

An Interior-Only Framework for Hessian Recovery: The KR Defect Method, with Uniqueness, Stability, and Boundary-Constrained Applications

RamaKrishna Pasupuleti

workisfun415@gmail.com

INDEPENDENT RESEARCHER <https://orcid.org/0009-0008-8418-1430>

Research Article

Keywords: Hessian recovery, derivative-free methods, interior-only sampling, generalised finite differences, uniqueness, stability, boundary-constrained problems, scattered data

Posted Date: July 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10399075/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

An Interior-Only Framework for Hessian Recovery: The KR Defect Method, with Uniqueness, Stability, and Boundary-Constrained Applications

RamaKrishna Pasupuleti

Independent Researcher, Suryapet, Telangana, India

ORCID: 0009-0008-8418-1430 Email: workisfun415@gmail.com

Abstract

Central finite-difference stencils for the Hessian require evaluations on both sides of each point and are therefore difficult to apply directly near a domain boundary, on scattered points, or when measurements are missing, without additional treatment. We develop an interior-only Hessian recovery framework—the KR defect framework—in which the three evaluation points used per direction, $\{x, x + Rh, x + h\}$, lie strictly inside the segment joining two interior points, so that no evaluation outside the domain is ever required. For $f \in C^2(\Omega)$ on a convex domain $\Omega \subset \mathbb{R}^n$, the normalised KR defect recovers directional curvature from these interior evaluations. We present eight theorem statements with accompanying proofs—covering curvature recovery, a local limit, a reconstruction formula using the minimal $n(n + 1)/2$ measurements, a uniqueness principle, a stability estimate, a boundary-interior property, and rotational invariance—together with an entrywise dimensional consistency proposition and an explicit complexity statement. The method is *first-order consistent*, $O(h)$; the design objective is not to improve asymptotic order on structured interior grids, where central differences already achieve $O(h^2)$, but to enable interior-only recovery where central differences are not directly applicable. We evaluate the framework on boundary-constrained problems (non-convex domains and an Euler–Bernoulli beam near its supports), scattered data, and missing-sensor scenarios, with full statistical reporting, and we introduce an adaptive step-selection rule that substantially reduces boundary error in low-noise settings while exposing an explicit truncation–noise trade-off. An open-source implementation reproducing all experiments is publicly available.

Keywords: Hessian recovery; derivative-free methods; interior-only sampling; generalised finite differences; uniqueness; stability; boundary-constrained problems; scattered data.

MSC 2020: 65D25; 65N06; 26B35; 49M37.

1 Introduction

1.1 The interior-only recovery problem

Estimating the Hessian of a function from point evaluations is a basic task in numerical analysis, arising in structural health monitoring, magnetic resonance elastography, geophysical inversion, and derivative-free optimisation. The standard tool—the central finite-difference (FD) stencil—estimates the (i, i) Hessian entry from the three points $\{x - he_i, x, x + he_i\}$. This stencil is second-order accurate and highly effective in the interior of a structured grid.

It carries, however, a structural restriction: it requires the point $x - he_i$. When x lies within h of the boundary in the $-e_i$ direction, that point falls outside Ω , and the stencil cannot be

evaluated without additional treatment such as one-sided stencils, ghost points, or extrapolation. The same restriction reappears, in different form, on scattered point sets, where no axis-aligned neighbour exists, and under sensor failure, where the required neighbour is missing. These are precisely the situations that dominate in practice: the most critical bending stresses occur at structural supports; the most diagnostically relevant stiffness contrasts occur at organ boundaries; and physical sensors are rarely placed on a perfect grid and sometimes fail.

This paper studies the following question: *can second-order derivative information be recovered using only evaluations guaranteed to lie inside the domain?* We give a constructive, affirmative answer—the KR defect framework—and analyse both its guarantees and its limitations. The framework samples three points per direction, all interior to the segment $[x, x + h] \subset \Omega$, and reconstructs the Hessian from the minimal number of such measurements.

1.2 Scope and honest positioning

We state at the outset what the method does and does not offer, so that its contribution is not mistaken for a claim it does not make. The framework is *first-order consistent*, $O(h)$, whereas central differences are $O(h^2)$; on smooth interior grids central differences are both more accurate and faster, and our experiments in Section 5.6 confirm this. The framework is likewise not designed to compete on accuracy with moving least squares (MLS) or radial-basis methods when evaluations are cheap: those methods use many neighbours and are typically more accurate. The KR framework is intended for the specific regime in which evaluations must remain interior, in which each evaluation is expensive, or in which a uniqueness or stability guarantee from a minimal measurement set is required. Section 6 gives concrete guidance on when the method is, and is not, the appropriate choice.

1.3 Contributions

The paper makes the following contributions.

