Monte Carlo Shapley Effects for Correlated Inputs: A Unified Framework with Surrogate Model Integration

Frederick R. Bennett
Department of Environment, Tourism, Science and Innovation
Queensland Government, Australia

Abstract

Shapley effects provide interpretable, game-theoretic measures of input importance that remain valid under arbitrary dependence structures, making them essential for sensitivity analysis of models with correlated inputs. However, their estimation requires conditional Monte Carlo sampling from the joint input distribution, which is computationally demanding and algorithmically subtle—particularly when inputs follow mixed distribution types (Normal, LogNormal, Uniform) with a Gaussian copula dependence. We present a unified Monte Carlo framework implemented in the **ShapleyX** Python package that addresses these challenges through three innovations: (i) a modular distribution class interface with exact analytical conditional sampling via latent Gaussian copula conditioning, enabling mixed marginal types to be handled by a single, reusable architecture; (ii) a coalition truncation strategy ($k_{\max}$) that exploits bounded interaction orders in surrogate models to reduce the subset enumeration from $O(2^d)$ to $O(d^{k_{\max}})$, yielding up to 5,000× speedups with zero accuracy loss; and (iii) tight integration with sparse Legendre polynomial chaos expansion (PCE) surrogates constructed via Automatic Relevance Determination (ARD) regression, enabling accurate Shapley estimation at negligible computational cost beyond surrogate training. The Owen–Prieur covariance-based Pick-Freeze estimator is used throughout, providing unbiased single-shot estimation without the inner-loop bias corrections required by double Monte Carlo. Three detailed case studies validate the framework: (a) the correlated Ishigami function ($d = 3$), where a sweep over $\rho \in [-0.9, 0.9]$ reveals how Shapley effects redistribute importance as correlation strength increases, while Sobol indices lose interpretability; (b) the cantilever beam structural reliability problem ($d = 6$, mixed LogNormal + Normal marginals, $\rho(l_X, l_Y) = -0.55$), where a full PCE surrogate of order 10 trained on 2,048 Sobol samples achieves $R^2 = 0.99987$ and reproduces target Shapley effects within ± 0.01; and (c) the wing weight function ($d = 10$, all Uniform inputs), where independent and correlated Shapley analyses are compared against OpenTURNS reference values.

Keywords: Shapley effects, correlated inputs, Monte Carlo, surrogate modelling, uncertainty quantification, Python.

Shapley effects provide interpretable, game-theoretic measures of input importance that remain valid under arbitrary dependence structures, making them essential for sensitivity analysis of models with correlated inputs. However, their estimation requires conditional Monte Carlo sampling from the joint input distribution, which is computationally demanding and algorithmically subtle—particularly when inputs follow mixed distribution types (Normal,

LogNormal, Uniform) with a Gaussian copula dependence. We present a unified Monte Carlo framework implemented in the **ShapleyX** Python package that addresses these challenges through three innovations: (i) a modular distribution class interface with exact analytical conditional sampling via latent Gaussian copula conditioning, enabling mixed marginal types to be handled by a single, reusable architecture; (ii) a coalition truncation strategy ($k_{\max}$) that exploits bounded interaction orders in surrogate models to reduce the subset enumeration from $O(2^d)$ to $O(d^{k_{\max}})$, yielding up to 5,000× speedups with zero accuracy loss; and (iii) tight integration with sparse Legendre polynomial chaos expansion (PCE) surrogates constructed via Automatic Relevance Determination (ARD) regression, enabling accurate Shapley estimation at negligible computational cost beyond surrogate training. The Owen–Prieur covariance-based Pick-Freeze estimator is used throughout, providing unbiased single-shot estimation without the inner-loop bias corrections required by double Monte Carlo. Three detailed case studies validate the framework: (a) the correlated Ishigami function ($d = 3$), where a sweep over $\rho \in [-0.9, 0.9]$ reveals how Shapley effects redistribute importance as correlation strength increases, while Sobol indices lose interpretability; (b) the cantilever beam structural reliability problem ($d = 6$, mixed LogNormal + Normal marginals, $\rho(l_X, l_Y) = -0.55$), where a full PCE surrogate of order 10 trained on 2,048 samples ($R^2 = 0.99987$) reproduces target Shapley effects within ± 0.01 of the Demange-Chryst et al. (2022) importance-sampling reference; and (c) the wing weight aerospace function ($d = 10$), where introducing four physically-motivated correlations fundamentally reorders the importance hierarchy, with wing area S_w rising from rank 4 (Sh = 0.125) to rank 1 (Sh = 0.337)—a finding invisible to classical Sobol indices. The framework supports both exhaustive and permutation-based estimation with built-in bootstrap confidence intervals, Sobol indices as a free by-product, and a documented interface for user-defined distribution classes.

1. Introduction

Global sensitivity analysis (GSA) quantifies how uncertainty in model inputs propagates to uncertainty in model outputs Saltelli et al. (2008). Among the many GSA methods developed over the past three decades—spanning variance-based approaches Sobol' (2001); Saltelli (2002), derivative-based measures Lamboni et al. (2013), moment-independent methods Borgonovo (2007); Pianosi and Wagener (2015), and dependence measures Da Veiga (2015)—variance-based Sobol indices have become the dominant paradigm for quantitative GSA owing to their rigorous decomposition of output variance into contributions from individual inputs and their interactions.

However, Sobol indices have a well-known limitation: when inputs are correlated, the variance decomposition is no longer unique, and the classical interpretation of S_i as "the fraction of variance explained by input i alone" breaks down Chastaing et al. (2012); Mara and Tarantola (2012); Kucherenko et al. (2012). This is not a marginal issue—correlated inputs are the norm in engineering design (geometric parameters co-vary), environmental modelling (physical parameters exhibit dependence), and financial risk analysis (market factors are correlated).

Shapley effects Owen (2014); Song et al. (2016), rooted in cooperative game theory Shapley (1953), provide a principled resolution. By treating each input variable as a "player" in a cooperative game whose "worth" is the variance explained by a coalition of variables, the Shapley value allocates output variance to each input in a way that is non-negative, sums to one, and—crucially—remains interpretable under arbitrary input dependence Owen and

Prieur (2017); Iooss and Prieur (2019). Shapley effects distribute shared variance contributions fairly among interacting and dependent variables, making them the appropriate tool for models with correlated inputs.

The computational challenge of estimating Shapley effects is substantial. For d inputs, the exact valuation of all $2^d - 1$ coalitions requires $O(2^d \cdot N)$ model evaluations, where N is the Monte Carlo sample size per coalition. Each coalition evaluation involves conditional sampling from $P(\mathbf{X}_{-u} \mid \mathbf{X}_u = \mathbf{x}_u)$, which is distribution-dependent and analytically non-trivial for non-Gaussian marginals. Song et al. Song et al. (2016) established the foundational computational framework, Owen & Prieur Owen and Prieur (2017) provided the covariance-based Pick-Freeze estimator that avoids inner-loop bias corrections, and subsequent work has improved estimation efficiency through variance reduction Broto et al. (2020), fast permutation algorithms Plischke et al. (2021), and importance sampling for rare-event reliability analysis Demange-Chryst et al. (2023); Il Idrissi et al. (2021).

