

An Interpretable Surrogate Model with an Enhanced Multivariate Product Representation Support and a Neural Residual: Dimension-Independent Accuracy on Product Structure and a Diagnostic for Real Data

ERCAN GÜRVİT

Department of Mathematics, Faculty of Sciences
Marmara University
İstanbul
TÜRKİYE

Abstract: - High-dimensional model representation (HDMR) gives a surrogate that is interpretable by construction: with an orthonormal basis, its component functions translate directly into variance-based Sobol' sensitivity indices. The established HDMR surrogates, however, expand the response as a sum of low-order additive components, so a genuinely multiplicative (product or power law) response is reached only through a cascade of interaction terms whose count grows combinatorially with the dimension. We attach to the additive expansion a single log-additive multiplicative support $S(x) = \exp(\sum_j g_j(x_j))$, in the spirit of the Enhanced Multivariate Product Representation (EMPR), and let a small neural network learn the remaining non-separable part. The support is used only when it lowers a held-out error, so the method guards itself. On clean product-structured benchmarks, the surrogate is dimension-independent in accuracy: on a Genz product peak, it holds a test-normalised mean squared error near 2×10^{-4} from five to thirty variables, while a Gaussian process and a multilayer perceptron drift from 10^{-3} up to about 0.65 over the same range. We then study two real systems. On Delft model basin yacht resistance data, the method is competitive and recovers the correct dominant variable (the Froude number). On hydrogen ignition delay over fifteen perturbed Arrhenius rate parameters, generated from a real kinetic mechanism, it again recovers the physically correct ranking (the chain branching reaction $H+O_2 \rightleftharpoons O+OH$) but is beaten on accuracy by the black box models. A simple linear-in-log fit quality explains and predicts this split: near one (clean products), the support brings accuracy dominance; well below one (real coupled physics), it brings only interpretability. A direct comparison with polynomial chaos expansion further locates this product structure advantage in a logarithmic transform combined with an additive truncation rather than in the choice of HDMR, and the method is positioned within the broader landscape of spectral surrogates, post hoc explainers such as SHAP, screening methods, and Shapley effects for dependent inputs. The contribution is therefore stated precisely rather than broadly.

Key-Words: - High-dimensional model representation, Enhanced multivariate product representation, Surrogate modelling, Sobol' indices, Global sensitivity analysis, Curse of dimensionality, Combustion kinetics, Neural residual, Polynomial chaos expansion, Shapley effects, Screening methods.

2010 Mathematics Subject Classification: 41A63, 62K20, 65D15, 62J02.

Received: April 23, 2026. Revised: June 11, 2026. Accepted: July 15, 2026. Published: October 2, 2026.

1 Introduction

A surrogate model replaces an expensive simulator with a cheap approximation, so that uncertainty propagation, optimisation and global sensitivity analysis can be carried out without running the simulator thousands of times. Two qualities are usually at odds. Flexible black box regressors, such as Gaussian processes [1], gradient boosted trees [2] and neural networks, fit almost any response but say nothing about which inputs are responsible for the

output. Variance decomposition surrogates, of which high-dimensional model representation (HDMR) is the central example [3], [4], expose exactly that information, yet they can be less accurate per sample when the response does not match their structure.

HDMR and its random sampling variant write a function as a mean, a sum of univariate terms, a sum of pairwise terms, and so on. With an orthonormal basis, the squared coefficients of each block are the variance contributions, so the first- and total-order Sobol' indices [5], [6] follow algebraically, without

any extra Monte Carlo integration. Most of the later work has improved the way these component functions are estimated from limited samples; recent examples fit them with relevance vector machine regression [7]. What these approaches share is that the expansion stays additive in low-order components. A multiplicative response such as $f = \prod_j h_j(x_j)$ is then represented only by adding main effects and a growing number of interaction terms, and this is exactly the situation in which an additive expansion is most expensive: power law drag, products of Arrhenius rate constants, ratios of products in flow models.

The Enhanced Multivariate Product Representation (EMPR) of Demiralp and co-workers was proposed to address this very point [8]. EMPR multiplies the expansion by univariate support functions and, as a result, represents functions with a high level of multiplicativity far better than plain HDMR. We take a numerically stable, log-additive form of that idea and combine it with two further ingredients. First, a small neural network absorbs whatever the decomposition leaves behind, which keeps the surrogate competitive with a black box while the decomposition keeps its exact Sobol' interpretability. Second, the support is applied only when a held-out test says it helps, so the method falls back to ordinary HDMR on responses that are not separable in log and never produces a numerical blow-up.

