Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Local QEF Displacement Bounds in Implicit CSG: Correspondence, Conservatism, and Marginal Value

Jeremy Pei

jypei89@gmail.com

Independent Researcher

Article

Keywords:

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11570117/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Local QEF Displacement Bounds in Implicit CSG: Correspondence, Conservatism, and Marginal Value

Jeremy Pei
Independent Researcher
jypei89@gmail.com

Abstract

We study how errors in approximate hard constructive solid geometry (CSG) affect a local quadratic error function (QEF). Each observed Hermite row is paired with a selected-primitive projection. Strict margins along its Boolean path establish that the reference point is visible on the exact boundary. Supplied value, gradient, Lipschitz and reach bounds control displacement between equally regularized, corresponding-row QEF solutions. Explicit arithmetic-error budgets extend the comparison to stored output. The result is a relative solution bound, not a bound on distance to the intended surface.

We distinguish tighter bounds, additional admissions and lower computation cost. In 11,977 historical CSG cells, spectral strengthening improves 27 denominators but adds one admission at tolerance $0.5h$. Nine locally predeclared confirmation tasks yield no selective benefit: the QEF tests reject every eligible cell, including 11,814 audit-safe candidates. A tiered rule preserves the full-spectral admission set and reduces median cold-cache verification time by 13.32% on 22 fixed tasks, but costs about 6.71 times the cheap-only rule for one additional admission. In a separate known-data pilot of 96 stored cells, a componentwise enclosure admits 71 declarations versus 27 and 22 for matched analytic and scalar controls, at half the largest stored cell width. These results locate conservatism in specific bound constructions rather than establish an intrinsic limit of QEF certification. Historical admissions remain numerical audits. The contribution is a well-defined local correspondence and an analysis of its guarantees and limited marginal value, not a general CAD speedup or complete surface-accuracy guarantee.

Keywords: implicit surfaces; constructive solid geometry; Hermite correspondence; quadratic error functions; interval arithmetic; Dual Contouring.

1 Introduction

Dual Contouring fits a vertex to planes obtained from points and normals on sign-changing cell edges [1]. The QEF fits the supplied planes; relating its vertex to an intended surface first requires relating those planes to that surface. An approximate Boolean expression may select a different operand, hide a projected point inside another operand, or have a large normal error despite a small value error.

Pointwise minimum and maximum failing to preserve exact signed distance is established background [6, 7], not the contribution of this paper. Our question is more specific:

> Given declared errors for approximate CSG operands, what exact visible boundary row corresponds to each observed Hermite row, and when does a worst-case QEF displacement bound provide a useful additional decision?

Our reference point is the nearest point on the exact primitive selected by the observed expression. Strict margins along the selected Boolean path establish that this point is exposed on the exact Boolean boundary. It need not be a root on the same grid edge. This correspondence defines the two QEFs being compared; it cannot be replaced by an unspecified true vertex.

The paper contributes (i) this query-local visible-row correspondence for finite signed min/max trees, (ii) conditional displacement bounds specialized to it using standard projection and perturbation arguments, and (iii) a mechanism, decision, and verification-cost study of their conservatism. The spectral characterization explains why a tighter bound need not add an admission. A locally predeclared confirmation design permits a negative practical outcome without changing the mathematical claim. An equal-admission proposition then isolates a different question: if the full-spectral service is already required, can its evaluation be made cheaper? The measured answer is conditional on the workload and cache regime, not a claim that the additional service is worth buying. We use these three contributions as the paper's scope, rather than

treating every later diagnostic as a new method. A fixed-output-mesh coverage study is retained only in Supplementary Section S6. It neither changes QEF rows nor supplies missing evidence of a QEF-guided improvement in surface reconstruction.

The distinction matters in interpreting the experiments. Historical data show overlapping protections: a normal threshold and the QEF gate can reject the same failures. Removing either one alone does not establish that it is useless. Conversely, a deterministic bound has no demonstrated practical advantage merely because it rejects many audit-safe cells. We measure both extra interceptions and extra safe rejections, with an empty admitted set having undefined conditional risk rather than zero risk. The confirmed outcome supports a local guarantee and a conservatism analysis, not a general recommendation to add a QEF gate to contouring pipelines.

1.1 Related work and scope

The comparisons concern different input information and output objects. Marching Cubes, feature-sensitive extraction and Dual Contouring produce meshes from scalar or Hermite data [2, 3, 1]; Manifold Dual Contouring addresses connectivity [9]. Learned variants predict extraction quantities or optimize mesh–field consistency [4, 5]. Distance-sample contouring, tangent-point, power-diagram and gradient-contact methods exploit SDF information for surface reconstruction [13, 16, 15, 14]. They are reconstruction comparators when input access matches, not weaker versions of a local QEF-displacement guarantee.