Despite these advances, a practical gap remains: most existing implementations are either research code tightly coupled to a specific estimation strategy (e.g., the Demange-Chryst IS framework) or lack a general distribution interface that handles the variety of input models encountered in practice. The surrogate-assisted workflow—construct a cheap-to-evaluate approximation of an expensive model, then run MC Shapley on the surrogate—has been explored by Stein & Singh Stein and Singh (2023) for polynomial chaos expansions, but the integration of surrogate metadata (specifically, the known interaction order of an RS-HDMR expansion) into the Shapley estimation procedure has not been systematically developed.

This paper presents the Monte Carlo Shapley estimation framework implemented in **ShapleyX** v0.5.1 Bennett (2026,b), a general-purpose Python package for GSA. The framework provides:

1. A **modular distribution class interface** that decouples correlation modelling (handled in latent normal space via a Gaussian copula) from marginal modelling (handled by per-variable inverse CDF transformations). Three built-in classes and a documented extensibility pattern support Normal, LogNormal, Uniform, and truncated marginals with arbitrary correlation structures.
2. **Coalition truncation** ($k_{\max}$) that exploits the bounded interaction order of surrogate models to reduce subset enumeration from $O(2^d)$ to $O(d^{k_{\max}})$, with **auto-detection** from surrogate metadata.
3. **Surrogate integration** that couples the RS-HDMR sparse Legendre PCE surrogate Bennett and Roberts (2025) with the MC Shapley machinery, enabling accurate sensitivity analysis at negligible computational cost beyond surrogate construction.
4. **Two estimation methods** (exhaustive and permutation-based) with built-in bootstrap confidence intervals and Sobol first-order/total-order indices extracted as a free by-product.

The framework is validated through three case studies spanning $d = 3$ to $d = 10$, independent and correlated inputs, mixed distribution types, and both variance-based and reliability-oriented (target) Shapley effects.

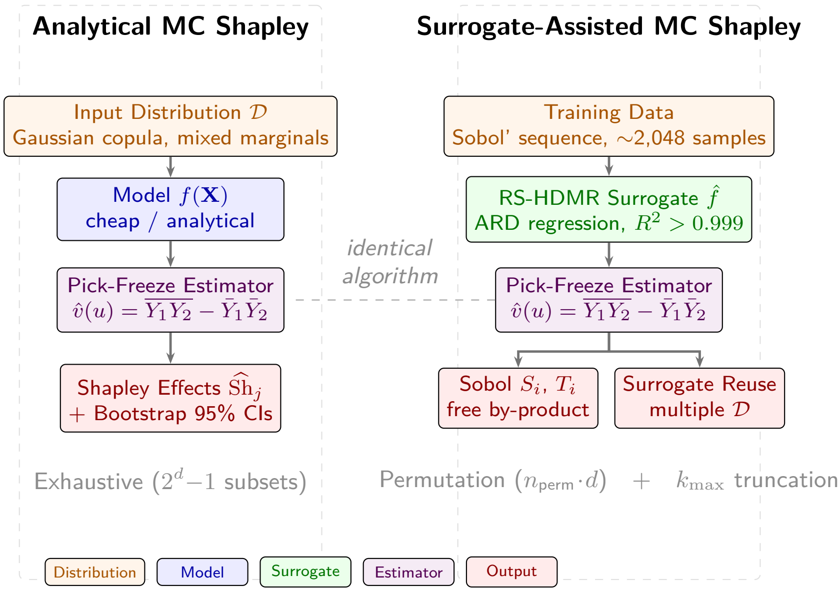

Figure 1: Two-path architecture of the MC Shapley estimation framework. The analytical path (left) evaluates the model f directly. The surrogate path (right) first trains an RS-HDMR surrogate $\hat{f}$ on Sobol' samples (∼2,048), then substitutes $\hat{f}$ for f in the estimation pipeline. Both paths use the same Pick-Freeze covariance estimator. The exhaustive method (left) enumerates all $2^d - 1$ subsets; the permutation method (right) scales to high dimensions, with $k_{\max}$ truncation exploiting bounded interaction orders in surrogate models. Shapley effects with bootstrap 95% confidence intervals are the primary output; Sobol S_i and T_i are obtained as a free by-product.

2. Theoretical Background

2.1. Shapley Effects

Let $\mathbf{X} = (X_1, \ldots, X_d)$ be a random vector of uncertain inputs with joint density $f_{\mathbf{X}}$, and let $f : \mathbb{R}^d \to \mathbb{R}$ be a square-integrable model. The value of a coalition $u \subseteq D = \{1, \ldots, d\}$ is the variance of the conditional expectation:

$$v(u) = \mathbb{V}[\mathbb{E}(f(\mathbf{X}) \mid \mathbf{X}_u)] \tag{1}$$

with $v(\emptyset) = 0$ and $v(D) = \mathbb{V}(f(\mathbf{X}))$. The Shapley effect for input i Owen (2014) is:

$$\mathrm{Sh}_i = \frac{1}{v(D)} \cdot \frac{1}{d} \sum_{u \subseteq D \setminus \{i\}} \binom{d-1}{|u|}^{-1} (v(u \cup \{i\}) - v(u)) \tag{2}$$

Shapley effects satisfy four axioms: efficiency ($\sum_i \mathrm{Sh}_i = 1$), symmetry, linearity, and the null-player property ($v(u \cup \{i\}) = v(u)$ for all $u \implies \mathrm{Sh}_i = 0$) Shapley (1953). Under input dependence, classical Sobol indices lose these properties, but Shapley effects retain them Iooss and Prieur (2019).

2.2. The Owen–Prieur Covariance Estimator

Owen & Prieur Owen and Prieur (2017) established that $v(u)$ can be expressed as a covariance:

$$v(u) = \mathrm{Cov}[f(\mathbf{X}), f(\mathbf{X}_u)] \tag{3}$$

where $\mathbf{X}_u$ is a random vector that shares the u-components with $\mathbf{X}$ but resamples the complement $\mathbf{X}_{-u}$ independently from the conditional distribution $P(\mathbf{X}_{-u} \mid \mathbf{X}_u)$. This leads to the **Pick-Freeze** (PF) estimator:

$$\hat{v}(u) = \frac{1}{N}\sum_{i=1}^{N} Y_1^{(i)} Y_2^{(i)} - \bar{Y}_1 \bar{Y}_2 \tag{4}$$

where $Y_1^{(i)} = f(\mathbf{X}^{(i)})$, $Y_2^{(i)} = f(\mathbf{X}_u^{(i)})$, $\bar{Y}_1 = \frac{1}{N}\sum Y_1^{(i)}$, and $\bar{Y}_2 = \frac{1}{N}\sum Y_2^{(i)}$. This requires $2N$ model evaluations per subset and is unbiased without the inner-loop bias corrections needed by double Monte Carlo estimators Demange-Chryst et al. (2023).

The PF estimator is particularly advantageous for binary failure indicators in reliability-oriented sensitivity analysis (ROSA). The double Monte Carlo (dMC) estimator applied to $\psi_t = \mathbf{1}(f(\mathbf{X}) > t)$ requires estimating $v(u) = \mathbb{V}[\mathbb{E}(\psi_t \mid \mathbf{X}_u)]$: when $p_f \ll 1$, the inner loop sees almost no failures, making $\mathbb{E}(\psi_t \mid \mathbf{X}_u) \approx 0$ and its variance ≈ 0, producing a Shapley effect of zero. The PF estimator estimates $\mathrm{Cov}[Y_1, Y_2]$ directly, and this covariance can be non-zero even when both Y_1 and Y_2 are almost always zero, provided rare failure events exhibit covariance structure.