It is useful to place this contribution within the wider landscape of spectral surrogates and sensitivity analysis. Polynomial chaos expansion (PCE) [9], [10] is the closest relative: like HDMR, it expands the response in an orthonormal polynomial basis and reads Sobol' indices analytically from the squared coefficients [11], and sparse variants selected by least angle regression [12] have made it a workhorse for global sensitivity analysis. Kriging and Gaussian process surrogates [1] predict flexibly but expose no such decomposition. On the interpretability side, model-agnostic explainers such as SHAP [13] extract post hoc feature importance from any trained model, including the neural networks we criticise, while Shapley effects [14], [15] extend variance-based sensitivity to correlated inputs, where Sobol' indices lose their uniqueness. For high-dimensional problems, a screening step, classically the elementary effects method of Morris [16], is used to discard non-influential factors before a costly variance-based analysis. We return to each of these connections in the discussion; here we note only that our method is not an alternative to this ecosystem but an additive in log support and a neural residual that can sit on top of any orthonormal additive core, and that we compare

directly against polynomial chaos expansion in Section 5.6.

The message of the paper is deliberately narrow. On clean product-structured functions, the surrogate becomes dimension-independent in accuracy, which is the opposite of the curse of dimensionality that degrades the black-box baselines. On real physical systems, the accuracy advantage need not appear, because real responses carry interactions that an additive support cannot hold; what does carry over to real data is the exact, physically meaningful sensitivity. We make this boundary precise through a cheap diagnostic, the quality of a linear-in-log fit, that tells in advance which of the two regimes a given problem belongs to. The whole study is reproducible from a single self-contained notebook that even regenerates the combustion data in place using a real kinetic mechanism [17].

2 HDMR and Sobol' Indices

Each input is rescaled to the unit interval and expanded in a normalised Legendre basis $\{P_1, \dots, P_m\}$, with the constant dropped so that every block has zero mean and unit variance. The additive surrogate is

$$f(x) \approx \beta_0 + \sum_i \sum_p \beta_{i,p} P_p(x_i) + \sum_{i < j} \sum_{p,q} \beta_{ij,p,q} P_p(x_i) P_q(x_j), \quad (1)$$

fitted by ridge-regularised least squares. Because the basis is orthonormal, the squared coefficient mass of each group is its Sobol' variance contribution: the univariate blocks give first-order indices, the pairwise blocks give second-order indices, and a total effect index is the sum over every block that contains the variable in question. No sampling-based integration is needed, which is the practical reason HDMR is attractive for sensitivity analysis [4], [5]. This additive orthonormal expansion is, in fact, a main effect polynomial chaos expansion, so the indices it yields are analytic in exactly the sense that polynomial chaos indices are; we return to this equivalence in Section 6.

3 The Multiplicative Support and the EMPR Surrogate

When the response is positive, we fit a log-additive support,

$$\log S(x) = \gamma_0 + \sum_j (\text{Legendre expansion of } x_j), \quad (2)$$

by least squares on $\log f$, and set $S(x) = \exp(\log S(x))$. For a separable product $f = \prod_j h_j(x_j)$, the logarithm is exactly additive, so S reproduces f from d univariate fits at any dimension and never incurs a combinatorial blow-up. The additive HDMR of Eq.

(1) is then fitted to the residual f/S , and the surrogate is

$$\hat{f}(x) = S(x) \cdot \text{HDMR}(f/S)(x). \quad (3)$$

The support exponent is clipped to its training range so that test predictions cannot extrapolate into an exponential blow-up. This is the log-additive realisation of the support function principle of EMPR [8], chosen here for numerical robustness.

The sensitivity indices are read from an additive HDMR of f itself, not from the factored residual; the support is an accuracy device, not a sensitivity device. The two roles are kept separate on purpose.

Self-guarding. The support is not applied blindly. On each fit, we hold out one fifth of the training data, fit the surrogate with and without the support on the rest, and keep the support only if it lowers the held-out error. This single rule makes the method abstain cleanly on responses that change sign, that are strongly interacting, or that are simply too data-starved to validate a support, and it removes the numerical failures that an unguarded log fit would otherwise produce.

Neural residual. A two-hidden-layer network with smooth activations is trained on the standardised EMPR residual $y - \hat{f}_{\text{EMPR}}(x)$ and added back, using the PyTorch library [18]. Pairwise HDMR interactions are turned on only when the sample size is large enough to estimate them, so that interaction coefficients are never fitted from too few points.

Baselines and metric. We compare against a Gaussian process with a squared exponential kernel [1], the same network applied directly to the data, gradient-boosted trees [2], and, in Section 5.6, polynomial chaos expansion; the tree and network baselines run through scikit-learn and PyTorch [19], [18]. Accuracy is the test normalised mean squared error, $\text{NMSE} = E[(y - \hat{y})^2] / \text{Var}[y]$, on a large independent test set, averaged over random seeds.

