energies

MDPI

Article

Py-GC/FID Screening Framework for Coal–Biomass–Solid Recovered Fuel Blends for Co-Firing: Additivity and Surrogate-Assisted Screening for Indonesian Fuels

Soleh Mochamad [1,2,*], Darmanto Prihadi Setyo [1], Pasek Ari Darmawan [1], Mukmin Gea Fardias [1] and Apriani Rachmawati [3]

[1] Faculty of Mechanical and Aerospace Engineering, Bandung Institute of Technology, Jalan Ganesa No. 10, Lebak Siliwangi, Coblong, Bandung 40132, West Java, Indonesia; prihadi.setyo@itb.ac.id (D.P.S.); aripasek@itb.ac.id (P.A.D.); gea_fardias@itb.ac.id (M.G.F.)
[2] PT PLN (Persero) Pusat Penelitian dan Pengembangan Ketenagalistrikan, Jalan Duren Tiga Raya No. 102, Duren Tiga, Kecamatan Pancoran, Kota Jakarta Selatan 12760, Daerah Khusus Ibukota Jakarta, Indonesia
[3] Faculty of Vocation, Department of Pulp and Paper Processing Technology, Bandung Institute of Technology and Science, Kota Deltamas Lot-A1 CBD, Ganesha Boulevard No. 1, Pasirranji, Central Cikarang District, Bekasi Regency 17530, West Java, Indonesia; rachmawati@itsb.ac.id
* Correspondence: 33121003@mahasiswa.itb.ac.id

Abstract

Screening coal–biomass–SRF blends requires a fast testing method. This study assessed whether Curie-point pyrolysis with gas chromatography and flame ionisation detection (Py-GC/FID) can detect non-additive volatile release in blends of an Indonesian sub-bituminous coal with sawdust, rice husk, and SRF BBJP. Twelve blends (20–80 wt.% non-coal) were pyrolysed at 590 °C, and a Shape Deviation Index (SDI, 0–2) was used to quantify the dissimilarity between area-normalised chromatograms. For blends containing sawdust, early peaks outside the 15.68 s analysis window are excluded, so their SDI is not directly comparable to that of blends without sawdust. Three variants are distinguished: SDI_rep, between duplicate runs (noise; median 0.89); SDI_add, the residual from the additive expectation with the coal profile estimated from the blends (1.00); and SDI_ref, the deviation with an external coal reference (1.59, largely reference mismatch). SDI_add exceeded SDI_rep in five of nine blends by a median of 0.05. With duplicate runs, no departure from additivity exceeding run-to-run variation was detected; interactions below about 0.3 in SDI cannot be ruled out. A surrogate with a genetic-algorithm search, trained on SDI_ref, described the measured shape space (highest near 0.8 coal with BBJP; not a recommendation). SDI captures shape, not yield. Screening needs replicated runs and retention-time alignment.

Keywords: co-firing; solid recovered fuel; biomass; pyrolysis-GC/FID; volatile release; additivity; repeatability; surrogate model; Indonesia

Check for updates

Academic Editor: F. Rubiera

Received: 29 August 2026
Revised: 19 October 2026
Accepted: 21 October 2026
Published: 30 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Co-firing biomass and waste-derived fuels in existing pulverised-coal boilers is one of the few options that can raise the share of renewable and recovered energy quickly without new plant [1]. In Indonesia, the state utility PT PLN (Persero) has adopted co-firing with biomass and solid recovered fuel (SRF) as part of its transition strategy [2,3]. The country's SRF is known as BBJP (*bahan bakar jumputan padat*), and it is produced from processed municipal and commercial waste. Agricultural and forestry residues are abundant in

Indonesia, and rice husk is among the residues considered for co-firing [4,5]. Coal supplied more than 64% of Indonesia's electricity generation in 2023 [6].

Selecting blends for such boilers is difficult. The fuels differ strongly in composition and reactivity, and blends do not necessarily behave as the weighted sum of their components. Thermogravimetric studies of coal–biomass blends have reported deviations from the weighted behaviour of the components during heating, with promoting or inhibiting effects on reactivity [7,8]. Blending refuse-derived fuel with coal has also been reported to shift the ignition, peak, and burnout temperatures [9,10]. Volatile release governs ignition and flame stability, and coal and biomass particles differ in how they devolatilise and ignite [11]. However, thermogravimetric analysis reports mass loss only, and full-scale or pilot boiler trials are costly and slow. A rapid method that examines the volatiles themselves with a very small sample would therefore be useful as a pre-screening step before those tests.

Curie-point pyrolysis coupled with gas chromatography and flame ionisation detection (Py-GC/FID) is a candidate for this role. The sample (a fraction of a milligram) is heated to a fixed temperature within milliseconds, and the released volatiles are separated by GC and detected by FID, giving a fingerprint chromatogram within minutes of signal acquisition. Curie-point pyrolysis with gas chromatography has been used to characterise coals and coal macerals [12,13], and pyrolysis-GC/MS has been applied to coal–biomass blends [14]. Interactions between fuels have been investigated in Py-GC/MS studies by comparing the yields of individual compounds with those expected from the components [14,15]. In one such study of pine wood with coals of different rank, pyrolysis-GC/MS and thermogravimetry showed additive behaviour, whereas larger-scale batch pyrolysis gave non-additive tar composition [14]. Such compound-level comparison is informative but demands identification and calibration for many peaks, and it is not easily used as a routine screening measure. A single number that summarises how far a blend chromatogram departs from the weighted sum of its components could be used earlier in the workflow, provided that it is repeatable and that the component chromatograms are available.

This study asks whether such a number, a Shape Deviation Index (SDI), can detect non-additive volatile release in blends of an Indonesian sub-bituminous coal with sawdust, rice husk, and BBJP. SDI is defined as the integrated absolute difference between two area-normalised chromatograms. It measures shape dissimilarity, not total volatile yield, and it is a descriptor, not an objective: A higher value is not treated as more desirable for co-firing. The aims were fourfold. Aim (i) was to quantify how repeatable Py-GC/FID chromatograms of duplicate runs are using the SDI between two runs of the same sample and the sensitivity of that measure to binning and retention-time alignment. Aim (ii) was to test the additivity of twelve binary, ternary, and quaternary blends with 20–80 wt.% non-coal material, far above the biomass pellet shares of about 5% assumed in current PLN planning for pulverised-coal and circulating-fluidised-bed boilers [2]. The coal component was taken from a reference chromatogram and, as a check independent of that reference, from a profile estimated from the blends themselves, and the deviation was compared with the replicate variation. Aim (iii) was to estimate how many replicate runs would be needed to resolve differences between blends. Aim (iv) was to use a Gaussian-process surrogate of SDI with a genetic-algorithm search and a bootstrap as a screening tool for blend compositions and to validate it before interpreting its optimum. The results show that, with duplicate runs, no departure from additivity exceeding run-to-run variation was detected, that interactions below about 0.3 in SDI cannot be ruled out, and that the default surrogate has no cross-validated skill. A surrogate with hyperparameters optimised in a nested validation has moderate skill with wide uncertainty, and its optimum (about 0.8 coal with BBJP) is a hypothesis for confirmation. This paper also sets out what is needed to make Py-GC/FID usable as a co-firing screening tool.

2. Materials and Methods

2.1. Fuel Materials and Sample Preparation

Four fuels were used, representing feedstocks relevant to PLN's co-firing programme [2,3]: an Indonesian sub-bituminous coal (C) from East Kalimantan; rice husk (RH), an agricultural residue; sawdust (SD), a lignocellulosic by-product of the timber and furniture industry; and BBJP (*bahan bakar jumputan padat*), a solid recovered fuel produced at a certified facility in West Java in accordance with SNI 8966:2021 [16]. BBJP is a mixed-waste fuel containing lignocellulosic fibres, polyolefins, polystyrene and textile fragments; the standard specifies a net calorific value of at least 15 MJ/kg and a chlorine content of at most 0.2 wt.% (dry basis) [16]. The fuels are the same materials as those of the authors' thermogravimetric co-combustion study (manuscript in preparation), but the blend compositions differ. Proximate and ultimate analyses of the fuels (as received) are given in Supplementary Table S1. The coal contained 35.9 wt.% volatile matter and 5.0 wt.% ash, with a gross calorific value of 5249 kcal/kg; the volatile matter of sawdust, rice husk, and BBJP was 58.5, 55.6, and 50.1 wt.% (82, 77 and 88 wt.% on a dry, ash-free basis, against 49 wt.% for the coal), and their ash content was 1.2, 19.4, and 36.2 wt.%.

All samples were air-dried at room temperature before milling. Coal, sawdust, rice husk, and BBJP were prepared in the same way: Each was ground in a mortar mill to a particle size below 100 µm. Blends were prepared gravimetrically on an analytical balance (repeatability ±0.1 mg) and mixed by tumbling for 20 min. Twelve blends were prepared: six binary blends (80 and 60 wt.% coal with sawdust, rice husk and BBJP; 80C_20SD, 80C_20RH, 80C_20BBJP, 60C_40SD, 60C_40RH, 60C_40BBJP), three ternary blends (60C_26RH_13BBJP, 60C_13RH_26BBJP, 60C_26SD_13RH), and three quaternary blends (60C_13SD_13RH_13BBJP, 40C_20each, 20C_26each). A blend name gives the nominal coal fraction followed by the fractions of the other components in wt.%; "each" denotes equal fractions of sawdust, rice husk, and BBJP. The design covers coal fractions of 0.2–0.8.

2.2. Curie-Point Py-GC/FID Analysis

The Py-GC/FID analyses were carried out at the Tonokura Laboratory, the University of Tokyo, Japan. Pyrolysis was performed with a Curie-point pyrolyser (JCI-22, Japan Analytical Industry, Tokyo, Japan) with an F590 pyrofoil, which sets the pyrolysis temperature to 590 °C. About 0.3–0.5 mg of sample was loaded into the ferromagnetic foil. The pyrolysis vapours were transferred through a deactivated stainless-steel transfer line held at 250 °C to a gas chromatograph (7890B, Agilent Technologies, Santa Clara, CA, USA). Separation was carried out on a DB-5 capillary column (30 m × 0.25 mm inner diameter, 0.25 µm film of 5%-phenyl methylpolysiloxane; Agilent J&W) with helium as carrier gas at a constant flow of 1.0 mL/min. The injector was operated in split mode (split ratio 1:50) at 250 °C. The flame ionisation detector was kept at 300 °C with hydrogen, air, and nitrogen make-up flows of 40, 400, and 25 mL/min, and the signal was acquired at 25 Hz (sampling interval 40 ms) for 900 s per run. Runs were repeated where material and instrument time allowed; the replicate structure is given in Section 2.4.1. Compounds were not identified; only the FID signal was recorded.

2.3. Data Export and Processing

Chromatograms were processed in OpenChrom and exported as retention time (ms) and detector signal (the column is named TIC by OpenChrom). The baseline was corrected with the SNIP (statistics-sensitive nonlinear iterative peak-clipping) algorithm with 100 iterations, and the exported files are the baseline-corrected traces. Each chromatogram was exported over the full acquisition time of 900 s (0–899.96 s; 22,500 points at 40 ms in a complete file); several files lack points near the start or contain short gaps, which the common analysis

window (Section 2.4.1) takes into account. The common analysis window, the resampling onto a 40 ms grid, and the area normalisation are described in Section 2.4.1. All calculations were carried out in Python 3.10 (Numpy 2.2.6,· Scipy 1.14.1, Pandas 2.3.3, Scikit-learn 1.7.2, Matplotlib 3.10.9, Seaborn 0.13.2, Pymoo 0.6.2, Numba 0.68.0, Ilvmlite 0.50.0). The notebook and output tables are deposited at https://github.com/mhdsoleh-data/TGA-CFD-MOO (accessed on 10 August 2026) and archived at https://doi.org/10.5281/zenodo.20550862.

