26th International Congress on Modelling and Simulation, Adelaide, SA, Australia, 30 Nov–4 Dec 2025
mssanz.org.au/modsim2025

ShapleyX – a Python package for the efficient estimation of Shapley effects and Sobol' indices from tabular data

FR Bennett

Queensland Government Department of the Environment, Tourism, Science and Innovation, Science Division, Brisbane, Queensland
Email: frederick.bennett@detsi.qld.gov.au

Abstract: The increasing complexity and dimensionality of computational models across various scientific and engineering disciplines necessitate robust and interpretable sensitivity analysis (SA) techniques. Traditional SA methods often fall short in adequately addressing models characterised by non-linear responses, strong interactions, and interdependencies among input variables. This paper introduces ShapleyX, a novel Python package designed to overcome these limitations by providing an efficient and accurate framework for global sensitivity analysis, with a particular emphasis on computing Shapley Effects through the strategic application of Random Sampling High-Dimensional Model Representations (RS-HDMR) and Relevance Vector Machine (RVM) regression.

The primary motivation behind ShapleyX stems from the critical need for a more nuanced understanding of how individual and interacting input variables contribute to model output variance. While widely used, Sobol' indices, for instance, struggle to fairly apportion variance in the presence of correlated inputs, often leading to ambiguous interpretations. Shapley Effects, rooted in cooperative game theory, offer a theoretically sound alternative by providing a unique and equitable distribution of the total output variance among all input variables. This is achieved by systematically evaluating the marginal contribution of each input across all possible coalitions of variables, thereby accounting for complex interaction effects without over- or under-attributing importance.

ShapleyX's methodological foundation is built upon two powerful techniques. Firstly, the package leverages RS-HDMR, a meta-modelling approach that approximates complex, high-dimensional functions as a sum of lower-dimensional component functions. This decomposition is particularly advantageous because the coefficients of the RS-HDMR expansion can be directly mapped to various sensitivity indices, including Sobol' and Shapley Effects, significantly reducing the computational burden associated with Monte Carlo-based SA. Unlike intrusive methods that require direct access to or modification of the underlying model, RS-HDMR operates non-intrusively, requiring only one set of input-output samples. Secondly, to efficiently and accurately estimate the coefficients of the RS-HDMR expansion, ShapleyX employs Relevance Vector Machine (RVM) regression. RVM, a Bayesian sparse learning technique, is adept at identifying and retaining only the most relevant basis functions, thereby promoting model sparsity and enhancing generalisation capabilities, especially in scenarios where the number of potential basis functions far exceeds the number of available samples. This addresses the ill-posed problem often encountered in high-dimensional regression, ensuring robust coefficient estimation even with relatively limited data.

Key innovations within ShapleyX include its seamless integration of these advanced computational methods into a user-friendly Python environment. The package allows for flexible basis function construction using orthonormal Legendre polynomials, and offers various regression strategies, including a cross-validated Automatic Relevance Determination (ARD) method (ard_cv) for optimal model selection. This combination ensures high accuracy in estimating sensitivity measures, even for models with strong non-linearities and parameter interdependencies. Furthermore, ShapleyX incorporates bootstrap resampling to provide robust confidence intervals for the computed sensitivity indices, enhancing the reliability and interpretability of the results.

The practical applications of ShapleyX are extensive, spanning various domains where understanding model behaviour and uncertainty is paramount. From environmental modelling, such as hydrological and gully erosion simulations, to complex engineering and scientific systems, ShapleyX facilitates informed decision-making by providing clear, quantitative insights into input variable importance.

Keywords: *Shapley Effects, Global Sensitivity Analysis, High Dimensional Model Representation (HDMR), Relevance Vector Machine (RVM)*

1. INTRODUCTION

In the courtroom of model interpretation, Shapley Effects are the impartial judge that weighs contributions fairly. Unlike classical variance-based measures that can misattribute importance when inputs interact or correlate, Shapley Effects compute each variable's marginal contribution averaged across all coalitions, producing a unique, axiomatic allocation of output variance.

The proliferation of sophisticated computational models across diverse scientific and engineering disciplines, from climate prognostics to financial forecasting and environmental management, has underscored the critical importance of understanding their underlying dynamics and uncertainties. As these models grow in complexity, often encompassing numerous interacting parameters and non-linear relationships, the task of robustly assessing the influence of input variables on model outputs becomes increasingly challenging. Sensitivity analysis (SA) serves as an valuable tool in this endeavour, quantifying the relative importance of model inputs in determining output variability and informing decisions related to model development, calibration, and application (Wang et al. 2023; Jakeman et al. 2006).