Table 1. Benchmark functions and their roles. Here d is the input dimension

Function	d	Role
Borehole	8	product, win
Wing-weight	10	powers, win
Genz peak	5–30	scaling test
Environmental	4	sum of products
Piston	7	mixed
OTL circuit	6	additive control
Ishigami	3	interaction
Sobol'-g	15	high order
Yacht	6	real, power law
H ₂ ignition	15	real, ~85% mult.

Source: created by the authors.

4 Benchmarks and Baselines

The benchmark functions, listed in Table 1, span the structural regimes that matter here. The engineering models and their input ranges are the standard ones from the Virtual Library of Simulation Experiments [20]; the Genz product peak is a classical test integrand with a known analytic decomposition [21]; the environmental model is the pollutant spill problem of [22]. The controls are reported rather than hidden, because the cases where the method is inert or loses are what make the positive claims credible.

5 Results

5.1 Dimension scaling

Table 2. Test NMSE versus dimension on the Genz product peak, at a fixed training size $N = 600$ (mean over seeds)

d	EMPR+NN	GP	MLP
5	0.0000	0.0008	0.0052
10	0.0001	0.0572	0.0436
20	0.0002	0.4411	0.5146
30	0.0002	0.6424	0.6532

Source: created by the authors.

Table 2. is the central result. For the Genz product peak with a fixed budget of 600 training points, the proposed surrogate maintains almost constant accuracy as the dimension grows from five to thirty, whereas the Gaussian process and the network deteriorate by three orders of magnitude. The log-additive support represents the product exactly, so accuracy is essentially independent of d , while the black-box models pay the full price of the curse of dimensionality

5.2 Data efficiency and honest controls

Table 3 collects the data efficiency results for all eight analytic functions at three representative training sizes. On the product structured functions, the proposed method matches or beats the Gaussian process, with the largest margins at small samples: it leads on wing-weight and borehole at $N = 75$, and on the twenty-dimensional Genz peak it is already near 10^{-4} at $N = 300$ while the baselines remain above 0.3. The controls behave as the mechanism predicts. On the additive OTL circuit, the support is inert, and the method reduces to HDMR, where the Gaussian process and the network are stronger. On Ishigami, the sign-changing x_1 – x_3 interaction is best captured by the Gaussian process. On the high-order, near-singular Sobol'-g function, no low-order method wins, and

gradient boosting is preferable; with the self-guard, the proposed method stays bounded there rather than diverging

Table 3. Test NMSE versus training size N on the analytic benchmarks (mean over seeds). The best entry in each row is in bold. EMPR+NN is the proposed method

Function	N	EMPR	EMPR+NN	GP	MLP	GBoost
Borehole	75	0.0019	0.0015	0.0038	0.0025	0.0771
	300	0.0010	0.0002	0.0000	0.0003	0.0250
	1200	0.0000	0.0000	0.0000	0.0002	0.0092
Wing-weight	75	0.0017	0.0016	0.0018	0.0047	0.1924
	300	0.0002	0.0000	0.0002	0.0004	0.0558
	1200	0.0000	0.0000	0.0000	0.0002	0.0213
Genz peak (d=20)	75	0.2939	0.2937	0.7497	0.9863	1.6369
	300	0.0001	0.0001	0.5501	0.4945	0.8355
	1200	0.0001	0.0001	0.3070	0.1764	0.6502
Environmental	75	0.1096	0.0457	0.0500	0.0327	0.2082
	300	0.0051	0.0020	0.0039	0.0046	0.0577
	1200	0.0044	0.0004	0.0001	0.0017	0.0206
Piston	75	0.0898	0.0592	0.0289	0.0144	0.1089
	300	0.0593	0.0078	0.0067	0.0045	0.0272
	1200	0.0062	0.0005	0.0013	0.0034	0.0137
OTL (additive)	75	0.0127	0.0089	0.0017	0.0014	0.0946
	300	0.0092	0.0013	0.0001	0.0005	0.0228
	1200	0.0000	0.0000	0.0000	0.0004	0.0073
Ishigami (interaction)	75	0.6416	0.6728	0.2940	0.3380	0.2731
	300	0.2758	0.1584	0.0032	0.0396	0.0947
	1200	0.2575	0.1633	0.0000	0.0254	0.0461

Function	N	EMPR	EMPR+NN	GP	MLP	GBoost
Sobol'-g (high order)	75	1.8158	1.8201	1.0130	1.1592	0.7574
	300	0.4951	0.7158	0.7149	0.6667	0.4763
	1200	0.4168	0.5991	0.3872	0.3232	0.3138

Source: created by the authors.