2.3. Gaussian Copula Dependence Modelling

Sklar's theorem Sklar (1959) states that any d-dimensional joint distribution $F_{\mathbf{X}}$ with continuous marginals F_j can be expressed as:

$$F_{\mathbf{X}}(x_1, \ldots, x_d) = C(F_1(x_1), \ldots, F_d(x_d)) \tag{5}$$

where $C : [0,1]^d \to [0,1]$ is a copula. The Gaussian copula parameterises C through a latent multivariate normal:

$$C(u_1, \ldots, u_d) = \Phi_\Sigma(\Phi^{-1}(u_1), \ldots, \Phi^{-1}(u_d)) \tag{6}$$

where Φ_Σ is the joint CDF of $\mathcal{N}_d(\mathbf{0}, \Sigma)$ and Φ^{-1} is the inverse standard normal CDF. This decouples the problem into two independently tractable parts:

1. **Correlation modelling**: The $d \times d$ correlation matrix Σ controls all dependence structure.
2. **Marginal modelling**: Each variable j has its own distribution F_j, applied via $X_j = F_j^{-1}(\Phi(Z_j))$, where $\mathbf{Z} \sim \mathcal{N}_d(\mathbf{0}, \Sigma)$.

For LogNormal and Normal marginals—the most common cases in engineering applications—$\rho_{\text{physical}} \approx \rho_{\text{latent}}$, so the latent correlation matrix can be populated directly with the desired physical correlations.

2.4. Conditional Sampling in Latent Normal Space

The core computational primitive for MC Shapley estimation is drawing from $P(\mathbf{X}_{-u} \mid \mathbf{X}_u = \mathbf{x}_u)$. Under the Gaussian copula, this is performed analytically:

1. Map fixed variables to latent space: $Z_j = \Phi^{-1}(F_j(x_j))$ for $j \in u$.
2. Condition in latent normal space using the closed-form conditional multivariate normal distribution Eaton (1983):
$$\mathbf{Z}_{-u} \mid \mathbf{Z}_u = \mathbf{z}_u \sim \mathcal{N}(\boldsymbol{\mu}_{-u|u}, \boldsymbol{\Sigma}_{-u|u}) \tag{7}$$
where $\boldsymbol{\mu}_{-u|u} = \boldsymbol{\Sigma}_{-u,u}\boldsymbol{\Sigma}_{u,u}^{-1}\mathbf{z}_u$ and $\boldsymbol{\Sigma}_{-u|u} = \boldsymbol{\Sigma}_{-u,-u} - \boldsymbol{\Sigma}_{-u,u}\boldsymbol{\Sigma}_{u,u}^{-1}\boldsymbol{\Sigma}_{u,-u}$.
3. Map back: $X_j = F_j^{-1}(\Phi(Z_j))$ for $j \notin u$.

This is exact—no nearest-neighbour approximation, no Gibbs sampling. A Cholesky decomposition of $\boldsymbol{\Sigma}_{-u|u}$ is pre-computed once per subset u and reused for all N conditional draws.

3. The Distribution Class Interface

The MC Shapley machinery is distribution-agnostic. Any input model that implements a four-method contract can be plugged into the estimation pipeline:

- `self.d` (int) — number of input dimensions.
- `sample_joint(n)` $\to (n, d)$ array — n i.i.d. samples from $f_{\mathbf{X}}$.
- `sample_conditional(u, x)` $\to (d,)$ array — one sample from $P(\mathbf{X} \mid \mathbf{X}_u = \mathbf{x})$.
- `sample_conditional_batch(u, X)` $\to (N, d)$ array — N conditional samples, vectorised. This is the performance-critical method; `sample_conditional` typically delegates to it.

This interface decouples the estimation algorithm (which only calls these methods) from the distribution implementation (which handles the specifics of each marginal type and dependence structure).

3.1. Built-in Classes

Three distribution classes are provided with the package (Table 1):

Table 1: Built-in distribution classes with exact conditional sampling.

Class	Marginals	Conditional sampling method
`MultivariateNormal`	$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	Closed-form Cholesky (Eq. 7)
`GaussianCopulaUniform`	$\mathcal{U}[a_j, b_j]$	Latent normal conditioning, then Φ mapping
`TruncatedMVNormal`	$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with bounds	Gibbs sampler (tunable burn-in)

`MultivariateNormal` is the fastest, requiring a single Cholesky decomposition per subset. `GaussianCopulaUniform` maps uniform variables through Φ^{-1} into latent normal space, conditions analytically, then maps back via Φ. `TruncatedMultivariateNormal` uses Gibbs sampling (tunable via `joint_burn_in` and `cond_burn_in`), trading exactness for generality when truncation bounds are active.

3.2. Extensibility: The Custom Class Pattern

For input models beyond the built-in classes, users write a custom class following the same four-method contract. The Gaussian copula pattern from the case studies demonstrates the approach:

- **Init**: Store marginal specifications as (type, params) tuples; construct a `MultivariateNormal` in latent space with the desired correlation matrix Σ.
- **sample_joint**: Draw $\mathbf{Z} \sim \mathcal{N}_d(\mathbf{0}, \Sigma)$, then map $X_j = F_j^{-1}(\Phi(Z_j))$ with per-column dispatch on marginal type. Apply rejection rules if needed.
- **sample_conditional_batch**: Map fixed X to Z via $Z_j = \Phi^{-1}(F_j(X_j))$, call `MultivariateNormal.sa` in latent space, map result back.

Adding a new marginal type (e.g., Weibull) requires only a new branch in the forward ($X \to Z$) and backward ($Z \to X$) transform methods. The fire spread case study Bennett (2026c) implements this pattern for a 10-variable mixed Normal + LogNormal + scaled LogNormal distribution with truncation rules and $\rho(m_d, U) = -0.8$, achieving 100% acceptance rate and empirical correlation $\hat{\rho} = -0.799$.

4. MC Shapley Estimation Methods

4.1. Exhaustive Method

The exhaustive method evaluates every non-empty subset $u \subseteq D$, computing $\hat{v}(u)$ via the PF estimator (Eq. 4) and aggregating into Shapley effects via the binomial-weighted sum (Eq. 2).

For d inputs, there are $2^d - 1$ subsets. Each subset requires $2N$ model evaluations (one for the joint sample, one for the conditional sample), plus a full-variance evaluation, yielding a total of approximately $2N \cdot 2^d$ evaluations. For $d \leq 8$, this is tractable on modern hardware (e.g., $d = 6$ with $N = 10{,}000$ requires $\sim 1.3 \times 10^6$ evaluations, completed in under a minute for cheap analytical functions).