Traditional SA approaches, such as local methods based on partial derivatives, offer computational efficiency but provide only a localised understanding of parameter influence, often failing to capture global behaviours, especially in the presence of strong non-linearities or interactions (Cacuci et al. 2003). Global SA methods, in contrast, explore the entire parameter space, yielding more reliable sensitivity measures. Variance-based techniques, notably those derived from Sobol' indices (Sobol 2001), have gained widespread acceptance for their ability to decompose output variance into contributions from individual inputs and their interactions. However, a significant limitation of classical Sobol' indices arises when input variables are interdependent or correlated; in such scenarios, the interpretability of individual contributions becomes ambiguous, and interaction effects can be misattributed or overcounted (Il Idrissi et al. 2021). This challenge highlights a fundamental gap in existing SA methodologies: the fair and unambiguous attribution of input variable importance in complex, high-dimensional settings with correlated inputs.

This paper introduces ShapleyX, a novel Python package developed to address these limitations by offering a robust, model-agnostic approach to global sensitivity analysis. ShapleyX leverages the powerful concept of Shapley Effects, derived from cooperative game theory, to fairly distribute output variance among inputs, even in the presence of strong interactions and dependencies. This is achieved by systematically evaluating the marginal contribution of each input variable across all possible permutations of variables. To ensure computational efficiency and scalability for high-dimensional model representations, ShapleyX integrates Shapley Effects with Random Sampling High-Dimensional Model Representations (RS-HDMR) and employs Relevance Vector Machine (RVM) regression for sparse and accurate coefficient estimation. The package provides accurate estimates of sensitivity measures without requiring intrusive modifications to the underlying model, making it a versatile tool for a broad scientific audience.

In summary ShapleyX offers a robust and computationally efficient framework for global sensitivity analysis. Its core contributions lie in its ability to:

- **Fairly Attribute Variance:** By leveraging Shapley Effects, ShapleyX provides a unique and equitable distribution of the total output variance among input variables, overcoming the limitations of traditional methods like Sobol' indices in handling correlated inputs and complex interactions.
- **Handle High Dimensionality:** The combination of RS-HDMR for efficient meta-modelling and RVM for sparse coefficient estimation enables ShapleyX to scale effectively to high-dimensional problems, reducing the need for extensive model evaluations.
- **Operate Non-Intrusively:** As a model-agnostic tool, ShapleyX requires only input-output data, making it applicable to a wide range of "black-box" models without requiring internal modifications.
- **Provide Robust Estimates:** The integration of bootstrap resampling allows for the quantification of uncertainty in sensitivity measures, enhancing the reliability and interpretability of the results.

2. THEORETICAL BACKGROUND

2.1. Variance-Based Sensitivity Indices

Let $Y = f(\mathbf{x})$ be a function with input vector $\mathbf{x} \in [0,1]^k$. The HDMR decomposition expresses Y as a sum of functions of increasing order:

$$Y = f_0 + \sum_{i} f_i\,(x_i) + \sum_{i<j} f_{ij}\,(x_i, x_j) + \cdots \tag{1}$$

The partial variances are defined as:

$$D_i = \int f_i^2(x_i)\, dx_i, \quad D_{ij} = \int f_{ij}^2(x_i, x_j)\, dx_i dx_j, \quad \text{etc} \tag{2}$$

The total variance $D = \sum_u D_u$. Sobol' indices are normalized partial variances:

$$S_i = \frac{D_i}{D}, \quad S_{ij} = \frac{D_{ij}}{D}, \quad S_i^T = \frac{D_i^T}{D}, \quad D_i^T = \sum_{u \ni i} D_u \tag{3}$$

2.2. Shapley Effects and Interaction Quantification

Shapley effects generalise the idea of "contribution" by averaging marginal contributions over all coalitions:

$$S_h^i = \sum_{u \subseteq A\{i\}} \frac{(|u|)!\,(k - |u| - 1)!}{k!} \big(c(u \cup \{i\}) - c(u)\big) \tag{4}$$

where $c(u)$ is the variance of the conditional expectation given input set u, $|u|$ is the cardinality of u and A is the full set of variables. When linked to Sobol' indices:

$$S_h^i = \sum_{u \subseteq A\{i\}} \frac{S_{i,u}}{|u| + 1} \tag{5}$$

where the numerator term $S_{i,u} = S_{\{i\} \cup u}$, is the normalised partial variance for the exact combination of i and the variables in u. This ensures $S_i \le S_h^i \le S_i^T$, with fair distribution of interaction contributions.