5.3 Interpretability against ground truth

Where an analytic answer exists, the recovered Sobol' indices are correct. For Ishigami, the method returns $S_1 = 0.314$, $S_2 = 0.443$, $S_3 = 0.000$ and a leading (x_1, x_3) interaction of 0.243, against the analytic values 0.314, 0.442, 0 and 0.244. For the borehole model, the total effect ranking is led by the radius r_w at 0.867, in agreement with the value reported in the literature. For an eight-dimensional Genz peak, whose main effect indices are available in closed form, the recovered indices match the analytic ones to a maximum absolute error of 0.006. These come for free from the same fit that produces the prediction; none of the black-box baselines offers a comparable decomposition.

5.4 Real data I: yacht hydrodynamics

The Delft model basin data set has 308 experiments with six inputs (five hull geometry coefficients and the Froude number) and a residuary resistance target [23]; it is fetched live in the notebook. The resistance is power law in the Froude number, and the recovered total effect ranking reflects this exactly: the Froude number sits at 0.949, and every geometric coefficient is below 0.05. On accuracy, the proposed method is competitive, winning at $N = 160$, but gradient boosting and the network are strong here because the problem is effectively one-dimensional, a single variable carrying about ninety-five per cent of the variance. This is the regime in which a high-dimensional product advantage simply does not apply.

5.5 Real data II: hydrogen ignition delay

Table 4. Real data accuracy (test NMSE). Yacht and hydrogen ignition delay; best entry of each row in bold

Data	N	EMPR+NN	GP	MLP/GB
Yacht	120	0.0050	0.0101	0.0048
	160	0.0021	0.0067	0.0041
Ignition τ	250	0.0598	0.0088	0.0045
	900	0.0127	0.0020	0.0016

Source: created by the authors.

To test a genuinely high dimensional real problem, we generate data in place from a real kinetic mechanism with Cantera [17]: for a constant-volume hydrogen–air reactor, we perturb the Arrhenius rate multipliers of the fifteen most ignition-sensitive reactions and record the ignition delay τ . The result is a fifteen-dimensional problem with τ spanning about two orders of magnitude. Table 4 reports a deliberately mixed outcome. The total effect ranking is again physically correct: the chain-branching reaction $\text{H}+\text{O}_2 \rightleftharpoons \text{O}+\text{OH}$ leads at 0.40, followed by the peroxide reactions $\text{H}+\text{H}_2\text{O}_2$ and $2\text{OH}(\text{+M})$, exactly the reactions that combustion theory identifies as rate-controlling. On accuracy, however, the black box models win, and the proposed method is competitive but not dominant.

The reason is measurable. A linear in log fit of τ has a coefficient of determination of about 0.85: the response is roughly eighty-five per cent multiplicative and fifteen per cent reaction coupling interactions. Adding pairwise terms to an HDMR of log τ cuts its error from 0.13 to 0.04, which confirms that the interactions are real. The additive support captures the separable part; the black box models additionally capture the coupling, which at feasible sample sizes in fifteen dimensions the decomposition cannot.

5.6 Comparison with polynomial chaos expansion

Polynomial chaos expansion (PCE) is the natural spectral competitor to HDMR: it expands the response in an orthonormal polynomial basis and, like an orthonormal HDMR, yields Sobol' indices analytically from the squared coefficients [11]. It is therefore reasonable to ask whether the dimension-independent accuracy we report on product functions survives once PCE is given the same advantage we exploit, namely a logarithmic transform of a positive product response. Using an orthonormal Legendre basis in the manner of standard polynomial chaos software [24], we ran PCE on the Genz product peak under the identical budget of Table 2 ($N = 600$ training points, an independent test set of 20 000 points, five seeds). Three variants are compared: a standard total degree PCE fitted to f ; the same total degree PCE fitted to log f and exponentiated; and a PCE fitted to log f but truncated to its additive, main effect terms only, which is exactly the log additive support at the core of our method. In every case, the truncation degree is the largest that the sample budget supports, so no variant is handicapped. Results are in Table 5.

Table 5. Test NMSE on the Genz product peak at $N = 600$ (mean over five seeds). Standard PCE fitted to f ; log-PCE fitted to log f under total-degree truncation; log-PCE under additive (main-effect) truncation, which coincides with our log-additive support. EMPR+NN values are reproduced from Table 2; the PCE columns are from an independent implementation released with the notebook

d	PCE on f	log-PCE (total degree)	log-PCE (additive = support)	EMPR+NN
5	0.1221	0.0024	0.0004	0.0000
10	0.8315	0.0840	0.0009	0.0001
20	0.9853	0.2020	0.0021	0.0002
30	1.0104	1.0290	0.0029	0.0002

Source: created by the authors.