1. An interior-only Hessian recovery framework (Section 2), whose three sampling points are provably interior on convex domains (Theorem 15).
2. Eight theorem statements with accompanying proofs, including a uniqueness principle (Theorem 11) and a stability estimate (Theorem 14), presented as central theoretical contributions and discussed relative to prior work in Section 3.
3. A dimensional analysis (Section 4): an entrywise consistency proposition whose bound does not depend explicitly on the ambient dimension, an explicit complexity statement, and an explanation of an aggregate-metric artifact.
4. An adaptive step-selection rule (Section 5.2) that substantially reduces boundary error in low-noise settings, together with an explicit and experimentally supported truncation–noise trade-off.

5. Ten numerical experiments with full statistical reporting, and a publicly available open-source implementation reproducing them (Section 7).

1.4 Related work

The central difference formula was formalised through the work of Gregory, Taylor, and Euler and is standard. One-sided stencils avoid the backward evaluation but place a point outside the segment $[x, x + h]$, which may violate domain constraints on annular or multiply-connected regions. Generalised finite differences (GFD) and mesh-free finite differences handle scattered geometry through local weighted least-squares fits over node neighbourhoods. MLS and radial-basis-function differentiation are standard on irregular domains. Finite-element Hessian recovery, such as superconvergent patch recovery and L^2 projection, post-processes finite-element solutions on meshes. The derivative-free optimisation literature constructs implicit Hessian models for minimisation rather than explicit recovery. The KR defect specialises, for the weight K , to the classical Jensen gap studied as a function of the interpolation parameter; a companion paper established the one-dimensional case. Section 3 compares the framework with GFD in detail.

2 The KR Defect Framework

2.1 Definitions and standing assumptions


Assumption 1. (Standing assumptions) (A1) $\Omega \subset \mathbb{R}^n$ is bounded, open, convex, with Lipschitz boundary. (A2) $f \in C^2(\Omega)$. (A3) $(K, R) \in \mathcal{A} := \{(K, R): K > 0, R > 0, K + R = 1\}$; for uniform results, $(K, R) \in \mathcal{A}_\delta := \{K \geq \delta > 0\}$. (A4) $x \in \Omega$, $h = sv \neq 0$, $v \in S^{n-1}$, $s > 0$, $x + h \in \Omega$.

Definition 2. (KR defect and normalised KR defect) Under Assumption 1, the KR defect is $D_f(x, h; K, R) := Kf(x) + Rf(x + h) - f(x + Rh)$, and the normalised KR defect is

$$\Phi_f(x, h; K, R) := \frac{D_f(x, h; K, R)}{KR\|h\|^2} = \frac{Kf(x) + Rf(x + sv) - f(x + Rsv)}{KR s^2}.$$

The three evaluation points are x , the interior point $z = x + Rh$, and $x + h$. Because $z = Kx + R(x + h)$ with $R \in (0, 1)$, it is a strict convex combination of x and $x + h$; Equation [eq:defect] is thus the Jensen gap evaluated at z , and Equation [eq:phi] normalises it into a dimensionless curvature estimator. The denominator $KR\|h\|^2$ is strictly positive under Assumption 1.

Interior-only sampling near a boundary



The interior-only sampling geometry. The three KR evaluation points x , $x + Rh$, and $x + h$ (black) all lie inside the domain, with $x + Rh$ strictly between x and $x + h$ (Lemma 3). The central finite-difference point $x - h$ (red) falls outside the domain near the boundary and cannot be evaluated (Theorem 15).

2.2 Preliminary lemmas

Lemma 3. (Interior sampling) Under Assumption 1, $z := x + Rh = Kx + R(x + h) \in \Omega$ and lies strictly between x and $x + h$.

Proof. Convexity of Ω and $R \in (0,1)$ give $(1 - R)x + R(x + h) = x + Rh \in \Omega$; strictness follows from $R \in (0,1)$. $\square$

Corollary 4. (Star-shaped extension) If Ω is star-shaped with respect to x and $x + h \in \Omega$, then $x + Rh \in \Omega$.

Lemma 5. (Taylor form) For $g \in C^2$ and $a, a + s \in \Omega$, $g(a + s) = g(a) + \nabla g(a)^\top s + \int_0^1 (1 - t) s^\top H_g(a + ts) s dt$.

Lemma 6. (Hessian difference estimate) Under Assumption 1, for all $t \in [0,1]$, $|v^\top [H_f(x + th) - H_f(x + tRh)]v| \leq \omega_{H_f}(s)$, where ω_{H_f} is the modulus of continuity of H_f ; consequently $|\int_0^1 (1 - t) v^\top [H_f(x + th) - H_f(x + tRh)]v dt| \leq 1/2 \omega_{H_f}(s)$.

