Iranian Journal of Numerical Analysis and Optimization
Vol. ??, No. ??, ??, pp ??
https://doi.org/??
https://ijnao.um.ac.ir/

Research Article

A stable and high-order accurate D-RBF-PU method via hybrid kernels

M. Esmaeilbeigi*, and R. Cavoretto

Abstract

The direct radial basis function partition of unity (D-RBF-PU) method is a powerful mesh-free technique for solving partial differential equations (PDEs). However, its performance is inherently constrained by a critical trade-off between accuracy and numerical stability. Finite-order smoothness kernels, such as thin plate splines and Matérn kernels, ensure well-conditioned systems but severely limit the convergence rate. Conversely, infinitely smooth kernels, such as Gaussians, yield high accuracy but suffer from severe ill-conditioning, leading to computational breakdown at finer discretizations. In this paper, we propose a novel approach using hybrid kernels, strategically combining Gaussian kernels with finitely smooth kernels to effectively bridge this gap.

Numerical experiments on 2D benchmark elliptic PDEs demonstrate that the proposed hybrid formulations achieve an optimal balance. Specifically, hybrid kernels exhibit robust, high-order convergence rates that remain highly competitive with, and in certain structural combinations significantly surpass, the accuracy of pure infinitely smooth bases. Simultaneously, they maintain manageable local and global condition numbers, allowing the method to

*Corresponding author

Received ??? ; revised ??? ; accepted ???

Mohsen Esmaeilbeigi
Faculty of Mathematical and Statistical Sciences, Malayer University, Malayer, Iran. e-mail: m.esmaeilbeigi@malayeru.ac.ir

Roberto Cavoretto
Department of Mathematics "Giuseppe Peano", University of Torino, Torino, Italy. e-mail: roberto.cavoretto@unito.it

How to cite this article
Esmaeilbeigi, M. and Cavoretto, R., A stable and high-order accurate D-RBF-PU method via hybrid kernels. *Iran. J. Numer. Anal. Optim.*, ??; ??(??): ??-??. ??

successfully scale to finer spatial resolutions where traditional infinitely smooth kernels fail. These results confirm that the hybrid kernel approach significantly enhances the robustness and accuracy of the D-RBF-PU method, offering a highly practical solution for large-scale simulations.

AMS subject classifications (2020): Primary 65N35; Secondary 65D05, 41A30, 65F35.

Keywords: Radial basis functions; Partition of unity; Hybrid kernels; Accuracy-stability trade-off.

1 Introduction

Radial basis function (RBF) methods offer a flexible framework for accurate approximation and have been successfully applied in numerous computational settings. These methods have proven particularly valuable in solving partial differential equations (PDEs), which are essential for modeling physical phenomena involving spatial and temporal variations. PDEs describe quantities that fluctuate with respect to geographical and temporal factors, enabling precise encapsulation of system behavior and prediction from initial conditions and external influences [1]. Despite their numerous advantages, global RBF approximations often become inefficient for large-scale problems due to their high computational costs and severe stability challenges.

This limitation has prompted the development of specialized localized techniques, based on the foundational partition of unity method (PUM) pioneered by Babuška and Melenk [16], which led to the development of RBF partition of unity (RBF-PU) methods [20, 7, 2, 3, 13].

In recent years, RBF methods, particularly their localized variants like RBF-PU [15], have transitioned from experimental approaches for small-scale PDE problems to prominent techniques for large-scale PDE simulations. To further enhance the computational efficiency of these localized methods, the direct RBF-PU (D-RBF-PU) method was introduced [17]. This innovative approach utilizes a direct discretization strategy that completely avoids computing the derivatives of PU weights and the lower derivatives of local RBF approximants. By effectively embedding the RBF Finite Difference (RBF-FD) method within a PU framework, the D-RBF-PU method offers a remarkably simpler and more efficient scheme for solving PDEs. Furthermore, it permits the use of discontinuous PU weight functions, which significantly reduces the overall computational expense.

A critical consideration in the D-RBF-PU method is the choice of the basis kernel. In its initial formulation [17], finitely smooth kernels without a shape parameter, such as polyharmonic

splines, were employed to mitigate computational instability. Subsequent advances incorporated finitely smooth kernels with a shape parameter [4]. However, these choices inherently compromise the approximation accuracy. Conversely, utilizing infinitely smooth kernels with a shape parameter (e.g., Gaussian) promises exceptionally high accuracy but introduces severe ill-conditioning and stability breakdowns, especially at finer resolutions.

To circumvent this accuracy-stability dilemma, a rich literature on stabilizing smooth RBF computations has evolved, encompassing advanced methodologies such as the RBF-QR algorithm [12, 11], generalized PDE-solving frameworks [10], the Contour–Padé method [11], and stable basis expansions via Hilbert–Schmidt SVD [8]. While highly powerful, these change-of-basis or complex-plane transformation strategies often induce nontrivial computational overhead or present structural complexities when integrated into highly localized, irregularly bounded partition patches.

In this article, we focus on the D-RBF-PU method for solving benchmark elliptic PDEs, specifically addressing this critical trade-off between accuracy and numerical stability. Departing from the traditional reliance on pure finitely or infinitely smooth kernels, we introduce *hybrid kernels* as a novel and robust solution.

While kernel blending concepts have established antecedents in Gaussian Process regression and machine learning for multi-scale feature representation—including our prior work on hybrid kernels in statistical learning [6]—as well as recent applications in localized meshfree variants [18, 14], their structural mechanism in localized numerical PDE analysis is fundamentally distinct. Specifically, the present work introduces the hybrid formulation not merely as a multi-scale fitting tool, but as a mathematically grounded, localized spectral floor regularizer designed to mitigate severe numerical ill-conditioning in asymmetric PDE collocation systems within the D-RBF-PU architecture. Furthermore, the positioning of this contribution must be explicitly distinguished from previous work in [4]. Although [4] enhanced the flexibility and accuracy of the direct RBF-PU method by transitioning from parameter-free polyharmonic spline kernels to local radial kernels dependent on a shape parameter augmented with local polynomial expansions, it still operated within the framework of pure single-kernel representations. Consequently, it remained susceptible to the classic accuracy-stability trade-off when smooth components are scaled. In contrast, the current study relies on a fundamental structural hybridization of the kernel itself ($\phi_h = \alpha\phi_{\text{inf}} + \beta\phi_{\text{fin}}$). This ensures robust computational resilience even under fixed shape parameters and in the flat limit ($\varepsilon \to 0$), completely bypassing the severe ill-conditioning inherent to pure smooth bases without requiring intricate polynomial augmentations for stability.

These hybrid formulations strategically blend the mathematical strengths of finite-smoothness and infinitely smooth kernels, aiming to achieve an optimal balance that guarantees both high-order accuracy and well-conditioned linear systems.

The remainder of this paper is organized as follows: Section 2 outlines the mathematical formulation of the D-RBF-PU method and discusses the profound impact of kernel selection on its performance. Section 3 introduces the proposed hybrid kernels, detailing how they provide a balanced approach by combining the properties of finite- and infinitely smooth kernels. Section 4 presents comprehensive numerical experiments on 2D benchmark elliptic PDEs, demonstrating significant improvements in both accuracy and stability. Finally, section 5 provides concluding remarks and describes perspectives for future research.

2 The direct RBF-PU method

The D-RBF-PU method is a highly efficient mesh-free local approximation scheme for solving PDEs. This methodology builds upon the standard RBF-PU framework, but introduces a paradigm shift by directly approximating the differential operators. This direct approach substantially simplifies the computational procedures and improves overall efficiency [17]. In this section, we review the mathematical foundations of the D-RBF-PU method, detail its step-by-step construction, and highlight the crucial role of kernel selection in mediating the fundamental trade-off between approximation accuracy and numerical stability.