Three things are visible. First, standard PCE on f suffers the full curse of dimensionality, its error climbing from 0.12 at five variables to essentially the variance of the data at thirty, exactly like the Gaussian process and the network of Table 2. Second, the logarithmic transform is genuinely powerful: at five and ten variables it recovers two to three orders of magnitude of accuracy, confirming that the multiplicativity, not the polynomial machinery, is what matters. But under the standard total degree truncation, the shared degree budget must fall as the dimension rises — at thirty variables, $N = 600$ points support only a first-degree expansion — and the log transform alone then collapses back to failure. Third, when the log space PCE is truncated additively instead, so that each variable keeps its own degree and the term count grows only as $O(d \cdot p)$ rather than combinatorially, dimension-independent accuracy returns: the error stays near 10^{-3} from five to thirty variables. That additive in log expansion is precisely our multiplicative support; the full method reaches the smaller value of Table 2 by adding the neural residual on top.

The experiment sharpens rather than weakens the paper's claim. The advantage on products is not a property of HDMR versus PCE — the two share the same orthonormal spectral foundation — but of two design choices that a practitioner must make: transforming to log space, and truncating additively rather than by total degree. A PCE user who does both recovers our behaviour; our contribution is to make both automatic and safe, through the self-guard that decides when the log support helps and the neural residual that captures what the additive log model leaves behind. Because these numbers use an independent PCE implementation and our own draw of Genz parameters, they should be read as evidence

of the mechanism rather than as digit-for-digit comparable with Table 2.

6 Discussion

The result is not a uniformly best surrogate, and the controls and the combustion case show why such a claim would be false. It is sharper: a single log-additive multiplicative factor on top of additive HDMR buys dimension-independent accuracy when, and only when, the response is close to separable in log, and it buys exact Sobol' interpretability always. The two real systems mark the boundary cleanly. The yacht is effectively low dimensional, and the hydrogen mechanism is genuinely fifteen dimensional but about fifteen percent coupled; both give the interpretability win and decline the accuracy win. This rules out the comfortable reading that the yacht result was merely an effect of low dimension, since combustion is high dimensional and still does not yield accuracy dominance, simply because real physics is not a clean product.

The practical tool that comes out of this is the linear-in-log fit quality. It is computable from the training data alone and tells, before any surrogate is built, which regime a problem belongs to: near one for clean products predicts accuracy dominance, well below one for coupled physics predicts interpretability only. Relative to recent work that improves the coefficient estimator of an additive expansion with relevance vector machines [7], our contribution is orthogonal: the multiplicative support and the neural residual could sit on top of any such additive core. For global sensitivity analysis specifically, the method is useful on both real systems, returning the correct rate-controlling reactions for ignition and the correct dominant variable for yacht resistance as a reusable algebraic by-product, at an accuracy that is at worst competitive.

6.1 Relation to polynomial chaos expansion

Beyond accuracy, PCE returns Sobol' indices directly from its coefficients, whereas we read indices from a separate additive HDMR off, and it is fair to ask whether this second fit is an avoidable complication. The answer clarifies the design. Our sensitivity model is itself an orthonormal spectral expansion — a main effect polynomial chaos in the Legendre basis — so its indices are equally analytic, obtained from squared coefficients with no Monte Carlo integration; we do not use a different mechanism from PCE; we use a second, deliberately additive, instance of the same one. It is kept separate from the predictive surrogate for a structural reason: the multiplicative support and

the neural residual, which are what buy predictive accuracy, destroy the orthogonal variance partition that makes the indices exact, so folding them into the index computation would forfeit precisely the interpretability we want to guarantee. PCE avoids the second fit only because it never introduces a support in the first place, and Table 5 shows the price of that in high dimension, where the single predictive and interpretive PCE fit is the one that succumbs to the curse, while our additive index model, being cheap and low order, is untouched. Sparse and adaptive PCE variants such as least angle regression PCE [12] narrow the accuracy gap by selecting a parsimonious basis and are a natural alternative core; as with the relevance vector machine HDMR of [7], our support and residual are orthogonal additions that could sit on top of a sparse PCE core just as readily as on additive HDMR.