2.4. Shape Deviation Index and Additivity Test

2.4.1. Chromatogram Processing

Each Py-GC/FID chromatogram was exported as a detector signal against retention time (OpenChrom). The chromatograms of sawdust, rice husk, BBJP, the coal reference, and the twelve blends were interpolated linearly onto a common retention-time grid with a step of $\Delta t = 40$ ms, covering 15.68–790.00 s (19,359 points), which is the window shared by all first-run files and the coal reference. The window starts at 15.68 s, which is both the first retention time of the coal reference chromatogram and the start of the first run of 80C_20RH; before this time, the coal component of the additive expectation with the reference (Section 2.4.3) is not defined, and not every blend was recorded. The first peak of the coal reference lies at 89.6 s. Peaks that elute earlier, such as those of sawdust at about 7 s and 15 s, therefore lie outside the window (Section 3.1). Negative interpolated values were set to zero. Each profile was normalised to unit area:

$$\begin{gathered} y_k(t) = \frac{s_k(t)}{\int s_k(t)dt}, \\ \int y_k(t)dt = 1 \end{gathered} \tag{1}$$

where $s_k(t)$ is the interpolated signal of sample k. Area normalisation removes differences in absolute detector response and injected mass so that only the distribution of the signal over retention time is compared. It follows that the analysis carries no information on total volatile yield.

The coal reference chromatogram was processed like the other chromatograms and is continuous on the 40 ms grid.

Nine of the twelve blends and the three pure biomass fuels and SRFs were measured twice. The second run of 60C_40BBJP was excluded because its file begins at 128 s and does not cover the analysis window. The ternary blends 60C_13RH_26BBJP and 60C_26SD_13RH were measured once. The nominal fractions of five blends (60C_26RH_13BBJP, 60C_13RH_26BBJP, 60C_26SD_13RH, 60C_13SD_13RH_13BBJP, and 20C_26each) summed to 0.98–0.99 and were renormalised to sum to 1.

2.4.2. Shape Deviation Index

The dissimilarity between two normalised chromatograms $a(t)$ and $b(t)$ is quantified as

$$SDI(a,b) = \int |a(t) - b(t)|dt \approx \Delta t.\sum_{j=1}^{N}\left|a\left(t_j\right) - b\left(t_j\right)\right| \tag{2}$$

Because both profiles have unit area, SDI equals twice the total-variation distance between them and ranges from 0 (identical shapes) to 2 (no overlap). A value of 1 corresponds to 50% shared area. SDI measures shape dissimilarity. It does not identify which compounds differ; it is not a measure of yield, and it is not by itself evidence of a chemical interaction.

2.4.3. Additive Model: Coal Reference and Coal Profile Estimated from the Blends

If a blend j with mass fractions w_j,i of coal (C), sawdust (SD), rice husk (RH), and BBJP behaved additively, its chromatogram would be

$$y_j(t) = w_{j,c}c(t) + \sum_{i\in\{SD,RH,BBJP\}} w_{j,i}y_i(t) \quad (3)$$

where c(t) is the normalised chromatogram of the coal component. Two versions of c(t) were used. In the first, c(t) is the coal reference chromatogram (Section 2.4.1); the SDI of blend j from Equation (3) with this c(t) is denoted SDI_ref,j. In the second, c(t) was estimated at every retention time by least squares over the twelve blends,

$$r_j(t) = y_{j,meas}(t) + \sum_{i\in\{SD,RH,BBJP\}} w_{j,i}y_i(t)$$
$$c(t) = max\left\{0, \frac{\sum_j w_{j,C}r_j(t)}{\sum_j w_{j,c}^2}\right\} \quad (4)$$

and renormalised to unit area. This estimate does not depend on the reference chromatogram. The additivity residual of blend j is then

$$SDI_{add,j} = SDI(y_{j,meas}(t), y_{j,model}(t)) \quad (5)$$

with y_j,model from Equation (3) using the estimated c(t). A leave-one-out version, in which blend j was omitted when c(t) was estimated, was also calculated. SDI_add is an optimistic bound on non-additivity, because c(t) is fitted to minimise the residual (SDI_add is minimised by construction): Any interaction that scales with the coal fraction is absorbed into the estimated coal profile and does not appear in the residual, and c(t) is fitted from the same data. It cannot detect such interactions. To assess how far the reference agrees with the coal in the blends, the SDI between each blend and the reference was compared with the largest value an ideal mixture of the same coal could reach, Σ_{i≠C} w_j,i · SDI(y_i, c_ref), which follows from the convexity of the total-variation distance.

Three SDI variants are therefore used and kept apart throughout: SDI_rep (Equation (2) between two runs of the same sample; the noise level), SDI_add (Equation (5); residual of the additive model with the coal profile estimated from the blends; the only variant used as evidence on additivity, always read against SDI_rep), and SDI_ref (additive model with the external coal reference; a large part of its level is reference mismatch, and it is not evidence of interaction). The surrogate and genetic-algorithm analysis (Sections 2.5–2.7) uses SDI_ref as its target.

2.4.4. Replicate Variation and Its Robustness

The SDI between the two runs of the same sample (SDI_rep, Equation (2)) was calculated for the nine blends and three pure biomass fuels and SRFs with a valid second run. SDI_rep is the smallest dissimilarity the method can resolve with the present sampling and injection repeatability, and it was compared with SDI_add of the same blend. Its sensitivity to the resolution and alignment of the chromatograms was tested by (i) rebinning the profiles to widths of 0.2–5 s and (ii) an optimistic window-wise alignment, in which the best lag (±2 s in 20 s windows or ±5 s in 10 s windows) was chosen separately in every window so that the resulting SDI_rep is a lower bound on what alignment can achieve. Alignment algorithms such as correlation-optimised warping or dynamic time warping were not applied.

2.4.5. Repeatability of the Estimated Coal Profile and of the Residual

The nine blends with a second run were analysed separately for run 1 and run 2: In each run, c(t) was estimated from those nine blends and their pure-fuel profiles (Equation (4)),

and SDI_add was calculated (Equation (5)). The SDI between the two estimated coal profiles, the difference in SDI_add between runs (mean, standard deviation, and maximum), and the Spearman rank correlation between runs were reported. The number of runs per blend needed to detect a difference δ in SDI between two blends was estimated as $n = 2(1.96 + 0.84)^2\ \sigma^2/\delta^2$, with σ^2 denoting the single-run variance estimated from the run-to-run differences in SDI_add (two-sided 5% level, 80% power), assuming equal variance for all blends and normally distributed errors. A 95% range on n was obtained from the chi-square interval of the variance. Because it is based on nine pairs, the estimate is illustrative.

2.4.6. Subset-Based Consistency Test of the Estimated Coal Profile

Under additive mixing, the coal profile estimated from any subset of blends must be the same. Coal was mixed at different fractions (0.2–0.8) and with different biomass partners, so interactions that depend on the partner or that are not proportional to the coal fraction make the estimates from different subsets disagree; this check does not depend on the coal reference. Equation (4) was applied separately to (i) the binary blends grouped by biomass partner (a 0.8 and a 0.6 coal blend with sawdust, with rice husk and BBJP) and (ii) the 0.8 and the 0.6 coal binary blends. The SDI between these estimates was compared with (a) the SDI between the coal profiles estimated from run 1 and run 2 of the same single blend and (b) 2000 random regroupings of the same six binary blends into three pairs, for which their mean pairwise SDI gives the spread expected when the grouping is unrelated to the biomass partner.

2.5. Surrogate Model of the SDI

The SDI of each blend against the additive expectation with the coal reference (Section 2.4.3) was modelled as a function of the four mass fractions. Every run was one observation: The nine blends with a valid second run contributed two, the other three contributed one, giving 21 observations from twelve blends (the second run of 60C_40BBJP was excluded; Section 2.4.1). The prediction target for validation was the mean SDI of the available runs of a blend. The surrogate was a weighted combination of four regressors: the mean of the training data, a linear Scheffe mixture model with ridge penalty (penalty chosen by leave-one-out), and two Gaussian processes (constant × ARD squared-exponential kernel, and constant × isotropic Matern-5/2 kernel, each plus a white-noise term; targets normalised; fixed jitter 10^{-6}; hyperparameters optimised with ten restarts). The non-negative weights (summing to one) were obtained by minimising the leave-one-out squared error, and the prediction standard deviation was multiplied by a scale factor so that the leave-one-out 90% intervals reached nominal coverage. Pure fuels were not used as SDI = 0 training points, because in earlier variants, they lowered the leave-one-out skill.

Validation used leave-one-out cross-validation grouped by blend (all runs of a blend held out together) with the mean SDI of the remaining blends as the baseline predictor. Because the model weights and the scale factor were chosen on the same folds, the whole procedure was repeated inside every outer fold (weights and scale factor estimated only from the other eleven blends; nested cross-validation). Skill is reported as $1 - \text{RMSE_model}/\text{RMSE_baseline}$. The run-to-run standard deviation of SDI was estimated from the nine replicate pairs and taken as the noise level of the target.

As an alternative to the fixed stack, the hyperparameters of a single Gaussian process were optimised. The search space comprised the kernel (ARD or isotropic squared-exponential, Matern-3/2 and-5/2, rational quadratic, or squared-exponential plus linear), the bounds of the length scales and of the amplitude, the treatment of noise (white noise learned within bounds, or fixed at the replicate-based run-to-run variance of the training blends multiplied by 0.5, 1, or 2), and the composition transform (mass fractions, their

square root, or the isometric log-ratio of the fractions shifted by 0.02). Configurations were scored by leave-one-out RMSE grouped by blend and searched by random search (48 configurations) and Bayesian optimisation (tree-structured Parzen estimator, 40 trials). To include the cost of the search in the reported accuracy, the search was repeated inside every outer leave-one-out fold using only the other eleven blends, and the held-out blend was predicted with the configuration chosen there (nested cross-validation); the 90% interval of the nested skill was obtained by bootstrapping the folds. As controls, the whole procedure was repeated 20 times with the compositions shuffled among the blends and with two further random-search seeds. The final configuration was searched on all twelve blends, its prediction standard deviation was calibrated on leave-one-out residuals, and the tuned model replaced the stack in the genetic algorithm and the bootstrap (Sections 2.6 and 2.7).

2.6. Genetic Algorithm

A genetic algorithm maximised a lower confidence bound of the predicted SDI (mean minus κ prediction standard deviations, $\kappa = 1$) within seven composition categories (three binary, three ternary, and the quaternary blend), with a minimum component fraction of 0.005 and mass fractions summing to one. Compositions farther than a trust radius from every measured blend (Euclidean distance in mass-fraction space; radius of 0.190, the median distance between neighbouring measured blends) were penalised by ten standard deviations of the target per radius so that the search could not report optima far outside the measured compositions. The following parameters were used: population 200, 300 generations, tournament selection (size 3), BLX-α crossover ($\alpha = 0.5$, probability 0.75), Gaussian mutation (probability 0.20, $\sigma = 0.05$), elitism, five restarts. The sensitivity of the optima to κ (0, 1, 2) and to the trust radius (0.5, 1, and 1.5 times the reference radius and unlimited) was tabulated, and the effect of the tuning grid (64 combinations of population, generations, crossover, mutation, and tournament size) and of the minimum fraction was checked.

2.7. Bootstrap

The stability of the optimum was assessed with N = 500 bootstrap replicates. In each replicate, the twelve blends were resampled with replacement (with all runs of a resampled blend), the four components of the surrogate were refitted (weights and scale factor kept at their full-data values; two restarts), and the genetic algorithm (population 100, 100 generations, one restart, trust region defined by the resampled blends) was run for each category. As a control, 20 repetitions with different GA random seeds were run on the full-data surrogate. The standard deviation of the optimal SDI, the median and percentiles of the optimal coal fraction, the share of replicates with an optimal coal fraction of 0.75–0.85, and the share of replicates in which each category had the highest predicted SDI were reported.

2.8. Reporting

SDI values are reported to two decimals. Software and repository are given in Section 2.3.