Lemma 7. (Polarisation) For symmetric A , $u^\top Av = 1/2 [(u + v)^\top A(u + v) - u^\top Au - v^\top Av]$, and the $n(n + 1)/2$ forms $\{e_i e_i^\top, 1/2 (e_i e_j^\top + e_j e_i^\top)\}_{i \leq j}$ span $\text{Sym}_n(\mathbb{R})$.

Lemma 8. (Poincaré–Wirtinger) For bounded connected Lipschitz Ω there is $C_P(\Omega) > 0$ with $\|u\|_{L^2} \leq C_P \|\nabla u\|_{L^2}$ for $u \in H_0^1(\Omega)$; for convex Ω of diameter d , $C_P \leq d/\pi$.

2.3 Main theorems

Theorem 9. (Curvature recovery) Under Assumption 1 with $(K, R) \in \mathcal{A}_\delta$, there exist $\xi \in (x, x + h)$ and a residual E with $|E| \leq 1/2\delta \omega_{H_f}(s)$ such that $\Phi_f(x, h; K, R) = 1/2 v^\top H_f(\xi) v + E$.

Proof. Taylor-expand $f(x + h)$ and $f(x + Rh)$ about x (Lemma 5) and substitute into [eq:defect]. Using $K + R = 1$, the value and gradient terms cancel, leaving $D_f = R \int_0^1 (1-t) h^\top H_f(x + th) h dt - R^2 s^2 \int_0^1 (1-t) v^\top H_f(x + tRh) v dt$. Writing $H_f(x + th) - RH_f(x + tRh) = KH_f(x + th) + R[H_f(x + th) - H_f(x + tRh)]$ and dividing by KRs^2 gives $\Phi_f = \int_0^1 (1-t) v^\top H_f(x + th) v dt + E$ with $E = R/K \int_0^1 (1-t) v^\top [H_f(x + th) - H_f(x + tRh)] v dt$. The weighted mean value theorem applied to the first integral yields $1/2 v^\top H_f(\xi) v$ for some $\xi = x + \tau h$, $\tau \in (0,1)$; the bound on E follows from Lemma 6. $\square$

Theorem 10. (Local limit) Under Assumption 1 with $(K, R) \in \mathcal{A}_\delta$, $\lim_{s \rightarrow 0^+} \Phi_f(x, sv; K, R) = 1/2 v^\top H_f(x) v$, uniformly in $(K, R) \in \mathcal{A}_\delta$ and $v \in S^{n-1}$, with error at most $1/2 \omega_{H_f}(s)$.

Proof. By Theorem 9, $\xi_s \in (x, x + sv)$ so $\|\xi_s - x\| < s \rightarrow 0$; uniform continuity of H_f bounds both the main term deviation and $|E|$ by multiples of $\omega_{H_f}(s)$, independent of (K, R) on $\mathcal{A}_\delta$. $\square$

Theorem 11. (Uniqueness principle) Under Assumption 1, if $\Phi_{f_1} = \Phi_{f_2}$ for all admissible (x, h, K, R) , then $f_1 - f_2$ is affine.

Proof. Let $g = f_1 - f_2$; then $\Phi_g \equiv 0$, and [eq:ll] gives $v^\top H_g(x) v = 0$ for all v, x . By Lemma 7, $H_g \equiv 0$, hence g is affine. $\square$

Theorem 12. (Reconstruction) Under Assumption 1, $H_f(x)$ is determined by $n(n+1)/2$ KR measurements ($= \dim \text{Sym}_n(\mathbb{R})$, necessary and sufficient): $(H_f)_{ii} = 2 \lim_{s \rightarrow 0} \Phi_f(x, se_i; K, R)$, and for $i < j$ with $v_{ij} = (e_i + e_j)/\sqrt{2}$, $(H_f)_{ij} = \lim_{s \rightarrow 0} 2 \Phi_f(x, sv_{ij}; K, R) - 1/2 (H_f)_{ii} - 1/2 (H_f)_{jj}$.

Theorem 13. (Consistency) Under Assumption 1 with $(K, R) \in \mathcal{A}_\delta$, $|\Phi_f(x, sv; K, R) - 1/2 v^\top H_f(x) v| \leq 1/2 \omega_{H_f}(s) = O(\omega_{H_f}(s))$; for $f \in C^3$ this is $O(s)$. Consistency alone does not imply convergence, which requires the stability of Theorem 14.

Theorem 14. (Stability) Under Assumption 1 with $n \leq 3$, $f \in H^2(\Omega) \cap C^2(\Omega)$, $f|_{\partial\Omega} = 0$, and $|\Phi_f| \leq \varepsilon$ everywhere, there exists $C(\Omega, n) > 0$ (for $\Omega = (0, d) \subset \mathbb{R}$ with normalised constants, $C = (d/\pi)^2$) such that $\|f\|_{L^\infty(\Omega)} \leq C(\Omega, n) \varepsilon$.