Algorithm 1 Exhaustive MC Shapley estimation

Require: Model f, distribution $\mathcal{D}$, sample size N, dimension d
Ensure: Shapley effects $\mathrm{Sh}_1, \ldots, \mathrm{Sh}_d$, bootstrap CIs

$\mathcal{V} \leftarrow \{\}$ ▷ Cache of (Y_1, Y_2) arrays per subset
for $k \leftarrow 1$ **to**d **do**
  for each subset u of size k **do**
    $\mathbf{X}_1 \leftarrow \mathcal{D}$.sample_joint$(N)$
    $\mathbf{Y}_1 \leftarrow f(\mathbf{X}_1)$
    $\mathbf{X}_2 \leftarrow \mathcal{D}$.sample_conditional_batch$(u, \mathbf{X}_1[:, u])$
    $\mathbf{Y}_2 \leftarrow f(\mathbf{X}_2)$
    $\mathcal{V}[u] \leftarrow (\mathbf{Y}_1, \mathbf{Y}_2)$
    $\hat{v}(u) \leftarrow \overline{\mathbf{Y}_1 \mathbf{Y}_2} - \bar{\mathbf{Y}}_1 \bar{\mathbf{Y}}_2$
  end for
end for
$\mathbf{X}_{\mathrm{full}} \leftarrow \mathcal{D}$.sample_joint$(N)$
$v(D) \leftarrow \mathbb{V}[f(\mathbf{X}_{\mathrm{full}})]$
for $i \leftarrow 1$ **to**d **do**
  $\mathrm{Sh}_i \leftarrow \frac{1}{d \cdot v(D)} \sum_{u \subseteq D \setminus \{i\}} \binom{d-1}{|u|}^{-1} (\hat{v}(u \cup \{i\}) - \hat{v}(u))$
end for
Bootstrap: Resample $\mathcal{V}$ with replacement B times; compute $\mathrm{Sh}_i^{(b)}$; return empirical $\alpha/2$ and $1 - \alpha/2$ quantiles.

Algorithm 1 stores the raw (Y_1, Y_2) arrays for each subset in a cache $\mathcal{V}$. This cache enables bootstrap confidence intervals to be computed by resampling within each subset's paired observations, preserving the covariance structure that the PF estimator depends on. By default, $B = 200$ bootstrap replications with $\alpha = 0.05$ for 95% confidence intervals.

4.2. Permutation Method

For $d > 8$, the exponential subset count of the exhaustive method becomes prohibitive. The permutation method Plischke et al. (2021) evaluates only those subsets encountered along n_{perm} random permutations of $\{1, \ldots, d\}$, with lazy caching: if a subset u appears in multiple permutations, its $v(u)$ value is evaluated once and reused. The number of unique subsets grows as $O(n_{\mathrm{perm}} \cdot d)$ rather than $O(2^d)$, making this method suitable for d up to ~ 50 at the cost of introducing small Monte Carlo error in the Shapley aggregation. Empirically, 100–500 permutations are sufficient for stable estimates.

Algorithm 2 Permutation-based MC Shapley (single permutation)

Require: Model f, distribution $\mathcal{D}$, sample size N, dimension d, permutation π
Ensure: Coalition values $\hat{v}(u)$ for all u appearing in π

$\mathbf{X}_1 \leftarrow \mathcal{D}.\text{sample\_joint}(N), \quad \mathbf{Y}_1 \leftarrow f(\mathbf{X}_1)$
for $k \leftarrow 1$ to$d$ **do**
  $u \leftarrow \{\pi_1, \ldots, \pi_k\}$ ▷ prefix of permutation
  if $u \notin \mathcal{V}$ **then**
    $\mathbf{X}_2 \leftarrow \mathcal{D}.\text{sample\_conditional\_batch}(u, \mathbf{X}_1[:, u])$
    $\mathbf{Y}_2 \leftarrow f(\mathbf{X}_2)$
    $\mathcal{V}[u] \leftarrow (\mathbf{Y}_1, \mathbf{Y}_2)$
    $\hat{v}(u) \leftarrow \overline{\mathbf{Y}_1\mathbf{Y}_2} - \bar{\mathbf{Y}}_1\bar{\mathbf{Y}}_2$
  end if
end for

Each permutation evaluates d new subsets, reusing the same joint sample $\mathbf{X}_1$ for all prefixes, reducing evaluations to $(n_{\text{perm}} + 1) \cdot d \cdot N$ rather than $2N \cdot d \cdot n_{\text{perm}}$.

4.3. Coalition Truncation ($k_{\max}$)

When the model's HDMR decomposition Rabitz and Al. (1999) has bounded interaction order—as in RS-HDMR surrogates where `polys=[10, 5]` implies only first- and second-order terms—subsets larger than the expansion order contribute nothing beyond what is already captured by their sub-coalitions. ShapleyX exploits this via a **coalition truncation** parameter $k_{\max}$:

- Only subsets of size $\leq k_{\max}$ are evaluated (plus the full set D, always needed for total variance).
- The Shapley accumulation skips coalition pairs where $v(u)$ or $v(u \cup \{i\})$ is unavailable, then renormalises to distribute the missing variance proportionally.
- This is **exact** when the model truly has no interactions above order $k_{\max}$.

Table 2: Coalition truncation savings. Subset count for exhaustive method with and without truncation.

d	Full ($2^d - 1$)	$k_{\max} = 2$	Factor	$k_{\max} = 3$	Factor
6	63	22	2.9×	42	1.5×
10	1,023	56	18×	176	5.8×
15	32,767	121	270×	576	57×
20	1,048,575	211	4,970×	1,351	780×
30	1.07×10^9	466	$2.3 \times 10^6\times$	4,526	$2.4 \times 10^5\times$

For a second-order RS-HDMR model at $d = 15$, coalition truncation evaluates 121 subsets instead of 32,767—a 270× reduction with zero accuracy loss. When using the high-level `model.get_mc_shapley()` interface, $k_{\max}$ is auto-detected from the surrogate's `polys` parameter ($k_{\max} = \text{len(polys)}$), requiring no user configuration.

4.4. Sobol Indices as a By-Product

A key computational advantage is that the $v(u)$ values are exactly the closed Sobol indices: $v(u) = \mathbb{V}[\mathbb{E}(f(\mathbf{X}) \mid \mathbf{X}_u)]$. First-order ($S_i$) and total-order ($T_i$) Sobol indices are therefore obtained without additional evaluations Sobol' (2001):

$$S_i = \frac{v(\{i\})}{v(D)}, \qquad T_i = 1 - \frac{v(D \setminus \{i\})}{v(D)} \tag{8}$$

These are extracted from the same $\mathcal{V}$ cache used for Shapley estimation at zero additional cost.

5. Surrogate Model Integration

5.1. RS-HDMR Surrogate Construction

ShapleyX's surrogate module Bennett and Roberts (2025) constructs a sparse Legendre polynomial chaos expansion (PCE) using the RS-HDMR framework Li et al. (2002). Input data are first transformed to the unit hypercube $[0,1]^d$ via min-max scaling. The function is approximated as:

$$f(\mathbf{x}) \approx \sum_{\boldsymbol{\alpha} \in \mathcal{A}} c_{\boldsymbol{\alpha}} \Psi_{\boldsymbol{\alpha}}(\mathbf{x}) \tag{9}$$

where $\Psi_{\boldsymbol{\alpha}}(\mathbf{x}) = \prod_{j=1}^{d} \psi_{\alpha_j}(x_j)$ are products of shifted Legendre polynomials Horner (1965), orthonormal on $[0,1]$. The `polys` parameter controls the truncation: e.g., `polys=[10, 5]` includes up to 10th-degree univariate terms and 5th-degree bivariate interactions. A single-element list `polys=[10]` constructs a full PCE with total-degree truncation $\|\boldsymbol{\alpha}\|_1 \leq 10$, yielding $\binom{d+10}{10}$ candidate basis functions spanning all interaction orders up to d.