3. Results

3.1. Feedstock Volatile Fingerprints

Figure 1 shows the chromatograms of sawdust, rice husk and BBJP, and the coal reference chromatogram. Each is dominated by a few peaks. Within the analysis window (15.68–790.00 s, Section 2.4.1), the sawdust chromatogram has nine peaks with a prominence of at least 2% of its maximum. The largest is at 30.8 s, followed by peaks at 156.9 s (34% of the largest), 54.0 s, and 417.2 s (8% each). Sawdust has larger peaks before the window at about 7 s and 15 s, which contain 51% of the area of its exported chromatogram and

are not part of the SDI analysis. Rice husk has six peaks in the window, the largest at 76.6 s, followed by 153.8 s (40%), 104.3 s (16%), and 438.6 s (14%); BBJP has five, the largest at 91.4 s, followed by 167.4 s (46%), 118.8 s (19%), and 449.8 s (15%). The coal reference chromatogram has six peaks. One peak at 155.8 s dominates, and there are smaller peaks at 235.7 s (10% of the largest), 165.8 s (9%), and 89.6 s (8%), with the last being the earliest peak of the coal reference.

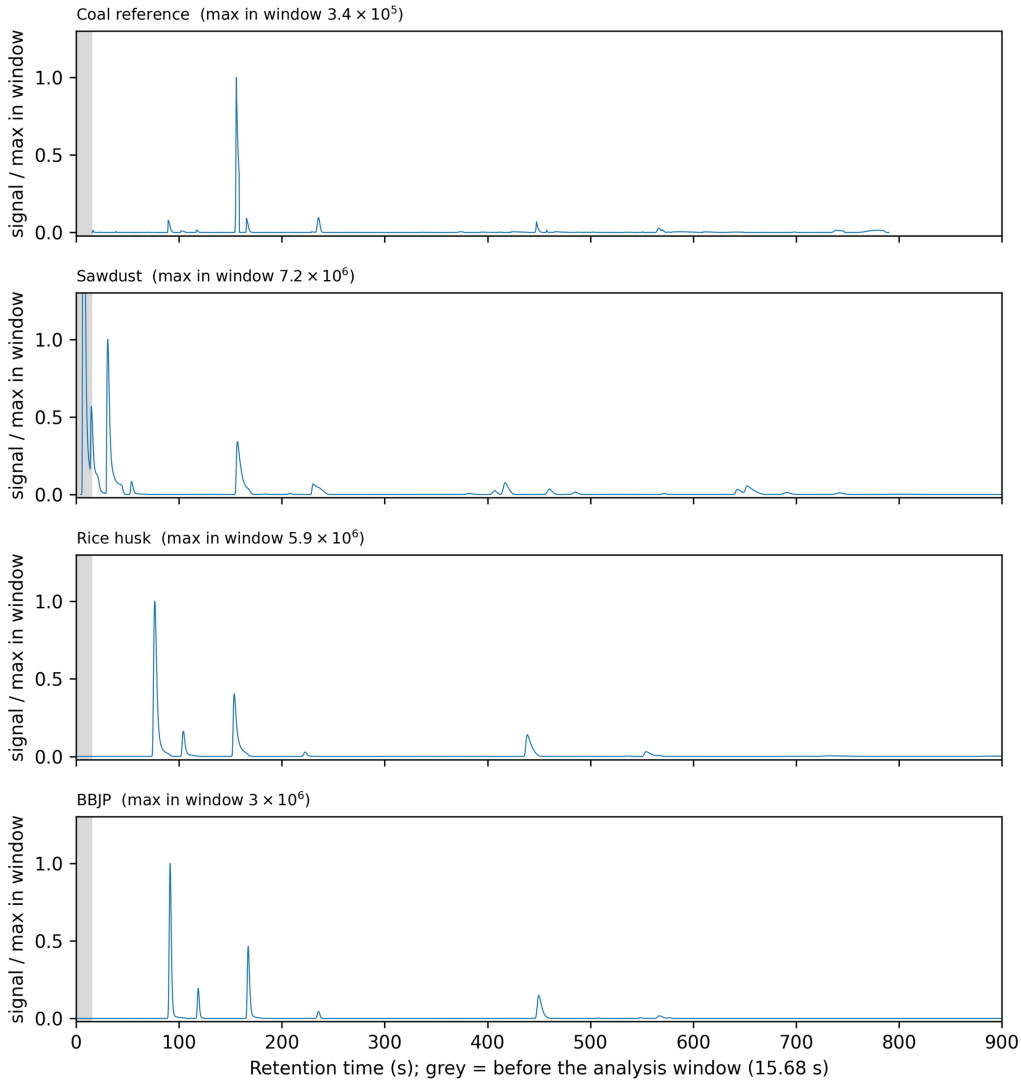


Figure 1. Chromatograms of sawdust, rice husk, and BBJP and the coal reference chromatogram, each scaled to its maximum inside the analysis window. Grey: retention times before the window (below 15.68 s); the sawdust peaks at about 7 s and 15 s exceed the scale.

The share of the area inside the window found before 100 s was 0.57 for sawdust, 0.51 for rice husk, 0.47 for BBJP, and 0.05 for the coal reference; between 100 and 300 s, it was 0.25, 0.33, 0.36, and 0.67, and after 300 s, it was 0.18, 0.16, 0.17, and 0.27. The main peaks of rice husk (76.6 s) and BBJP (91.4 s) do not coincide. The SDI between the fuels, from 1.78 (sawdust and rice husk) to 1.95 (rice husk and BBJP), shows that the chromatograms of the fuels overlap little in shape (Section 3.3). Peaks of sawdust, rice husk, and BBJP lie at similar retention times near 150–170 s (156.9, 153.8, and 167.4 s) and 440–450 s (438.6 and 449.8 s), so peaks at these positions in a blend chromatogram cannot be attributed to a single fuel. The absolute signal levels differ between samples and files; because each profile

is normalised to unit area, only its shape enters the analysis. No peak was assigned to a compound because only FID data were recorded (Section 4.1).

3.2. Repeatability of Duplicate Runs

Nine of the twelve blends and the three pure biomass fuels and SRFs have a valid second run (Section 2.4.1). The SDI between the two runs of the same blend (SDI_rep) had a median of 0.89 (range 0.68–1.71, n = 9; Table 1; Figure 2). For the pure biomass fuel and SRF, it was 1.21 (sawdust), 1.73 (rice husk), and 1.89 (BBJP) (Figure S1). An SDI of 0.89 corresponds to about 55% shared normalised area between the two runs of a blend, and SDI values of 1.2–1.9 correspond to 5–40% for the pure fuels.

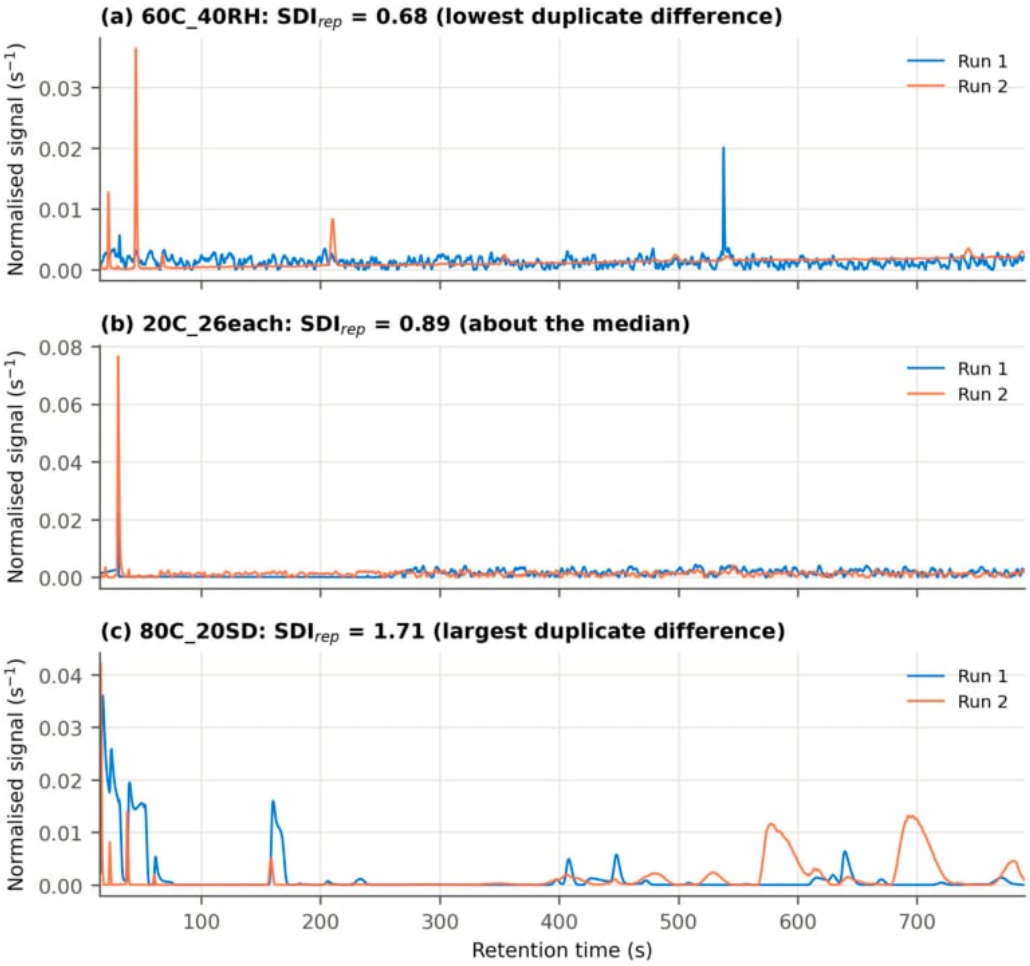


Figure 2. Area-normalised Py-GC/FID chromatograms of duplicate runs of three blends: (**a**) 60C_40RH, the smallest difference between runs (SDI_rep = 0.68); (**b**) 20C_26each, close to the median (0.89); (**c**) 80C_20SD, the largest difference (1.71). Analysis window 15.68–790 s; gaps in the exported files are bridged by linear interpolation, as in the analysis.

The result did not depend on the time resolution or on a simple retention-time correction (Table 1, Figure S2). Rebinning the chromatograms from 0.04 s to 5 s lowered the median SDI_rep only from 0.89 to 0.73. An optimistic alignment, with the best lag chosen separately in every window, lowered it to 0.80 (20 s windows, ±2 s) and 0.68 (10 s windows, ±5 s), which still corresponds to only about 66% shared area. The variation between duplicate runs is therefore not fine-scale peak jitter or a retention-time offset of a few seconds.

Table 1. Additivity residual and duplicate-run variation of the twelve blends.

Blend	Type	Coal Fraction	SDI_add (All Blends)	SDI_add (Leave-One-Out)	SDI_rep	SDI_add − SDI_rep
80C_20SD	binary	0.80	1.28	1.40	1.71	−0.44
80C_20RH	binary	0.80	1.26	1.40	1.38	−0.11
80C_20BBJP	binary	0.80	1.06	1.17	1.46	−0.40
60C_26RH_13BBJP	ternary	0.61	0.99	1.03	0.79	+0.21
60C_13RH_26BBJP	ternary	0.61	0.89	0.93	not run	–
60C_26SD_13RH	ternary	0.61	0.83	0.87	not run	–
60C_13SD_13RH_13BBJP	quaternary	0.61	0.90	0.94	0.85	+0.05
60C_40SD	binary	0.60	1.06	1.12	1.43	−0.37
60C_40RH	binary	0.60	0.85	0.88	0.68	+0.17
60C_40BBJP	binary	0.60	1.00	1.04	excluded (file starts at 128 s)	–
40C_20each	quaternary	0.40	0.98	0.99	0.76	+0.22
20C_26each	quaternary	0.20	1.38	1.39	0.89	+0.50

3.3. Coal Reference and Estimated Coal Profile

The coal reference chromatogram is continuous on the 40 ms grid (19,359 points, 15.68–790.00 s) and differed from the chromatograms of sawdust, rice husk, and BBJP by SDI 1.63, 1.73, and 1.67, similarly to the differences between the biomass and SRF themselves (1.78–1.95), so the four fuel chromatograms overlap little in shape.

The coal chromatogram estimated from the twelve blends (Equation (4)) differed from the reference by SDI 1.57 (Figure 3c) and from the chromatograms of sawdust, rice husk, and BBJP by SDI 1.64, 1.89, and 1.93. The coal profile estimated separately from run 1 and from run 2 of the nine replicated blends differed by SDI 0.68, less than the replicate SDI of the blends but far from identical.