2.1 Partition of unity framework

Let $\{\Omega_\ell\}_{\ell=1}^{N_c}$ be an open and bounded overlapping covering of a given domain $\Omega \subset \mathbb{R}^d$. A PU subordinate to this covering is a family of compactly supported nonnegative functions $\{w_\ell\}_{\ell=1}^{N_c}$ that satisfy two primary conditions [21]:

1. The support of each weight function is strictly contained within its corresponding patch:

$$\operatorname{supp}(w_\ell) \subseteq \Omega_\ell, \quad \text{for } \ell = 1, \ldots, N_c.$$

2. The sum of the weight functions equals exactly one everywhere in the domain:

$$\sum_{\ell=1}^{N_c} w_\ell(x) = 1, \quad \text{for all } x \in \Omega.$$

Given local approximants s_ℓ of an unknown function u defined in individual patches Ω_ℓ, a globally smooth approximation $s(x)$ is constructed by blending these local functions using the

PU weights:

$$s(x) = \sum_{\ell=1}^{N_c} w_\ell(x) s_\ell(x).$$

In the context of RBF approximations, the local function $s_\ell(x)$ is typically represented as a linear combination of RBFs and polynomial terms:

$$s_\ell(x) = \sum_{j \in J_\ell} c_j \phi(\|x - x_j\|_2) + \sum_{k=1}^{Q} b_k p_k(x),$$

where $J_\ell = \{j : x_j \in X_\ell\}$ represents the indices of the nodes that reside in patch Ω_ℓ, ϕ is the selected radial kernel, and the polynomial terms are necessary to ensure a well-posed system if ϕ is conditionally positive definite (CPD).

2.2 RBF-PU Construction

Let $\phi : \mathbb{R}^d \to \mathbb{R}$ be a CPD RBF of order n. We denote the set of global trial points by $X = \{x_1, \ldots, x_N\} \subset \Omega$, and the localized trial points within a specific patch Ω_ℓ by $X_\ell = X \cap \Omega_\ell$. The local approximation space is formally defined as follows:

$$V_\ell = \text{span}\{\phi(\cdot - x_j) : x_j \in X_\ell\} \oplus \mathbb{P}_{n-1}(\mathbb{R}^d),$$

where $\mathbb{P}_{n-1}(\mathbb{R}^d)$ denotes the space of d-variable polynomials of total degree at most $n-1$, and Q represents its dimension. It is worth noting that if a strictly positive definite (SPD) kernel is employed, the polynomial augmentation can be completely omitted.

Local approximants $s_\ell \in V_\ell$ are determined by interpolating the data values at the local nodes X_ℓ. Enforcing the interpolation conditions, along with the polynomial constraint

$$\sum_{j \in J_\ell} c_j p_k(x_j) = 0,$$

leads to a classic saddle-point linear system:

$$\begin{bmatrix} \Phi & P \\ P^T & 0 \end{bmatrix} \begin{bmatrix} \mathbf{c} \\ \mathbf{b} \end{bmatrix} = \begin{bmatrix} \mathbf{u}|_{X_\ell} \\ \mathbf{0} \end{bmatrix},$$

where the interpolation matrix entries are given by $\Phi_{ij} = \phi(\|x_i - x_j\|_2)$, and the entries of the polynomial matrix are $P_{jk} = p_k(x_j)$ for a chosen basis $\{p_1, \ldots, p_Q\}$ of $\mathbb{P}_{n-1}$.

2.3 The direct RBF-PU method

Consider a general boundary value PDE problem defined as

$$\begin{aligned} Lu(x) &= f(x), \qquad x \in \Omega, \\ Bu(x) &= g(x), \qquad x \in \Gamma, \end{aligned}$$

where $\Omega \subset \mathbb{R}^d$ is the computational domain, $\Gamma = \partial\Omega$ is its boundary, and L and B are linear differential operators governing the interior and boundary physics, respectively.

In traditional PU methods, evaluating the action of a differential operator L on the global approximation $s(x)$ requires applying the operator to the product $w_\ell(x)s_\ell(x)$. By the product rule, this generates higher-order cross-derivatives of both the PU weights w_ℓ and the local approximants s_ℓ, significantly increasing the mathematical complexity and computational overhead.

The D-RBF-PU method circumvents this bottleneck by proposing a *direct* approximation of the PDE operators. Instead of differentiating the final global function, the method constructs global approximations for the derivatives directly:

$$Lu(x) \approx s^L(x) := \sum_{\ell=1}^{N_c} w_\ell(x) s_\ell^L(x),$$

$$Bu(x) \approx s^B(x) := \sum_{\ell=1}^{N_c} w_\ell(x) s_\ell^B(x).$$

Here, s_ℓ^L and s_ℓ^B are localized RBF approximants of the functions Lu and Bu, respectively. This structure completely bypasses the need to compute any derivatives of the PU weights w_ℓ, thereby significantly reducing the computational cost and simplifying the implementation.

2.4 Accuracy-stability trade-off

A central and persistent challenge in the implementation of the D-RBF-PU method is the inherent trade-off between numerical accuracy and matrix stability, a balance that is dictated entirely by the choice of the base kernel ϕ.

Finite-order smoothness kernels, such as the thin plate spline (TPS), offer robust numerical stability and well-conditioned local matrices. This stability arises largely because they do not rely on a tunable shape parameter. However, their limited algebraic smoothness inherently restricts the maximum achievable convergence rate, making them less suitable for problems demanding extreme high-resolution accuracy [7].

In contrast, infinitely smooth kernels equipped with a shape parameter ε, such as the Gaussian RBF defined by

$$\phi(r) = e^{-\varepsilon^2 r^2},$$

have the potential to achieve spectral or exponential convergence rates under optimal conditions. However, this exceptional accuracy comes at a severe cost driven by two distinct ill-conditioning mechanisms. On one hand, under the flat limit ($\varepsilon \to 0$) at a fixed node distribution, the basis functions physically flatten out, causing the columns of the local interpolation matrices to approach a constant vector and inducing a structural rank deficiency. On the other hand, under the refinement limit ($h \to 0$) at a fixed shape parameter, severe clustering of nodes leads to highly correlated matrix rows due to the dense oversampling of smooth profiles. Both scenarios result in an exponential growth of the condition number, often causing catastrophic numerical breakdowns before the high-order accuracy potential can be practically realized [9].

2.5 Direct approximation and system construction

To compute the local approximants without differentiating the weights, the D-RBF-PU method constructs $s_\ell^L(x)$ and $s_\ell^B(x)$ using generalized Lagrange basis functions:

$$s_\ell^L(x) = \sum_{j \in J_\ell} \mathcal{L}u_j^*(\ell; x)\, u(x_j), \quad \text{and} \quad s_\ell^B(x) = \sum_{j \in J_\ell} \mathcal{B}u_j^*(\ell; x)\, u(x_j).$$

The basis functions $\mathcal{L}u_j^*(\ell; x)$ are obtained by solving a modified local interpolation system in which the target functional is replaced by the action of the operator L on the basis functions. For a given evaluation point $x \in \Omega_\ell$, the system is

$$\begin{bmatrix} \Phi & P \\ P^T & 0 \end{bmatrix} \begin{bmatrix} \mathcal{L}u^*(\ell;x) \\ \mathcal{L}v^*(\ell;x) \end{bmatrix} = \begin{bmatrix} L\phi(\cdot - x)|_{X_\ell} \\ Lp(x) \end{bmatrix}. \tag{1}$$

A strictly analogous linear system is constructed and solved to obtain $\mathcal{B}u^*(\ell;x)$, simply by replacing the differential operator L with the boundary operator B.

