

A Surrogate-Informed Framework for Sparse Grid Interpolation

M. Rosellini, F. Fruzza, A. Mariotti, M. V. Salvetti, L. Tamellini

November 26, 2025

Abstract

Approximating complex, high-dimensional, and computationally expensive functions is a central problem in science and engineering. Standard sparse grids offer a powerful solution by mitigating the curse of dimensionality compared to full tensor grids. However, they treat all regions of the domain isotropically, which may not be efficient for functions with localized or anisotropic behavior. This work presents a surrogate-informed framework for constructing sparse grid interpolants, which is guided by an error indicator that serves as a zero-cost estimate for the hierarchical surplus. This indicator is calculated for all candidate points, defined as those in the next-level grid $w + 1$ not already present in the base grid w . It quantifies the local approximation error by measuring the relative difference between the predictions of two consecutive interpolants of level w and $w - 1$. The candidates are then ranked by this metric to select the most impactful points for refinement up to a given budget or following another criterion, as, e.g., a given threshold in the error indicator. The final higher-order model is then constructed using a surrogate-informed approach: the objective function is evaluated only at the selected high-priority points, while for the remaining nodes of the $w + 1$ grid, we assign the values predicted by the initial w -level surrogate. This strategy significantly reduces the required number of expensive evaluations, yielding a final model that closely approximates the accuracy of a fully-resolved $w + 1$ grid at a fraction of the computational cost. The accuracy and efficiency of the proposed surrogate-informed refinement criterion is demonstrated for several analytic function and for a real engineering problem, i.e., the analysis of sensitivity to geometrical parameters of numerically predicted flashback phenomenon in hydrogen-fueled perforated burners.

Contents

1	Introduction	3
2	Methodology	4
2.1	Selection of the refinement points	4
2.2	Design of the final surrogate model	5
2.2.1	Surrogate-informed approach	5
2.2.2	Error-correction superposition	5
2.2.3	Equivalence of the Surrogate-informed and Error-correction approaches	6
3	Results	8
3.1	Validation on analytical benchmark functions	8
3.1.1	The Ishigami Function	8
3.1.2	The Sobol G-function	12
3.1.3	The obscillatory Function	15
3.2	Application to the hydrogen-fueled perforated burners case study	19
4	Conclusions	23

1 Introduction

The analysis and design of complex engineering systems frequently rely on high-fidelity computational models. While these models provide accurate predictions, their significant computational cost often makes them impractical for applications requiring repeated evaluations.

To overcome this computational limitation, surrogate modeling has emerged as a powerful and widely adopted strategy. A surrogate model is a mathematical approximation of the original complex, high-order function at a significantly reduced computational cost. It is employed to map input parameters to output responses in situations where the true relationship is either unknown or too computationally expensive to evaluate directly.

When a high number of evaluations is needed, as is typically the case in tasks that need exploring a large parameter space like, e.g., uncertainty quantification or optimization problems, surrogate models become essential. By replacing the original computationally intensive simulations with efficient approximations, surrogate models enable extensive exploration of the parameter space without incurring prohibitive computational costs.

For instance, stochastic-collocation methods provide a polynomial approximation of a quantity of interest as a function of parameters characterized by probability density functions (PDFs). However, when the number of uncertain parameters is large, the computational cost of standard tensor-product collocation methods grows exponentially, becoming prohibitive for practical applications. In this context, sparse grids have emerged as a powerful and widely adopted technique to address this so-called curse of dimensionality [Bungartz and Griebel, 2004]. Originally introduced by Smolyak [1963], sparse grids are built through a hierarchical combination of lower-dimensional tensor product grids, selectively discarding grid points that contribute least to the accuracy of the approximation. This construction dramatically reduces the number of required evaluations while preserving high approximation quality for functions with sufficient regularity. As a result, sparse grids significantly mitigate the curse of dimensionality and make high-dimensional problems more tractable.

However, standard sparse grids are inherently isotropic; they refine the approximation space uniformly across all dimensions and regions of the domain. This approach is suboptimal for a large class of problems characterized by anisotropic or localized features, where the function exhibits different behavior or importance across different input parameters or in specific subregions. This limitation has motivated the development of adaptive sparse grid methods, which aim to allocate computational resources more efficiently by refining the grid non-uniformly [Gerstner and Griebel, 2003]. A common strategy involves using the hierarchical surplus [Ma and Zabaras, 2009], defined as the difference in value between a new grid point and the interpolation from the previous coarser grid, as an a posteriori error indicator to guide the refinement process. In this context, significant contributions have been made to develop both dimension-adaptive and spatially-adaptive schemes.

A major drawback of most adaptive sparse grid strategies lies in the computation of the hierarchical surplus, which requires evaluating the original, expensive function at all candidate refinement points. This process is computationally demanding and may prove inefficient, particularly for points that are ultimately discarded during the grid construction. As a result, the cost of adaptivity can offset its intended benefits. In this work, we introduce a novel surrogate-informed adaptive framework that circumvents this issue. Our main contribution is a zero-cost error indicator that estimates the hierarchical surplus without requiring any new evaluations of the expensive objective function. This indicator quantifies the local approximation error by measuring the relative discrepancy between the predictions of two existing, nested surrogates of consecutive levels. More specifically, if we wish to refine the accuracy obtained at level w , we compute the relative discrepancy between the predictions given by the level w and $w - 1$ grids on all the additional points of the grid at level $w + 1$. When a baseline sparse grid of level w is considered, the nested structure of sparse grids guarantees that all points from level $w - 1$ are already available. This property enables efficient reuse of surrogate evaluations across levels and allows the error indicator to be constructed without incurring any further computational cost.