The reference and the coal profile estimated from the blends lead to different levels of SDI. With the reference chromatogram as the coal component, the twelve blends had an SDI of 1.43–1.68 (median 1.59) from their additive expectation, and the SDI exceeded the blend's own SDI_rep in eight of nine blends (median of 1.60 against 0.89 for the nine blends with a replicate). With the coal profile estimated from the blends, the median fell to 1.00 (Section 3.4). The distance of the blends from the reference coal itself was 1.62–1.70 for the three 0.8-coal blends, whereas a blend composed additively of the same coal could lie at most 0.40 from it (excess 1.22–1.30); for the 0.6-coal blends, it was 1.55–1.66 against a bound of about 0.80 (excess 0.76–0.86), for 0.4 coal, it was 1.61 against 1.20, and for 0.2 coal, it was 1.61 against 1.59 (Figure S4). For 80C_20BBJP, for example, the SDI from the additive expectation was 1.64 with the reference and 1.06 with the estimated coal profile (Figure 3a,b). Part of the SDI level obtained with the reference therefore reflects the difference between the reference and the coal in the blends, particularly for the blends with a high coal fraction, and only differences between blends, or SDI relative to SDI_rep, are interpreted below. The choice of the coal component can dominate an additivity test of this kind.

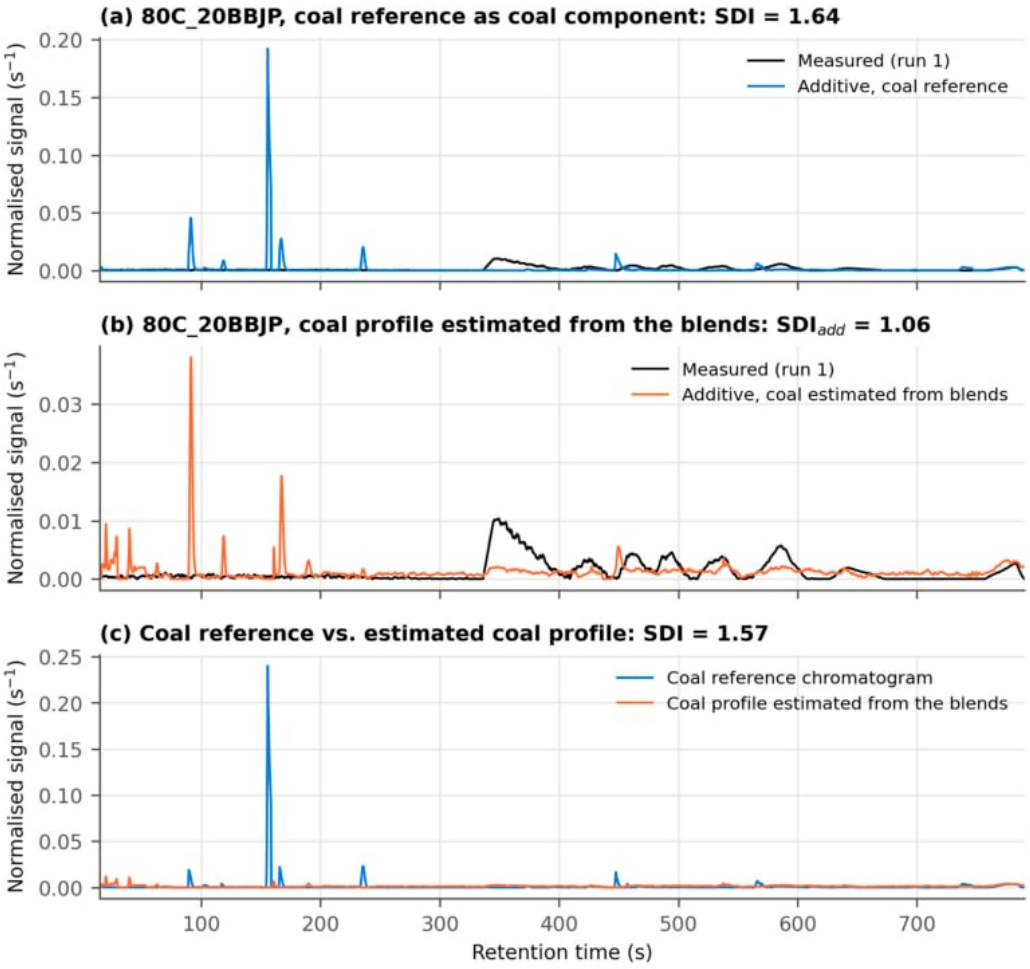


Figure 3. Measured chromatogram of 80C_20BBJP (run 1) against its additive expectation (Equation (3)): (**a**) with the coal reference chromatogram as the coal component (SDI = 1.64); (**b**) with the coal profile estimated from the twelve blends (Equation (4); SDI_add = 1.06); (**c**) the coal reference chromatogram and the estimated coal profile (SDI = 1.57).

3.4. Additivity Residual Compared with Replicate Variation

With the coal profile estimated from the blends, the additivity residual SDI_add of the twelve blends had a median of 1.00 (range 0.83–1.38); the leave-one-out version had a median of 1.04 (Table 2). By coal fraction, SDI_add ranged from 1.06 to 1.28 at 80 wt.% coal and 0.83 to 1.06 at about 60 wt.%, and it was 0.98 at 40 wt.% and 1.38 at 20 wt.%. The 20 wt.% coal blend has the highest residual because coal contributes little to it, so its residual mainly reflects the variation in the three biomass and SRF chromatograms.

For the nine blends with a second run, SDI_add had a median of 1.00 compared with a median SDI_rep of 0.89 for the same blends. SDI_add exceeded the blend's own SDI_rep in five of nine blends (median difference +0.05; range −0.44 to +0.50; Figure 4). It was lower than SDI_rep for the four blends with the highest replicate variation (80C_20SD, 80C_20BBJP, 60C_40SD, 80C_20RH) and higher by 0.17–0.50 for the blends at 20–60 wt.% coal with the smallest replicate variation (60C_26RH_13BBJP, 60C_40RH, 40C_20each, 20C_26each). The deviation of the blends from the additive model is therefore of the same size as the variation between duplicate runs of the same blends. This is a result of limited statistical power, not a demonstration of additivity: With duplicate runs, no departure from additivity exceeding run-to-run variation was detected, and interactions of magnitude below about 0.3 in SDI cannot be ruled out (the run-to-run SD of SDI_add is 0.18, and about 24 runs per blend would be needed to resolve a difference of 0.1; Section 4.5). SDI_add is

also an optimistic bound (Section 2.4.3), so the result does not rule out interactions that scale with the coal fraction.

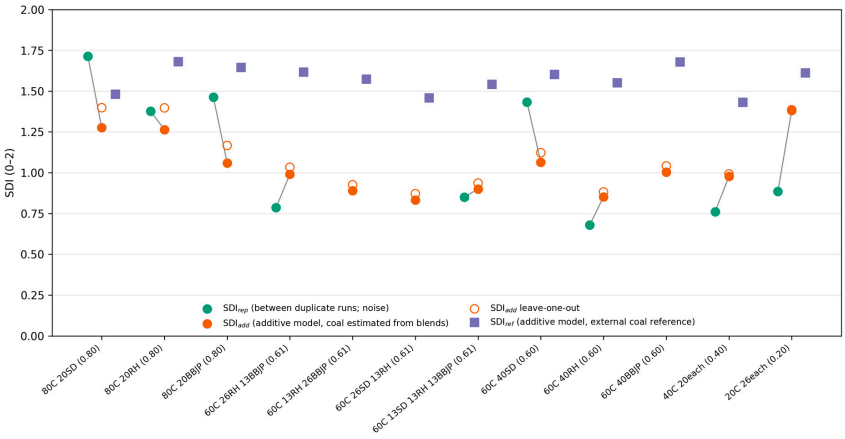


Figure 4. The three SDI variants for the twelve blends, ordered by coal fraction (in brackets); *y*-axis: SDI (0–2). SDI_rep: SDI between the two runs of the same blend (noise level). SDI_add: residual from the additive model with the coal profile estimated from all blends (filled) and leave-one-out (open). SDI_ref: deviation from the additive model with the external coal reference chromatogram. Lines join SDI_rep and SDI_add of the same blend; blends without a valid second run have no SDI_rep.

Table 2. Median SDI between duplicate runs of nine blends under rebinning and optimistic alignment.

Processing	Median SDI_rep
Native grid (0.04 s)	0.89
Rebinned to 0.2 s/0.5 s/1 s	0.88/0.88/0.87
Rebinned to 2 s/5 s	0.84/0.73
Best lag per 20 s window, ±2 s	0.80
Best lag per 10 s window, ±5 s	0.68

3.5. Repeatability of the Residual and Resolvable Differences

Repeating the additivity analysis separately with the run-1 and run-2 chromatograms of the nine replicated blends gave the same median SDI_add in both runs (about 1.06), but the individual values differed between runs by a mean of +0.09 with a standard deviation of 0.18 and a maximum of 0.48 (80C_20RH, 1.23 in run 1 and 0.75 in run 2). The rank agreement between runs was weak (Spearman 0.47, $p = 0.21$, n = 9; Figure S3). Blends therefore cannot be ranked by SDI_add with the present data.

The run-to-run standard deviation corresponds to a single-run standard deviation of about 0.12. Under the assumptions in Section 2.4.5, about 24 runs per blend would be needed to detect a difference of 0.1 in SDI between two blends (95% range 11–89), about six runs for 0.2 (3–22), and about three runs for 0.3 (1–10; Figure S9). Duplicate runs, as used here, can therefore resolve only differences well above 0.3, and the estimates are illustrative because they rest on nine pairs.

3.6. Consistency of the Estimated Coal Profile Across Blend Subsets

The coal profiles estimated from the binary blends of each biomass partner (a 0.8 and a 0.6 coal blend each) differed from each other by SDI 1.32 (sawdust against rice husk), 1.32 (sawdust against BBJP), and 0.95 (rice husk against BBJP), and the profiles estimated from the 0.8 and the 0.6 coal binary blends differed by SDI 1.06 (Figure S5). These differences are of the same size as the noise in a single-blend estimate: The coal profile estimated from run 1 and from run 2 of the same blend differed by SDI 1.47 (median of five blends;

range 0.72–1.81), and estimates from any two of the twelve single blends differed by a median of 1.12 (66 pairs). Grouping the six binary blends by biomass partner did not make the estimates more different from each other than random regrouping did: The mean pairwise SDI between partner groups was 1.20 against a median of 1.18 (5th–95th percentile 1.13–1.24) for random regroupings, and 72% of the random regroupings gave a spread at or below the observed one. The coal profile estimated from the blends therefore does not depend on the biomass partner or on the coal fraction beyond the noise, which is consistent with additive mixing, but given the size of the noise, it also has little power to reject it.

3.7. Gaussian-Process Surrogate of the SDI

3.7.1. Data and Noise Level

The surrogate was trained on the twelve blends (six binary, three ternary, and three quaternary). Nine blends have a valid second run, so 21 observations (one per run) entered the training set; the run of 60C_40BBJP that begins at 128 s was excluded (Section 2.4.1). The SDI of a blend, computed against the additive reference (Section 2.4.2), ranged from 1.447 (40C_20each) to 1.685 (80C_20BBJP) when averaged over the available runs (1.431–1.680 in run 1). The standard deviation of a single SDI between two runs of the same blend was 0.063 (nine pairs), that of a two-run mean was 0.045, and the standard deviation of the blend means was 0.076. Removing the run-to-run variance leaves a between-blend standard deviation of about 0.061. All model comparisons below are made against this level.

3.7.2. Model and Validation

The SDI was modelled by a weighted combination of four regressors: a constant (mean of the training data), a linear Scheffe mixture model with ridge penalty, and two Gaussian processes (ARD squared-exponential kernel and isotropic Matern-5/2 kernel, each with a white-noise term). The non-negative weights (summing to one) were obtained by minimising the leave-one-out squared error; they were 0.70 (ARD-GP), 0.29 (Matern-GP), 0.01 (constant), and 0 (linear model). The prediction standard deviation was rescaled so that leave-one-out 90% intervals reached their nominal coverage; the scale factor was 1.01. No pure-fuel SDI = 0 anchors were used: In an earlier variant, they lowered the leave-one-out skill in every test (Section 3.7.4).