2.3. RS-HDMR and Non-Intrusive Regression

In RS-HDMR, the HDMR functions are approximated using a Legendre polynomial basis. The model becomes:

$$Y \approx \sum_d c_d\, \phi_d(\mathbf{x}) \tag{6}$$

where ϕ_d are uni- and multivariate orthogonal basis functions and the expansion coefficients $\{c_\alpha\}$ are estimated via regression analysis. Traditional least squares fail when the number of basis functions is greater than the number of training samples necessitating regularisation. RVM, a Bayesian sparse regression technique built on the principles of automatic relevance determination (ARD) (MacKay 1996), naturally solves this by inducing sparsity through precision hyperparameters ($\boldsymbol{\alpha}$) introduced by placing a zero-mean Gaussian prior on each of the expansion coefficients.

$$p(\, c_d \mid \alpha_d\,) = \mathcal{N}(\, c_d \mid 0, \alpha_d^{-1}\,) \tag{7}$$

The RegressionARD class in the ShapleyX implements an iterative, fast marginal likelihood maximisation algorithm to optimise the precision hyperparameters (Tipping and Faul 2003). During optimisation, $\alpha_d \to \infty$ for irrelevant features forcing $c_d \to 0$ while relevant features retain finite α_d, keeping their weights non-zero. Basis functions for which $c_d = 0$ are essentially discarded from the basis set.

The resulting metamodel can be thought of as a special type of sparse polynomial chaos expansion with expansion coefficients that can be mapped directly to Sobol' indices and Shapley effects for sensitivity analysis.

3. PACKAGE DESIGN AND IMPLEMENTATION (SHAPLEYX)

The methodology described is implemented in ShapleyX, an open-source Python package. ShapleyX is designed for ease of use and integrates the entire workflow from data input to sensitivity analysis results.

3.1. Architecture and Core Functionalities

ShapleyX is implemented in Python with modular components. The transformation module rescales primitives to $[0,1]$ and records ranges for surrogate use. The legendre module builds primitive and product basis functions and constructs a design matrix. The regression module offers several solvers: a custom RegressionARD (RVM-style) implementation optimised for RS-HDMR bases, OMP and OMP_CV wrappers from sklearn, and sklearn's ARDRegression. The indicies module converts coefficients to Sobol' indices, computes Shapley and

Shapley-Owen interaction measures, and packages results. The predictor module fits a fast Ridge surrogate on retained basis terms for use in PAWN/delta/h indices and in bootstrapping. The resampling module implements bootstrap quantile computations. Utilities provide plotting and diagnostics.

3.2. API and user workflow

The primary user entry point is the rshdmr class. The user supplies a path to a CSV file or DataFrame with columns for input variables and a column 'Y' for responses and specifies polynomial orders for first and second (and optionally higher) orders via the polys argument. Calling run_all() executes: input transformation, basis construction, sparse regression, model diagnostics (explained variance, MAE, MSE), Sobol' and Shapley computation, surrogate fitting, and optional bootstrap resampling to produce confidence. The package also supplies methods to compute PAWN (Pianosi and Wagener 2015, 2018), Delta Moment-Independent Measure (Borgonovo 2007; Plischke et al. 2013) and Kullback-Leibler (KL) divergence-based H indices using the surrogate model, and to export the pruned basis for external predictions.

3.3. Algorithmic scaling and practical considerations

The basis size grows combinatorially with input dimension and polynomial orders. For many practical systems, truncation at second order with modest polynomial degrees (e.g., 5–10 for first order, 3–6 for second order) captures most variance. The RVM/ARD solver typically selects a small active set s of basis functions (sparse), and the computational cost becomes dominated by operations on s rather than the full basis. In principle, memory can be conserved by constructing basis columns on demand and by pruning via a variance-threshold prior to fitting. For very high dimensional problems, pre-screening (e.g. using the built in rapid PAWN analysis tool (Pianosi and Wagener 2018)) is recommended to reduce input set before full RS-HDMR expansion (Ziehn & Tomlin, 2009).