The global approximations are then assembled via straightforward summation:

$$s^L(x) = \sum_{\ell=1}^{N_c} \sum_{j\in J_\ell} w_\ell(x)\, \mathcal{L}u_j^*(\ell;x)\, u(x_j),$$

$$s^B(x) = \sum_{\ell=1}^{N_c} \sum_{j\in J_\ell} w_\ell(x)\, \mathcal{B}u_j^*(\ell;x)\, u(x_j).$$

2.6 Final collocation system

Let $Y = Y_\Omega \cup Y_\Gamma$ denote the total set of collocation points, partitioned into interior points Y_Ω and boundary points Y_Γ. The governing PDE and its boundary conditions are enforced discretely at these locations:

$$\begin{aligned} (Lu)(y_k) &= f(y_k), \qquad \text{for } y_k \in Y_\Omega, \\ (Bu)(y_k) &= g(y_k), \qquad \text{for } y_k \in Y_\Gamma. \end{aligned}$$

Substituting the global expansions s^L and s^B into these collocation conditions yields the final global sparse discrete system:

$$\begin{bmatrix} A_L \\ A_B \end{bmatrix} \mathbf{u} = \begin{bmatrix} f|_{Y_\Omega} \\ g|_{Y_\Gamma} \end{bmatrix}, \tag{2}$$

where $\mathbf{u} = [u(x_1), \ldots, u(x_N)]^T$ is the vector of unknown nodal values. The elements of the global matrix are assembled efficiently by accumulating the localized Lagrange weights. Specifically, for an interior point $y_k \in Y_\Omega$:

$$A_L(k,j) = \sum_{\ell\in I(y_k)} w_\ell(y_k)\, \mathcal{L}u_j^*(\ell;y_k),$$

and for a boundary point $y_k \in Y_\Gamma$:

$$A_B(k,j) = \sum_{\ell \in I(y_k)} w_\ell(y_k)\, \mathcal{B}u_j^*(\ell; y_k),$$

where $I(y_k)$ represents the index set of all patches Ω_ℓ that cover the specific evaluation point y_k. Thus

$$\mathbf{A}_{\text{global}}\mathbf{U} = \mathbf{F}. \tag{3}$$

Remark 1. While Proposition 1 in the next section rigorously guarantees the nonsingularity and unisolvence of the localized interpolation systems constructed on each individual partition patch Ω_j, the well-posedness of the final assembled global asymmetric collocation matrix $\mathbf{A}_{\text{global}}$ in (3) warrants careful qualification. In asymmetric Kansa-type RBF collocation methods applied to elliptic operator equations, establishing a universal proof for global unisolvence is a famously difficult open problem. Even for standard smooth kernels, counterexamples involving specialized node configurations where the global matrix becomes singular can theoretically be engineered.

Recently, significant progress has been achieved in this domain; for instance, recent mathematical frameworks (e.g., [19]) have successfully established formal global unisolvence criteria for Kansa-type collocation under elliptic setups specifically using CPD polyharmonic splines (PHS). For the hybrid formulation proposed herein, since the global matrix is a blended sparse assembly of an infinitely smooth component and a PHS-like or Matérn component, we intuitively expect the global system to inherit the robust well-posedness floor of its finitely smooth constituent. However, a formal, mathematically rigorous proof of global unisolvence for hybrid-kernel asymmetric collocation remains a recognized theoretical gap, which we defer to dedicated future investigations in approximation theory.

In summary, the implementation of the D-RBF-PU framework inherently requires the solution of two distinct categories of linear systems:

1. **Local small-scale systems** (1), defined in each overlapping patch Ω_ℓ, to compute operator-dependent Lagrange functions. These dense systems are of size $(|J_\ell| + Q) \times (|J_\ell| + Q)$.

2. **A global collocation system** (2), which is sparse and of size $N \times N$, assembled directly from the local components to produce the final PDE solution.

Crucially, both types of systems are highly susceptible to numerical ill-conditioning, making the strategic mitigation of this issue—such as through the proposed hybrid kernels—vital for the practical viability of the method.

3 Hybrid kernels: A balanced approach

In this section, we propose the integration of *hybrid kernels* as a novel and strategic mechanism to govern the critical trade-off between approximation accuracy and numerical stability in the D-RBF-PU method. By intelligently blending the mathematical strengths of infinitely smooth and finitely smooth RBFs, hybrid kernels offer a highly flexible and robust framework to enhance the performance of mesh-free solvers, particularly for complex elliptic PDEs.

3.1 Formulation of hybrid kernels

A hybrid kernel is fundamentally defined as a positive linear combination of two distinct types of RBFs: One possessing infinite smoothness and the other characterized by finite smoothness. Formally, we construct the hybrid kernel ϕ_h as follows:

$$\phi_h(r) = \alpha\phi_{\text{inf}}(r,\varepsilon) + \beta\phi_{\text{fin}}(r),$$

where

- $\phi_{\text{inf}}(r,\varepsilon)$ is an infinitely smooth RBF equipped with a shape parameter $\varepsilon > 0$.
- $\phi_{\text{fin}}(r)$ is a finitely smooth RBF, which typically does not depend on a shape parameter.
- $\alpha > 0$ and $\beta > 0$ are tunable weight parameters that control the relative contribution of each constituent kernel. Throughout this study, we fix $\alpha = 1$ as a standard normalization choice, since spectral properties and matrix conditioning are strictly governed by the relative weight ratio β/α; altering α merely scales the entire matrix uniformly without affecting its intrinsic condition number.
- $r = \|x - y\|_2$ denotes the standard Euclidean distance between any two points $x, y \in \mathbb{R}^d$.

The core mathematical motivation for this formulation lies in the use of the complementary spectral properties of the two types of kernels. Infinitely smooth kernels provide exponential or spectral convergence rates; however, their associated interpolation matrices exhibit exponentially decaying eigenvalues, leading to rapid ill-conditioning as the grid is refined or as $\varepsilon \to 0$. Conversely, finitely smooth kernels exhibit algebraically decaying eigenvalues, which inherently prevents extreme condition number explosion but limits the maximum achievable accuracy. The hybrid kernel ϕ_h essentially uses ϕ_{fin} to provide a stabilizing "floor" for the matrix spectrum, while allowing ϕ_{inf} to drive the high-order accuracy of the approximation.

3.2 Definiteness, kernel families, and system augmentation

To guarantee the well-posedness and unique solvability of localized collocation systems within the D-RBF-PU framework, it is crucial to establish the definiteness properties of the proposed hybrid kernel ϕ_h. While the sum of two SPD kernels is trivially SPD, the case in which an SPD kernel is combined with a CPD kernel requires careful analysis.

We formalize the characterization of the hybrid formulation by the following proposition:

Proposition 1. Let $\phi_{\text{inf}}(r)$ be an SPD kernel in $\mathbb{R}^d$, and let $\phi_{\text{fin}}(r)$ be a CPD kernel of order n in $\mathbb{R}^d$. Then, for any positive scaling parameters $\alpha > 0$ and $\beta > 0$, the hybrid kernel $\phi_h(r) = \alpha\phi_{\text{inf}}(r) + \beta\phi_{\text{fin}}(r)$ is strictly CPD of order n on $\mathbb{R}^d$.

Proof. Let $X = \{x_1, x_2, \ldots, x_N\} \subset \mathbb{R}^d$ be a set of distinct collocation points. Let A_{inf}, A_{fin}, and A_h denote the evaluation matrices associated with ϕ_{inf}, ϕ_{fin}, and ϕ_h on X, respectively, such that $(A_h)_{ij} = \alpha(A_{\text{inf}})_{ij} + \beta(A_{\text{fin}})_{ij}$.