By ranking all candidate points according to this metric, we can identify and prioritize the most impactful regions for refinement before committing computational resources. The final, higher-order model is then constructed using a hybrid approach: the true function is evaluated only at the selected high-priority points, while the values for the remaining nodes of the new grid level are imputed using the predictions from the last available surrogate. This strategy dramatically reduces the number of expensive function evaluations, yielding a final model that closely approximates the accuracy of a

fully-resolved, higher-level grid at a fraction of the computational cost.

The proposed strategy has been implemented in a tool developed by [Piazzola and Tamellini \[2022\]](#), which constructs a sparse grid interpolant using the Smolyak combination technique, where the final surrogate model is a weighted linear combination of simpler, anisotropic tensor-product grids. The selection of these component grids is governed by an isotropic total degree rule, which systematically determines the multi-indices and corresponding coefficients for the combination formula. To ensure numerical stability and efficiency, each tensor grid is built upon Clenshaw-Curtis nodes arranged according to a nested doubling rule, which guarantees that points evaluated at coarser levels are reused in finer ones. The final approximation at any given point is then obtained by summing the values of the individual Lagrange interpolants from each grid in the combination, weighted by their specific Smolyak coefficient.

The structure of this work is as follows. In Section 2, we present and formally derive the theoretical foundations of our novel surrogate model, specifically describing the proposed error indicator in section 2.1 and designing the final surrogate-informed model in 2.2. Section 3 is devoted to assessing the performance of the proposed framework. We begin by testing it on a set of analytical benchmark functions commonly used in uncertainty quantification, including the Ishigami function (Section 3.1.1), the Sobol G-function (Section 3.1.2), and a highly oscillatory function (Section 3.1.3). Finally, in Section 3.2, we demonstrate the applicability of our method to a real-world scenario involving hydrogen-fueled perforated burners.

2 Methodology

The core of our proposed methodology lies in the efficient construction of a high-fidelity surrogate model, designed to approximate a full $w + 1$ level sparse grid while reducing the number of required function evaluations. This is achieved through two key steps: first, a selection of a small subset of the most impactful refinement points, and second, the design of the final surrogate model by evaluating the function at these points and using the w level sparse grid surrogate model for the remaining ones.

2.1 Selection of the refinement points

Instead of uniformly enriching the grid, which can be inefficient, our approach identifies and prioritizes regions of the domain where the current surrogate model is least accurate. This is accomplished using a computationally efficient, a-posteriori error indicator which gives an estimate of the hierarchical surplus. The core principle is to estimate the local interpolation error at a candidate point $\mathbf{x}$ by quantifying the model’s sensitivity to a change in its complexity. This concept is inspired by the hierarchical surplus in sparse grids, but it is implemented here as a zero-cost proxy. We measure the discrepancy between the predictions of two surrogate models of adjacent levels, both constructed using only the already available data from the base grid. A large discrepancy suggests that the lower-level model was insufficient at that point, indicating a high local error and a high potential for improvement. The discrepancy criterion is calculated for every candidate point $\mathbf{x} \in \mathcal{X}_{\text{cand}}$, where the candidate set is defined as the difference between the point sets of the target (level $w + 1$) and base (level w) grids:

$$\mathcal{X}_{\text{cand}} = \mathcal{X}_{w+1} \setminus \mathcal{X}_w \quad (1)$$

Let $D_w = \{(\mathbf{x}_i, f_{\text{true}}(\mathbf{x}_i)) \mid \mathbf{x}_i \in \mathcal{X}_w\}$ be the dataset containing the known function values on the base grid. We define two auxiliary surrogate models using $\mathcal{A}_k$, the sparse grid interpolation operator of level k , with data D_w :

1. A w -level interpolant: $f_w(\mathbf{x}) = \mathcal{A}_w[D_w](\mathbf{x})$
2. A $(w - 1)$ -level interpolant: $f_{w-1}(\mathbf{x}) = \mathcal{A}_{w-1}[D_w](\mathbf{x})$

Crucially, both models are constructed using the same limited dataset D_w . The calculation of their predictions at any new point $\mathbf{x}$ therefore incurs no additional calls to the expensive true function f_{true} . The point-wise discrepancy, $\Delta(\mathbf{x})$, is the absolute difference between these two predictions:

$$\Delta(\mathbf{x}) = |f_w(\mathbf{x}) - f_{w-1}(\mathbf{x})| \quad (2)$$

To provide a normalized and numerically stable metric, the final discrepancy criterion, $\eta(\mathbf{x})$, is formulated as a hybrid relative-absolute error. This approach avoids instabilities when the predicted value $f_w(\mathbf{x})$ is close to zero. The criterion is defined as:

$$\eta(\mathbf{x}) = \begin{cases} \frac{|\Delta(\mathbf{x})|}{|f_w(\mathbf{x})|} & \text{if } |f_w(\mathbf{x})| > \epsilon \\ |\Delta(\mathbf{x})| & \text{if } |f_w(\mathbf{x})| \leq \epsilon \end{cases} \quad (3)$$

where ϵ is a small numerical tolerance (e.g., 10^{-12}). Once the indicator $\eta(x)$ has been computed for each candidate point $x \in \mathcal{X}_{\text{cand}}$, the points are sorted in descending order according to their indicator values. To determine which points should be refined, several selection strategies can be considered. In this work, we focus on three representative approaches:

1. Fixed budget: The top B candidates are selected, where B is a predefined computational budget.
2. Relative threshold: All candidates whose indicator value exceeds a fraction τ of the maximum observed indicator, η_{max} , are selected. The set is thus defined as $\mathcal{X}_{\text{sel}} = \{x \in \mathcal{X}_{\text{cand}} \mid \eta(x) \geq \tau \cdot \eta_{\text{max}}\}$.
3. Elbow identification: A geometric heuristic identifies the *elbow* point of the sorted indicator curve as the point with the maximum perpendicular distance to a line connecting the first and last points. The number of points in $\mathcal{X}_{\text{sel}}$ is set to the rank of this elbow point, thus establishing a cost-benefit optimal budget.

