carlo uncertainty study and a 10,240-evaluation Sobol' global sensitivity analysis. After calibration of its C/N and pH factors against 19 stable steady-state periods from two independent laboratory CSTR studies, the kinetic model reproduced the measured specific methane yield with a leave-one-out cross-validated R² of 0.89 (NRMSE 9%). At the nominal operating point, the plant achieves an 82.5% first-law (energy) efficiency but only a 40.7% second-law (exergy) efficiency (36.0–42.7% across realistic hot-water supply and return temperatures), the difference arising almost entirely from the low exergy content of the recovered heat, underscoring the value of exergy-based, composition-resolved thermodynamic accounting over simple heating-value bookkeeping.

Gradient Boosting surrogate models reproduced the mechanistic-model predictions of specific methane yield and CHP exergy efficiency with held-out R² of 0.990 and 0.995, respectively, and Sobol' global sensitivity analysis showed that digester temperature and organic loading rate govern methane yield while feedstock composition and organic loading rate govern exergy efficiency - a finding consistently reproduced by permutation feature importance extracted from the trained ML surrogates. Organic loading rate is thus the only operating variable that strongly affects both objectives, lowering both as it rises, whereas temperature acts mainly on yield and feedstock blend mainly on exergy; plant operation therefore involves a trade-off between throughput and per-unit conversion quality that a combined mechanistic/machine-learning decision-support tool of the type demonstrated here can help resolve.

The advantages of the proposed framework are its full reproducibility (an open-source Python implementation with all scripts and data), its thermodynamic rigour (every enthalpy, entropy, adiabatic flame temperature and exergy value is computed from first principles via PYroMat rather than looked up), and its dual physical/data-driven sensitivity accounting; its principal limitations are the reduced-order biochemical kinetics,

which were benchmarked only for manure/maize-silage feeds at mesophilic temperature, the untested temperature factor and food-waste parameters, and the empirical CHP electrical-efficiency correlation, all of which are natural targets for future refinement and for validation against field-metered industrial plant data.

ACKNOWLEDGMENT

The authors gratefully acknowledge the institutional support of the Department of Structural, Geotechnical and Building Engineering, Politecnico di Torino, Italy, and of the Institute for Nanomaterials, Advanced Technologies and Innovation, Technical University of Liberec, Czech Republic.

CONFLICT OF INTERESTS

The authors declare that there is no conflict of interest associated with this publication.

DATA AND CODE AVAILABILITY

The complete Python implementation, the initial and full machine-learning datasets, the industrial-run time series, the Monte Carlo results and summary statistics, the Sobol' sensitivity indices, the machine-learning model-performance and feature-importance tables, the experimental reference dataset extracted from [26, 27], the calibration and ADM1 benchmarking scripts with their outputs, the composition-diagnostics and heat-exergy sensitivity scripts with their outputs and all figures presented in this manuscript are provided as supplementary files accompanying this submission and are fully reproducible from a single fixed random seed.

REFERENCES

1. Outlook for Biogas and Biomethane: Prospects for Organic Growth – Analysis. Available online: <https://www.iea.org/reports/outlook-for-biogas-and-biomethane-prospects-for-organic-growth> (accessed on 15 August 2026).
2. Kougias, P.G., Angelidaki, I. Biogas and Its Opportunities—A Review. *Front. Environ. Sci. Eng.* **2018**, *12*, 14.
3. Achinas, S., Achinas, V., Euverink, G.J.W. A Technological Overview of Biogas Production from Biowaste. *Engineering* **2017**, *3*, 299–307.
4. Batstone, D.J., Keller, J., Angelidaki, I., Kalyuzhnyi, S.V., Pavlostathis, S.G., Rozzi, A., Sanders, W.T.M., Siegrist, H., Vavilin, V.A. The IWA Anaerobic Digestion Model No 1 (ADM1). *Water Science and Technology* **2002**, *45*, 65–73.
5. Zwietering, M.H., Jongenburger, I., Rombouts, F.M., Van 't Riet, K. Modeling of the Bacterial Growth Curve. *Appl Environ Microbiol* **1990**, *56*, 1875–1881.
6. Lay, J.-J., Li, Y.-Y., Noike, T. Influences of pH and Moisture Content on the Methane Production in High-Solids Sludge Digestion. *Water Research* **1997**, *31*, 1518–1524.