Let $\Pi_{n-1}(\mathbb{R}^d)$ denote the space of d-variate polynomials of total degree up to $n-1$. Consider any nonzero coefficient vector $c = [c_1, c_2, \ldots, c_N]^T \in \mathbb{R}^N \setminus \{0\}$ that satisfies the polynomial vanishing constraints:

$$\sum_{i=1}^{N} c_i p(x_i) = 0, \quad \forall p \in \Pi_{n-1}(\mathbb{R}^d).$$

We examine the quadratic form associated with the hybrid matrix A_h on this constrained subspace:

$$c^T A_h c = \alpha \left(c^T A_{\text{inf}} c\right) + \beta \left(c^T A_{\text{fin}} c\right).$$

Since ϕ_{inf} is SPD, its interpolation matrix A_{inf} is SPD over the entire Euclidean space, meaning $c^T A_{\text{inf}} c > 0$ holds universally for any nonzero vector c. Moreover, since ϕ_{fin} is CPD of order n, its matrix satisfies $c^T A_{\text{fin}} c \geq 0$ for all nonzero vectors c in $\ker(P^T)$. Given $\alpha > 0$ and $\beta > 0$, the linear combination yields $c^T A_h c > 0$. Consequently, A_h inherits strict conditional positive definiteness of order n, guaranteeing the unisolvence and nonsingularity of the augmented local interpolation systems [7, 21]. □

Remark 2. It is essential to clarify that Proposition 1 strictly guarantees the unisolvence and nonsingularity (invertibility) of the local linear systems. It does not, however, provide a quantitative bound on the spectral condition number. The practical stabilization benefit—namely, capping the growth of condition numbers under spatial refinement—is a distinct physical prop-

erty derived from the heuristic spectral floor analysis on the constrained subspace $V = \ker(P^T)$ presented in Section 3.3.

Based on Proposition 1, the requirement for polynomial augmentation and the resulting size of the local saddle-point system matrix $(|J_\ell| + Q) \times (|J_\ell| + Q)$ on a patch Ω_ℓ are systematically classified as follows:

1. **SPD Constituent Pairing (GA + Matérn C^4):** Since both components are SPD ($n = 0$), the hybrid kernel ϕ_h is unconditionally SPD. The local system is invertible without polynomial augmentation ($Q = 0$), yielding an unaugmented local system matrix of size $|J_\ell| \times |J_\ell|$.

2. **CPD Constituent Pairing of Order $n = 2$ (GA + TPS / GA + Cubic):** Since ϕ_{fin} is CPD of order $n = 2$, ϕ_h is CPD of order 2. Ensuring well-posedness requires augmenting the local basis with degree-1 polynomials ($p \in \Pi_1(\mathbb{R}^d)$). In $d = 2$ spatial dimensions, this introduces $Q = \binom{1+2}{2} = 3$ polynomial terms $(1, x, y)$, expanding the local augmented saddle-point matrix to size $(|J_\ell| + 3) \times (|J_\ell| + 3)$.

To construct this framework under rigorous guidelines, we utilize standard radial families. Commonly employed infinitely smooth kernels include:

$$\begin{aligned} \text{Gaussian (GA):} \quad & \phi_{\text{inf}}(r) = e^{-(\varepsilon r)^2}, \\ \text{Multiquadric (MQ):} \quad & \phi_{\text{inf}}(r) = \sqrt{1 + (\varepsilon r)^2}, \\ \text{Inverse MultiQuadric (IMQ):} \quad & \phi_{\text{inf}}(r) = \frac{1}{\sqrt{1 + (\varepsilon r)^2}}. \end{aligned}$$

In our formulation, the finitely smooth kernels $\phi_{\text{fin}}(r)$ are deliberately kept parameter-free, meaning that $r = \|x - x_j\|_2$ represents the unscaled Euclidean distance without the shape parameter ε. While employing a fixed length scale (or omitting shape parameter tuning) for $\phi_{\text{fin}}(r)$ may be sub-optimal in terms of approximation accuracy compared to an independently scaled constituent, this represents a deliberate stability-oriented trade-off. By decoupling the finitely smooth component from ε, we ensure that $\phi_{\text{fin}}(r)$ provides an invariant, numerically stable spectral baseline that is entirely unaffected by ill-conditioning in the flat limit ($\varepsilon \to 0$) of the infinitely smooth component $\phi_{\text{inf}}(r)$. Furthermore, the Matérn kernel belongs to the global smoothness class $C^4(\mathbb{R}^d)$ for spatial dimensions $d \leq 3$ [7].

3.3 Heuristic error analysis and accuracy-stability trade-off

While a rigorous derivation of native space error bounds for the hybrid formulation within the localized PU framework is deferred to future research, we provide a heuristic argument to illustrate how the blending parameter β systematically governs the trade-off between approximation accuracy and numerical stability.

In standard RBF collocation, theoretical error bounds are governed by the power function $P_{\phi,X}(x)$, whereas numerical stability is dictated by the conditioning of the collocation linear system. To account for both SPD (e.g., Matérn) and CPD (e.g., TPS or Cubic of order m) kernels, the spectral behavior must be evaluated on the polynomial-constrained subspace:

$$V = \ker(P^T) = \left\{c \in \mathbb{R}^N \setminus \{0\} : P^T c = 0\right\},$$

where $P \in \mathbb{R}^{N\times Q}$ represents the local polynomial basis matrix (with $V = \mathbb{R}^N$ for unaugmented SPD kernels). On this subspace V, the constrained quadratic form $c^T Ac$ is strictly positive, matching the augmented saddle-point formulation analyzed in Proposition 1.

We define the constrained minimum eigenvalue (or minimal Rayleigh quotient on V) as follows:

$$\lambda^V_{\min}(A) := \min_{c\in V,\, \|c\|_2=1} c^T Ac > 0.$$

For the infinitely smooth component ϕ_{inf} (e.g., GA), $\lambda^V_{\min}(A_{\text{inf}})$ decays exponentially ($\sim e^{-c/\varepsilon^2}$), leading to severe ill-conditioning and computational collapse as $h \to 0$ or $\varepsilon \to 0$. Conversely, the finitely smooth stabilizer ϕ_{fin} exhibits an algebraic spectral decay, maintaining a strictly positive constrained lower bound $\lambda^V_{\min}(A_{\text{fin}}) \sim \mathcal{O}(h^{2k})$ for a positive integer k determined by the kernel smoothness order.

By construction of the hybrid collocation system $A_h = \alpha A_{\text{inf}} + \beta A_{\text{fin}}$, applying the Rayleigh-Ritz principle restricted to the constrained subspace V yields

$$\lambda^V_{\min}(A_h) \geq \alpha\lambda^V_{\min}(A_{\text{inf}}) + \beta\lambda^V_{\min}(A_{\text{fin}}).$$

When the spatial discretization is highly refined ($h \to 0$), $\lambda^V_{\min}(A_{\text{inf}})$ rapidly drops below the machine precision ($\approx 10^{-16}$). At this critical threshold, the constrained term $\beta\lambda^V_{\min}(A_{\text{fin}})$ becomes dominant, establishing a strict, positive spectral "floor" on V that scales linearly with the blending parameter β. This effectively guarantees the well-posedness of the augmented saddle-

point system solved in Equation (1). Consequently, β functions as a continuous regularization parameter governing the trade-off as follows:

- **Accuracy-Dominated Regime ($\beta \to 0$):** The hybrid kernel behaves almost purely as an infinitely smooth RBF, recapturing high-order convergence rates in the pre-asymptotic regime, while remaining susceptible to numerical saturation on ultra-fine grids.

- **Stability-Dominated Regime (Larger β):** The algebraic stability floor $\beta\lambda^V_{\min}(A_{\text{fin}})$ elevates the constrained spectrum, cushioning the system against spectral collapse. This bounds the condition number of the augmented matrix and enables the D-RBF-PU solver to operate stably on fine spatial resolutions, transitioning the asymptotic convergence rate to the algebraic order of ϕ_{fin}.