Automatic Relevance Determination (ARD) regression Tipping (2001) prunes irrelevant basis terms via Bayesian evidence maximisation. Starting from an overcomplete basis set, ARD iteratively adds basis functions and removes those whose posterior precision tends to infinity (indicating zero contribution), yielding a sparse representation from typically thousands of candidates down to tens or hundreds of active terms. Bootstrap resampling of the training data provides confidence intervals for surrogate-derived sensitivity indices.

5.2. The Surrogate-to-MC-Shapley Pipeline

The surrogate integrates with the MC Shapley machinery as a drop-in replacement for the expensive model:

1. **Train**: Construct the RS-HDMR surrogate $\hat{f}$ on N_{train} independent Sobol' samples (typically 512–4,096).
2. **Validate**: Check R^2, MAE, and slope on held-out test data. $R^2 \gtrsim 0.999$ is typically required for reliable Shapley estimation.
3. **Estimate**: Call the standard MC Shapley routine with $\hat{f}$ as the model and the target (possibly correlated) distribution $\mathcal{D}$ as the input model. The surrogate evaluations are $O(10^4$–$10^6\times)$ faster than the original model.

4. **Auto-truncate**: $k_{\max}$ is auto-detected from the surrogate's `polys` parameter, eliminating the exponential subset cost for high-dimensional problems.

This decoupling of surrogate training and Shapley estimation has a key practical advantage: the same surrogate can be interrogated under **multiple** distribution scenarios (independent, correlated with different correlation structures) at negligible additional cost, enabling exploration of how dependence affects the sensitivity structure.

Under independent inputs, Sobol indices can be extracted directly from the squared Legendre coefficients grouped by variable label Sudret (2008), providing an independent validation path. Under correlated inputs, the coefficient-based extraction breaks down because the variance of interaction terms no longer factorises—hence the need for MC Shapley estimation on the surrogate.

6. Case Studies

We present three case studies that collectively validate the framework across different dimensions, distribution types, correlation structures, and quantities of interest. All analyses were performed using **ShapleyX** v0.5.1. Complete code and data are available in the **ShapleyX** repository Bennett (2026b).

6.1. Case Study 1: Correlated Ishigami Function ($d = 3$)

The Ishigami function Ishigami and Homma (1991) is the archetypical benchmark for sensitivity analysis:

$$f(\mathbf{X}) = \sin X_1 + a \sin^2 X_2 + b X_3^4 \sin X_1 \tag{10}$$

with $X_j \sim \mathcal{U}[-\pi, \pi]$, $a = 7$, $b = 0.1$. Under independent inputs, the analytical Shapley effects are $\mathrm{Sh} = (0.436, 0.442, 0.122)$ and first-order Sobol indices $S = (0.314, 0.442, 0)$ with one non-zero interaction $S_{1,3} = 0.244$ Sudret (2008).

We extend the standard benchmark by imposing a Gaussian copula correlation ρ between X_1 and X_3 and computing Shapley effects via the exhaustive method ($N = 5{,}000$, $2^3 - 1 = 7$ subsets, $2N{\cdot}7 = 70{,}000$ evaluations per ρ value). The sweep covers $\rho \in \{-0.9, -0.7, -0.5, -0.2, 0.0, +0.2, +0.5, +0$ reproducing the analysis of Iooss & Prieur Iooss and Prieur (2019).

RS-HDMR surrogate. A surrogate was trained on 256 independent Sobol' samples with `polys=[10, 5]` (105 candidate basis functions), achieving $R^2 = 0.999994$ and $\mathrm{MAE} = 0.007$. Under independent inputs, the surrogate-derived Shapley effects are $(0.430, 0.451, 0.119)$, within 0.006 of the analytical values and consistent with the analytical reference values from Table 3 of Bennett Bennett and Roberts (2025).

Table 3: Shapley effects and first-order Sobol indices for the Ishigami function at selected values of $\rho(X_1, X_3)$. Exhaustive MC Shapley with $N = 5{,}000$.

ρ	Sh_1	Sh_2	Sh_3	S_1	S_3
-0.9	0.196	0.587	0.217	0.299	0.342
-0.5	0.338	0.486	0.176	0.295	0.173
0.0	0.425	0.453	0.122	0.290	0.014
$+0.9$	0.199	0.582	0.220	0.310	0.341

Table 3 reveals the core insight that motivates Shapley effects for correlated inputs. At $\rho = 0$ (independence), $S_3 = 0.014$ correctly indicates that X_3 has negligible first-order influence, while $\mathrm{Sh}_3 = 0.122$ reflects its contribution through the $S_{1,3}$ interaction—the Shapley allocation distributes the interaction variance fairly between X_1 and X_3. At $|\rho| = 0.9$, however, S_1 and S_3 remain essentially unchanged ($S_1 \approx 0.30$, $S_3 \approx 0.34$) while the Shapley effects redistribute dramatically: Sh_1 drops from 0.425 to 0.196, Sh_2 rises from 0.453 to 0.587, and Sh_3 rises from 0.122 to 0.217. The Sobol indices fail to reflect the correlation entirely—$S_3 = 0.341$ at $\rho = -0.9$ gives the misleading impression that X_3 is the most important input, when in reality X_2 (which is uncorrelated with the others and contributes genuine first-order variance) dominates the Shapley allocation.

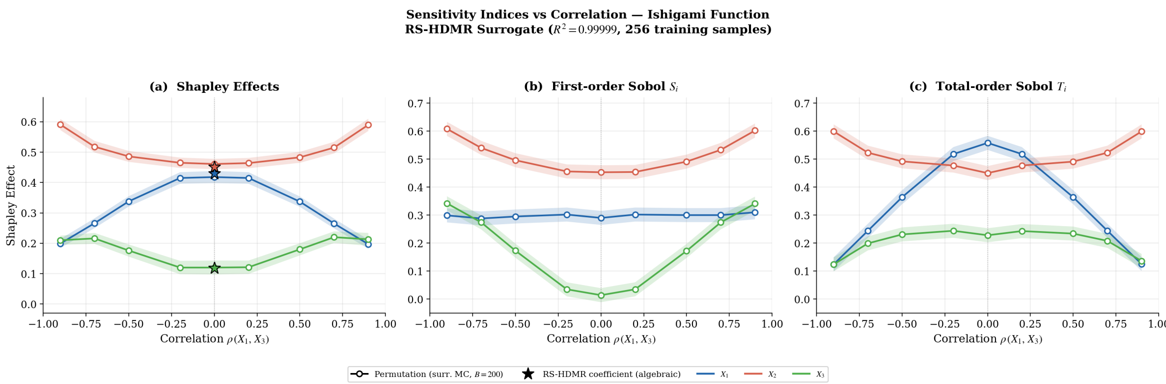

Figure 2: Shapley effects and Sobol indices for the Ishigami function as a function of correlation $\rho(X_1, X_3)$, computed using the RS-HDMR surrogate ($R^2 = 0.99999$, 256 training samples, `polys=[10, 5]`). (a) **Shapley effects**: solid lines with markers show permutation-based MC Shapley estimates ($n_{\mathrm{perm}} = 50$, $N = 5{,}000$, $B = 200$ bootstrap), with shaded bands indicating 95% confidence intervals. Star markers at $\rho = 0$ show RS-HDMR coefficient-derived (algebraic) Shapley effects. The classical U-shaped pattern observed by Iooss & Prieur (2019) is reproduced: Sh_2 dominates at the correlation extremes while Sh_1 dominates near independence. (b) **First-order Sobol indices** S_i: nearly constant across ρ, failing to reflect the dependence structure. (c) **Total-order Sobol indices** T_i: wider spread but still fail to fairly redistribute shared variance, with $S_i > T_i$ appearing under strong correlation—a well-known artefact of applying Sobol decomposition to dependent inputs.