6.2 Intrinsic decomposition versus post hoc explanation

We criticise black-box regressors for opacity, and this opacity is not absolute: model-agnostic explainers, and SHAP in particular [13], can attribute feature importance to a trained network after the fact. It is worth stating clearly how the two notions of interpretability differ. Our indices are intrinsic and global: they are variance shares of the fitted function, defined before any prediction is made, exact against analytic ground truth (Section 5.3), and identical for every user because they are a property of the model rather than of a sampling scheme or a background dataset. SHAP values are post hoc and, in their common form, local: they explain individual predictions through Shapley attributions over features, are aggregated to a global importance only by averaging, and depend on the chosen background distribution and on approximations whose cost grows with dimension. The two are complementary rather than competing. A SHAP analysis of our neural residual, for instance, would describe the non-separable coupling that the residual absorbs and that the variance decomposition, by construction, does not localise; conversely, our decomposition gives SHAP a ground truth check on the separable part. For a surrogate whose purpose is global sensitivity analysis, an intrinsic variance decomposition is the more direct instrument, because it delivers exactly the Sobol' indices the analysis asks for without a second explanatory model; we therefore regard SHAP not as a substitute but as a useful post hoc probe of the residual, in line with work that advocates combining intrinsic and post hoc views for design exploration.

6.3 The prediction–interpretability trade-off

The whole paper can be read as a statement about this trade-off, and it is worth making explicit. The classical position is that flexibility and transparency pull against each other: a Gaussian process or a network predicts well and explains nothing; an additive decomposition explains everything and predicts less well off its structure. Our results refine this into two regimes governed by a single measurable quantity: the quality of a linear-in-log fit. When that quality is near one — clean product structure — the trade-off vanishes: the log additive support is simultaneously the most accurate model and the exactly interpretable one, and there is nothing to trade (Tables 2 and 5). When it is well below one — real coupled physics — the trade-off returns in a controlled form: the black box captures coupling the additive support cannot, so it predicts better, while our method concedes some accuracy but keeps the exact, physically correct sensitivity (Section 5.5). The contribution is therefore not a claim to have abolished the trade-off but a diagnostic that says, in advance and from the training data alone, which regime a problem is in and hence how much accuracy, if any, interpretability will cost.

6.4 Screening and a two-step sensitivity workflow

For the fifteen-dimensional combustion problem, engineering practice would ordinarily begin with a screening method — the elementary effects method of Morris [16] — to isolate the few influential factors before committing to a more expensive variance-based decomposition, and it is natural to ask how our method fits such a workflow. It fits as the second stage, and our own results are evidence for the first. The total effect ranking we recover concentrates almost all of the variance in three reactions, led by the chain branching $\text{H}+\text{O}_2 \rightleftharpoons \text{O}+\text{OH}$; a Morris screen, which is far cheaper than a full decomposition, would be expected to flag the same handful and to discard the remaining rate constants as non-influential. A natural and recommended pipeline is therefore to run a Morris screening first, reduce the fifteen parameters to the four or five that survive, and then apply the EMPR surrogate on that reduced set, where both the additive support and any affordable second-order interaction terms become much cheaper to estimate, and the residual coupling that defeats the method in full dimension may become tractable. Screening and our decomposition are complementary — the former a cheap filter on which variables matter, the latter a precise account of how much and through what structure — and aligning them matches standard sensitivity analysis practice [6].

6.5 Correlated inputs, Sobol' indices, and Shapley effects

Our interpretability rests on Sobol' indices, and their central limitation should be acknowledged: the variance decomposition on which they stand assumes independent inputs, and under correlation the decomposition is no longer unique, so first- and total-order indices can double-count shared variance and distort a ranking [25]. Shapley effects, which distribute each portion of the output variance fairly among inputs using the Shapley value from cooperative game theory, remain well defined under dependence and are the more robust choice when inputs are correlated [14], [15]. Two points justify our use of Sobol' indices here. First, in every case we study, the inputs are independent by construction: the benchmark variables are sampled independently, the yacht design parameters vary over an experimental design, and the perturbed Arrhenius multipliers are drawn independently, so the variance partition is unique and the indices are exactly the quantities validated against analytic ground truth in Section 5.3. Second, the orthonormal component functions our method produces are precisely the objects a Shapley effect computation consumes, so extending the method to correlated inputs is a change of aggregation, not of surrogate: the same fitted decomposition that yields Sobol' indices under independence can feed a Shapley effect estimator under dependence, as the relevance vector machine HDMR line of work also notes [7]. We therefore present Sobol' indices as correct for the independent input problems studied, and flag Shapley effects as the appropriate instrument once real, correlated inputs are in play.