2.2 Design of the final surrogate model

Let $\mathcal{X}_w$ be the sparse grid point set at level w . The corresponding dataset, D_w , is obtained by evaluating the true function f_{true} at each point in this grid:

$$D_w = \{(\mathbf{x}_i, f_{\text{true}}(\mathbf{x}_i)) \mid \mathbf{x}_i \in \mathcal{X}_w\} \quad (4)$$

The w -level base model, $f_b(\mathbf{x})$, is then constructed by applying the w -level sparse grid interpolation operator, $\mathcal{A}_w$, to this discrete dataset:

$$f_b(\mathbf{x}) = \mathcal{A}_w[D_w](\mathbf{x}) \quad (5)$$

Let $\mathcal{X}_{w+1}$ be the set of nodes for the target grid of level $w + 1$, and let $\mathcal{X}_{\text{sel}} \subset \mathcal{X}_{w+1}$ be the subset of nodes selected for refinement using one of the approaches described in Section 2.1.

Building on the refined point set $\mathcal{X}_{\text{sel}}$, we propose two alternative strategies for constructing the final surrogate model. Each method offers a different perspective on how to integrate new information while preserving the structure and efficiency of sparse grid interpolation.

2.2.1 Surrogate-informed approach

The surrogate model on the sparse grid at level $w + 1$ is obtained by additional calls of the function only at the nodes selected through the criterion described in Section 2.1, while for the remaining nodes the surrogate model at level w is used. The function $g(\mathbf{x})$ is defined at the nodes of the $w + 1$ grid as:

$$g(\mathbf{x}) = \begin{cases} f_{\text{true}}(\mathbf{x}) & \text{if } \mathbf{x} \in \mathcal{X}_{\text{sel}} \\ f_b(\mathbf{x}) & \text{if } \mathbf{x} \in \mathcal{X}_{w+1} \setminus \mathcal{X}_{\text{sel}} \end{cases} \quad (6)$$

The final model is then given by:

$$f_{\text{sur}}(\mathbf{x}) = \mathcal{A}_{w+1}[g](\mathbf{x}) \quad (7)$$

2.2.2 Error-correction superposition

The final surrogate model can also be built through an error-correction superimposition as the point-wise sum of two continuous functions, the base model and an error-correction term $E(\mathbf{x})$:

$$f_{\text{corr}}(\mathbf{x}) = f_b(\mathbf{x}) + E(\mathbf{x}) \quad (8)$$

The error-correction term $E(\mathbf{x})$ is itself a sparse grid interpolant of an error function, $e(\mathbf{x})$, which is defined as:

$$e(\mathbf{x}) = \begin{cases} f_{true}(\mathbf{x}) - f_b(\mathbf{x}) & \text{if } \mathbf{x} \in \mathcal{X}_{sel} \\ 0 & \text{if } \mathbf{x} \in \mathcal{X}_{w+1} \setminus \mathcal{X}_{sel} \end{cases} \quad (9)$$

The error surrogate is thus:

$$E(\mathbf{x}) = \mathcal{A}_{w+1}[e](\mathbf{x}) \quad (10)$$

2.2.3 Equivalence of the Surrogate-informed and Error-correction approaches

Let $f_{surr}(\mathbf{x})$ be the surrogate model constructed via the surrogate-informed approach, and let $f_{corr}(\mathbf{x})$ be the model from the Error-correction approach. The first method builds the refined surrogate by populating the target grid with a mix of true function values and predicted values from the base model, while the second one first builds a surrogate for the approximation error and then adds it to the base model. For a nested, polynomial sparse grid framework, the two models are proved to be identical:

$$f_{surr}(\mathbf{x}) \equiv f_{corr}(\mathbf{x})$$

By construction, a sparse grid interpolant, denoted as $\mathcal{A}[f](\mathbf{x})$, is exact at its own nodes. That is, for any node $\mathbf{x}_i$ within the grid's point set, the interpolant satisfies the property:

$$\mathcal{A}[f](\mathbf{x}_i) = f(\mathbf{x}_i) \quad (11)$$

where $f(\mathbf{x}_i)$ is the function value supplied at that node.

Let $\mathcal{A}_w$ be the sparse grid interpolation operator for a grid of level w . Let $f_b(\mathbf{x}) = \mathcal{A}_w[f_{true}](\mathbf{x})$ be the baseline surrogate model at level w . Let $\mathcal{X}_{target}$ be the set of nodes for the target grid of level $w + 1$, and let $\mathcal{X}_{sel} \subset \mathcal{X}_{target}$ be the subset of nodes selected for refinement. The equivalence between the surrogate-informed approach and the error-correction approach can be demonstrated by recalling the linearity of the sparse grid interpolation operator. The operator $\mathcal{A}_w$ is linear, meaning $\mathcal{A}_w[f + g] = \mathcal{A}_w[f] + \mathcal{A}_w[g]$ where f and g are two different functions and $\mathcal{A}_w[c \cdot f] = c \cdot \mathcal{A}_w[f]$ where f is a function and c a number.