The permutation method was validated against the exhaustive method at three anchor ρ

values (-0.9, 0.0, $+0.9$) with $n_{\text{perm}} = 50$. All Shapley effects agree within ± 0.009, confirming the permutation method's reliability for this problem.

6.2. Case Study 2: Cantilever Beam Structural Reliability ($d = 6$)

The cantilever beam problem from Demange-Chryst et al. Demange-Chryst et al. (2023) models a rectangular beam with two orthogonal tip forces. The maximum tip displacement is Zhou et al. (2014); Li et al. (2017):

$$\varphi(\mathbf{X}) = \frac{4L^3}{E l_X l_Y} \sqrt{\left(\frac{F_X}{l_X^2}\right)^2 + \left(\frac{F_Y}{l_Y^2}\right)^2} \tag{11}$$

with six inputs: three LogNormal (forces F_X, F_Y and elastic modulus E) and three correlated Normal (cross-section dimensions l_X, l_Y and beam length L). The correlation structure is $\rho(l_X, l_Y) = -0.55$, $\rho(L, l_X) = \rho(L, l_Y) = 0.45$. Failure occurs when $\varphi > t = 0.066$ m, with $p_f \approx 0.015$.

This problem exercises the framework on mixed marginals, a non-trivial correlation structure, and both variance-based and target (reliability-oriented) Shapley effects.

Analytical MC Shapley. Using a custom `GaussianCopulaMixed` class—Normal + LogNormal marginals with latent Gaussian copula conditioning—variance-based Shapley effects were computed exhaustively ($N = 10{,}000$, $2^6 - 1 = 63$ subsets, 1.26M evaluations). Target Shapley effects used $N = 100{,}000$ for higher precision on the binary failure indicator (estimated $p_f = 0.0153$).

PCE Surrogate. A full PCE of order 10 (`polys=[10]`) was trained on 2,048 independent uniform Sobol' samples. The total-degree truncation $\|\boldsymbol{\alpha}\|_1 \leq 10$ with $d = 6$ produced 8,007 candidate basis functions (all interaction orders up to 6-way). ARD regression achieved $R^2 = 0.99987$ and MAE $= 0.003$ m. The surrogate was then evaluated under the correlated mixed-distribution specification for MC Shapley estimation.

Table 4: Cantilever beam: variance-based and target Shapley effects. Analytical and surrogate estimates compared with Demange-Chryst et al. (2022) IS reference. All estimates include bootstrap 95% CIs ($B = 200$).

Variable	Sh_D (analyt.)	Sh_D (surr.)	T-Sh (analyt.)	T-Sh (surr.)	T-Sh (ref. IS)
F_X	0.109	0.111	0.150	0.144	0.146
F_Y	−0.007	−0.006	0.006	0.009	0.001
E	0.076	0.083	0.093	0.102	0.103
l_X	0.205	0.211	0.282	0.275	0.282
l_Y	0.364	0.360	0.262	0.263	0.254
L	0.254	0.242	0.207	0.208	0.214

The results (Table 4) demonstrate strong agreement across all three sources. Variance-based Shapley effects identify l_Y (0.364) and L (0.254) as the dominant inputs for deflection variability, consistent with their cubic and quartic roles in Eq. (11). The PCE surrogate reproduces all six target Shapley effects within ± 0.01 of the analytical values: for the most influential input l_X, the surrogate estimate is 0.275 vs. 0.282 analytical and 0.282 IS reference.

This close agreement, achieved with only 2,048 training samples, demonstrates that the full PCE—spanning all interaction orders through its 8,007 candidate basis functions—faithfully captures both the bulk sensitivity structure and the tail behaviour relevant to the binary failure indicator.

The PCE surrogate also yields Sobol first-order indices from its coefficient decomposition. These show $S_{l_Y} = 0.617$, $S_L = 0.243$, and $S_{l_X} = 0.208$. The substantial gap between Shapley and Sobol allocations for l_Y (0.360 vs. 0.617) reflects the correlation structure: when l_X and l_Y are negatively correlated, the Sobol first-order index over-allocates variance to l_Y, whereas the Shapley effect fairly distributes the shared component.

6.3. Case Study 3: Wing Weight Function ($d = 10$)

The wing weight function Forrester et al. (2008) models the weight W (lb) of a light aircraft wing:

$$W = 0.036 S_w^{0.758} W_{fw}^{0.0035} \left(\frac{A}{\cos^2 \Lambda}\right)^{0.6} q^{0.006} \lambda^{0.04} \left(\frac{100 t_c}{\cos \Lambda}\right)^{-0.3} (N_z W_{dg})^{0.49} + S_w W_p \tag{12}$$

with 10 independent Uniform inputs spanning five orders of magnitude (from $W_p \sim 0.05$ to $W_{dg} \sim 2{,}100$ lb). This case study demonstrates the framework at larger dimension, validates the surrogate against published OpenTURNS Sobol indices Baudin et al. (2015), and explores how correlated inputs fundamentally reshape the importance ranking.

Analytical MC Shapley. Exhaustive estimation with $N = 10{,}000$ ($2^{10} - 1 = 1{,}023$ subsets, 20.46M evaluations) using `GaussianCopulaUniform`.

RS-HDMR Surrogate. Trained on 2,048 Sobol' samples with `polys=[8, 6, 4]` (9,380 candidate basis functions—80 first-order, 1,620 second-order, 7,680 third-order). Achieved $R^2 = 0.9999998$, MAE $= 0.014$ lb. Sobol indices from the surrogate coefficients match the published OpenTURNS reference values, with exact agreement for the dominant variable N_z ($S_{N_z} = 0.413$ vs. 0.413 reference).

Correlated Extension. Four physically-motivated correlation pairs were introduced: $\rho(S_w, A) = 0.6$, $\rho(S_w, W_{dg}) = 0.7$, $\rho(A, \lambda) = -0.4$, $\rho(\Lambda, q) = 0.3$. Surrogate-based MC Shapley was then run under this correlated specification.

Table 5: Wing weight: Shapley effects under independent and correlated inputs. All estimates from surrogate-based MC Shapley ($N = 10{,}000$, exhaustive). Published OpenTURNS Sobol indices for validation.

Variable	**Sh (indep.)**	**Sh (corr.)**	**Rank (indep. → corr.)**	S_i **(OT ref.)**
N_z	0.415	0.228	1 → 2	0.413
A	0.223	0.214	2 → 3	0.228
t_c	0.143	0.075	3 → 5	0.135
S_w	0.132	**0.337**	4 → **1**	0.130
W_{dg}	0.085	0.137	5 → 4	0.088

The introduction of correlations (Table 5) fundamentally redistributes the Shapley allocation. Wing area S_w rises from rank 4 (Sh = 0.132) to rank 1 (Sh = 0.337)—a 2.6× increase—because it participates in two strong correlations ($\rho = 0.6$ with A, $\rho = 0.7$ with W_{dg}). The Shapley allocation fairly distributes the shared variance contributions, and S_w, as the "hub" through which multiple correlated effects propagate, absorbs a larger share. Meanwhile, N_z drops from 0.415 to 0.228 despite remaining uncorrelated with any variable: the total variance increases from 2,341 to 3,021 lb^2 under the correlated specification, diluting N_z's relative share.