Validation used leave-one-out cross-validation grouped by blend (all runs of a blend held out together), with the predictor "mean SDI of the remaining blends" as baseline (RMSE 0.079). Table 3 and Figure 5a give the results. When the stack weights and the scale factor were fitted on the same leave-one-out predictions, the stack reached an RMSE of 0.065 (skill 0.17 relative to the baseline). That figure includes the selection of the model on the data on which it is scored. Repeating the whole procedure inside every outer fold (weights and scale factor estimated only from the eleven remaining blends; nested cross-validation) gave an RMSE of 0.082, larger than the baseline (skill −0.04, $R^2 = -0.29$); 83% and 92% of the held-out blend means fell inside the 90% and 95% intervals. The surrogate therefore showed no predictive skill on an unseen blend. The two blends with the highest measured SDI were underpredicted by 0.12–0.13 (60C_40BBJP: 1.552 predicted against 1.678; 80C_20BBJP: 1.561 against 1.685), and the model returned predictions between 1.47 and 1.63 for measured values between 1.45 and 1.69; that is, it regressed towards the mean.

Table 3. Leave-one-out cross-validation of the SDI surrogate (12 blends, grouped by blend; baseline RMSE 0.079).

Predictor	RMSE	R^2	Skill vs. Baseline	95% Coverage
Constant (training mean)	0.079	−0.18	0.01	1.00
Scheffé linear model	0.092	−0.60	−0.16	0.92
GP, ARD squared-exponential	0.066	0.17	0.16	0.83
GP, Matern-5/2	0.071	0.04	0.10	1.00
Stack, weights fitted on the same folds	0.065	0.19	0.17	0.92
Stack, nested (used for the conclusions)	0.082	−0.29	−0.04	0.92

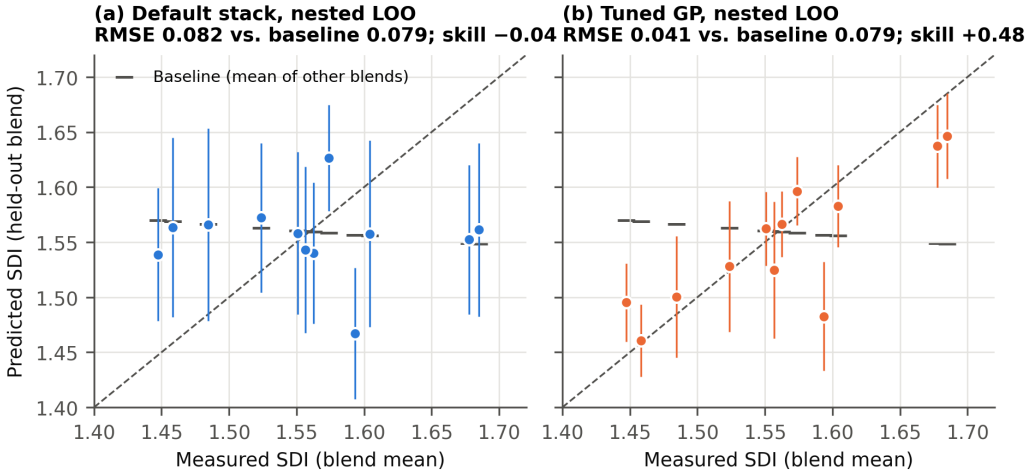


Figure 5. Leave-one-out predictions of the blend-mean SDI (held-out blend, ± prediction SD) against the measured value: (**a**) default stacked surrogate with nested selection of the weights (skill −0.04); (**b**) Gaussian process with hyperparameters selected by nested Bayesian search (skill +0.48; Section 3.7.5). Dashes: baseline prediction (mean of the other blends); dashed line: 1:1.

3.7.3. Shape of the Fitted Response

Along a coal-BBJP line, the fitted SDI rose from 1.57 at 0.2 coal to 1.60, 1.66, 1.67, and 1.68 at 0.4, 0.6, 0.7, and 0.8 coal and decreased to 1.66 at 0.9 coal (prediction standard deviation 0.04–0.07). Along the coal-rice husk line, it was almost flat between 0.2 and 0.7 (1.55–1.57) and reached 1.60 at 0.8 and 1.63 at 0.9. Along the coal–sawdust line, it decreased from 1.56 at 0.2–0.4 to 1.51–1.53 at 0.7–0.8 and rose to 1.60 at 0.9 (Figure 6). The differences between the three lines (0.05–0.15) are of the order of the run-to-run standard deviation (0.063) and of the prediction standard deviation, so the surrogate provides at most a weak indication that BBJP raises the SDI more than rice husk or sawdust at 0.6–0.8 coal.

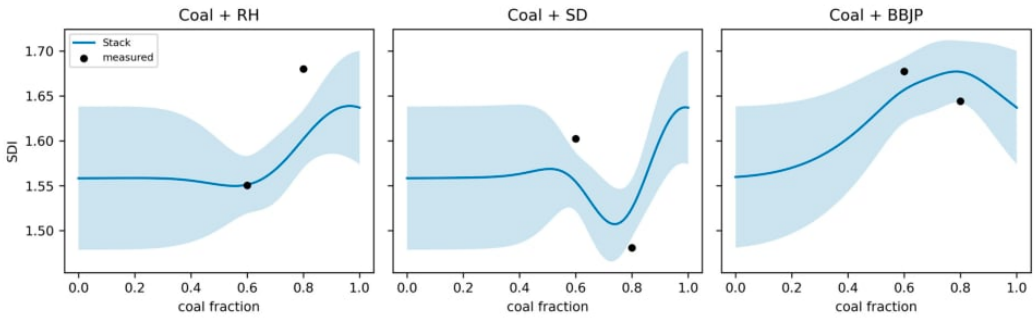


Figure 6. SDI predicted by the surrogate (stack) along the coal + rice husk, coal + sawdust, and coal + BBJP lines, with the prediction interval (shaded) and the measured binary blends (points, run 1).

3.7.4. Effect of the SDI = 0 Anchors and Interpretation

Pure fuels have SDI = 0 by definition and were tried as training points at the corners of the composition simplex. With the anchors, the leave-one-out RMSE of the best single model rose from 0.14 to 0.29–0.41 (skill −0.86 to −1.68) in a preliminary analysis with an earlier version of the coal reference chromatogram, and the anchors were therefore not used in the analysis reported here. The blend chromatograms were 1.62–1.70 (80 wt.% coal blends) and 1.55–1.66 (60 wt.% blends) from this reference, whereas an ideal mixture of the same coal could not exceed 0.40 and about 0.80, respectively (excess 0.76–1.30). Part of the SDI level (1.45–1.69) therefore reflects a difference between the reference coal and the coal in the blends, and only the *differences* between blends carry information about interaction. These differences are small compared with the noise, which is why the surrogate has no skill.

3.7.5. Alternative Surrogate: Nested Hyperparameter Optimisation

The result of Section 3.7.2 refers to a fixed model family. As an alternative, the hyperparameters of a single Gaussian process were searched (Section 2.5), with the search repeated inside every outer fold so that the reported accuracy includes the cost of the search. The nested RMSE was 0.041 for the random search and 0.041 for the Bayesian search against 0.079 for the mean of the other blends—that is, a skill of +0.48 for both (Figure 5b), with a fold-bootstrap 90% interval of +0.11 to +0.71 (Bayesian search). For comparison, a single Gaussian process in its default configuration (ARD kernel, learned noise, mass fractions as input) had a skill of +0.02, and the stack in Section 3.7.2 had a skill of −0.04. The configuration chosen inside the folds was stable: the isometric log-ratio input in all twelve folds, an ARD kernel of squared-exponential or Matern type in eleven (squared-exponential 5, Matern-3/2 3, Matern-5/2 3), and a squared-exponential-plus-linear kernel in one, with learned white noise in nine. Of 48 fixed configurations drawn at random (Figure S8), 45 had a lower nested RMSE than the baseline (median skill +0.12; best +0.48, which cannot be selected without hindsight); the median skill was +0.25 for ARD kernels and +0.11 for isotropic kernels, and it was +0.04 for mass fractions as input against +0.16 for their square root and +0.12 for the log-ratio transform. The final configuration, searched on all twelve blends, was a squared-exponential kernel with a linear term on the log-ratio input with learned noise; its leave-one-out RMSE was 0.041.

Three checks were made. With the compositions shuffled among the blends (20 permutations so that no composition-SDI relation remained), the same nested procedure gave a mean skill of −0.12 (5th-95th percentile −0.42 to +0.11; maximum +0.20), and none of the permutations reached the observed value (+0.42 in the corresponding run). Two further random-search seeds gave +0.44 and +0.48. The nested RMSE (0.041) is close to the standard deviation expected for a two-run blend mean (0.045, from the nine replicate pairs), so little further improvement is possible with the present measurements, and the wide interval of the skill reflects that it rests on twelve blends. The search space and the 0.02 shift of the log-ratio transform were defined by the authors and are not part of the nested selection.

3.8. Genetic-Algorithm Search, Bootstrap and Candidate Blends

3.8.1. Search on the Surrogate

A genetic algorithm (population 200, 300 generations, elitist, tournament selection, BLX-alpha crossover, five restarts; Section 2.6) maximised a lower confidence bound of the predicted SDI (mean minus one prediction standard deviation) within seven composition categories (three binary, three ternary, and the quaternary blend), with a minimum component fraction of 0.005. The search was limited to a trust region around the measured blends: Compositions farther than 0.190 (the median distance between neighbouring measured

blends, Euclidean distance in mass-fraction space) were penalised. All seven optima were inside this region, and none exceeded the highest measured SDI (1.685; Table 4, Figure S7).

Table 4. GA optima on the SDI surrogate (lower-confidence-bound fitness, trust radius of 0.190).

Category	Coal	Second Component	Third/Fourth	Predicted SDI ± SD	Nearest Measured Blend (Distance)
Coal + sawdust	0.934	SD 0.066	-	1.624 ± 0.058	80C_20SD (0.190)
Coal + rice husk	0.902	RH 0.098	-	1.633 ± 0.047	80C_20RH (0.144)
Coal + BBJP	0.789	BBJP 0.211	-	1.677 ± 0.035	80C_20BBJP (0.016)
Coal + SD + RH	0.902	RH 0.094	SD 0.005	1.633 ± 0.048	80C_20RH (0.147)
Coal + SD + BBJP	0.786	BBJP 0.210	SD 0.004	1.676 ± 0.035	80C_20BBJP (0.017)
Coal + RH + BBJP	0.788	BBJP 0.208	RH 0.004	1.676 ± 0.034	80C_20BBJP (0.016)
Quaternary	0.785	BBJP 0.207	RH 0.004, SD 0.004	1.675 ± 0.034	80C_20BBJP (0.017)

The optimum was near 0.79 coal for every category that contained BBJP (0.785–0.789), with BBJP at 0.207–0.211 (predicted SDI 1.675–1.677, standard deviation 0.034–0.035) and at 0.016–0.017 from the measured blend 80C_20BBJP. For coal + rice husk and coal + sawdust, the optima lay at higher coal fractions (0.90 and 0.93) with predicted SDI 1.633 and 1.624. In the ternary and quaternary categories, the additional components fell to about 0.005 (pseudo-binary compositions). This does not show that a third or fourth component has no effect: the surrogate has no skill and contains no interaction term that the data could support. The optimiser settings were immaterial: the mean predicted optimum of 64 combinations of population, generations, crossover, mutation, and tournament size spread over 0.00015, removal of the minimum fraction changed the fitness by at most 0.002, and 20 repetitions with different random seeds gave standard deviations of at most 0.0001.

3.8.2. Bootstrap of the Optimum

The stability of the optimum was assessed with a bootstrap of N = 500 replicates. In each replicate, the twelve blends were resampled with replacement (with all runs of a resampled blend), the surrogate was refitted, and the genetic algorithm (population 100, 100 generations, one restart) was run on the refitted surrogate (Section 2.7). Keeping the GA random seed as the only source of variation gave a standard deviation of the optimum of at most 0.0001, whereas resampling the data gave 0.03–0.07 (Table 5), so the variation of the optimum was caused by the data and not by the search. No bootstrap optimum exceeded the highest measured SDI (0 of 3500 category results).