Proof sketch. $|\Phi_f| \leq \varepsilon$ with Theorem 10 gives $\|H_f\|_{\text{op}} \leq 2\varepsilon$ pointwise, hence $|\Delta f| \leq 2n\varepsilon$. Elliptic H^1 and H^2 estimates, Poincaré (Lemma 8), and the Sobolev embedding $H^2 \hookrightarrow L^\infty$ for $n \leq 3$ combine to the stated bound. The restriction $n \leq 3$ is exactly the embedding condition $2 - n/2 > 0$. $\square$

Theorem 15. (Boundary-interior property) Under Assumption 1, $x, x + Rh, x + h \in \Omega$ with $x + Rh \in (x, x + h)$ strictly, whereas central FD requires $x - h$ (outside Ω near the boundary) and one-sided FD requires $x + 2h \notin [x, x + h]$. Corollary 4 extends this to star-shaped domains.

Proof. Immediate from Lemma 3. $\square$

Theorem 16. (Rotational invariance) For orthogonal Q and $g(x) = f(Qx)$, $\Phi_g(x, h; K, R) = \Phi_f(Qx, Qh; K, R)$ and $H_g(x) = Q^\top H_f(Qx)Q$.

2.4 Reconstruction algorithms

$f, x \in \Omega$, step $s > 0$, weights $\{K_1, \dots, K_m\}$ symmetric $\hat{H} \in \text{Sym}_n(\mathbb{R})$ $\hat{H}_{ii} \leftarrow 2 \bar{\Phi}_f(x, se_i) v_{ij} \leftarrow (e_i + e_j)/\sqrt{2}$ $\hat{H}_{ij} \leftarrow 2 \bar{\Phi}_f(x, sv_{ij}) - 1/2 \hat{H}_{ii} - 1/2 \hat{H}_{jj}$; $\hat{H}_{ji} \leftarrow \hat{H}_{ij}$ $\hat{H}$

x , baseline h_0 , factor $\gamma \in (0, 1/2]$, floor $h_{\min} \geq 0$, domain Ω $d \leftarrow \text{dist}(x, \partial\Omega)$ $h(x) \leftarrow \max(\min(h_0, \gamma d), h_{\min})$ $\hat{H}$ from Algorithm [alg:hessian] with step $h(x)$

3 Relation to Generalised Finite Differences

Because the KR stencil samples interior points that need not be grid-aligned, its relationship to generalised finite differences (GFD) and mesh-free finite differences requires careful statement. We compare them in Table 1 along five axes. The two approaches are not strict competitors in every regime: GFD is a flexible weighted least-squares fitting framework, whereas the KR method is a critically-determined operator carrying uniqueness and stability statements. They are suited to different questions, and we do not claim that KR supersedes GFD in general.

Comparison of the KR framework with generalised finite differences. Entries describe structural differences in formulation and guarantees, not claims of superiority. Statements about the absence of a result in the GFD literature are made to the author's knowledge.

Aspect	GFD / mesh-free FD	KR (this work)
Formulation	Weighted least-squares fit of a local polynomial over a node neighbourhood, then analytic differentiation	Single normalised defect from three interior evaluations per direction; no polynomial fit
Reconstruction	Over-determined; local system size grows with the neighbourhood (typically 15–30 nodes in 2D)	Critically determined; exactly $n(n+1)/2 = \dim \text{Sym}_n$ measurements, the minimum possible
Assumptions	Sufficiently rich neighbourhood; user-chosen weight function and basis affect the result	$x, x+h \in \Omega$ convex; interior sampling point then guaranteed inside (Lemma 3)
Guarantees	Consistency results established; a uniqueness theorem of the present type is not known to the author in the standard GFD literature	Uniqueness and stability under stated conditions (Theorems 11, 14)
Cost / order	Cost per point scales with neighbourhood size; order depends on	Fixed three evaluations per direction; first-order $O(h)$

Aspect	GFD / mesh-free FD basis	KR (this work)
--------	-----------------------------	----------------

4 Dimensional Analysis

Proposition 17. (Dimension-independent entrywise consistency) *Under the assumptions of Theorem 13, each recovered entry satisfies $|H_{ij}^{KR}(x) - H_{ij}(x)| = O\left(\omega_{H_f}(h)\right)$, with a constant that does not depend explicitly on the ambient dimension n ; for $f \in C^3$ the bound is $O(h)$.*

Proof. Each diagonal entry is formed from one directional defect and each off-diagonal entry from three (Algorithm [alg:hessian]). Every defect obeys the n -independent bound of Theorem 13; a fixed number of them per entry preserves the bound with an n -independent constant. $\square$

Remark 18. *Proposition 17 concerns individual entries; it does not assert dimension-independence of the full-matrix error in a matrix norm, and we make no such claim (see Section 5.5).*