This constrained interplay demonstrates how tuning β provides a mathematically sound mechanism to bypass the standard RBF uncertainty principle while preserving full compatibility with Proposition 1.

3.4 Architectural comparison with alternative stabilization techniques

To properly contextualize the proposed hybrid kernel framework within the broader spectrum of RBF stabilization literature, it is instructive to contrast its mathematical and implementation philosophy with established alternatives, such as the RBF-QR algorithm [12], the Contour–Padé method [11], stable bases via Hilbert–Schmidt SVD (HS-SVD) [8], and ε-adaptive strategies [5]. While these methods are highly effective at mitigating the ill-conditioning associated with the flat limit ($\varepsilon \to 0$), they fundamentally differ from the hybrid kernel approach in three major aspects:

1. **Algorithmic Structure and Pre-processing Complexity:** Methods like RBF-QR rely on complex analytical expansions (e.g., spherical harmonics or Chebyshev polynomials) to transition to a stable basis, whereas HS-SVD requires dense matrix decompositions to establish an orthogonal basis in the native space. These operations introduce non-trivial analytical and pre-processing transformations per collocation point. In contrast, the hybrid kernel modifies only the intrinsic entry definitions of the evaluation matrix: $(A_h)_{ij} = \alpha(A_{\text{inf}})_{ij} + \beta(A_{\text{fin}})_{ij}$. Because it operates entirely within the standard RBF

basis, it avoids additional change-of-basis factorizations or special function evaluations during linear system assembly.

2. **Algorithmic Simplicity and Implementation Ease:** Implementing change-of-basis techniques (RBF-QR/HS-SVD) or complex-plane contour integrations (Contour–Padé) requires a substantial rewrite of the standard RBF interpolation pipeline. This restricts their immediate use as "plug-and-play" components in pre-existing industrial PDE codes. Conversely, the hybrid kernel formulation requires no modifications to the overarching linear solver or the localized matrix assembly routines, offering an out-of-the-box, structurally straightforward stabilization tool for practitioners.

3. **Compatibility with Localized Formulations (D-RBF-PU):** Most classical stabilization techniques were originally formulated for global RBF interpolation or specialized geometries (e.g., spheres or torus domains). Extending RBF-QR or Contour–Padé into highly localized, irregularly bounded patches within a PU framework introduces severe mathematical and algorithmic hurdles. The hybrid kernel framework bypasses this geometric restriction entirely, as its stabilizing spectral floor acts symmetrically and isotropically across any local node configuration, regardless of the boundary geometry.

3.5 Integration into the D-RBF-PU Framework

The D-RBF-PU method faces a persistent operational challenge: traditional finitely smooth kernels ensure stable local matrices but yield suboptimal accuracy, while infinitely smooth kernels yield superior accuracy but suffer from computational breakdown at fine resolutions. By adopting ϕ_h, practitioners can actively navigate this trade-off. Increasing the weight α emphasizes accuracy by leveraging the rapid convergence of ϕ_{inf}, whereas increasing β injects robustness and limits the growth of the condition number through ϕ_{fin}.

To integrate hybrid kernels into the D-RBF-PU architecture, we seamlessly replace the standard RBF ϕ with the hybrid formulation ϕ_h. In each local patch Ω_ℓ, the basic approximant is redefined as:

$$s_\ell(x) = \sum_{j \in J_\ell} c_j \phi_h(\|x - x_j\|_2) + \sum_{k=1}^{Q} b_k p_k(x).$$

For the direct approximation of differential operators, local expansions take the form $s_\ell^L(x) = \sum_{j \in J_\ell} \mathcal{L}u_j^*(\ell; x)u(x_j)$ and $s_\ell^B(x) = \sum_{j \in J_\ell} \mathcal{B}u_j^*(\ell; x)u(x_j)$. The modified generalized Lagrange

functions are computed by solving the localized system:

$$\begin{bmatrix} \Phi_h & P \\ P^T & 0 \end{bmatrix} \begin{bmatrix} \mathcal{L}u^*(\ell;x) \\ \mathcal{L}v^*(\ell;x) \end{bmatrix} = \begin{bmatrix} L\phi_h(\cdot - x)|_{X_\ell} \\ Lp(x) \end{bmatrix}.$$

A significant computational advantage of this approach is that, due to the linearity of the differential operators L and B, the action on the hybrid kernel is computed trivially by superposition:

$$L\phi_h(r) = \alpha L\phi_{\text{inf}}(r,\varepsilon) + \beta L\phi_{\text{fin}}(r).$$

Therefore, evaluating the vectors on the right-hand side does not require complex derivations beyond those already established for standard kernels.

Finally, the globally approximated fields remain consistent with the original formulation:

$$s^L(x) = \sum_{\ell=1}^{N_c} w_\ell(x)\, s_\ell^L(x),$$
$$s^B(x) = \sum_{\ell=1}^{N_c} w_\ell(x)\, s_\ell^B(x),$$

and the global sparse collocation system is assembled precisely as detailed in Section 2. By transitioning to the operational matrices A_L and A_B derived from ϕ_h, the overall computational pipeline benefits from significantly improved conditioning at the local level, directly translating to enhanced stability and solvability of the final global system.

4 Numerical experiments and discussion

In this section, we present a comprehensive set of numerical experiments designed to systematically evaluate the performance of the proposed D-RBF-PU method. The primary objective of these experiments is to empirically demonstrate the inherent accuracy-stability trade-off associated with classical RBFs and to validate the efficacy of the proposed hybrid kernel approach in overcoming these limitations.

To this end, we consider a benchmark linear elliptic PDE in a two-dimensional domain $\Omega = [0,1]^2$, governed by

$$Lu = -\Delta u + u = f \quad \text{in } \Omega.$$

Mixed boundary conditions are imposed to rigorously test the method: Dirichlet conditions are prescribed along the vertical boundaries ($x = 0$ and $x = 1$), while Neumann conditions are applied on the horizontal edges ($y = 0$ and $y = 1$). The forcing term f and the boundary data are analytically extracted from the exact solutions chosen. We investigate two distinct test cases:

- **Test case 1 (The Franke Function):** This widely used benchmark is characterized by severe local variations, steep gradients, and multiple peaks, making it a challenging test for mesh-free interpolants. The exact solution is given by

$$\begin{aligned} u(x,y) = &\frac{3}{4}\exp\left(-\frac{1}{4}\left((9x-2)^2+(9y-2)^2\right)\right) \\ &+\frac{3}{4}\exp\left(-\frac{1}{49}(9x+1)^2-\frac{1}{10}(9y+1)^2\right) \\ &+\frac{1}{2}\exp\left(-\frac{1}{4}\left((9x-7)^2+(9y-3)^2\right)\right) \\ &-\frac{1}{5}\exp\left(-\left((9x-4)^2+(9y-7)^2\right)\right). \end{aligned}$$

- **Test case 2 (A Smooth Trigonometric Function):** To evaluate the spectral convergence properties on a globally smooth profile, we employ

$$u(x,y) = \sin(\pi x)\cos\left(\frac{\pi y}{2}\right).$$

Spatial discretization utilizes uniformly distributed collocation points with a nominal node spacing parameter $h = (c_d N)^{-1/d}$, which serves as a global refinement indicator for convergence analysis and plotting purposes. While the rigorous geometric fill distance is classically defined as $h_{X,\Omega} = \sup_{x\in\Omega}\min_{x_j\in X}\|x - x_j\|_2$, for our uniformly distributed trial points, the parameter h provides a computationally straightforward and asymptotically equivalent measure of the discretization scale. Here, N represents the total number of trial points, $d = 2$ is the spatial dimension, and $c_d = \mathrm{vol}(B)/\mathrm{vol}(\Omega)$, where B denotes the bounding box of Ω. For the standard test domain $\Omega = [0,1]^2$, the bounding box coincides with the domain itself ($B = \Omega$), yielding $c_d = 1$ and an effective spatial fill distance of $h = N^{-1/2} = 1/\sqrt{N}$. This explicit relation maps the nodal spacing h directly to the refinement parameter $\sqrt{N}$ depicted on the horizontal axes of our convergence figures. We progressively refine the grid using $N \in \{121, 225, 441, 841, 1681\}$, which corresponds to $\sqrt{N} \in \{11, 15, 21, 29, 41\}$ (or spatial resolution $h \in [1/11, 1/41]$).