Looking at Equation 6, the hybrid function $g(\mathbf{x})$ from the surrogate-informed approach can be expressed in terms of the base model $f_b(\mathbf{x})$ and a delta function $\delta(\mathbf{x})$:

$$g(\mathbf{x}) = f_b(\mathbf{x}) + \delta(\mathbf{x}) \quad (12)$$

where $\delta(\mathbf{x})$ can be expressed as:

$$\delta(\mathbf{x}) = \begin{cases} f_{true}(\mathbf{x}) - f_b(\mathbf{x}) & \text{if } \mathbf{x} \in \mathcal{X}_{sel} \\ 0 & \text{if } \mathbf{x} \in \mathcal{X}_{target} \setminus \mathcal{X}_{sel} \end{cases} \quad (13)$$

By inspection, this function $\delta(\mathbf{x})$ is identical to the error function $e(\mathbf{x})$ from Equation 9 used in the error-correction approach. Now, if we apply the interpolation operator to $g(\mathbf{x})$ and make use of its linearity we obtain:

$$\begin{aligned} f_{surr}(\mathbf{x}) &= \mathcal{A}_{w+1}[g](\mathbf{x}) \\ &= \mathcal{A}_{w+1}[f_b + \delta](\mathbf{x}) \\ &= \mathcal{A}_{w+1}[f_b](\mathbf{x}) + \mathcal{A}_{w+1}[\delta](\mathbf{x}) \end{aligned} \quad (14)$$

Analyzing the two terms on the right-hand side of Eq. 14 we can notice that:

1. The term $\mathcal{A}_{w+1}[\delta](\mathbf{x})$ is by definition the interpolation of the error function $\delta(\mathbf{x}) = e(\mathbf{x})$, which is precisely the error surrogate $E(\mathbf{x})$;
2. The term $\mathcal{A}_{w+1}[f_b](\mathbf{x})$ represents the interpolation of the base model $f_b(\mathbf{x})$ on the higher-level grid structure $w + 1$. The base model $f_b(\mathbf{x})$ is a polynomial of a degree determined by the w sparse grid. Since the Clenshaw-Curtis grids are nested and the w_{target} grid is richer than the w_{start} grid, the $w + 1$ interpolant has the capacity to exactly reproduce any polynomial that the w interpolant can represent. Therefore, the interpolation of a lower-order polynomial on a higher-order grid structure simply returns the polynomial itself.

$$\mathcal{A}_{w+1}[f_b](\mathbf{x}) = f_b(\mathbf{x}) \quad (15)$$

Substituting these results back into Eq. 14, we get:

$$f_{\text{surr}}(\mathbf{x}) = f_b(\mathbf{x}) + E(\mathbf{x}) = f_{\text{corr}}(\mathbf{x}), \quad (16)$$

proving the equivalence of the models obtained with the two approaches

3 Results

To validate the performance and robustness of the proposed adaptive refinement workflow, we first applied it to some analytical functions that are notoriously challenging for surrogate modeling. For all the tested functions, a sparse grid model was initially constructed using a low-fidelity grid ($w = 2$), with the target resolution set at level $w + 1 = 3$. Using our algorithm, we identified the most sensitive points, which were subsequently used to build a surrogate-informed model. Subsequently, the method was deployed on a practical engineering case study to demonstrate its real-world applicability.

3.1 Validation on analytical benchmark functions

The chosen benchmark functions, Ishigami function [Ishigami and Homma [1990]], Sobol’s G-function and an oboscillating function, test the algorithm’s ability to handle strong non-linearities, parameter interactions, and non-smooth features [Saltelli et al. [2010]].

3.1.1 The Ishigami Function

The Ishigami function is a non-linear, non-monotonic benchmark defined in three dimensions, characterized by a strong interaction between its first and third variables. It is formally expressed as:

$$f(\mathbf{x}) = \sin(x_1) + a \sin^2(x_2) + bx_3^4 \sin(x_1), \quad (17)$$

where the input domain is $\mathbf{x} \in [-\pi, \pi]^3$, and the standard coefficients are $a = 7$ and $b = 0.1$. The combination of monotonic and periodic components makes this function a comprehensive test case for assessing surrogate-based refinement strategies.

To identify candidate points for refinement, we apply the discrepancy criterion described in Equation 3, evaluating all nodes from level $w + 1$ that are not already present in level w . Given the sparsity of the lower-level grid (e.g., only 7 points in the case of the zero-level $w = 1$), the discrepancy values are generally expected to be high, as the coarse representation fails to capture the local variability of the function across most candidate locations. In this context, we adopt a threshold-based selection strategy, retaining all candidate points whose discrepancy exceeds 20%. This approach effectively isolates regions where the lower-level surrogate is expected to fail, resulting in the selection of 36 refinement points. These points are then used to construct the surrogate-informed model following the procedure detailed in Section 2. Figure 1 presents a pointwise comparison of the percentage error for the baseline model at level w , the surrogate-informed hierarchical model, and the target model at level $w + 1$, evaluated over 50 randomly selected test locations. The results demonstrate that the hierarchical model consistently outperforms the baseline, achieving significantly lower errors across nearly all test points. Furthermore, its predictions closely align with those of the target model, indicating that the surrogate-informed refinement effectively reproduces the accuracy of a higher-fidelity simulation.

An alternative perspective is provided in Figure 2, which displays the histogram of the normalized percentage error density for both the baseline and hierarchical models. While the pointwise comparison highlights individual performance differences, the histogram emphasizes the global error distribution. As shown, the baseline model yields a broader spread of errors, with values reaching up to 50%. In contrast, the hierarchical model exhibits a sharp concentration of errors around lower values, with no instances exceeding 10%. This distributional view reinforces the improved accuracy and robustness of the hierarchical approach in capturing the underlying function behavior.

Figure 3 depicts the response surfaces on the (x_1, x_2) plane for $x_3 = \pi$, obtained from the interpolants constructed on the baseline grid at level w , the surrogate-informed grid, and the target grid at level $w + 1$. These plots illustrate how the surrogate-informed model successfully approximates the target resolution, yielding high-quality results compared to the baseline configuration. Notably, the response surface derived from the baseline grid fails to capture the full amplitude of the sine-induced peaks, resulting in a visibly smoothed representation of the function. In contrast, the hierarchical model restores the accuracy observed in the target reference, effectively recovering the sharper features and local variations that are otherwise too much smoothed in the coarse baseline approximation.