This ranking reversal has direct aerospace design implications. Under independent inputs, the Sobol analysis identifies N_z as dominant ($S_{N_z} = 0.413$), suggesting that tightening the load factor tolerance is the most effective way to reduce wing weight uncertainty. Under realistic correlations, S_w emerges as the critical parameter—a conclusion invisible to classical Sobol indices.

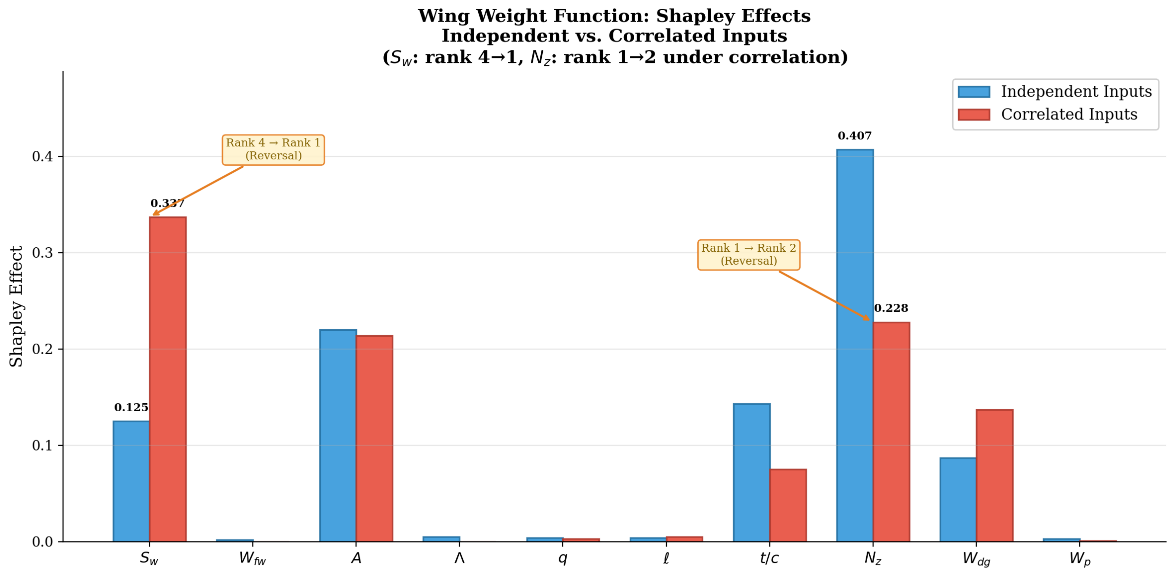

Figure 3: Shapley effects for the 10-variable wing weight function under independent (blue) and correlated (red) inputs. The introduction of four physically-motivated correlation pairs ($\rho(S_w, A) = 0.6$, $\rho(S_w, W_{dg}) = 0.7$, $\rho(A, \lambda) = -0.4$, $\rho(\Lambda, q) = 0.3$) fundamentally reorders the importance hierarchy: wing area S_w rises from rank 4 (Sh = 0.132) to rank 1 (Sh = 0.337), while load factor N_z drops from rank 1 (Sh = 0.415) to rank 2 (Sh = 0.228). All estimates from surrogate-based MC Shapley ($N = 10{,}000$, exhaustive, $B = 200$).

7. Computational Characteristics

Table 6 summarises the computational cost for the three case studies. For analytical functions, MC Shapley with $N = 10{,}000$ and exhaustive subsets completes in 5–20 seconds. For surrogate-based estimation, the dominant cost is surrogate training (2,048 evaluations, 130–170 seconds for ARD on 10^3–10^4 basis functions). Once trained, surrogate-based MC Shapley runs in comparable time to the analytical case, and the surrogate can be reused for multiple distribution scenarios at no additional training cost.

Table 6: Computational cost summary. All timings on a single CPU core.

Case Study	d	Method	Evaluations	Time	Key cost
Ishigami	3	Exhaustive, $N = 5{,}000$	70,000	~ 5 s	MC Shapley
Cantilever (analyt.)	6	Exhaustive, $N = 100{,}000$	12.6M	~ 2 min	MC Shapley
Cantilever (surr.)	6	PCE training	2,048	~ 170 s	ARD regression
		Surrogate MC Shapley	12.5M	~ 3 min	Surrogate evals (cheap)
Wing weight (analyt.)	10	Exhaustive, $N = 10{,}000$	20.5M	~ 20 s	MC Shapley
Wing weight (surr.)	10	RS-HDMR training	2,048	~ 130 s	ARD regression
		Surrogate MC Shapley	3.5M	~ 20 s	Surrogate evals (cheap)

For expensive models where a single evaluation takes minutes to hours, the surrogate-based workflow is transformative: 2,048 model evaluations for surrogate construction, followed by sensitivity analysis under multiple distribution scenarios for negligible additional cost. The auto-detected $k_{\max}$ from the surrogate's `polys` parameter eliminates the exponential subset cost: for the wing weight surrogate ($k_{\max} = 3$), the 1,023 full subsets are reduced to 176, though the full enumeration was used here for validation.

8. Limitations and Future Work

The current framework has several limitations that define directions for future development:

1. **No importance sampling.** For reliability-oriented sensitivity analysis with extreme failure probabilities ($p_f \lesssim 10^{-3}$), standard MC Shapley estimation is limited by the rarity of failure events. The surrogate workflow partially addresses this (once built, large N is cheap), but for models where surrogates are insufficiently accurate in the failure region, importance sampling—as developed by Demange-Chryst et al. Demange-Chryst et al. (2023)—is essential. Integrating IS estimators into the distribution class interface is a priority.
2. **No given-data mode.** The framework requires the ability to draw fresh conditional samples from the joint distribution. It cannot operate on a fixed pre-existing dataset, a capability provided by the nearest-neighbour approximation in the Demange-Chryst given-data framework Demange-Chryst et al. (2023).
3. **Gaussian copula only.** The current implementation assumes elliptical dependence through a Gaussian copula. For problems with tail dependence (e.g., financial risk, extreme environmental events), alternative copulas (Student-t, Gumbel, Clayton) or vine copula constructions would be needed. The distribution class interface is designed to accommodate these extensions.
4. **No double Monte Carlo.** Only the covariance-based PF estimator is implemented. The dMC estimator, while requiring bias correction under IS, provides an alternative estimation path that may be preferable in some regimes.
5. **Permutation method bias.** For small n_{perm}, the permutation method introduces Monte Carlo error in the Shapley aggregation. Systematic bias characterisation as a function of d and n_{perm} remains future work.
6. **Surrogate extrapolation risk.** Surrogates trained on independent uniform samples and deployed under correlated distributions may extrapolate into regions poorly represented in the training data. Quantifying this extrapolation error and developing training strategies that cover the correlated input space are open problems.

9. Conclusion

We have presented a unified Monte Carlo framework for Shapley effect estimation with correlated inputs, implemented in the **ShapleyX** Python package. The framework's key contributions are:

1. **A modular distribution class interface** that decouples correlation modelling (Gaussian copula in latent normal space) from marginal modelling (arbitrary inverse CDF transformations), with built-in support for Normal, LogNormal, Uniform, and truncated Normal marginals, and a documented extensibility pattern.
2. **Coalition truncation** ($k_{\max}$) that exploits bounded interaction orders in surrogate models to achieve 270×–5,000× reductions in subset enumeration, with auto-detection from surrogate metadata.
3. **Surrogate integration** that couples sparse Legendre PCE surrogates (RS-HDMR with ARD regression) with the MC Shapley machinery, enabling accurate sensitivity analysis at negligible cost for expensive models.
4. **Sobol indices as a free by-product** from the same coalition values $v(u)$.