Proposition 19. (Complexity) *Recovering the full Hessian by Algorithm [alg:hessian] over m weights requires $3m n(n+1)/2 = O(mn^2)$ function evaluations, $O(n^2)$ arithmetic operations beyond those evaluations, and $O(n^2)$ memory. Each entry is computed independently; the procedure assembles no global linear system.*

Remark 20. (Reconstruction structure) *The reconstruction computes each entry independently and therefore does not solve a dimension-dependent local linear system as part of the reconstruction procedure; the arithmetic conditioning of that step does not depend on the ambient dimension. Some scattered-data schemes, including certain GFD and MLS implementations, do form and solve a local least-squares system whose conditioning depends on the neighbourhood; the behaviour is implementation-dependent, and we make no general claim about all such methods.*

5 Numerical Experiments

All experiments use Python with NumPy on identical hardware. Unless stated, the KR estimator averages over $m = 9$ weights $K \in \{0.1, \dots, 0.9\}$. Statistical experiments report mean, standard deviation, and 95% confidence intervals (CI, computed as $\bar{x} \pm 1.96 \sigma / \sqrt{N}$). Experiments are ordered by the strength of the claim they support, beginning with the boundary-constrained setting that motivates the method.

5.1 Boundary-constrained recovery on non-convex domains

Table 2 recovers the Hessian entry H_{11} of $f(x, y) = \sin(2\pi x)\cos(2\pi y)$ near the boundaries of three non-convex domains. At each “near-boundary” position, central FD requires a point

outside the domain and is not applicable, whereas the KR stencil remains interior (Theorem 15) and returns an estimate.

Hessian H_{11} recovery near boundaries of non-convex domains ($h = 0.04$). “N/A” marks positions where central FD requires an out-of-domain evaluation. KR returns an estimate at every position; the larger errors at near-boundary points are consistent with the $O(h)$ bound of Theorem 13.


Domain	Position (x_1, x_2)	True H_{11}	KR error	FD status
Annulus (inner boundary)	(0.31,0.00)	1.235	0.773	N/A
Annulus (interior)	(0.60,0.00)	−5.933	0.701	OK
L-shape (corner)	(0.03,0.25)	−3.014	0.803	N/A
L-shape (interior)	(0.25,0.25)	−5.785	0.682	OK
Rectangle+hole (near)	(0.21,0.00)	0.000	0.001	N/A

5.2 Adaptive step selection near a beam support

We recover the signed second derivative $u''(x) = q/2EI x(x - 1)$ of a simply-supported beam ($q = 10^3$, $EI = 10^4$, $L = 1$), whose bending moment vanishes at both supports. Table 3 compares a fixed step $h = 0.04$ with the adaptive rule of Algorithm [alg:adaptive] ($\gamma = 0.05$, noiseless). Because the true value vanishes at the supports, the fixed-step relative error grows there (up to 95.5%), while the adaptive rule holds it below 2.5%, reducing the mean relative error from 28.8% to 1.9% (Figure 2). Central FD is not applicable at $x = 0.02$ or $x = 0.98$.

Adaptive versus fixed step for beam second-derivative recovery (noiseless). The absolute error is roughly uniform; the fixed-step relative error grows near the supports only because the true value vanishes there. Adaptive stepping restores a uniform relative accuracy.

x	True $u''(x)$	Fixed- h error (%)	Adaptive error (%)
0.02	−0.00098	95.5	2.4
0.05	−0.00238	36.9	2.4
0.10	−0.00450	17.2	2.2
0.30	−0.01050	3.6	1.4
0.50	−0.01250	0.2	0.1
0.70	−0.01050	4.0	1.5
0.90	−0.00450	18.3	2.2
0.95	−0.00238	38.9	2.4
0.98	−0.00098	44.6	2.5
Mean relative error		28.8	1.9




Relative error in the recovered beam second derivative along the span. With a fixed step (triangles) the relative error grows sharply near the supports, where the true value vanishes; the adaptive rule of Algorithm [\[alg:adaptive\]](#) (circles) holds it below 2.5% across the span. Central differences are not applicable at $x = 0.02$ or $x = 0.98$.

5.2.0.1 Truncation–noise trade-off.

The estimator divides by $KR\|h\|^2$, so measurement noise of standard deviation σ induces a Hessian-entry error scaling as σ/h^2 , while truncation scales as h . This yields the heuristic model $E(h) \approx Ch + C'\sigma h^{-2}$, minimised at $h^* \approx (2C'\sigma/C)^{1/3}$; the noise term was verified to scale as h^{-2} to within 3%. Table 4 reports the mean relative beam error against γ at several noise levels: smaller γ is best in the noiseless limit, but the optimal γ increases with noise. The method is therefore not tuned to a single γ ; the appropriate value depends predictably on the noise level, and a floor $h_{\min}$ or a noise-aware $\gamma(\sigma)$ (Section 6) prevents step collapse.