Figure 4 presents a cross-sectional view of the function behavior along the slice defined by $x_1 = x_3 = \pi$. This visualization enables a direct comparison between the true reference function, the interpolated baseline model at level w , and the surrogate-informed hierarchical model. The plot clearly illustrates




Figure 1: Comparison of the error on 50 testing points for the $w = 2$ and surrogate informed sparse grid for the Ishigami's function




Figure 2: Histogram of the error between the baseline and the informed model




Figure 3: Response surface comparison for the Ishigami's function at $x_3 = \pi$.

how the hierarchical refinement improves the approximation, particularly in regions where the baseline model fails to capture the underlying non-linear structure of the function.

This cross-section offers a detailed perspective on the trends already observed in Figure 3, where the response surfaces were visualized over the (x_1, x_2) plane for $x_3 = \pi$. Specifically, it confirms that the baseline model produces significantly attenuated and smoothed peaks compared to the target reference, failing to reproduce the full amplitude of the sine-driven oscillations. In contrast, the hierarchical model closely matches the target behavior, successfully recovering both the sharpness and magnitude of the peaks.




Figure 4: Comparison between the baseline sparse grid, the surrogate-informed one and the true values for the Ishigami's function at $x_1 = x_3 = \pi$




Figure 5: Convergence of maximum absolute error and RMSE for the Ishigami's function

Convergence study

In our initial setup, candidate points were selected using a threshold criterion with $\tau = 0.2$, which included all nodes exhibiting unit relative error. For the three-dimensional Ishigami function, this resulted in a refinement set of 36 points. Beyond the static selection, our methodology also enables a dynamic convergence analysis by incrementally enriching the surrogate model with additional re-

finement points. This process allows us to monitor how the approximation quality evolves as the computational budget increases.

Figure 5 illustrates this convergence behavior with respect to both the maximum absolute error and the Root Mean Square Error (RMSE). The plot is constructed by generating a sequence of surrogate-informed models, each incorporating an increasing number of refinement points—starting from zero (corresponding to the baseline sparse grid at level w) and progressing toward the full target grid at level $w + 1$. For each configuration, an interpolant is built and evaluated on a set of 200 randomly distributed test points, yielding pointwise error estimates against the true function values.

This incremental refinement strategy enables us to track how the approximation accuracy improves as additional points are incorporated. Both the maximum absolute error and the RMSE—decrease steadily, indicating a monotonic convergence behavior. As the refinement set grows, we observe a convergence plateau beyond which further point additions result in negligible error reduction. This trend provides practical guidance for determining an effective refinement budget: once the error stabilizes, adding further refinement points yields only marginal improvements in accuracy. Importantly, even with a limited number of refinement points, the surrogate-informed model achieves an accuracy comparable to that of the full sparse grid interpolant at level $w + 1$, confirming the efficiency and robustness of the proposed strategy in capturing the essential features of the target function.

3.1.2 The Sobol G-function

The Sobol G-function is a scalable, high-dimensional benchmark characterized by discontinuities in its derivatives. For our analysis, we used a $d = 4$ dimensional implementation, defined as:

$$g(\mathbf{x}) = \prod_{i=1}^d \frac{|4x_i - 2| + a_i}{1 + a_i} \quad (18)$$

where the domain is $\mathbf{x} \in [0, 1]^d$. The coefficients a_i are chosen to control the importance of each variable and are set to $a_i = (i - 1)/2$. The function is most sensitive to the first dimensions, testing the ability of the adaptive algorithm to correctly allocate refinement points. Figure 6 shows the trend of the discrepancy function computed between levels w and $w - 1$ across candidate points at level $w + 1$. In this case, each point corresponds to a different error magnitude. Applying the 20% threshold criterion successfully excluding from refinement those points with a relative error below the threshold, leads to the selection of 50 points and the exclusion of the remaining 46. If, instead, the elbow selection method had been used, in this particular configuration the algorithm would have selected all points with non-zero error, that would have resulted in 52 points being chosen and 44 discarded.




Figure 6: Discrepancy criterion for the nodes of the $w = 3$ sparse grid for the Sobol’s function

Figure 7 reports the percentage error of the hierarchical surrogate-informed model and the one of the baseline model. For reference, the figure also includes model at level $w + 1$ (target model). The results show that the hierarchical surrogate-informed model achieves an error nearly indistinguishable from the surrogate at level $w + 1$, confirming its ability to replicate high-fidelity behavior while leveraging surrogate information.

Figure 8 shows the histogram of the normalized percentage error for both the baseline and the hierarchical surrogate models. As observed in the previous benchmark, the hierarchical model performs better than the baseline, with the distribution being shifted toward smaller errors.

Figure 9 presents the response surfaces obtained from the three models, baseline at level w , hierarchical surrogate-informed, and target at level $w + 1$, on the (x_1, x_2) plane, fixing $x_3 = x_4 = 0$. While the baseline model provides a coarse approximation of the overall trend, it fails to accurately capture the localized features that emerge near the boundaries of the domain. In the reference configuration, these edge regions exhibit sharp peaks and rapid variations that are entirely missed by the baseline model, resulting in a smoothed response surface which is not able to represent the true complexity of the system. In contrast, the hierarchical surrogate-informed model is able to recover these localized phenomena with remarkable fidelity. It reconstructs a response surface that closely resembles the reference, both in terms of amplitude and spatial distribution of the peaks. This improved resolution near the domain boundaries highlights the model’s ability to integrate multi-level information and adapt to features that are otherwise inaccessible at the baseline resolution.

Convergence study

As in the case of the Ishigami function, Figure 10 illustrates the convergence behavior of the refinement method as the number of budget points increases. The trend is clearly monotonic: both the maximum absolute error and the RMSE decrease steadily as more refinement points are added. In particular, it can be seen that the hierarchical surrogate reaches an accuracy comparable to that of the full sparse grid model at level $w + 1$, with only 50 additional function evaluations instead of 96. This confirms the efficiency of the refinement strategy in balancing computational cost and approximation quality.