Trettner and Kobbelt generate samples directly on quadric-based CSG surfaces with geometric guarantees [30]. Their specialized primitive representation provides an existing route to valid CSG surface samples. Our question is different: given observed Hermite rows from an approximate field and declared errors, which exact rows correspond to them, and how do those errors affect a regularized QEF? We do not claim to introduce the general problem of producing guaranteed CSG surface points.

CSG estimate semantics and pseudo-SDF repair are established topics [7, 6]. Lipschitz pruning preserves a regional expression while reducing evaluation cost [11]; fuzzy Boolean operators serve differentiable modeling [12]. Conservative neural fields supply producer-side non-overestimation properties [17]. None of those properties alone supplies the visible row correspondence and complete uncertainty declarations used here. Probabilistic quadrics optimize expected quadratic error under plane uncertainty [8]; our comparison instead uses deterministic errors between equally regularized corresponding QEFs.

Relation to the author's witness-contract manuscript. The author has a separate related manuscript, *Witnessed Boolean Contracts*, currently under review and not yet accepted for publication. It studies visible projected witnesses and local distance deficiency for approximate Boolean unions, using value/gradient errors, reach, and branch/visibility conditions. That framework is inherited background here, not a second claim of a new witness principle. The present work develops a different comparison object: observed Hermite rows paired with selected-primitive projection rows, finite mixed signed min/max paths, and displacement between equally regularized corresponding QEFs. Its additional analysis concerns row-to-QEF propagation, geometric conservatism, decision changes, and fixed-service verification cost. The relevant projection and path arguments are given in full, so the QEF claims do not depend on an unavailable proof from that manuscript.

Interval verification and local check/refine procedures are mature methods [20, 21]. Isotopic meshing, representation-aware RBF enclosures and interval Krawczyk surface tests establish geometric or zero-set guarantees under their own input assumptions [10, 24, 23, 22]. Our interval control begins later, at the declared uncertain QEF rows; it is not a reproduction of those surface algorithms. Verified least-squares bounds and matrix decompositions are direct numerical precedents [26, 27, 28]. Indeed, the augmented regularized problem $A_\lambda = [N; \sqrt{\lambda}I]$, $b_\lambda = [h; 0]$ has full column rank for $\lambda > 0$ even if the normal rows are rank deficient. Those solvers remain applicable alternatives. Rump's componentwise bounds use error identities and verified residual computations rather than treating a numerical solution as exact [27]. Our new common-box control uses a preconditioned componentwise enclosure with exact certificate checking; it is not an implementation of Rump's QR-based algorithm or INTLAB, and it does not represent the strongest possible verified least-squares solver.

Positive reach provides the projection tube [18]; a statistical reach estimator is not automatically a deterministic producer certificate [25]. Standard perturbation algebra [19] supplies the QEF estimates. Likewise, cheap checks followed by more expensive verified computation are prior art in adaptive geometric predicates [29]. Our tiered rule uses two nested displacement bounds on the same fixed QEF, not an exact determinant sign. We claim neither a new general filtering principle nor a new matrix-perturbation theory. The specific contribution is the visible-row comparison object, its conditional displacement bounds, and the mechanism and cost evidence for using those bounds.

The auxiliary mesh audit addresses another established problem: directed surface-distance verification. Near-zero mesh-to-mesh methods exploit geometric upper bounds and local refinement [31]; cascading triangle-soup bounds provide an especially close algorithmic precedent [32]. We use elementary member-witness, same-triangle and partition bounds with exact acceptance checks. We do not claim new Hausdorff theory, a reimplementation of those complete algorithms, or a speed ranking against them. This auxiliary track is retained in Supplementary Section S6 as a traceable coverage–resource study on fixed downstream meshes, outside the paper's three core claims.

2 Information model and correspondence

Let Ψ define the intended solid by $\{x : \Psi(x) \le 0\}$, and let $\widehat{\Psi}$ be the approximate expression queried by the contourer. Both use the same finite CSG tree. With the positive-outside convention, union, intersection, and difference use $\min(a, b)$, $\max(a, b)$, and $\max(a, -b)$, respectively. The modeled set is defined by the nonpositive values of this expression. In particular, a binary difference represents $\{x : a(x) \le 0,\ b(x) \ge 0\}$, which includes the cut boundary rather than strictly subtracting all points of the closed second solid. The expression Ψ is generally not the resulting solid's Euclidean SDF.

We distinguish three objects: the intended boundary S, the approximate zero set $\widehat{S} = \{x : \widehat{\Psi}(x) = 0\}$, and the extracted triangle mesh M. A field residual, distance to S, and distance between two mesh vertices are different quantities. None is silently substituted for another below.