Table 5. Bootstrap (N = 500, surrogate refitted in every replicate) of the GA optimum.

Category	Optimal SDI, Mean ± SD	5th–95th Percentile	Coal Fraction, Median (5th–95th)	Share Within 0.75–0.85 Coal	Highest Category
Coal + sawdust	1.519 ± 0.063	1.459–1.573	0.61 (0.58–0.93)	9%	0%
Coal + rice husk	1.552 ± 0.053	1.504–1.601	0.80 (0.60–0.93)	44%	7.0%
Coal + BBJP	1.622 ± 0.069	1.503–1.651	0.80 (0.60–0.82)	66%	88.4%
Coal + SD + RH	1.557 ± 0.027	1.513–1.599	0.80 (0.59–0.93)	42%	0%
Coal + SD + BBJP	1.623 ± 0.034	1.527–1.649	0.79 (0.59–0.81)	63%	0%
Coal + RH + BBJP	1.627 ± 0.027	1.561–1.649	0.80 (0.60–0.81)	69%	3.2%
Quaternary	1.615 ± 0.032	1.556–1.647	0.78 (0.20–0.81)	57%	1.4%

The median optimal coal fraction was 0.78–0.80 in six of the seven categories (0.61 for coal + sawdust); the 5th-95th percentile ranged from 0.58 to 0.93 (0.20–0.81 for the quaternary category; Figure S6). The share of replicates with an optimal coal fraction

between 0.75 and 0.85 was 63–69% for coal + BBJP and the two ternary categories that contain BBJP, 57% for the quaternary blend, 42–44% for coal + rice husk and coal + sawdust + rice husk, and 9% for coal + sawdust. The category with the highest predicted SDI in each replicate was coal + BBJP in 88.4% of the replicates, coal + rice husk in 7.0%, coal + rice husk + BBJP in 3.2%, and the quaternary blend in 1.4%; the three categories that contain sawdust without the quaternary blend (coal + sawdust, coal + sawdust + rice husk, coal + sawdust + BBJP) were never highest. The optimum near 0.8 coal and the higher SDI of the BBJP-containing categories were therefore repeated in most replicates, but the spread of the optimal coal fraction (standard deviation 0.09–0.20) does not allow an optimum to be stated to better than about ±0.1 in coal fraction.

3.8.3. Sensitivity to the Confidence Bound and the Trust Radius

The optimum of the categories containing BBJP was insensitive to the choices made (coal fraction of 0.78–0.79 for confidence-bound weights of 0, 1 and 2 and for trust radii from 0.5 to unlimited times the reference radius). The coal + rice husk and coal + sawdust optima depended on them: 0.85–0.97 and 0.87–0.99 for coal, respectively. In no setting did a predicted optimum exceed the highest measured SDI (Supplementary Table S2).

3.8.4. Blends Proposed in the Submitted Version

The submitted version recommended GA-Opt-2 (80C_12RH_8BBJP, predicted SDI 1.745) and GA-Opt-3 (75C_15RH_10SD, 1.692). Those predictions were higher than the highest value predicted at the training blends (1.643); that is, they were extrapolations of an overfitted model. Scored with the present surrogate, GA-Opt-2 had a predicted SDI of 1.631 ± 0.037 and GA-Opt-3 of 1.545 ± 0.041 (Table 6); both lie inside the trust region (0.113 and 0.122 from 80C_20RH). The predicted values are 0.11–0.15 lower than in the submitted version and below that of the measured blend 80C_20BBJP (predicted 1.677 ± 0.035; measured 1.685 as the mean of two runs, 1.644 and 1.725). The present analysis does not support GA-Opt-2 and GA-Opt-3 as priority candidates.

Table 6. Blends of the submitted version scored by the present surrogate.

Blend	SDI, Submitted Version	SDI, Present Surrogate (± SD)	Lower Bound	Nearest Measured Blend (Distance)
GA-Opt-2, 80C_12RH_8BBJP	1.745	1.631 ± 0.037	1.594	80C_20RH (0.113)
GA-Opt-3, 75C_15RH_10SD	1.692	1.545 ± 0.041	1.504	80C_20RH (0.122)
80C_20RH (measured, mean of two runs 1.604)	-	1.602 ± 0.032	1.570	0.000
80C_20BBJP (measured, mean of two runs 1.685)	-	1.677 ± 0.035	1.642	0.000

3.8.5. Interpretation

The surrogate was trained on SDI_ref, the deviation of each blend from the additive expectation formed with the coal reference chromatogram (Section 2.5). Section 3.3 showed that a large part of this quantity is reference mismatch and not interaction (the 0.8-coal blends were 1.62–1.70 from the reference, although an additive blend of the same coal could lie at most 0.40 from it) and that SDI_ref against the reference grows with the coal fraction for that reason. The search and its bootstrap therefore describe the shape space of the measured data, namely at which coal fraction and with which partner the chromatogram lies farthest from the additive expectation built on the reference coal. They do not identify a blend with a chemical synergy, and a higher SDI_ref is not treated as desirable. The search and its bootstrap point to a coal fraction of about 0.8, and BBJP as partner gives the highest predicted SDI_ref (88% of the bootstrap replicates); this is a description of where the measured SDI_ref

values are largest, and the data do not show why. Because the coal profile estimated from the blends (SDI_add) is minimised by construction, it cannot be used to confirm or refute this ordering either. The support for the description is limited: the default surrogate had no cross-validated skill (nested RMSE of 0.082 against 0.079 for the mean; the tuned alternative of Section 3.7.5 had a skill of +0.48 with a 90% interval of +0.11 to +0.71), the SDI_ref values of the categories (1.52–1.68) differ by about the run-to-run standard deviation (0.063), and the optimum falls on or beside measured blends (0.016–0.017 for the BBJP categories), so the search largely returns the location of the best measured compositions. The result is not a blend recommendation. Replicated measurements (Section 4.5) of 80C_20BBJP and 80C_20RH, of the intermediate compositions the search suggests (about 0.79 coal with 0.21 BBJP; about 0.90 coal with 0.10 rice husk), and of the coal in the blends themselves would show whether the ordering is repeatable and how much of it is reference mismatch. No third or fourth component can be excluded or recommended on the present data.

3.8.6. Genetic Algorithm and Bootstrap with the Tuned Surrogate

The search of Section 3.8.1 and the bootstrap of Section 3.8.2 were repeated with the tuned surrogate (same trust region, confidence-bound weight, and settings; bootstrap N = 500 with the tuned configuration refitted in every replicate; Table 7). The optimum of coal + BBJP was 0.80 coal and 0.20 BBJP (predicted SDI 1.680 ± 0.030), 0.005 from the measured blend 80C_20BBJP, and that of coal + rice husk + BBJP was 0.81 coal (1.677 ± 0.030). For coal + rice husk, coal + sawdust + rice husk, coal + sawdust + BBJP, and the quaternary blend, the optima lay at 0.93 coal, at the edge of the trust region (0.187–0.190 from the nearest measured blend), because the tuned model predicts a further increase in SDI with coal above 0.8 for these partners; these are extrapolations to the boundary and not findings. Coal + sawdust gave 0.49 coal (1.583 ± 0.046), 0.161 from 60C_40SD. No optimum exceeded the highest measured SDI (1.685). In the bootstrap, the standard deviation of the optimal SDI was 0.025–0.069 (median over categories 0.036; 0.000 when only the GA seed was varied), and a category containing BBJP had the highest predicted SDI in 94% of the resamples (coal + BBJP 60%, coal + sawdust + BBJP 28%, coal + rice husk + BBJP 3%, quaternary 3%). The share of resamples with an optimal coal fraction of 0.75–0.85 was 32% for coal + BBJP and 35% for coal + rice husk + BBJP. The tuned surrogate therefore agrees with the stack on the ordering of the categories in the data (BBJP highest SDI_ref; a description, not an explanation) and on the location of the coal + BBJP optimum at about 0.8 coal, and it differs in that it predicts a further rise in SDI with coal for the other partners.

Table 7. GA optima and bootstrap (N = 500) with the tuned surrogate.

Category	Coal	Partner(s)	Predicted SDI ± SD	Nearest Measured Blend (Distance)	Bootstrap Coal, Median (5th–95th)	Highest Category
Coal + sawdust	0.486	SD 0.514	1.583 ± 0.046	60C_40SD (0.161)	0.60 (0.48–0.93)	0%
Coal + rice husk	0.934	RH 0.066	1.643 ± 0.033	80C_20RH (0.190)	0.93 (0.63–0.93)	5%
Coal + BBJP	0.803	BBJP 0.197	1.680 ± 0.030	80C_20BBJP (0.005)	0.77 (0.60–0.93)	60%
Coal + SD + RH	0.932	RH 0.063, SD 0.005	1.642 ± 0.030	80C_20RH (0.190)	0.93 (0.58–0.93)	1%
Coal + SD + BBJP	0.932	BBJP 0.063, SD 0.005	1.681 ± 0.030	80C_20BBJP (0.190)	0.91 (0.52–0.93)	28%
Coal + RH + BBJP	0.810	BBJP 0.186, RH 0.004	1.677 ± 0.030	80C_20BBJP (0.017)	0.80 (0.60–0.93)	3%
Quaternary	0.928	BBJP 0.063, RH 0.005, SD 0.005	1.678 ± 0.030	80C_20BBJP (0.187)	0.85 (0.16–0.93)	3%

3.9. Sensitivity to the Start of the Analysis Window

The analysis window starts at 15.68 s, which is the first retention time of the coal reference and of the first run of 80C_20RH. To test how much the exclusion of the early sawdust peaks matters, the additivity analysis (coal profile estimated from the blends, Equation (4)) was repeated with windows starting at 7.0 s and 6.0 s, using only the blends for which their first run covers the window ('08_code/window_sensitivity.py'; Table S3). The first-run files of the other blends begin at 8.6–15.7 s, and the second-run files of 80C_20SD and 20C_26each begin at 14.6 s and 9.4 s, so only a few blends and replicate pairs remain. The window from 7.0 s retains 90% of the exported sawdust area (49% for the 15.68 s window; 99.9% from 6.0 s), but cuts the peak at about 7 s.

With the 7.0 s window (six blends, three replicate pairs) the median SDI_add was 1.04 (1.03 for the same six blends in the 15.68 s window); SDI_add of the sawdust blends was 1.19 (80C_20SD; 1.11 in the 15.68 s window of the same subset), 0.98 (40C_20each; 1.02), and 1.33 (20C_26each; 1.40). With the 6.0 s window (four blends), it was 1.12 and 1.02 for 80C_20SD and 40C_20each (1.03 and 1.05 in the 15.68 s window). The SDI between the two runs of pure sawdust rose from 1.21 to 1.39 (7.0 s) and 1.44 (6.0 s), so the early peaks add replicate variation rather than remove it. SDI_add moved by at most 0.09, and the comparison with SDI_rep (two of three pairs with SDI_add above SDI_rep) is unchanged, but only one replicated sawdust blend (40C_20each) remains, so the test has very little power. SDI involving sawdust is therefore on a different effective basis from that of the other blends, and its dependence on the early peaks cannot be settled with the present files. SDI, as defined here, neither contains total volatile yield (area normalisation) nor, without compound identification, total volatile behaviour; this requires TGA or Py-GC/MS.

4. Discussion

4.1. What the Results Show

Two runs of the same blend shared only about 55% of their normalised chromatogram area (median SDI_rep 0.89), and two runs of the same pure biomass fuel or SRF shared 5–40%. With the coal profile estimated from the blends, the deviation of the blends from an additive model (median SDI_add 1.00) was of the same size as this run-to-run variation, and in five of nine blends, it exceeded the blend's own replicate SDI by a median of only 0.05. The present data therefore do not show any departure from additive mixing of the fuels, and they do not show additivity either: The variation between duplicate runs is too large to detect an interaction of moderate size. The result is power-limited: With duplicate runs, no departure from additivity exceeding run-to-run variation was detected, and interactions below about 0.3 in SDI cannot be ruled out. Reports of non-additive volatile behaviour in coal–biomass and coal–SRF blends [14,17,18] are neither confirmed nor contradicted by this study.