Figure 7: Comparison of the error on 50 testing points for the $w = 2$ and surrogate informed sparse grid for the Sobol's function




Figure 8: Histogram of the error between the baseline and the hierarchical model




Figure 9: Response surface comparison for the Sobol's function at $x_3 = x_4 = 0$




Figure 10: Convergence of maximum absolute error and RMSE for the Sobol's function

3.1.3 The obscillatory Function

Oscillatory 1 Function The Oscillatory 1 function is a smooth, periodic benchmark designed to evaluate the sensitivity of numerical algorithms to heterogeneous frequency content across input dimensions. In our study, we consider a $d = 4$ dimensional formulation defined as:

$$f(\mathbf{x}) = \sum_{i=1}^d \cos(\theta_i x_i) \quad (19)$$

where the domain is $\mathbf{x} \in [0, 1]^d$. The frequency coefficients θ_i are chosen to induce anisotropic oscillatory behavior:

$$\boldsymbol{\theta} = [1.5\pi, 3.0\pi, 0.5\pi, 4.5\pi] \quad (20)$$

This configuration results in higher-frequency oscillations along x_2 and x_4 , and lower-frequency variation along x_3 , introducing directional smoothness disparities. For general dimensionality $d \neq 4$, the function defaults to a scalable frequency assignment $\theta_i = (i + 1)\pi$, ensuring extensibility. The function is continuously differentiable and free of discontinuities, making it suitable for testing surrogate modeling techniques and adaptive sampling strategies in smooth regimes. Its varying frequency content challenges algorithms to resolve fine-scale features selectively across dimensions. Figure 11 illustrates the behavior of the discrepancy function computed between levels w and $w - 1$ across candidate points at level $w + 1$. Each point corresponds to a distinct error magnitude. Applying the threshold-based selection criterion with a cutoff value of 5% results in the selection of 26 points and the exclusion of the remaining candidates.

Figure 12 reports the percentage error relative to the true function of the surrogate-informed model and the one of the baseline model showing once again as the surrogate-informed model has errors much lower than the baseline one, very close to those of the more expensive $w + 1$ level reference model.




Figure 11: Discrepancy criterion for the nodes of the $w = 3$ sparse grid for the obscillatory function

Figure 13 presents the histogram of the percentage error for both the baseline and the surrogate-informed models. Indeed, the baseline model exhibits a broader error distribution, with a significant number of points reaching higher error magnitudes relative to the true function values. In contrast, the histogram corresponding to the surrogate-informed model demonstrates a more concentrated distribution, successfully maintaining error levels across the evaluated points lower than 7%. Figure 14 shows the response surface projected onto the x_1 - x_2 plane, obtained by fixing $x_3 = x_4 = 0$. As in the previous cases, the baseline model at level w provides only a coarse approximation of the true behavior of the system. In particular, it fails to capture the sharp localized peaks that characterize the reference response surface. The hierarchical surrogate-informed model, on the other hand, succeeds in reconstructing these localized structures with high fidelity, as it produces a response surface

that closely matches the target both in amplitude and spatial distribution. This demonstrates the model’s ability to adapt to oscillatory behavior and resolve fine-scale features that are inaccessible at lower resolution levels. The improved agreement with the reference highlights the effectiveness of the hierarchical strategy in capturing complex patterns, especially in regions where the function exhibits rapid variations. Figure 15 shows the error between the baseline w model, target $w + 1$ model and the true values of the function, calculated along a one-dimensional slice defined by $x_1 = x_2 = x_3 = 0.5$, while varying the x_4 dimension. As evident from the plot, the baseline model exhibits significant deviations from the reference, particularly in regions where the function undergoes rapid oscillations. In contrast, the hierarchical surrogate-informed model demonstrates a markedly improved accuracy, with error levels consistently lower and more stable across the domain. The error profile thus reinforces the conclusions drawn from the response surface analysis, highlighting the benefits of our hierarchical surrogate modeling approach in capturing localized and high-frequency phenomena.

Convergence study

As for the previous benchmarks, Figure 16 illustrates the convergence behavior of the refinement strategy as the number of budget points increases. The plot shows a progressive reduction in both the maximum absolute error and the RMSE, confirming the monotonic improvement of the surrogate model. Notably, the curve reveals the existence of a saturation threshold beyond which the hierarchical surrogate achieves an accuracy comparable to that of the full sparse grid model at level $w + 1$, while requiring significantly fewer evaluations of the true objective function.

In this specific case, the hierarchical strategy reaches the same level of accuracy as the target model by adding only 26 refinement points, compared to the 96 evaluations required by the full sparse grid at level $w + 1$. This results in a substantial reduction in computational cost, demonstrating the efficiency of the proposed refinement procedure. The ability to achieve near-target accuracy with a fraction of the evaluations highlights the potential of the hierarchical surrogate for scalable uncertainty quantification and optimization tasks, especially in high-dimensional settings where function evaluations are expensive.




Figure 12: Comparison of the error on 50 testing points for the $w = 2$ and surrogate informed sparse grid for the oscillatory function




Figure 13: Histogram of the error between the baseline and the informed model




Figure 14: Response surface comparison for the oscillatory function at $x_3 = x_4 = 0$




Figure 15: Comparison between the baseline sparse grid, the surrogate-informed one and the true values for the oscillatory function at $x_1 = x_2 = x_3 = 0.5$.