To quantitatively assess performance, we monitor three key metrics:

1. The relative ∞-norm error: $E_\infty = \|\mathbf{s} - \mathbf{u}\|_\infty / \|\mathbf{u}\|_\infty$, where $\mathbf{s}$ and $\mathbf{u}$ are the numerical and exact solution vectors, respectively.

2. The numerical convergence order, derived from the slope of the linear regression in the log-log plot of the error versus h.

3. The condition numbers. In the subsequent tables, "Local Cond." denotes the condition number of the local interpolation matrices (averaged across all PU patches), while "Global Cond." reflects the condition number of the final assembled sparse collocation matrix.

To quantify the asymptotic accuracy of the numerical schemes, the global convergence orders p reported in the performance profiles are systematically estimated via a linear least-squares curve fitting applied to the log-log error data against the grid refinement indicator $\sqrt{N}$. Formally, the convergence rate is determined as $p = -a_1$, where a_1 is the optimal slope satisfying the first-degree polynomial relationship $\log(E_\infty) = a_1 \log(\sqrt{N}) + a_2$. For the stable finite-smoothness and proposed hybrid kernel frameworks, this linear regression spans the complete sequence of grid refinements up to $N = 1681$. Conversely, for the pure infinitely smooth configurations, the regression interval is truncated to the stable, pre-breakdown range ($N \leq 441$) to prevent numerical saturation from biasing the estimated convergence rate. All simulations were executed in MATLAB on a workstation equipped with an Intel Core i7 CPU (4.00 GHz) and 16 GB of RAM.

4.1 Baseline: Finitely smooth kernels

Our first set of experiments investigates finitely smooth kernels (Matérn, TPS, and Cubic). Theoretical expectations suggest that these kernels should provide robust numerical stability at the cost of algebraic (lower-order) convergence rates.

Table 1: Performance profile of finitely smooth kernels for Test Case 1 in 2D space ($d = 2$). Matrix size indicates the dimension $(|J_\ell| + Q) \times (|J_\ell| + Q)$ of local collocation systems.

Kernel	Poly. Deg.	N range	Conv.	Local Cond.	Global Cond.
Matérn (C^4)	None ($n=0$)	121–1681	2.27	5e8–5e11	5e4–3e6
TPS	Deg. 1 ($n=2$)	121–1681	1.32	2e6–1e9	8e4–5e6
Cubic	Deg. 1 ($n=2$)	121–1681	2.25	6e7–6e10	2e4–3e6

As illustrated in Table 1 for Test case 1, the numerical results align perfectly with theoretical predictions. The Matérn and Cubic kernels exhibit moderate convergence orders of approximately 2.27 and 2.25, respectively, while the TPS kernel yields a slower rate of 1.32. Notably, the method successfully computes solutions up to the finest grid ($N = 1681$). Although the local

condition numbers increase with grid refinement, they remain bounded within a manageable algebraic limit (up to $\sim 10^{11}$). Furthermore, the PU blending ensures that the global sparse system remains exceptionally well-conditioned (peaking at 10^6), allowing for efficient iterative or direct solutions.

4.2 The stability bottleneck: Infinitely smooth kernels

Next, we evaluate infinitely smooth kernels (GA, MQ, and IMQ), which are historically favored for their spectral convergence but are notorious for severe ill-conditioning as $h \to 0$.

Table 2: Performance profile of pure infinitely smooth kernels. Note: Empirical convergence orders are evaluated over the pre-breakdown regression window $N \in [121, 441]$, as finer grids suffer from severe numerical ill-conditioning.

Kernel	N range	Conv. Order	Local Cond.	Global Cond.
GA ($\varepsilon=2$)	121–441	4.38	2e12–3e16	1e4–9e5
MQ ($\varepsilon=2$)	121–441	3.95	2e8–3e12	4e4–5e5
IMQ ($\varepsilon=2$)	121–441	3.95	3e7–1e11	2e4–4e5

Table 2 reveals the stark contrast between accuracy and stability. For Test case 1, these kernels deliver impressive convergence orders (4.38 for GA, 3.95 for MQ/IMQ). However, this rapid convergence is accompanied by an exponential deterioration in the local condition numbers. For the GA kernel at merely $N = 441$, the local condition number hits the machine precision limit ($\sim 10^{16}$). For any $N > 441$, the local interpolation matrices become numerically singular, preventing the extraction of the generalized Lagrange functions and causing the entire D-RBF-PU algorithm to fail. This completely prohibits the use of pure infinitely smooth kernels on fine resolutions.

4.3 Breaking the barrier: The hybrid kernel approach

To resolve the impasse demonstrated in the previous subsections, we deploy the proposed hybrid kernels, combining the highly accurate GA kernel with stabilizing finitely smooth kernels. The weighting parameters α and β, together with the shape parameter ε, are optimally tuned to restrict the condition number growth while preserving the influence of the GA kernel.

Table 3 summarizes the main numerical results obtained in this study. To ensure a fair and transparent comparison, we note that the reported empirical convergence orders are computed

Table 3: Performance profile of hybrid kernels for Test Case 1.

Kernel Formulation	N range	Conv.	Local Cond.	Global Cond.
GA + Matérn ($\varepsilon=5$, $\alpha=1$, $\beta=10^{-8}$)	121–1681	3.68	1e11–5e15	3e4–2e6
GA + TPS ($\varepsilon=2.5$, $\alpha=1$, $\beta=10^{-6}$)	121–1681	5.10	3e10–1e15	2e4–2e6
GA + Cubic ($\varepsilon=2.5$, $\alpha=1$, $\beta=10^{-8}$)	121–1681	4.24	2e10–3e15	2e4–1e6

Note: Empirical convergence orders are evaluated over the full grid refinement window $N \in [121, 1681]$, made possible by the structural stability of the hybrid spectral floor.

via linear regression over the stable regime of each configuration: for pure smooth kernels, the regression window is truncated to $N \leq 441$ due to catastrophic ill-conditioning at finer resolutions, whereas for hybrid formulations, the window encompasses the full resolution range up to $N = 1681$. By introducing a small fractional weight (β) of a finite-smoothness kernel, we successfully extend the solvability of the method up to the finest grid ($N = 1681$), a feat impossible with the pure GA kernel. Remarkably, this structural stability does not degrade accuracy; instead, convergence orders remain exceptionally high. The GA+TPS combination achieves an outstanding convergence order of 5.10, outperforming both of its individual constituents. Local condition numbers are strictly capped below the threshold of total numerical singularity (up to 10^{15}), and global sparsity maintains the overarching system highly tractable.

To further substantiate these findings across different mathematical landscapes, we analyzed the hybrid kernels using Test case 2. Figure 1 visually maps the convergence trajectories, showing relative ∞-norm error versus $\sqrt{N}$.