7. Al Seadi, T., Rutz, D., Prassl, H., Köttner, M., Finsterwalder, T., Volk, S., Janssen, R. *Biogas Handbook*. University of Southern Denmark. Esbjerg, Denmark, **2008**.
8. Deublein, D.; Steinhauser, A. *Biogas from Waste and Renewable Resources: An Introduction*, 2nd ed.; Wiley-VCH: Weinheim, Germany, **2008**.
9. Weiland, P. Biogas Production: Current State and Perspectives. *Appl. Microbiol. Biotechnol.* **2010**, *85*, 849–860.
10. Chen, Y., Cheng, J.J., Creamer, K.S. Inhibition of Anaerobic Digestion Process: A Review. *Bioresource Technology* **2008**, *99*, 4044–4064.
11. Yenigün, O., Demirel, B. Ammonia Inhibition in Anaerobic Digestion: A Review. *Process Biochemistry* **2013**, *48*, 901–911.
12. Rittmann, B.E., McCarty, P.L. *Environmental Biotechnology: Principles and Applications*, McGraw-Hill New York, USA. **2001**.
13. Szargut, J., Morris, D.R., Steward, F.R. *Exergy Analysis of Thermal, Chemical, and Metallurgical Processes*. University of Michigan, USA, **1988**.
14. Rasi, S., Veijanen, A., Rintala, J. Trace Compounds of Biogas from Different Biogas Production Plants. *Energy* **2007**, *32*, 1375–1380.
15. Martin, C., Ranalli, J., Moore, J. PYroMat: A Python Package for Thermodynamic properties. *JOSS* **2022**, *7*, 4757.
16. Breiman, L. Random Forests. *Machine Learning* **2001**, *45*, 5–32.
17. Friedman, J.H. Greedy Function Approximation: A Gradient Boosting Machine. *Annals of statistics* **2001**, *29*, 1189–1232.
18. Lv, L., Niu, Z., Hao, P., Wang, X., Zheng, M., Ren, Z. Machine Learning Prediction and Process Optimization for Enhanced Methane Production from Straw Anaerobic Digestion with Biochar. *Sustainability* **2026**, *18*, 609.
19. Sobol', I.M. Global Sensitivity Indices for Nonlinear Mathematical Models and Their Monte Carlo Estimates. *Mathematics and Computers in Simulation* **2001**, *55*, 271–280.
20. Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., Tarantola, S. *Global Sensitivity Analysis: The Primer*, John Wiley & Sons, **2008**.
21. Herman, J., Usher, W. SALib: An Open-Source Python Library for Sensitivity Analysis. *JOSS* **2017**, *2*, 97.
22. Angelidaki, I., Alves, M., Bolzonella, D., Borzacconi, L., Campos, J.L., Guwy, A.J., Kalyuzhnyi, S., Jenicek, P., Van Lier, J.B. Defining the Biomethane Potential (BMP) of Solid Organic Wastes and Energy Crops: A Proposed Protocol for Batch Assays. *Water Science and Technology* **2009**, *59*, 927–934.
23. Angelidaki, I., Sanders, W. Assessment of the Anaerobic Biodegradability of Macropollutants. *Rev Environ Sci Biotechnol* **2004**, *3*, 117–129.
24. McKay, M.D., Beckman, R.J., Conover, W.J. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics* **2000**, *42*, 55–61.
25. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. Scikit-Learn: Machine Learning in Python. *The Journal of machine Learning Research* **2011**, *12*, 2825–2830.

26. Varol, A., Ugurlu, A. Comparative Evaluation of Biogas Production from Dairy Manure and Co-digestion with Maize Silage by CSTR and New Anaerobic Hybrid Reactor. *Engineering in Life Sciences* **2017**, 17, 402–412.
27. Seppälä, M., Pyykkönen, V., Väisänen, A., Rintala, J. Biomethane Production from Maize and Liquid Cow Manure – Effect of Share of Maize, Post-Methanation Potential and Digestate Characteristics. *Fuel* **2013**, 107, 209–216.
28. Github. dgaida Dgaida/PyADM1ODE. Available online: <https://github.com/dgaida/PyADM1ODE> (accessed on 11 August 2026).
29. Sadrimajd, P., Mannion, P., Howley, E., Lens, P.N.L. *PyADM1: A Python Implementation of Anaerobic Digestion Model No. 1*, **2021**.
30. Harris, C.R., Millman, K.J., Van Der Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., et al. Array Programming with NumPy. *Nature* **2020**, 585, 357–362.
31. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat Methods* **2020**, 17, 261–272.
32. McKinney, W. Data Structures for Statistical Computing in Python. *SciPy Proceedings* **2010**, 445, 51–56.