Sensitivity to γ : mean relative beam error (%). The best γ per row is in bold; the optimum shifts to larger γ as noise grows, consistent with the h^{-2} noise scaling.

Noise σ	$\gamma = 0.1$	$\gamma = 0.3$	$\gamma = 0.5$
0 (noiseless)	3.8	10.3	15.5
10^{-9}	33	10	16
10^{-8}	278	38	21
10^{-7}	2632	270	120



Mean relative beam error as a function of the boundary shrink factor γ at three noise levels (logarithmic vertical axis). In the noiseless limit smaller γ is better; as noise grows the optimal γ increases, reflecting the h^{-2} noise amplification of the truncation–noise trade-off.

5.3 Scattered data

Table 5 recovers the Hessian of $f(x, y) = e^x \sin y + \cos x e^y$ at 500 random points in $[0.1, 0.9]^2$. KR uses three evaluations per point; MLS uses all neighbours within bandwidth $r = 0.2$. MLS attains a lower mean RMSE because it uses far more information; KR attains lower variance and uses a fixed three evaluations, which is advantageous when evaluations are expensive. Central FD is not applicable without a structured grid. We do not claim KR improves on MLS accuracy; it does not.

Hessian RMSE on 500 scattered points (mean $\pm$ std, 500 estimates). MLS is more accurate but uses many more evaluations; KR uses exactly three per point.

Method	Mean RMSE	Std RMSE	Evaluations/point
KR (3 evaluations)	0.0520	0.0261	3
MLS (bandwidth $r = 0.2$)	0.0338	0.0353	many (neighbourhood)
Central FD	not applicable		—

5.4 Missing data

Table 6 reports KR recovery of $f = \sin(\pi x)\cos(\pi y)$ on a 25×25 grid under four sensor-failure patterns. KR uses the nearest available neighbour as sampling partner and degrades gracefully; central FD requires a complete grid and is not applicable for any pattern.

KR error under four missing-data patterns ($h = 0.05$; 95% CI over 80 evaluation points). Central FD is not applicable for any pattern.

Pattern	Missing (%)	Mean error	Std	95% CI
Random removal	40.0	1.109	1.540	[0.772, 1.446]


Pattern	Missing (%)	Mean error	Std	95% CI
Central block	21.9	0.340	0.164	[0.304,0.376]
Boundary sensors	42.2	0.216	0.115	[0.191,0.241]
Clustered failure	16.0	0.340	0.164	[0.304,0.376]

5.5 Behaviour with ambient dimension

Table 7 recovers the Hessian of $f(x) = \sum_i \cos(x_i)$ (diagonal Hessian) at $x = (1/2, \dots, 1/2)$. We report the per-entry diagonal error—the honest accuracy measure—alongside the aggregate matrix RMSE and the evaluation count.

Dimensional behaviour ($f = \sum_i \cos x_i$, $h = 0.05$, $m = 9$). Per-entry diagonal error is constant in n (Proposition 17); evaluations grow as $O(n^2)$ (Proposition 19). The aggregate RMSE decreases as a metric-averaging effect, explained below, and must not be read as improved accuracy.

n	Per-entry diagonal error	Aggregate matrix RMSE	Evaluations
2	0.01232	0.00909	27
5	0.01232	0.00641	135
10	0.01232	0.00523	495
20	0.01232	0.00452	1,890
50	0.01232	0.00403	11,475
80	0.01232	0.00390	29,160



Dimensional behaviour for $f = \sum_i \cos x_i$. The per-entry diagonal error (circles) is essentially constant in n , consistent with Proposition 17. The aggregate matrix RMSE (triangles) decreases, but only because the fixed per-entry errors are averaged over n^2 entries (Equation [eq:rmse]); this is a metric-averaging effect, not improved accuracy.

5.5.0.1 Why the aggregate RMSE decreases.

If the per-entry error is E_{diag} on each of the n diagonal entries and E_{off} on each of the $n^2 - n$ off-diagonal entries (true value zero here), then

$$\text{RMSE} = \sqrt{\frac{nE_{\text{diag}}^2 + (n^2 - n)E_{\text{off}}^2}{n^2}} = \sqrt{\frac{E_{\text{diag}}^2}{n} + (1 - 1/n)E_{\text{off}}^2}.$$


As n grows, the large diagonal errors are averaged over n^2 entries and the RMSE tends to the small off-diagonal floor E_{off} , not to zero; Equation [eq:rmse] matches the measured values in Table 7 to six significant figures. The per-entry diagonal error is unchanged across all tested n , consistent with Proposition 17 (Figure 4). Apparent reductions in aggregate RMSE are thus attributable to averaging over an increasing number of small off-diagonal entries and should not be interpreted as improved reconstruction accuracy.