Figure 16: Convergence of maximum absolute error and RMSE for the oscillatory function

3.2 Application to the hydrogen-fueled perforated burners case study

The primary objective of the proposed surrogate-informed hierarchical sparse grid refinement methodology is to enable efficient and accurate surrogate modeling of complex engineering systems, where direct evaluations of the objective function are computationally expensive. To demonstrate its practical applicability, we consider a representative case study involving hydrogen-fueled perforated burners. In this context, flashback phenomena pose critical risks due to hydrogen's high flame speed and diffusivity, making reliable prediction of flame behavior essential for design and control.

The physical model considers a representative unit cell with a single slit, defined by four geometric parameters: slit length (L), slit width (W), inter-slit spacing (D), and plate thickness (H), as shown in Figure 17. The simulations describe a hydrogen-air mixture at equivalence ratio $\phi = 0.6$ with uniform inlet conditions ($T_u = 300$ K) and ambient outlet pressure ($p = 1$ atm). The solid domain is modeled as stainless steel with thermally-coupled interfaces via Conjugate Heat Transfer. Detailed chemistry is captured through a 9-species, 22-reaction mechanism Bilger et al. [1990]. A steady-state approach is adopted to determine the critical flashback velocity V_{FB} , defined as the cold-flow bulk velocity at the slit entry:

$$V_{FB} = \frac{A_{\text{tot}}}{A_{\text{slit}}} V_{\text{in}} = V_{\text{in}} / \psi, \quad (21)$$

where $\psi = A_{\text{slit}}/A_{\text{tot}}$ is the porosity of the slit. This definition provides an intrinsic measure of the slit's resistance to flashback, independent of the trivial scaling with ψ . Under all investigated conditions, the jet Reynolds number remains within the laminar regime. For a comprehensive description of the numerical model, validation studies, and computational setup, the reader is referred to Fruzza et al. [2023, 2024b,a, 2025a].




Figure 17: Left panel: Slit pattern on the burner plate, highlighting the section of the computational domain. Center panel: Slit geometry (boundary conditions in red) and solid plate mesh. Right panel: Computational domain (fluid zone in light blue, solid zone in gray, boundary conditions in red). Reproduced from Fruzza et al. [2025a].

In a previous study Fruzza et al. [2025b], a surrogate-informed sparse-grid methodology was developed to map critical flashback velocities and maximum burner temperatures across the four-dimensional design space. Since each sparse grid point required a costly three-dimensional simulation, an initial sparse grid at level $w = 2$ was constructed, combining high-fidelity simulations with surrogate modeling to reduce computational cost. However, the response surfaces generated at this level were found to lack sufficient accuracy, particularly in regions of high sensitivity ($L \approx 0$). To address this limitation, the grid was manually refined by evaluating 33 additional simulated points at level $w = 3$ near $L = 0$, while the remaining points were populated using interpolated values from the initial surrogate model. Although this physics-informed manual refinement improved accuracy, it relied on prior domain knowledge and lacked a systematic criterion for point selection.

In the current study, we aim to generalize and optimize this refinement strategy by automatically identifying high-impact grid points for additional evaluations. The proposed surrogate-informed hierarchical methodology provides a systematic approach to targeted refinement while respecting computational constraints, eliminating the need for manual intervention based on physical intuition.

The two scalar output quantities of interest are flashback velocity and maximum burner temperature. Thus, the refinement process can be guided by either quantity. We observed that the resulting performance exhibits no significant differences depending on the variable chosen to drive the refinement, so in what follows, we refer to the process using the flashback velocity as the guiding quantity. Figure 18 presents the discrepancy criterion computed for candidate nodes at level $w + 1$, ranked in decreasing order of importance. To enable a direct comparison between the manually refined physics-based approach from [Fruzza et al. \[2025b\]](#) and the surrogate-informed hierarchical strategy, the budget in this study was fixed at 33 points, matching the number of simulations used during the previous empirical refinement phase.




Figure 18: Discrepancy criterion for the nodes of the $w = 3$ sparse grid for our case study

Figure 19 illustrates the three-dimensional refinement points selected by both the physics-based and the hierarchically-informed models, fixing the dimension $W = 0$. The figure displays the 33 points chosen according to the discrepancy-based criterion introduced in this study, allowing for a direct visual comparison with those obtained through manual empirical refinement. While the two sets are not identical, the hierarchical model successfully recovers the majority of the points identified by the physics-based approach, particularly in regions where the function exhibits strong gradients or localized features. This overlap confirms the effectiveness of the discrepancy criterion in targeting the most informative regions of the domain. In addition to this core agreement, the hierarchical strategy introduces a few supplementary points in more peripheral or less active areas of the input space. These additional samples, although not selected by the manual procedure, contribute to regularizing the surrogate model and improving its global accuracy. Overall, the comparison highlights the surrogate’s capacity to emulate expert-driven refinement while offering a more systematic and scalable alternative.

Figure 20 shows the histogram of the normalized percentage error for three models: the baseline, the physics-informed, and the hierarchically-informed surrogate. The results clearly demonstrate that the hierarchically-informed model achieves error distributions closely aligned with those of the physics-informed model. In particular, it successfully shifts the error peak toward lower values, indicating a substantial improvement over the baseline and confirming its effectiveness in capturing the underlying physical behavior with minimal computational cost.

Figure 21 compares the response surfaces for flame velocity across three configurations: the baseline sparse grid model, the proposed surrogate-informed hierarchical model, and the knowledge-based model. The hierarchical and knowledge-based models produce nearly identical predictions, exhibiting a more accurate representation of the velocity gradient across the domain, especially in regions where the baseline model tends to oversmooth the underlying physical behavior.