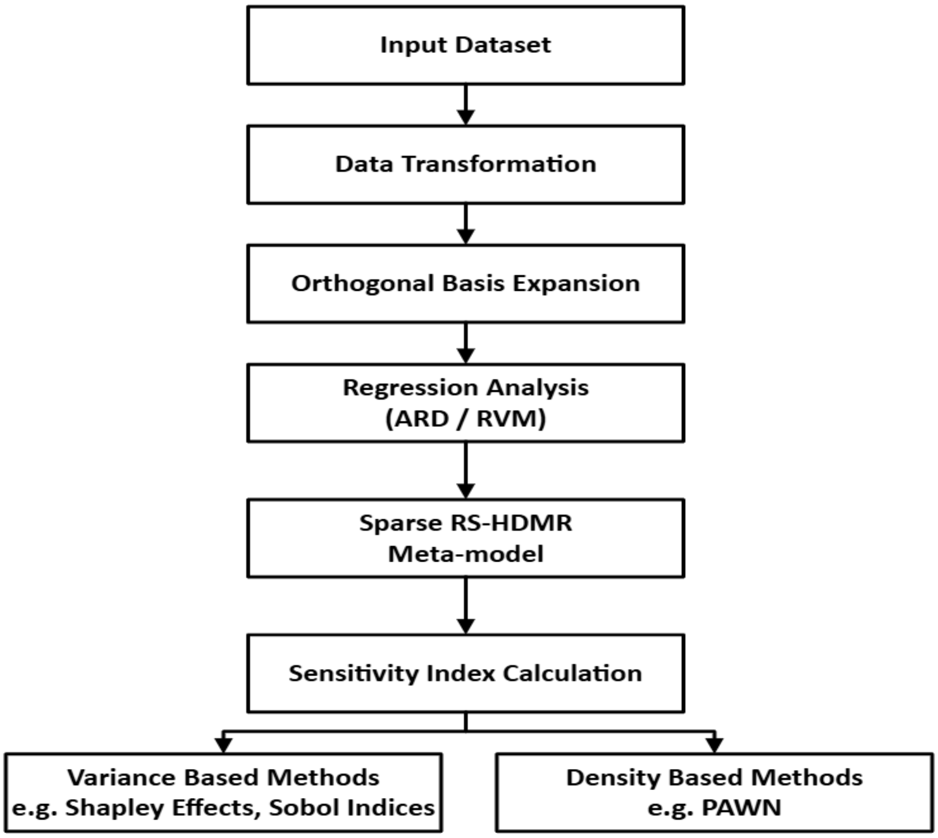

Figure 1 Data flow through the ShapleyX architecture.

4. CASE STUDY – OWEN'S PRODUCT FUNCTION

This section presents a case study demonstrating the application of the ShapleyX framework for estimating Shapley values using the Owen product function. This benchmark function, as detailed by (Owen 2014), is particularly valuable for stress-testing sensitivity estimators in high-dimensional settings due to its hierarchical variable importance and analytically tractable Sobol' indices.

The Owen product function is defined as:

$$f(\mathbf{x}) = \prod_{j=1}^{d} \left(\mu_j + \tau_j g_j(x_j)\right), \quad x_j \in [0,1],$$

where $g_j(x_j) = \sqrt{12}\left(x_j - \frac{1}{2}\right)$, $\mu_j = 1$ for all j, and $\tau_j = [4,4,2,2,1,1]/4$ for $d = 6$. This configuration creates a hierarchy of variable importance, with diminishing contributions from inputs x_1 to x_6.

A key characteristic of this function is that its analytical partial variances are given by $\sigma_u^2 = \prod_{j\in u} \tau_j^2 \prod_{j\notin u} \mu_j^2$. This allows for direct comparison between estimated and true Shapley values. The function is also notable because its first and second-order partial functions of the HDMR decomposition capture only about 83% of the variance indicating that higher order interactions contribute nontrivially to the functions total variance making it a challenging test case for sensitivity analysis methods.