Three case studies validate the framework across $d = 3$ to $d = 10$, independent and correlated inputs, mixed distribution types, and both variance-based and target Shapley effects. The surrogate-based workflow reproduces analytical and published reference values within confidence intervals, and the correlated-input analyses reveal importance ranking reversals invisible to classical Sobol indices—demonstrating the practical value of game-theoretic sensitivity measures for models with dependent inputs.

Code Availability

All code and data are available in the **ShapleyX** repository Bennett (2026b) (`https://github.com/frbennett/shapleyx`), under the `Examples/` directory. The RS-HDMR surrogate methodology is described in detail in Bennett Bennett and Roberts (2025).

Acknowledgments

The author thanks Art B. Owen and Clémentine Prieur for the foundational Pick-Freeze estimator, and Julien Demange-Chryst for making the target Shapley reference code publicly available.

References

A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. *Global Sensitivity Analysis: The Primer*. Wiley, 2008.

I. M. Sobol'. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55(1–3):271–280, 2001.

A. Saltelli. Making best use of model evaluations to compute sensitivity indices. *Computer Physics Communications*, 145(2):280–297, 2002.

M. Lamboni, B. Iooss, A.-L. Popelin, and F. Gamboa. Derivative-based global sensitivity measures: General links with Sobol' indices and numerical tests. *Mathematics and Computers in Simulation*, 87:45–54, 2013.

E. Borgonovo. A new uncertainty importance measure. *Reliability Engineering & System Safety*, 92(6):771–784, 2007.

F. Pianosi and T. Wagener. A simple and efficient method for global sensitivity analysis based on cumulative distribution functions. *Environmental Modelling & Software*, 67:1–11, 2015.

S. Da Veiga. Global sensitivity analysis with dependence measures. *Journal of Statistical Computation and Simulation*, 85(7):1283–1305, 2015.

G. Chastaing, F. Gamboa, and C. Prieur. Generalized Hoeffding-Sobol decomposition for dependent variables—application to sensitivity analysis. *Electronic Journal of Statistics*, 6:2420–2448, 2012.

T. A. Mara and S. Tarantola. Variance-based sensitivity indices for models with dependent inputs. *Reliability Engineering & System Safety*, 107:115–121, 2012.

S. Kucherenko, S. Tarantola, and P. Annoni. Estimation of global sensitivity indices for models with dependent variables. *Computer Physics Communications*, 183(4):937–946, 2012.

A. B. Owen. Sobol' indices and Shapley value. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):245–251, 2014.

E. Song, B. L. Nelson, and J. Staum. Shapley effects for global sensitivity analysis: Theory and computation. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):1060–1083, 2016.

L. S. Shapley. A value for n-person games. In *Contributions to the Theory of Games II*, pp. 307–317. Princeton University Press, 1953.

A. B. Owen and C. Prieur. On Shapley value for measuring importance of dependent inputs. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):986–1002, 2017.

B. Iooss and C. Prieur. Shapley effects for sensitivity analysis with correlated inputs: Comparisons with Sobol' indices, numerical estimation and applications. *International Journal for Uncertainty Quantification*, 9(5):493–514, 2019.

B. Broto, F. Bachoc, and M. Depecker. Variance reduction for estimation of Shapley effects and adaptation to unknown input distribution. *SIAM/ASA Journal on Uncertainty Quantification*, 8(2):693–716, 2020.

E. Plischke, G. Rabitti, and E. Borgonovo. Computing Shapley effects for sensitivity analysis. *SIAM/ASA Journal on Uncertainty Quantification*, 9(4):1411–1437, 2021.

J. Demange-Chryst, F. Bachoc, and J. Morio. Shapley effect estimation in reliability-oriented sensitivity analysis with correlated inputs by importance sampling. *International Journal for Uncertainty Quantification*, 13(3), 2023. arXiv:2202.12679v2.

M. Il Idrissi, V. Chabridon, and B. Iooss. Developments and applications of Shapley effects to reliability-oriented sensitivity analysis with correlated inputs. *Environmental Modelling & Software*, 143:105115, 2021.

A. Stein and T. Singh. Shapley effect estimation using polynomial chaos. arXiv:2301.04315, 2023.

F. Bennett. **ShapleyX**: A Python package for global sensitivity analysis with RS-HDMR. *Journal of Open Source Software* (submitted), 2026.

F. Bennett. **ShapleyX** repository, v0.5.1. `https://github.com/frbennett/shapleyx`, 2026.

F. R. Bennett and M. E. 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* (submitted), 2025.

F. Bennett. Shapley effects for rare-event sensitivity analysis: A case study on the Rothermel fire spread model. Technical report, Department of Environment, Tourism, Science and Innovation, Queensland Government, 2026.

A. Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8:229–231, 1959.

M. L. Eaton. *Multivariate Statistics: A Vector Space Approach.* Wiley, 1983.

H. Rabitz, Ö. F. Aliş, J. Shorter, and K. Shim. Efficient input-output model representations. *Computer Physics Communications*, 117(1–2):11–20, 1999.

G. Li, S.-W. Wang, and H. Rabitz. Practical approaches to construct RS-HDMR component functions. *Journal of Physical Chemistry A*, 106(37):8721–8733, 2002.

J. M. Horner. Generalized Rodrigues formula solutions for certain linear differential equations. *Transactions of the American Mathematical Society*, 115:31–42, 1965.

M. E. Tipping. Sparse Bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1:211–244, 2001.

B. Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93(7):964–979, 2008.

T. Ishigami and T. Homma. An importance quantification technique in uncertainty analysis for computer models. In *Proceedings of the First International Symposium on Uncertainty Modeling and Analysis*, pp. 398–403, 1991.

C. Zhou, Z. Lu, L. Zhang, and J. Hu. Moment independent sensitivity analysis with correlations. *Applied Mathematical Modelling*, 38(19–20):4885–4896, 2014.

B. Li, L. Zhang, X. Zhu, X. Yu, and X. Ma. Reliability analysis based on a novel density estimation method for structures with correlations. *Chinese Journal of Aeronautics*, 30(3):1021–1030, 2017.

A. Forrester, A. Sóbester, and A. Keane. *Engineering Design via Surrogate Modelling: A Practical Guide.* Wiley, 2008.

OpenTURNS Consortium. OpenTURNS: Reference Guide. `https://openturns.github.io/openturns/latest`, 2024.

T. Goda. A simple algorithm for global sensitivity analysis with Shapley effects. *Reliability Engineering & System Safety*, 213:107702, 2021.

M. Herin, M. Il Idrissi, V. Chabridon, and B. Iooss. Proportional marginal effects for global sensitivity analysis. *SIAM/ASA Journal on Uncertainty Quantification*, 12(2):667–692, 2024.

G. Blatman and B. Sudret. Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, 230(6):2345–2367, 2011.

Affiliation:

Frederick R. Bennett
Department of Environment, Tourism, Science and Innovation
Queensland Government
E-mail: frederick.bennett@detsi.qld.gov.au