A second check that does not depend on the coal reference supports this reading: Coal profiles estimated from blends with different biomass partners or coal fractions differed from each other no more than profiles between run 1 and run 2 of the same blend or those from random regroupings of the blends (Section 3.6). The additivity test is optimistic. The coal profile was estimated from the same blends, so any interaction that scales with the coal fraction is absorbed into it. In addition, SDI describes the distribution of the signal over retention time and not the amount of volatiles released, so a blend could release more or less total volatile matter than expected without changing SDI. The equality of the positive and negative parts of the difference between two unit-area chromatograms is a consequence of normalisation and carries no information on redistribution or mechanism. This study makes no claim about reaction pathways; candidate explanations from the literature, such as catalytic effects of alkali and alkaline-earth species in biomass ash, hydrogen transfer between

coal and biomass-derived fragments, and secondary gas-phase reactions [19–21], could only be tested with compound-level identification (Py-GC/MS) and with the individual components isolated. Only FID data were collected, so peaks were not identified.

4.2. Dependence of SDI on the Coal Component

Against the coal reference chromatogram, the twelve blends had an SDI of 1.43–1.68 (median 1.59) from their additive expectation, and the deviation exceeded the blend's own replicate SDI in eight of nine blends; with the coal profile estimated from the blends, the median was 1.00. The estimated profile differed from the reference by SDI 1.57, and the 0.8-coal blends differed by 1.62–1.70 from the reference, although a blend composed additively of the same coal could lie at most 0.40 from it (excess of 1.22–1.30; 0.76–0.86 for the 0.6-coal blends). Because the weight of the coal component in the additive prediction grows with the coal fraction, a difference between the reference and the coal in the blends also produces an apparent dependence of SDI on the coal fraction. The level of SDI against the reference (and therefore any statement of the type "all blends are strongly non-additive") should not be interpreted in isolation; differences between blends, or SDI relative to the duplicate-run SDI, are more informative. Estimating the coal component from the blends, as done here, avoids the dependence on the reference but cannot detect interactions that scale with that component (the consistency check across partners and coal fractions has too little power to do so given the noise), and it uses up the blend data (twelve blends for a profile of 19,359 points), which is why the leave-one-out value is also reported (1.04).

4.3. What the Surrogate and the Genetic Algorithm Can and Cannot Tell

The Gaussian-process surrogate (Section 3.7) predicted the SDI of an unseen blend no better than the mean of the other blends (nested RMSE 0.082 against 0.079; skill −0.04, R^2 −0.29), although the stack was better than the baseline when its weights were selected on the same folds (skill +0.17). The difference between the two figures is a measure of how much a model chosen on twelve blends can look better than it is. The noise level explains why: The standard deviation of a single SDI between two runs (0.063) is close to the standard deviation between blends (0.076), leaving a between-blend standard deviation of about 0.06 to be explained. The genetic-algorithm search (Section 3.8) therefore largely returned the location of the best measured blends: The optimum of the categories with BBJP was near 0.79 coal, 0.016 from 80C_20BBJP, and a bootstrap placed the optimal coal fraction at 0.58–0.93 (5th–95th percentile; 0.20–0.81 for the quaternary blend), with BBJP as the partner with the highest predicted SDI_ref in 88% of the resamples (a description of the measured data, not evidence of synergy, because SDI_ref contains reference mismatch, Section 3.3). This agrees in location (about 0.8 coal) with the optimum reported in the submitted version, but for a different reason: The search was confined to the neighbourhood of measured blends, and none of the predicted values exceeded the highest measured SDI. The two blends proposed in the submitted version, GA-Opt-2 and GA-Opt-3, are within this neighbourhood and are predicted at 1.63 and 1.55 (against 1.68 for 80C_20BBJP), which is not better than a measured blend.

The verdict on predictive skill depends on how the surrogate is configured. The stack and the default Gaussian process had no skill, whereas a Gaussian process for which its hyperparameters were optimised inside the validation (Section 3.7.5) reached a nested skill of +0.48 (90% interval +0.11 to +0.71), which none of 20 shuffled-composition controls reached. Because the search was nested, this is not the usual optimism of choosing a model on the same data; two limitations remain. First, the search space and the log-ratio shift of 0.02 were defined after the default analysis, and with twelve blends, the interval is wide. Second, the log-ratio transform with a shift turns an absent component into a distinct

level so that the model can separate blends by the components they contain (for example, blends with BBJP); this is a plausible structure for these data but is learnt from very few compositions. The nested RMSE (0.041) is at the noise level of a two-run blend mean (0.045), so a better model could not be recognised with these measurements. The tuned surrogate places the optimum of coal + BBJP at about 0.8 coal, next to the measured blend, and moves the other categories to the edge of the trust region (Section 3.8.6); it describes the ordering of the categories in the data more than any specific composition, and, because its target is SDI_ref, that ordering is not evidence of interaction. Measuring the candidate blends (Section 4.9) remains the decisive test.

4.4. Possible Sources of Run-to-Run Variation

The cause of the variation between duplicate runs was not isolated. It is not fine-scale peak jitter or a retention-time offset of a few seconds, because rebinning to 5 s or an optimistic window-wise alignment (±5 s) lowered the replicate SDI only from 0.89 to 0.73 and 0.68. Remaining candidates are (i) heterogeneity of the 0.3–0.5 mg Curie-point sample, which is small compared with the particle size and inhomogeneity of solid fuels and blends; (ii) differences between the acquisition or export of the first and second runs (the second-run files start at 0.16 s or later, and one starts at 128 s, against 15.7 s or earlier for the first runs); and (iii) nonlinear retention-time drift that a window-wise lag does not correct. The pure biomass fuel and SRF varied more between runs (SDI_rep 1.21–1.89) than the blends (0.68–1.71), which suggests that part of the variation arises in the second runs of the pure fuels. Correlation-optimised warping or dynamic time warping and more replicates are the next steps to test.

4.5. Sample Size Needed to Resolve Differences

The run-to-run standard deviation of SDI_add was 0.18, corresponding to a single-run standard deviation of about 0.12. Under the simplifying assumptions of equal variance for all blends and normally distributed errors, about six runs per blend would be needed to detect a difference of 0.2 in SDI between two blends (two-sided 5% level, 80% power), about three runs for 0.3, and about 24 runs for 0.1; the 95% ranges (3–22, 1–10 and 11–89 runs) are wide because only nine pairs were available (Figure S9). The estimates are illustrative, but they show that duplicate runs cannot resolve differences in the size seen between blends (roughly 0.1–0.3) and that triplicate runs, as is common in thermal analysis, would resolve only differences above about 0.3.

4.6. Relevance for Co-Firing Practice

The motivation for this work is a rapid, low-mass pre-screening step before pilot- or boiler-scale co-firing trials. In its present form, Py-GC/FID with SDI cannot serve that purpose, because it did not detect any departure from additivity and cannot rank blends. This does not show that the technique is unsuitable in principle; it shows that the repeatability of the chromatograms and the reference for the coal component limit what SDI can resolve.

Three further aspects should be considered before the method is used for co-firing decisions. First, the blends contain 20–80 wt.% non-coal material, whereas current PLN planning assumes about 5% biomass pellets in pulverised-coal and circulating-fluidised-bed boilers (30% in stoker boilers) and about 1% waste pellets [2,3]. At those shares, any non-additive signal is expected to be smaller, and the run-to-run variation would matter more; the method should be tested in that range. Second, Curie-point pyrolysis is performed at one temperature in an inert atmosphere with a heating profile that differs from that in a boiler, so its relation to ignition and burnout is unproven; it should be compared with thermogravimetric reactivity indices [7,8] and with pilot-scale data for the same fuels.

Third, a boiler decision also depends on ash behaviour, emissions, fuel handling and cost, which SDI does not address.

4.7. Cost and Time of Systematic Application

The Curie-point pyrolyser, the gas chromatograph, and the other laboratory equipment were made available free of charge for this study, so a per-sample cost cannot be derived from it, and the total cycle time per run (foil loading, pyrolysis, acquisition, and cool-down) was not recorded. The requirements can nevertheless be stated in terms of material and instrument time. The sample mass (0.3–0.5 mg) is small compared with the gram quantities used for proximate analysis. Signal acquisition takes 15 min per run (900 s, Section 2.2); at the three to six runs per blend indicated in Section 4.5, this is 45–90 min of acquisition per blend, excluding sample loading, cool-down, and data processing. Because more replicate runs are needed than for thermogravimetric analysis, the small sample mass does not by itself make Py-GC/FID the cheaper screening method, and a cost comparison should be made with the instrument rates of the laboratory that applies it.

4.8. Limitations

This study used a single coal, three pure biomass, and SRF samples with duplicate runs; twelve blends (six binary, three ternary, and three quaternary); at most two runs per sample (two ternary blends were run once and one second run was excluded because it did not cover the analysis window); a single pyrolysis temperature; and FID detection without peak identification. No combustion, emission, or ash data were collected. The blend fractions of five samples summed to 0.98–0.99 as prepared and were renormalised. The additivity test uses coal estimated from the same blends and is an optimistic bound. The sample-size estimates rest on nine replicate pairs. The analysis window (15.68–790.00 s) covers the volatiles eluting within the 900 s acquisition and starts after the leading peaks of the sawdust chromatogram (51% of its exported area), so the SDI of blends containing sawdust does not include them and is on a different effective basis from that of the other blends (Section 3.9). Because chromatograms are area-normalised and FID response factors differ between compound classes (no MS data), SDI detects a shape deviation of the retained window only; it does not capture total volatile yield or total volatile behaviour, which would need TGA or Py-GC/MS. The surrogate and genetic-algorithm analysis rests on twelve compositions; has no cross-validated skill in its default configuration and moderate skill, with a wide interval, after nested hyperparameter optimisation; and was confined to the neighbourhood of measured blends.

4.9. Recommendations for Future Work

1. Run at least three replicates of every blend and pure fuel with fixed acquisition settings and a stated retention-time window; use six or more if differences near 0.2 are of interest, particularly for the candidate blends near 0.8 coal with BBJP.
2. Apply retention-time alignment (for example, correlation-optimised warping or dynamic time warping) before SDI is calculated, and compare SDI with the replicate distribution.
3. Use Py-GC/MS to identify the peaks that contribute most to SDI so that the metric can be connected to specific volatiles.
4. Add blends at coal fractions between 0.2 and 0.6 and at 1–5 wt.% biomass or SRF, matched to current PLN planning, so that a surrogate can be trained on more than twelve compositions, and compare SDI with TGA reactivity and, where possible, pilot combustion results.