6.6 Limitations

The limitations follow from the same analysis. The support needs positivity, and sign-changing responses fall back to additive HDMR. A single global log additive factor cannot represent interactions, so capturing the residual coupling would require either affordable second-order HDMR, which is costly in fifteen dimensions, or a richer support; this is left for future work. The accuracy dominance claim is, by the diagnostic itself, restricted to near-separable responses, and we make no broader claim.

7 Conclusion

We presented an interpretable surrogate that augments additive HDMR with an EMPR-style log-additive multiplicative support and a neural residual, gated by held-out model selection. On clean product-structured

benchmarks, it is dimension-independent in accuracy while the black-box baselines succumb to the curse of dimensionality, and it recovers Sobol' indices exactly against analytic ground truth. On two real physical systems, yacht residuary resistance and hydrogen ignition delay over fifteen rate parameters, it returns the physically correct sensitivity structure while being competitive but not accuracy dominant, and a simple linear-in-log diagnostic predicts in advance which regime applies. A direct comparison with polynomial chaos expansion locates the source of the product structure advantage in a logarithmic transform together with an additive truncation rather than in HDMR itself, and we position the method within the wider landscape of spectral surrogates, post hoc explainers such as SHAP, screening methods, and Shapley effects for dependent inputs. The contribution is precise rather than sweeping: dimension-independent accuracy under near separability, exact interpretability always. Every experiment, including the live data fetch and the in-place combustion data generation, is reproducible from a single self-contained notebook.

Acknowledgement:

The author is grateful to Prof. Metin Demiralp, whose Enhanced Multivariate Product Representation provided the starting point for the multiplicative support studied here, and to the three anonymous reviewers, whose comments on polynomial chaos expansion, post hoc explainability, screening, and correlated input sensitivity substantially broadened the discussion.

Declaration of Generative AI and AI-assisted Technologies in the Writing Process

The authors wrote, reviewed and edited the content as needed and they have not utilised artificial intelligence (AI) tools. The authors take full responsibility for the content of the publication.

References:

- [1] C. E. Rasmussen and C. K. I. Williams, Gaussian Processes for Machine Learning, MIT Press, Cambridge, MA, USA, 2006, ISBN 978-0-262-18253-9.
- [2] J. H. Friedman, Greedy function approximation: a gradient boosting machine, The Annals of Statistics, Vol.29, No.5, 2001, pp. 1189–1232, doi: [10.1214/aos/1013203451](https://doi.org/10.1214/aos/1013203451).
- [3] H. Rabitz and Ö. F. Aliş, General foundations of high-dimensional model representations, Journal of Mathematical Chemistry, Vol.25, 1999, pp. 197–233, doi: [10.1023/A:1019188517934](https://doi.org/10.1023/A:1019188517934).
- [4] G. Li, S.-W. Wang, and H. Rabitz, Practical approaches to construct RS-HDMR component functions, Journal of Physical Chemistry A, Vol.106, No.37, 2002, pp. 8721–8733, doi: [10.1021/jp014567t](https://doi.org/10.1021/jp014567t).
- [5] I. M. Sobol', Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, Mathematics and Computers in Simulation, Vol.55, No.1–3, 2001, pp. 271–280, doi: [10.1016/S0378-4754\(00\)00270-6](https://doi.org/10.1016/S0378-4754(00)00270-6).
- [6] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola, Global Sensitivity Analysis: The Primer, John Wiley & Sons, Chichester, UK, 2008, ISBN 978-0-470-05997-5.
- [7] F. R. Bennett and M. Roberts, The efficient estimation of Sobol' indices and Shapley effects using random sampling high dimensional model representations and relevance vector machine regression, Environmental Modelling & Software, 2026 (in press).
- [8] B. Tunga and M. Demiralp, The influence of the support functions on the quality of enhanced multivariate product representation, Journal of Mathematical Chemistry, Vol.48, 2010, pp. 827–840, doi: [10.1007/s10910-010-9714-2](https://doi.org/10.1007/s10910-010-9714-2).
- [9] R. G. Ghanem and P. D. Spanos, Stochastic Finite Elements: A Spectral Approach, Springer-Verlag, New York, USA, 1991, ISBN 978-0-387-97456-9.
- [10] D. Xiu and G. E. Karniadakis, The Wiener–Askey polynomial chaos for stochastic differential equations, SIAM Journal on Scientific Computing, Vol.24, No.2, 2002, pp. 619–644, doi: [10.1137/S1064827501387826](https://doi.org/10.1137/S1064827501387826).
- [11] B. Sudret, Global sensitivity analysis using polynomial chaos expansions, Reliability Engineering & System Safety, Vol.93, No.7, 2008, pp. 964–979, doi: [10.1016/j.ress.2007.04.002](https://doi.org/10.1016/j.ress.2007.04.002).
- [12] G. Blatman and B. Sudret, Adaptive sparse polynomial chaos expansion based on least angle regression, Journal of Computational Physics, Vol.230, No.6, 2011, pp. 2345–2367, doi: [10.1016/j.jcp.2010.12.021](https://doi.org/10.1016/j.jcp.2010.12.021).
- [13] S. M. Lundberg and S.-I. Lee, A unified approach to interpreting model predictions, in Advances in Neural Information Processing