The subplots in Figure 1 uniformly confirm that the hybrid mechanism is robust and independent of the specific underlying exact solution. Convergence orders consistently exceed 4.0, verifying that the hybrid kernels seamlessly blend the spectral decay of the GA component with the numerical resilience of the finite-smoothness component. Note that the local and global condition numbers for Test case 2 are omitted here for brevity, as they follow a nearly identical growth trajectory and remain within the same orders of magnitude as those reported for Test case 1 in Table 3 due to the identical spatial node configurations. Ultimately, these results provide strong empirical evidence that the hybrid D-RBF-PU framework is a highly practical, stable, and accurate tool for large-scale PDE simulations.

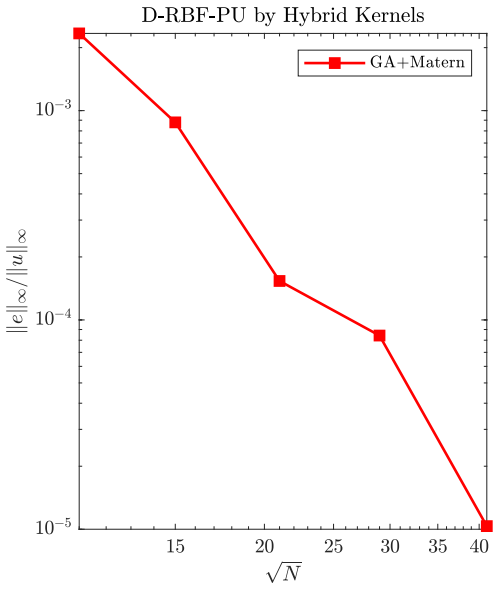

(a) GA + Matérn ($\varepsilon = 2$, $\alpha = 1$, $\beta = 10^{-6}$)
Convergence Order = 4.01

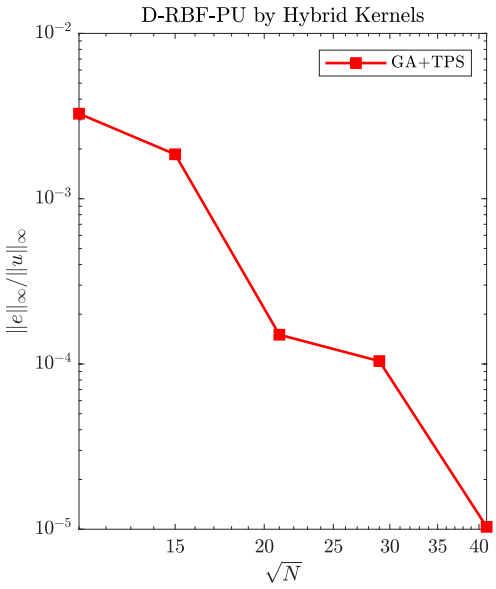

(b) GA + TPS ($\varepsilon = 2$, $\alpha = 1$, $\beta = 10^{-8}$)
Convergence Order = 4.38

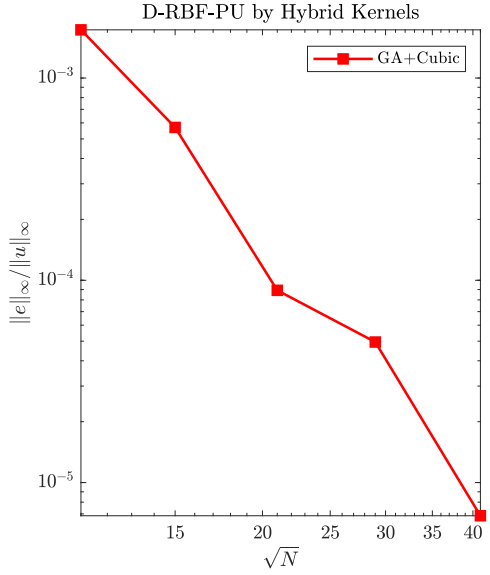

(c) GA + Cubic ($\varepsilon = 1.6$, $\alpha = 1$, $\beta = 10^{-6}$)
Convergence Order = 4.12

Figure 1: Convergence behavior of hybrid kernels for Test Case 2. The plots track the decay of the relative ∞-norm error ($\|e_h\|_\infty/\|u\|_\infty$) against grid refinement ($\sqrt{N}$). All hybrid combinations reliably maintain high-order convergence without succumbing to ill-conditioning. Note that the local and global condition numbers for Test Case 2 are omitted here for brevity, as they follow a nearly identical growth trajectory and remain within the same orders of magnitude as those reported for Test Case 1 in Table 3 due to the identical spatial node configurations.

4.4 Discussion on convergence and implementation aspects

To better appreciate the significance of the results presented in Tables 1–3, it is instructive to contrast the characteristics of the proposed hybrid kernel framework with existing literature in the domain of RBF-PUMs and D-RBF-PU formulations. Traditional RBF-PUM applications utilizing pure finitely smooth profiles, such as the Matérn C^4 or Cubic families, exhibit highly predictable but strictly limited algebraic convergence rates, seldom exceeding order 3.0 in practical spatial discretizations [3]. On the other hand, achieving high-order convergence rates exceeding 4.0 typically mandates the deployment of infinitely smooth kernels like GAs or MQs. However, as demonstrated in Table 2 and widely documented in classical mesh-free theory [7], these pure infinitely smooth configurations inevitably succumb to catastrophic ill-conditioning in the flat limit or under progressive grid refinement.

The hybrid kernel strategy introduced in this study effectively bridges this long-standing gap. By embedding a parameter-free finite-smoothness baseline, the formulation yields convergence rates as high as 5.10 (Table 3) while successfully maintaining complete numerical stability up to the finest grid ($N = 1681$). From an implementation perspective, this framework offers a distinct structural advantage over specialized stabilization paradigms such as the RBF-QR or RBF-GA algorithms. While those techniques successfully bypass ill-conditioning, they require nontrivial modifications to compute expanded basis transformations, particularly when extending to multi-dimensional boundary value problems. Because the proposed hybrid framework keeps matrix condition numbers within manageable bounds via a structural stability floor, it operates seamlessly with standard, unaugmented linear system solvers. This eliminates the need for change-of-basis algorithms, rendering the proposed scheme both highly accurate and remarkably straightforward to implement compared to standard stabilized RBF-PUM variants.

4.5 Parameter sensitivity and practical guidance

To address the crucial question of parameter selection and to ensure full reproducibility of the reported results, we first clarify the procedure used to select the shape parameter ε and blending weight β (fixing $\alpha = 1$ as a standard system normalization). To ensure a fair comparison based on the best achievable performance of each architectural paradigm, each method (both pure baseline kernels and hybrid configurations) was evaluated using its individually optimized shape parameter.

Specifically, for each kernel combination and test problem, the reported pairs (ε, β)—or ε alone for pure baselines—were systematically identified via an exhaustive grid search over $\varepsilon \in [1, 10]$ and $\beta \in \{10^{-4}, 10^{-5}, \ldots, 10^{-10}\}$, selecting the combination that minimized the Root Mean Square (RMS) error.

It is important to emphasize that this comparative strategy evaluates each formulation under its optimal operational regime. For the pure GA baseline, reducing the shape parameter on fine spatial discretizations causes catastrophic numerical breakdown due to severe ill-conditioning, forcing its optimal operational value to settle around $\varepsilon = 2$. Conversely, the spectral floor provided by the hybrid formulation allows the solver to operate stably at smaller shape parameters without matrix singularity. Consequently, the performance reported in Tables 2–3 reflects a comparison between separately optimized configurations, highlighting how hybrid stabilization enables practitioners to safely access smaller shape parameter regimes that remain computationally intractable for pure infinitely smooth RBFs.