Figure 19: Collocation of refinement points from the physics-based model and the hierarchical one




Figure 20: Histogram plot of the error from the baseline, the hierarchical and the physics-informed model

Figure 22 reports the percentage error on the flashback velocity, evaluated over 17 test points for three models: the baseline sparse grid model, the surrogate-informed hierarchical model proposed in this study, and the physics-based model. Since no analytical reference values are available in this experimental setting, the comparison is performed against empirical measurements taken at 17 locations specifically selected outside the collocation point of all sparse grid models. This ensures an unbiased evaluation of each model’s generalization capability. The hierarchical surrogate exhibits a performance closely aligned with that of the knowledge-based approach, confirming its ability to emulate expert-driven modeling decisions. Although the error profiles vary across individual test locations—reflecting localized differences in how each model captures the underlying physics—the overall trend is consistent, and both the hierarchical and physics-based models yield comparable global accuracy.




Figure 21: Response surface comparison for the flashback velocity in our case study at $L = 0.0$ and $H = 0.6$




Figure 22: Comparison of the error on 17 testing points for flame velocity in our case study

4 Conclusions

In this work, we introduced a novel surrogate-informed adaptive framework for constructing sparse grid interpolants, designed to overcome the computational limitations of traditional adaptive methods. The core of our contribution is a zero-cost error indicator that leverages two nested, pre-existing surrogates to estimate the hierarchical surplus without requiring any expensive evaluations of the true objective function. This allows for an efficient, a priori ranking of candidate refinement points, enabling a hybrid model construction where the computational budget is allocated exclusively to the most impactful regions of the parameter space.

The effectiveness of the proposed methodology was validated through a series of numerical investigations. On a set of analytical benchmark functions, the framework demonstrated its robustness and efficiency. The results consistently showed that our hybrid model achieved an accuracy closely approaching that of a fully-resolved, higher-level ($w+1$) grid, while requiring only a smaller fraction of the computational budget. This confirms that the proposed zero-cost indicator is highly effective at identifying the most salient points for grid refinement.

The practical applicability of the framework was then demonstrated on an experimental dataset from a hydrogen-fueled perforated burner. In this real-world scenario, our method yielded excellent results, comparable to those obtained from a specialized, physics-informed refinement of the surrogate model. A remarkable finding was that our algorithm automatically identified the majority of the refinement points that had been manually selected in the physics-informed approach based on specific knowledge. However, a key distinction emerged: our data-driven indicator also introduced a few additional points in regions far from the primary area of interest. This behavior highlights a significant advantage of our approach: by not being constrained by physical assumptions, the algorithm works to regularize the entire response surface, ensuring global accuracy rather than focusing solely on a localized, known region of high error.

In conclusion, the surrogate-informed adaptive sparse grid framework presents a powerful and efficient alternative to standard adaptive techniques. It effectively balances accuracy and computational cost, making high-dimensional problems more tractable. Future work will focus on applying this framework to a broader range of complex engineering problems, particularly in the context of optimization and uncertainty quantification. Further research could also explore the formulation of more sophisticated error indicators and the integration of this adaptive strategy within multi-fidelity modeling workflows.

References

- RW Bilger, SH Stårner, and RJ Kee. On reduced mechanisms for methane-air combustion in non-premixed flames. *Combustion and Flame*, 80:135–149, 1990.
- Hans-Joachim Bungartz and Michael Griebel. Sparse grids. *Acta numerica*, 13:147–269, 2004.
- Filippo Fruzza, Rachele Lamioni, Leonardo Tognotti, and Chiara Galletti. Flashback of H₂-enriched premixed flames in perforated burners: Numerical prediction of critical velocity. *International Journal of Hydrogen Energy*, 48:31790–31801, 2023. ISSN 0360-3199.
- Filippo Fruzza, Hongchao Chu, Rachele Lamioni, Temistocle Grenga, Chiara Galletti, and Heinz Pitsch. The importance of soot effect, preferential diffusion, and conjugate heat transfer for flashback limits of hydrogen-fueled perforated burners. *Proceedings of the Combustion Institute*, 40(1):105581, 2024a. ISSN 1540-7489.
- Filippo Fruzza, Rachele Lamioni, Alessandro Mariotti, Maria Vittoria Salvetti, and Chiara Galletti. Flashback propensity due to hydrogen blending in natural gas: Sensitivity to operating and geometrical parameters. *Fuel*, 362:130838, 2024b. ISSN 0016-2361.
- Filippo Fruzza, Hongchao Chu, Rachele Lamioni, Temistocle Grenga, Chiara Galletti, and Heinz Pitsch. Three-dimensional numerical investigation of flashback in premixed hydrogen flames within perforated burners. *Combustion and Flame*, 274:113987, 2025a. ISSN 0010-2180.
- Filippo Fruzza, Rachele Lamioni, Temistocle Grenga, , Alessandro Mariotti, Maria Vittoria Salvetti, and Chiara Galletti. Mapping flashback velocity in hydrogen-fueled perforated burners over a broad geometric design space. *International Journal of Hydrogen Energy*, 2025b. ISSN submitted.
- Thomas Gerstner and Michael Griebel. Dimension-adaptive tensor-product quadrature. *Computing*, 71(1):65–87, 2003.
- Tsutomu Ishigami and Toshimitsu Homma. An importance quantification technique in uncertainty analysis for computer models. In *[1990] Proceedings. First international symposium on uncertainty modeling and analysis*, pages 398–403. IEEE, 1990.
- Xiang Ma and Nicholas Zabarar. An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations. *Journal of Computational Physics*, 228(8):3084–3113, 2009.
- Chiara Piazzola and Lorenzo Tamellini. The sparse grids matlab kit—a matlab implementation of sparse grids for high-dimensional function approximation and uncertainty quantification. *arXiv preprint arXiv:2203.09314*, 2022.
- Andrea Saltelli, Paola Annoni, Ivano Azzini, Francesca Campolongo, Marco Ratto, and Stefano Tarantola. Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. *Computer physics communications*, 181(2):259–270, 2010.
- Sergey Smolyak. Quadrature and interpolation formulas for tensor products of certain classes of functions. In *Soviet Math. Dokl.*, volume 4, pages 240–243, 1963.