The following code block shows the basic Python syntax for running a ShapleyX job.

```python
# Step 1
# Load the ShapleyX package
from shapleyx import rshdmr

# Step 2
# Create an instance of the class rshdmr
model = rshdmr('owen_data.csv',
               polys = [8, 6, 4, 4],
               method = 'ard_cv'
              )
#Step 2
# Call the method run_all()
sobol, shapley, total = model.run_all()
```

After completing the first step of loading the `shapleyx` package, the second step involves instantiating the `rshdmr` constructor that configures the `model` object. Key arguments in this call are:

- `'owen_data.csv'` — the input dataset (features and target) to be modelled as described above. 256 samples have been prepared using scrambled and shifted Sobol' sequence sampling (Matoušek 1998; Sobol' 1967).
- `polys = [8, 6, 4, 4]` — a list of polynomial degrees / basis sizes used for the model. The four integers in the list denote the maximum order Legendre polynomial used to estimate the corresponding HDMR partial function terms. For first order terms, up to 8$^{\text{th}}$ order Legendre polynomials will be used, up to 6$^{\text{th}}$ order Legendre polynomials for second order partial functions etc. The length of the list and the maximum Legendre polynomial degree can be set arbitrarily, but the size of the basis set will scale combinatorially with these factors.
- `method = 'ard_cv'`— the `method` variable assigns the fitting/regularisation strategy; in this case "Automatic Relevance Determination" (ARD) with cross-validation. ARD is requested.

The `model.run_all()` method provides a single-call, reproducible workflow that sequences data transform → basis build → selection → evaluation → uncertainty estimation and reporting of Shapley and Sobol' indices.

Results of the *N*=256 Owen product function calculation are provided in Table 1 along with bootstrap uncertainty estimates and analytically determined values for direct comparison. It shows that with a modest initial computational investment of 256 Owen function evaluations, quite accurate results have been achieved despite the complexity of the non-linear function being emulated.

Table 1 Dominant Shapley effects and Shapley–Owen interaction effects for Owen's product function. Numerical estimates from ShapleyX with 95% bootstrap confidence intervals (1000 resamples). Analytical values shown for comparison.

Label	ShapleyX	Lower 95% CI	Upper 95% CI	Analytical
Sh_1	0.3596	0.3493	0.3699	0.3469
Sh_2	0.3451	0.3340	0.3562	0.3469
Sh_3	0.1222	0.1161	0.1282	0.1198
Sh_4	0.1128	0.1066	0.1189	0.1198
Sh_5	0.0246	0.0219	0.0272	0.0333
Sh_6	0.0330	0.0301	0.0360	0.0333
$Sh_{1,2}$	0.2293	0.2194	0.2393	0.2242
$Sh_{1,3}$	0.0791	0.0730	0.0853	0.0757
$Sh_{1,4}$	0.0779	0.0719	0.0839	0.0757
$Sh_{3,2}$	0.0722	0.0660	0.0783	0.0757
$Sh_{4,2}$	0.0690	0.0626	0.0754	0.0757
$Sh_{1,2,3}$	0.0479	0.0417	0.0540	0.0495
$Sh_{1,2,4}$	0.0471	0.0415	0.0527	0.0495
$Sh_{1,2,3,4}$	0.0118	0.0077	0.0159	0.0110

5. DISCUSSION AND CONCLUSION

This work presents a practical framework and open Python implementation—ShapleyX—that estimates Sobol' indices and Shapley Effects efficiently and non-intrusively from a single tabulated dataset. The combination of RS-HDMR expansions and sparse Bayesian regression yields compact and interpretable surrogates whose coefficients map directly to variance contributions. Shapley Effects, computed algebraically from these coefficients, provide dependence-aware, fair attributions of output variance and allocate interactions equitably. Case studies demonstrate accuracy and computational efficiency versus permutation-based Monte Carlo, particularly in moderate to high dimensions where naive Shapley estimators become costly.

Limitations include reliance on low-order HDMR truncation; when significant high-order interactions exist beyond the expansion order, additional terms or hybrid methods (e.g., RS-HDMR with Gaussian process regression for components) may be needed (Ren et al. 2022). The current workflow handles non-uniform input distributions via min–max scaling to [0,1]; future work could extend copula-aware transformations and direct modeling under dependence. Algorithmic extensions may include adaptive per-component order selection, improved parallel bootstrap, and GPU-accelerated linear algebra for very large bases.

ShapleyX is intended as a practical tool for researchers and practitioners in computational science, environmental modelling, and uncertainty quantification who require fair, dependence-robust sensitivity measures without prohibitive sampling costs. The package, sample notebooks reproducing the results, and documentation are provided with the submission.

The case study presented—Owen's product function—demonstrates ShapleyX's accuracy, scalability, and practical utility. It has been shown to correctly rank variable importance, capture intricate interaction effects, and provide nuanced insights into model behaviour that might be obscured by less comprehensive SA methods.