5.6 Convergence order

Table 8 measures the convergence order of KR and central FD for $f(x, y) = e^{x+y}$ at $(0.5, 0.5)$, with $p = \log(E_h/E_{h/2})/\log 2$. KR attains $p \approx 0.98 \approx O(h)$, as predicted by Theorem 13; central FD attains $p \approx 1.88 \approx O(h^2)$ and is more accurate on this smooth interior grid (Figure 5). This is an explicit and expected limitation of KR, stated here so that it is not discovered only in the results.

Convergence order for $f = e^{x+y}$. KR is first-order; central FD is second-order and more accurate on smooth interior grids—an explicit limitation of the present method.

h	KR error	p_{K-R}	FD error	p_{FD}
0.200	0.0888	—	0.0167	—
0.100	0.0456	0.96	0.0042	1.99
0.050	0.0233	0.97	0.0011	1.94
0.025	0.0118	0.98	0.0003	1.87
	Mean p	0.98		1.88



Convergence order for $f = e^{x+y}$ at $(0.5, 0.5)$. The KR error (circles) follows the $O(h)$ reference slope; the central-difference error (triangles) follows the $O(h^2)$ reference slope and is smaller on this smooth interior grid. This first-order behaviour is a deliberate trade-off for interior-only applicability, not an implementation deficiency.

5.7 Noise robustness

Table 9 reports H_{11} estimation error for $f = e^{x+y}$ under three noise models ($\sigma = 10^{-3}$, 1000 Monte Carlo trials). On this smooth interior grid, central FD is more stable under all three models—a further explicit limitation, consistent with the h^{-2} noise scaling discussed in Section 5.2.

Estimation error under measurement noise ($\sigma = 10^{-3}$, 1000 trials). Central FD is more noise-stable on this smooth interior grid.

Noise model	KR mean	KR std	FD mean	FD std
Gaussian	0.474	0.594	0.187	0.236
Impulsive (10% outliers)	1.454	2.028	0.421	0.734
Correlated	0.185	0.183	0.058	0.074

6 Discussion

6.0.0.1 When to use the framework.

The KR framework is the appropriate choice when evaluations must remain interior (structural supports, organ walls, holes), when data is scattered or partly missing, when the domain is non-convex, or when a uniqueness or stability certificate from a minimal measurement set is required. Central finite differences remain preferable on smooth interior grids, where they are more accurate ($O(h^2)$ versus $O(h)$; Table 8), more noise-

stable (Table 9), and faster; MLS or radial-basis methods are preferable when evaluations are cheap and accuracy is paramount (Table 5).

6.0.0.2 Adaptive stepping and the trade-off.

Fixed-step KR suffices in the interior. Near boundaries the adaptive rule of Algorithm [alg:adaptive] substantially reduces the relative error in low-noise settings (Table 3), but the h^{-2} noise amplification means the boundary benefit is bounded under significant noise (Table 4). We therefore do not claim boundary error is eliminated; rather, it is made controllable through γ , whose appropriate value grows with the noise level. A noise-aware rule $\gamma = \gamma(\sigma)$, or a floor $h_{\min}$, is a natural refinement.

6.0.0.3 Limitations.

The method is first-order; the stability estimate is established only for $n \leq 3$; the stability constant is not available in closed form for general domains; and the boundary-interior guarantee requires convexity (extended to star-shaped domains by Corollary 4). A magnetic resonance elastography application, in which the relevant stiffness contrasts occur at organ boundaries, is a plausible future use of the framework; we stress that this paper does not demonstrate an elastography result, and the demonstrated evidence is the structural and scattered-data setting.

7 Reproducibility

An open-source Python implementation reproducing every experiment in this paper is publicly available, installable via a standard package manager, and archived with a persistent identifier. It depends only on NumPy and includes a validation suite in which each test corresponds to a theorem of Section 2.

8 Conclusion

We have developed an interior-only framework for Hessian recovery in which all sampling points are provably inside the domain, and we have analysed its guarantees—curvature recovery, a local limit, a reconstruction formula on the minimal measurement set, a uniqueness principle, a stability estimate, a boundary-interior property, rotational invariance, and dimension-independent entrywise consistency—together with its explicit limitations, chiefly its first-order accuracy and its noise sensitivity relative to central differences on smooth grids. The framework is intended not as a general replacement for finite differences but as a tool for the boundary-constrained, scattered, and missing-data settings in which central differences are not directly applicable. An adaptive step-selection rule improves boundary accuracy in low-noise settings, subject to an explicit truncation–noise trade-off. Future work includes extending the stability estimate beyond $n \leq 3$, a noise-aware adaptive rule, and application-specific forward models.

Declarations

Funding: None. **Competing interests:** The author declares none. **Data availability:** The implementation and all experiment scripts are openly available and archived with a persistent identifier.