5. Conclusions

1. Py-GC/FID at 590 °C, using 0.3–0.5 mg of sample, was assessed as a rapid pre-screening step for coal–biomass–SRF blends. The test was whether the chromatogram of a blend equals the weighted sum of the chromatograms of its components, measured by a Shape Deviation Index (SDI) that is twice the total-variation distance between area-normalised chromatograms (range 0–2). The additive expectation was formed in two ways: with a coal reference chromatogram and with a coal profile estimated from the twelve blends under the hypothesis of additive mixing, which does not depend on the reference.
2. The SDI between two runs of the same blend had a median of 0.89 (range 0.68–1.71, n = 9), and between two runs of the same pure biomass or SRF, it was 1.21 (sawdust), 1.73 (rice husk), and 1.89 (BBJP). Rebinning to 5 s or an optimistic window-wise retention-time alignment lowered the replicate SDI only to 0.73 and 0.68, respectively. Two runs of the same blend therefore shared only about 55–66% of their normalised chromatogram area, and two runs of a pure biomass or SRF shared only 5–40%.
3. With the coal profile estimated from the blends, the residual SDI of the blends from the additive model had a median of 1.00 (range of 0.83–1.38; leave-one-out median of 1.04), close to the replicate SDI (0.89). The residual exceeded the blend's own replicate SDI in five of nine blends (median difference +0.05). The coal profile estimated from blends with different biomass partners or coal fractions varied no more than between two runs of the same blend or between random regroupings of the blends. This is a zero result of insufficient statistical power, not evidence of additivity: With duplicate runs, interactions below about 0.3 in SDI cannot be ruled out. It is also an optimistic bound, because any interaction that scales with the coal fraction is absorbed into the fitted coal profile. For blends containing sawdust, early peaks (about 7 s and 15 s, 51% of its exported area) outside the 15.68 s window are excluded, so their SDI is not directly comparable to that of blends without sawdust, for example, 80C_20SD and 80C_20BBJP.
4. The residual differed between run 1 and run 2 with a standard deviation of 0.18 (maximum 0.48), and the rank agreement between runs was weak (Spearman 0.47, $p = 0.21$, n = 9). With this scatter, about six runs per blend would be needed to detect a difference of 0.2 in SDI, and about three runs are needed for a difference of 0.3 (illustrative estimate from nine pairs). Duplicate runs cannot resolve the differences between blends that were observed.
5. Against the coal reference chromatogram, the twelve blends had an SDI of 1.43–1.68 (median 1.59), and the 0.8-coal blends lay 1.62–1.70 from the reference, although an additive blend of the same coal could lie at most 0.40 from it. With the coal profile estimated from the blends, the median was 1.00, and the estimated profile differed from the reference by SDI 1.57. Only differences between blends, or SDI relative to the duplicate-run SDI, should therefore be interpreted.
6. A Gaussian-process surrogate of SDI (12 blends and 21 runs) had no cross-validated skill in its default configuration (nested RMSE of 0.082 against 0.079 for the mean of the other blends; skill −0.04), whereas hyperparameters optimised inside a nested validation gave a skill of +0.48 (90% interval +0.11 to +0.71; the same procedure on shuffled compositions gave −0.12 on average). A genetic algorithm restricted to the neighbourhood of the measured blends placed the optimum near 0.79 coal with BBJP (predicted SDI of 1.68, 0.016 from the measured blend 80C_20BBJP), and a bootstrap of 500 resamples gave an optimal coal fraction of 0.58–0.93 (5th–95th percentile, 0.20–0.81 for the quaternary blend; median of 0.78–0.80 in six of seven categories) with BBJP as the partner with the highest predicted SDI_ref in 88% of the resamples. The surrogate

was trained on SDI_ref, which contains reference mismatch, so this is a description of the measured shape space and not evidence of chemical synergy. The blends recommended in the submitted version, GA-Opt-2 (80C_12RH_8BBJP) and GA-Opt-3 (75C_15RH_10SD), were predicted at 1.63 and 1.55 by the present surrogate, below the measured 80C_20BBJP (1.68), and are no longer proposed as priority candidates.

7. In its present form, Py-GC/FID with SDI does not replace thermogravimetric analysis, proximate analysis, or boiler trials for blend selection, and no synergistic or antagonistic blend is identified. The tested blends contain 20–80 wt.% non-coal material, whereas current PLN planning assumes about 5% biomass in pulverised-coal and circulating-fluidised-bed boilers (30% in stoker boilers) and about 1% waste [2]. The technique may become useful as a low-mass pre-screening tool, with 15 min of signal acquisition per run, once repeatability is improved.
8. This study used the following: a single coal, three pure biomass, and SRF samples with duplicate runs; twelve blends; at most two runs per sample (two ternary blends were run once, and one second run was excluded because it did not cover the analysis window; one pyrolysis temperature; FID detection without peak identification; and no combustion, emission, or ash data. The surrogate was trained on twelve compositions, and its optimum lies at or near measured blends; mechanisms and synergy were not addressed.
9. For next action, it is recommended to carry out the following: (i) Run at least three replicates of every blend and pure fuel with fixed acquisition settings and a stated retention-time window and at least six for the candidate blends near 0.8 coal with BBJP (80C_20BBJP, 80C_20RH). (ii) Apply retention-time alignment (for example, correlation-optimised warping) before calculating SDI. (iii) Identify the peaks that contribute most to SDI by Py-GC/MS. (iv) Test blends at biomass shares of 1–5 wt.%, as assumed in current PLN planning, and compare SDI with thermogravimetric reactivity and, where possible, pilot combustion results.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/en19215165/s1. Table S1: Proximate and ultimate analyses and gross calorific value of the pure fuels; Table S2: Sensitivity of the GA optima to the confidence-bound weight and the trust radius; Table S3: Sensitivity of the additivity analysis to the start of the analysis window; Figure S1: Duplicate runs of the pure fuels; Figure S2: Effect of rebinning and alignment on the duplicate-run SDI; Figure S3: Additivity residual from run 1 and run 2; Figure S4: SDI of the blends against the coal reference and the ideal-mixture bound; Figure S5: Consistency of the coal profile estimated from blend subsets; Figure S6: Bootstrap distribution of the optimal coal fraction; Figure S7: GA optima and nearest measured blends; Figure S8: Nested skill of fixed Gaussian-process configurations; Figure S9: Runs per blend needed to resolve a given SDI difference.

Author Contributions: Conceptualization, S.M.; methodology, S.M., P.A.D., and M.G.F.; software, S.M.; validation, P.A.D.; formal analysis, S.M.; investigation, S.M. and A.R.; resources, P.A.D.; data curation, S.M.; writing—original draft preparation, S.M.; writing—review and editing, D.P.S., P.A.D. and M.G.F.; visualization, S.M.; supervision, D.P.S. and P.A.D.; project administration, S.M.; funding acquisition, D.P.S. and P.A.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Institut Teknologi Bandung through the PUAPT 2024 international research programme.

Data Availability Statement: The Py-GC/FID chromatograms, the analysis notebook, and the output tables are available at https://github.com/mhdsoleh-data/TGA-CFD-MOO (accessed on 10 August 2026) and archived at https://doi.org/10.5281/zenodo.20550862.

Acknowledgments: The authors thank the Tonokura Laboratory, the University of Tokyo, for providing access to the Curie-point Py-GC/FID facilities used in this study and PT. PLN Indonesia Power UBP Suralaya for providing the coal and biomass used in this study.

Conflicts of Interest: Author Soleh Mochamad was employed by PT PLN (Persero) Pusat Penelitian dan Pengembangan Ketenagalistrikan at the time of the study. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Basu, P.; Butler, J.; Leon, M.A. Biomass co-firing options on the emission reduction and electricity generation costs in coal-fired power plants. *Renew. Energy* **2011**, *36*, 282–288. [CrossRef]
2. Direktorat Jenderal Energi Baru, Terbarukan dan Konservasi Energi (EBTKE), Kementerian Energi dan Sumber Daya Mineral (ESDM). Menuju Implementasi Cofiring Biomassa Pada PLTU Batubara [Towards the Implementation of Biomass Co-Firing at Coal Power Plants]. 5 November 2020. Available online: https://www.esdm.go.id/id/berita-unit/direktorat-jenderal-ebtke/menuju-implementasi-cofiring-biomassa-pada-pltu-batubara (accessed on 20 September 2026).
3. 3. Directorate General of Electricity, Ministry of Energy and Mineral Resources (ESDM). Bigger Share Given to Renewables in 2021–2030 Electricity Procurement Plan [RUPTL PT PLN 2021–2030, Decision of the Minister of Energy and Mineral Resources No. 188.K/HK.02/MEM.L/2021, 28 September 2021]. 5 October 2021. Available online: https://www.esdm.go.id/en/berita-unit/directorate-general-of-electricity/ruptl-2021-2030-diterbitkan-porsi-ebt-diperbesar (accessed on 20 September 2026).
4. Rhofita, E.I.; Rachmat, R.; Meyer, M.; Montastruc, L. Mapping analysis of biomass residue valorization as the future green energy generation in Indonesia. *J. Clean. Prod.* **2022**, *354*, 131667. [CrossRef]
5. Swastika, A.B.D.P.; Liyantono, L.; Wulandani, D. Rice Husk Availability Mapping as Biomass Cofiring Material at Power Plant in Indramayu. *J. Keteknikan Pertan.* **2024**, *12*, 284–300. [CrossRef]
6. U.S. Energy Information Administration. Indonesia (Country Analysis, International). Last Updated 25 August 2025. Available online: https://www.eia.gov/international/content/analysis/countries_long/Indonesia (accessed on 20 September 2026).
7. Guo, F.; He, Y.; Hassanpour, A.; Gardy, J.; Zhong, Z. Thermogravimetric analysis on the co-combustion of biomass pellets with lignite and bituminous coal. *Energy* **2020**, *197*, 117147. [CrossRef]
8. Yuan, Y.; He, Y.; Tan, J.; Wang, Y.; Kumar, S.; Wang, Z. Co-Combustion Characteristics of Typical Biomass and Coal Blends by Thermogravimetric Analysis. *Front. Energy Res.* **2021**, *9*, 753622. [CrossRef]
9. Azam, M.; Ashraf, A.; Setoodeh Jahromy, S.; Miran, S.; Raza, N.; Wesenauer, F.; Jordan, C.; Harasek, M.; Winter, F. Co-Combustion Studies of Low-Rank Coal and Refuse-Derived Fuel: Performance and Reaction Kinetics. *Energies* **2021**, *14*, 3796. [CrossRef]
10. Isaac, K.; Bada, S.O. The co-combustion performance and reaction kinetics of refuse derived fuels with South African high ash coal. *Heliyon* **2020**, *6*, e03309. [CrossRef] [PubMed]
11. Riaza, J.; Gibbins, J.; Chalmers, H. Ignition and combustion of single particles of coal and biomass. *Fuel* **2017**, *202*, 650–655. [CrossRef]
12. Nip, M.; de Leeuw, J.W.; Schenck, P.A. The characterization of eight maceral concentrates by means of Curie point pyrolysis-gas chromatography and Curie point pyrolysis-gas chromatography-mass spectrometry. *Geochim. Cosmochim. Acta* **1988**, *52*, 637–648. [CrossRef]
13. van Graas, G.; de Leeuw, J.W.; Schenck, P.A. Analysis of coals of different rank by Curie-point pyrolysis-mass spectrometry and Curie-point pyrolysis-gas chromatography-mass spectrometry. *Phys. Chem. Earth* **1980**, *12*, 485–494. [CrossRef]
14. Jones, J.M.; Kubacki, M.; Kubica, K.; Ross, A.B.; Williams, A. Devolatilisation characteristics of coal and biomass blends. *J. Anal. Appl. Pyrolysis* **2005**, *74*, 502–511. [CrossRef]
15. Li, T.; Song, F.; Bai, Y.; Wu, F.; Ruan, M.; Cao, Y.; Zhou, L.; Sun, F. Real-Time Emission, Chemical Properties, and Dynamic Evolution Mechanism of Volatile Organic Compounds during Co-Pyrolysis of Rice Straw and Semi-Bituminous Coal. *ACS ES&T Eng.* **2023**, *3*, 690–705. [CrossRef]
16. *SNI 8966:2021*; Bahan Bakar Jumputan Padat Untuk Pembangkit Listrik [Solid Recovered Fuel for Power Plants]. BSN (Badan Standardisasi Nasional): Jakarta, Indonesia, 2021.
17. Soncini, R.M.; Means, N.C.; Weiland, N.T. Co-pyrolysis of low rank coals and biomass: Product distributions. *Fuel* **2013**, *112*, 74–82. [CrossRef]
18. Liu, N.; Huang, H.; Huang, X.; Li, R.; Feng, J.; Wu, Y. Co-pyrolysis Behavior of Coal and Biomass: Synergistic Effect and Kinetic Analysis. *ACS Omega* **2024**, *9*, 31803–31813. [CrossRef] [PubMed]
19. Wang, W.; Lemaire, R.; Bensakhria, A.; Luart, D. Review on the catalytic effects of alkali and alkaline earth metals (AAEMs) including sodium, potassium, calcium and magnesium on the pyrolysis of lignocellulosic biomass and on the co-pyrolysis of coal with biomass. *J. Anal. Appl. Pyrolysis* **2022**, *163*, 105479. [CrossRef]

20. Tchapda, A.H.; Pisupati, S.V. A Review of Thermal Co-Conversion of Coal and Biomass/Waste. *Energies* **2014**, *7*, 1098–1148. [CrossRef]
21. Wu, L.; Guan, Y.; Li, C.; Shi, L.; Yang, S.; Reddy, B.R.; Ye, G.; Zhang, Q.; Liew, R.K.; Zhou, J.; et al. Free-radical behaviors of co-pyrolysis of low-rank coal and different solid hydrogen-rich donors: A critical review. *Chem. Eng. J.* **2023**, *474*, 145900. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.