- Systems 30 (NeurIPS 2017), Long Beach, CA, USA, 2017, pp. 4765–4774.
- [14] A. B. Owen, Sobol' indices and Shapley value, *SIAM/ASA Journal on Uncertainty Quantification*, Vol.2, No.1, 2014, pp. 245–251, doi: [10.1137/130936233](https://doi.org/10.1137/130936233).
- [15] E. Song, B. L. Nelson, and J. Staum, Shapley effects for global sensitivity analysis: theory and computation, *SIAM/ASA Journal on Uncertainty Quantification*, Vol.4, No.1, 2016, pp. 1060–1083, doi: [10.1137/15M1048070](https://doi.org/10.1137/15M1048070).
- [16] M. D. Morris, Factorial sampling plans for preliminary computational experiments, *Technometrics*, Vol.33, No.2, 1991, pp. 161–174, doi: [10.1080/00401706.1991.10484804](https://doi.org/10.1080/00401706.1991.10484804).
- [17] D. G. Goodwin, H. K. Moffat, I. Schoegl, R. L. Speth, and B. W. Weber, Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes, Version 3.0, 2023, <https://www.cantera.org> (Accessed Date: January 2, 2026).
- [18] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, et al., PyTorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, Vancouver, BC, Canada, 2019, pp. 8024–8035.
- [19] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, et al., Scikit-learn: Machine learning in Python, *Journal of Machine Learning Research*, Vol.12, 2011, pp. 2825–2830, <http://jmlr.org/papers/v12/pedregosa11a.html> (Accessed Date: January 2, 2026).
- [20] S. Surjanovic and D. Bingham, Virtual Library of Simulation Experiments: Test Functions and Datasets, Simon Fraser University, 2013, <http://www.sfu.ca/~ssurjano> (Accessed Date: January 2, 2026).
- [21] A. Genz, Testing multidimensional integration routines, in B. Ford, J.-C. Rault, and F. Thomasset (eds.), *Tools, Methods and Languages for Scientific and Engineering Computation*, North-Holland, Amsterdam, The Netherlands, 1984, pp. 81–94.
- [22] N. Bliznyuk, D. Ruppert, C. Shoemaker, R. Regis, S. Wild, and P. Mugunthan, Bayesian calibration and uncertainty analysis for computationally expensive models using optimization and radial basis function approximation, *Journal of Computational and Graphical Statistics*, Vol.17, No.2, 2008, pp. 270–294, doi: [10.1198/106186008X320681](https://doi.org/10.1198/106186008X320681).
- [23] I. Ortigosa, R. López, and J. García, A neural networks approach to residuary resistance of sailing yachts prediction, in *Proceedings of the International Conference on Marine Engineering (MARINE 2007)*, Barcelona, Spain, 2007.
- [24] J. Feinberg and H. P. Langtangen, Chaospy: An open source tool for designing methods of uncertainty quantification, *Journal of Computational Science*, Vol.11, 2015, pp. 46–57, doi: [10.1016/j.jocs.2015.08.008](https://doi.org/10.1016/j.jocs.2015.08.008).
- [25] S. Kucherenko, S. Tarantola, and P. Annoni, Estimation of global sensitivity indices for models with dependent variables, *Computer Physics Communications*, Vol.183, No.4, 2012, pp. 937–946, doi: [10.1016/j.cpc.2011.12.020](https://doi.org/10.1016/j.cpc.2011.12.020).

Contribution of Individual Authors to the Creation of a Scientific Article (Ghostwriting Policy)

Ercan Gürvit is the sole author of this article and carried out all stages of the research, from the formulation of the problem and the development of the method to the implementation, the experiments, and the writing of the manuscript.

Sources of Funding for Research Presented in a Scientific Article or Scientific Article Itself

No funding was received for conducting this study.

Conflict of Interest

The author has no conflicts of interest to declare that are relevant to the content of this article.

Creative Commons Attribution License 4.0 (Attribution 4.0 International, CC BY 4.0)

This article is published under the terms of the Creative Commons Attribution License 4.0 (Attribution 4.0 International, CC BY 4.0), https://creativecommons.org/licenses/by/4.0/deed.en_US