Despite its strengths, the current implementation of ShapleyX, like the underlying RS-HDMR formulation, primarily assumes uniform input distributions. This might limit its direct applicability in contexts where input variables follow significantly different, non-uniform distributions. Future research will explore extensions to accommodate arbitrary input distributions, potentially through adaptive sampling strategies or by integrating other machine learning algorithms capable of handling diverse data distributions. Further enhancements could also include optimising the selection of polynomial orders for RS-HDMR and exploring alternative sparse regression techniques to further improve computational performance for extremely large datasets.

In conclusion, ShapleyX represents a significant advancement in the field of sensitivity analysis, providing a powerful and accessible tool for researchers and practitioners. By combining computational efficiency with robust analytical capabilities, it facilitates a deeper understanding of complex model behaviours, improves uncertainty quantification, and ultimately supports more informed decision-making in scientific and engineering applications.

The ShapleyX package along with documentation and examples, including the Owen Product Function presented here is available at https://github.com/frbennett/shapleyx.

REFERENCES

Borgonovo E (2007) 'A new uncertainty importance measure', *Reliability Engineering & System Safety*, 92(6):771–784, doi:10.1016/j.ress.2006.04.015.

Cacuci DG, Ionescu-Bujor M and Navon IM (2003) *Sensitivity and uncertainty analysis*, Chapman & Hall/CRC Press, Boca Raton.

Il Idrissi M, Chabridon V and Iooss B (2021) 'Developments and applications of Shapley effects to reliability-oriented sensitivity analysis with correlated inputs', *Environmental Modelling & Software*, 143:105115, doi:10.1016/j.envsoft.2021.105115.

Jakeman AJ, Letcher RA and Norton JP (2006) 'Ten iterative steps in development and evaluation of environmental models', *Environmental Modelling & Software*, 21(5):602–614, doi:10.1016/j.envsoft.2006.01.004.

MacKay DJC (1996) 'Bayesian Non-Linear Modeling for the Prediction Competition', in GR Heidbreder (ed) *Maximum Entropy and Bayesian Methods: Santa Barbara, California, U.S.A., 1993*, Springer Netherlands, Dordrecht, doi:10.1007/978-94-015-8729-7_18.

Matoušek J (1998) 'On theL2-Discrepancy for Anchored Boxes', *Journal of Complexity*, 14(4):527–556, doi:10.1006/jcom.1998.0489.

Owen AB (2014) 'Sobol' Indices and Shapley Value', *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):245–251, doi:10.1137/130936233.

Pianosi F and Wagener T (2015) 'A simple and efficient method for global sensitivity analysis based on cumulative distribution functions', *Environmental Modelling & Software*, 67:1–11, doi:10.1016/j.envsoft.2015.01.004.

Pianosi F and Wagener T (2018) 'Distribution-based sensitivity analysis from a generic input-output sample', *Environmental Modelling & Software*, 108:197–207, doi:10.1016/j.envsoft.2018.07.019.

Plischke E, Borgonovo E and Smith CL (2013) 'Global sensitivity measures from given data', *European Journal of Operational Research*, 226(3):536–550, doi:10.1016/j.ejor.2012.11.047.

Ren O, Boussaidi MA, Voytsekhovsky D, Ihara M and Manzhos S (2022) 'Random Sampling High Dimensional Model Representation Gaussian Process Regression (RS-HDMR-GPR) for representing multidimensional functions with machine-learned lower-dimensional terms allowing insight with a general method', *Computer Physics Communications*, 271:108220, doi:10.1016/j.cpc.2021.108220.

Sobol' IM (1967) 'On the distribution of points in a cube and the approximate evaluation of integrals', *USSR Computational Mathematics and Mathematical Physics*, 7(4):86–112, doi:10.1016/0041-5553(67)90144-9.

Sobol' IM (2001) 'Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates', *Mathematics and Computers in Simulation*, 55(1–3):271–280, doi:10.1016/S0378-4754(00)00270-6.

Tipping ME and Faul AC (2003) 'Fast Marginal Likelihood Maximisation for Sparse Bayesian Models', in CM Bishop and BJ Frey (eds) *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, PMLR, https://proceedings.mlr.press/r4/tipping03a.html.

Wang Q, Guillaume JHA, Jakeman JD, Bennett FR, Croke BFW, Fu B, Yang T and Jakeman AJ (2023) 'A Decision-Relevant Factor-Fixing Framework: Application to Uncertainty Analysis of a High-Dimensional Water Quality Model', *Water Resources Research*, 59(8):e2022WR032194, doi:10.1029/2022WR032194.