99

R. Muthupillai et al., Magnetic resonance elastography by direct visualization of propagating acoustic strain waves, *Science* **269** (1995) 1854–1857.

R. Sinkus et al., Viscoelastic shear properties of in vivo breast lesions measured by MR elastography, *Magn. Reson. Imaging* **23** (2005) 159–165.

A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, 2005.

A.R. Conn, K. Scheinberg, L.N. Vicente, *Introduction to Derivative-Free Optimization*, SIAM, 2009.

J. Nocedal, S.J. Wright, *Numerical Optimization*, 2nd ed., Springer, 2006.

R.J. LeVeque, *Finite Difference Methods for Ordinary and Partial Differential Equations*, SIAM, 2007.

P. Lancaster, K. Šalkauskas, Surfaces generated by moving least squares methods, *Math. Comp.* **37** (1981) 141–158.

G.E. Fasshauer, *Meshfree Approximation Methods with MATLAB*, World Scientific, 2007.

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

B. Fornberg, N. Flyer, *A Primer on Radial Basis Functions with Applications to the Geosciences*, SIAM, 2015.

N. Flyer, G.A. Barnett, L.J. Wicker, Enhancing finite differences with radial basis functions, *J. Comput. Phys.* **316** (2016) 39–62.

T. Liszka, J. Orkisz, The finite difference method at arbitrary irregular grids, *Comput. Struct.* **11** (1980) 83–95.

J.J. Benito, F. Ureña, L. Gavete, Influence of several factors in the generalized finite difference method, *Appl. Math. Model.* **25** (2001) 1039–1053.

J.J. Benito, F. Ureña, L. Gavete, B. Alonso, Application of the generalized finite difference method to solve the advection–diffusion equation, *J. Comput. Appl. Math.* **235** (2007) 1849–1863.

F. Ureña, J.J. Benito, L. Gavete, Application of the generalized finite difference method to solve the three-dimensional wave equation, *J. Comput. Appl. Math.* **235** (2011) 3183–3195.

- J. Slak, G. Kosec, On generation of node distributions for meshless PDE discretizations, *SIAM J. Sci. Comput.* **41** (2019) A3202–A3229.
- O. Davydov, R. Schaback, Minimal numerical differentiation formulas, *Numer. Math.* **140** (2018) 555–592.
- O.C. Zienkiewicz, J.Z. Zhu, The superconvergent patch recovery and a posteriori error estimates, *Int. J. Numer. Methods Eng.* **33** (1992) 1331–1364.
- J.H. Bramble, A.H. Schatz, Higher order local accuracy by averaging in the finite element method, *Math. Comp.* **31** (1977) 94–111.
- S.C. Brenner, L.R. Scott, *The Mathematical Theory of Finite Element Methods*, 3rd ed., Springer, 2008.
- M.J.D. Powell, The NEWUOA software for unconstrained optimization without derivatives, in *Large-Scale Nonlinear Optimization*, Springer, 2006, pp. 255–297.
- M.J.D. Powell, The BOBYQA algorithm for bound constrained optimization without derivatives, Technical Report DAMTP 2009/NA06, University of Cambridge, 2009.
- J. Larson, M. Menickelly, S.M. Wild, Derivative-free optimization methods, *Acta Numer.* **28** (2019) 287–404.
- W. Hare, G. Jarry-Bolduc, Calculus identities for generalized simplex gradients, *SIAM J. Optim.* **30** (2020) 853–884.
- J.L.W.V. Jensen, Sur les fonctions convexes et les inégalités entre les valeurs moyennes, *Acta Math.* **30** (1906) 175–193.
- S.S. Dragomir, A refinement of Jensen’s inequality, *Bull. Math. Soc. Sci. Math. Roumanie* **34** (1990) 291–296.
- J.G. Liao, A. Berg, Sharpening Jensen’s inequality, *Amer. Statist.* **73** (2019) 278–281.
- S. Abramovich, J. Pečarić, S. Varošanec, Sharpening Jensen’s inequality via a linear interpolation, *J. Inequal. Pure Appl. Math.* **5** (2004) Article 91.
- R. Pasupuleti, A KR defect framework for classical inequalities, *Bound. Value Probl.* (2026), accepted.
- W. Rudin, *Principles of Mathematical Analysis*, 3rd ed., McGraw-Hill, 1976.
- L.E. Payne, H.F. Weinberger, An optimal Poincaré inequality for convex domains, *Arch. Rational Mech. Anal.* **5** (1960) 286–292.
- L.C. Evans, *Partial Differential Equations*, 2nd ed., American Mathematical Society, 2010.
- D. Gilbarg, N.S. Trudinger, *Elliptic Partial Differential Equations of Second Order*, 2nd ed., Springer, 1983.
- R.A. Adams, J.J.F. Fournier, *Sobolev Spaces*, 2nd ed., Academic Press, 2003.