This empirical analysis reveals that the D-RBF-PU method is highly robust and operates successfully over a remarkably wide parameter plateau rather than requiring a narrow, hypersensitive window. Because optimal parameters are selected independently for each kernel formulation and boundary value configuration, the interplay between ε and β is governed by three distinct operational regimes:

1. **The Under-Stabilized Regime ($\beta < 10^{-11}$):** If the blending parameter is chosen too small, the finite-smoothness kernel fails to inject a sufficient spectral floor. Consequently, the exponential decay of the GA component dominates, leading to severe ill-conditioning and early matrix singularity on fine resolutions ($N \geq 841$).

2. **The Optimal Blending Plateau ($10^{-9} \leq \beta \leq 10^{-5}$):** Within this wide range, the hybrid kernel achieves an ideal equilibrium. The local condition numbers are bounded safely below the computational breakdown threshold ($\sim 10^{15}$), while the high-order convergence properties driven by ϕ_{inf} are fully preserved. Across different test problems and kernel pairings (e.g., GA+Matérn, GA+TPS, GA+Cubic), the independently optimized blending weights consistently fall within $10^{-8} \leq \beta \leq 10^{-6}$. Minor shifts in optimal values are dictated by problem geometry and kernel smoothness: sharp, multi-peaked profiles favor slightly smaller β values to preserve maximum spectral accuracy, whereas globally smoother solutions accommodate larger stabilization floors.

3. **The Saturation Regime ($\beta > 10^{-3}$):** When β becomes overly dominant, the algebraic convergence behavior of the finitely smooth kernel overrides the high-order accuracy of the infinitely smooth component. In this case, the method remains perfectly stable, but

the convergence order drops asymptotically toward the baseline algebraic rates reported in Table 1.

For practical engineering implementations where the exact analytical solution is unavailable, we recommend a simple two-step heuristic workflow: first, choose a moderate shape parameter based on the local node density, and second, initialize the blending parameter at a safe baseline of $\beta = 10^{-7}$. For fully automated, data-driven applications, practitioners can seamlessly integrate a localized variant of Rippa's algorithm or the Leave-One-Out Cross-Validation (LOOCV) framework [7], minimizing the local residual norm across individual partition patches to systematically extract optimal (ε, β) pairs without manual intervention.

5 Conclusion

In this paper, we investigate the integration of hybrid kernels within the D-RBF-PU method to resolve the fundamental accuracy-stability trade-off in solving elliptic PDEs. By linearly combining infinitely smooth and finitely smooth kernels, the proposed approach successfully leverages the complementary properties of both classes: it achieves the high convergence rates characteristic of infinitely smooth kernels while utilizing the localized nature of finitely smooth kernels to keep the condition numbers of local interpolation matrices at strictly manageable levels.

Numerical experiments in benchmark test cases clearly demonstrated the distinct advantages of this hybrid formulation. While pure infinitely smooth kernels suffered from severe ill-conditioning and computational breakdown during grid refinement, the hybrid combinations reliably maintained high-order convergence without succumbing to numerical instability. These results confirm that hybrid kernels significantly enhance the robustness of the D-RBF-PU method, making it a practical and highly efficient tool for large-scale PDE simulations where both extreme accuracy and stability are required.

Beyond these foundational results, this study laid the groundwork for a more flexible and intelligent computational framework. Future work will focus on developing an adaptive hybrid strategy designed to dynamically employ different combinations of kernels and optimize the shape parameter (ε) alongside the weight parameters (α and β) based on local nodal distributions, problem dimensionality, or specific computational constraints. Furthermore, extending this methodology to other classes of PDEs (such as parabolic or hyperbolic equations) and conducting a rigorous theoretical analysis of the underlying stability and convergence bounds remain critical directions for future exploration.

Acknowledgements

The authors would like to express their sincere gratitude to the handling editor and the anonymous reviewers for their constructive comments and valuable suggestions, which have significantly improved the quality and presentation of this manuscript.

Funding

This research received no external funding.

Conflicts of Interest

The authors declare no conflicts of interest.

References

[1] Bleecker, D. and Csordas, G. *Basic partial differential equations*, Chapman and Hall/CRC, 1995.

[2] Cavoretto, R. *A numerical algorithm for multidimensional modeling of scattered data points,* Comput. Appl. Math. 34 (2015), 65–80.

[3] Cavoretto, R., De Marchi, S., De Rossi, A., Perracchione, E. and Santin, G. *Partition of unity interpolation using stable kernel-based techniques,* Appl. Numer. Math. 116 (2017), 95–107.

[4] Cavoretto, R., De Rossi, A. and Haider, A. *A shape-parameterized RBF-partition of unity technique for PDEs,* Appl. Math. Lett. 163 (2025), 109453.

[5] Cavoretto, R., De Rossi, A. and Lancellotti, S. *Bayesian approach for radial kernel parameter tuning,* J. Comput. Appl. Math. 441 (2024), 115716.

[6] Esmaeilbeigi, M. and Cheraghi, M. *Hybrid kernel approach to improving the numerical stability of machine learning for parametric equations with Gaussian processes in the noisy and noise-free data assumptions,* Eng. Comput. 40 (2024), 761–794.

[7] Fasshauer, G.E. *Meshfree approximation methods with MATLAB*, Interdisciplinary Mathematical Sciences, Vol. 6, World Scientific Publishing Co., 2007.

[8] Fasshauer, G.E. and McCourt, M. *Kernel-based approximation methods using MATLAB*, Interdisciplinary Mathematical Sciences, Vol. 19, World Scientific Publishing Co., Singapore, 2015.

[9] Fornberg, B. and Flyer, N. *A primer on radial basis functions with applications to the geosciences*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2015.

[10] Fornberg, B. and Flyer, N. *Solving PDEs with radial basis functions,* Acta Numer. 24 (2015), 215–258.

[11] Fornberg, B., Larsson, E. and Flyer, N. Stable computations with Gaussian radial basis functions, SIAM J. Sci. Comput. 33(2) (2011), 869–892.

[12] Fornberg, B. and Wright, G. *Stable computation of multiquadric interpolants for all values of the shape parameter,* Comput. Math. Appl. 48(5–6) (2004), 853–867.

[13] Garmanjani, G., Banei, S., Shanazari, K. and Azari, Y. *An RBF-PUM finite difference scheme for forward–backward heat equation,* Comput. Appl. Math. 42 (2023), 231.

[14] Hussain, M., Ghafoor, A., Hussain, A., Haq, S., Ali, I. and Ul Arifeen, S. *A hybrid kernel-based meshless method for numerical approximation of multidimensional Fisher's equation,* Math. Comput. Simul. 223 (2024), 130–157.

[15] Larsson, E., Shcherbakov, V. and Heryudono, A. *A least squares radial basis function partition of unity method for solving PDEs,* SIAM J. Sci. Comput. 39(6) (2017), A2539–A2564.

[16] Melenk, J.M. and Babuška, I. *The partition of unity finite element method: Basic theory and applications,* Comput. Methods Appl. Mech. Eng. 139(1) (1996), 289–314.

[17] Mirzaei, D. *The direct radial basis function partition of unity (D-RBF-PU) method for solving PDEs,* SIAM J. Sci. Comput. 43(1) (2021), A1223–A1249.

[18] Mishra, P.K., Fasshauer, G.E., Sen, M.K. and Ling, L. *A stabilized radial basis-finite difference (RBF-FD) method with hybrid kernels,* Comput. Math. Appl. 77(9) (2019), 2354–2368.

[19] Mohammadi, M., Sommariva, A. and Vianello, M. *Unisolvence of Kansa collocation for elliptic equations by polyharmonic splines with random fictitious centers,* Appl. Math. Lett. 168 (2025), 109571.

[20] Wendland, H. *Fast evaluation of radial basis functions: Methods based on partition of unity*, in Approximation Theory X: Wavelets, Splines, and Applications, Vanderbilt University Press, Nashville, 2002, 473–483.

[21] Wendland, H. *Scattered Data Approximation*, Cambridge University Press, 2005.