2.1 What the producer must supply

Each exact leaf ϕ_j is a smooth-surface SDF with known reach lower bound $\rho_j > 0$. The reach of a surface is the supremal radius for which every point closer than that radius has a unique nearest point on the surface [18]. Thus ρ_j specifies a tube in which the primitive projection is well-defined. The approximate leaf $\widehat{\phi}_j$ supplies:

- value-error bounds $\varepsilon_j(D)$ on every queried box D;
- gradient-error bounds $g_j(C)$ at the differentiable tube points of the source cell C;
- Lipschitz bounds for approximate subexpressions on the boxes enclosing the witness neighborhoods used in the visibility test.

Thus the gradient bound means $\|\nabla\widehat{\phi}_j(x) - \nabla\phi_j(x)\| \le g_j(C)$ wherever the query is in its declared tube domain, not that the exact SDF is differentiable everywhere in a cell containing its medial axis. The reach test below verifies the domain needed at the observed witness. Value and Lipschitz bounds remain valid even where the gradient is undefined.

An ordinary scalar grid does not supply all these quantities. This is an explicitly richer input model than scalar-only contouring. In the historical experiments the bounds come from a declared analytic perturbation; they are not measured from reference errors and then fed back into the admission decision.

2.2 Observed rows and the QEF comparison

An active cell in this prototype is one with at least one observed sign-changing edge, from which the extractor obtains Hermite rows. It provides numerical edge roots $\widehat{p}_i$ of $\widehat{\Psi}$, stored observed normals $\widehat{n}_i$, and stored plane offsets $\widehat{h}_i$. The observed normals need not have exact real unit length. For the same stored cell center c, the implementation forms an approximation to $\widehat{n}_i^{\mathsf{T}}(\widehat{p}_i - c)$; its arithmetic defect is retained in Lemma 5. Set

$$\widehat{N} = \begin{bmatrix} \widehat{n}_1^{\mathsf{T}} \\ \vdots \\ \widehat{n}_m^{\mathsf{T}} \end{bmatrix}.$$

The centered observed QEF minimizes

$$\|\widehat{N}y - \widehat{h}\|_2^2 + \lambda\|y\|_2^2, \qquad \lambda > 0.$$

Reference row i uses the visible primitive projection corresponding to observed row i. The reference system uses the same row count, center, regularization, and final box projection. Our bound compares these two systems. It does not bound regularization bias against an unregularized feature, or the error of a separate exact-field contouring run.

The executable output is either FLOAT-ADMITTED with a numerically evaluated displacement bound or UNKNOWN with rejection reasons. Invalid policy parameters are rejected before extraction. UNKNOWN means that the requested condition has not been established; it does not mean that the observed geometry is wrong. FLOAT-ADMITTED is deliberately not called a machine certificate: the historical implementation uses ordinary binary64 arithmetic. The corrected stored-data interface below does not retroactively change those decisions or certify their upstream error declarations.

3 Conditional bounds and decision mechanism

The assumption roles and proof dependencies are summarized in Table 1. All implications in this section are mathematical real-arithmetic statements, distinguishing exact reference rows from finite stored observations. The historical implementation evaluates its original formulas in ordinary floating point; the interval comparison verifies only the downstream QEF enclosure from supplied binary64 endpoints. Supplementary Section S7 records the additional obligations for reliable application to stored fields, rows and returned vertices.

3.1 Reference constraints for a finite hard-CSG tree

Let Ψ and $\widehat{\Psi}$ be exact and approximate scalar expressions on the same finite tree. Its selectors are finite-arity minimum or maximum, and unary negation is allowed. Thus union, intersection, and difference are represented by $\min(a,b)$, $\max(a,b)$, and $\max(a,-b)$. Each exact leaf ϕ_j is a signed-distance function for a smooth zero surface S_j, with a supplied reach lower bound $\rho_j > 0$. The composed expression Ψ need not be an exact SDF. Throughout, $B(x,r) = \{z : \|z-x\|_2 \le r\}$ denotes a closed ball, including $B(x,0) = \{x\}$.

The producer supplies value-error and approximate-field Lipschitz bounds on every queried box. The familiar inequalities for minimum, maximum, and negation propagate these bounds through a finite tree: a selector's value error and Lipschitz constant are bounded by the maxima of the corresponding child bounds, and negation leaves the bounds unchanged. No new distance semantics follow from this propagation.

At an approximate edge root $\widehat{p}_i$, the observed tree selects a leaf j_i and an effective sign $\sigma_i \in \{-1,1\}$. The sign is the product of the negations on its path. A reference Hermite row will use the closest point on this selected exact leaf, provided the following tests prove that the point is unique and visible on the exact Boolean boundary.

Lemma 1 (Selected-primitive projection)**.** *Suppose $\widehat{p} \in C$, the selected leaf obeys*

$$|\widehat{\phi}_j(x) - \phi_j(x)| \le \varepsilon_j \quad \text{for all } x \in C,$$

and define

$$r_{\text{solve}} = |\widehat{\phi}_j(\widehat{p})|, \qquad \eta_p = \varepsilon_j + r_{\text{solve}}.$$

If $\eta_p < \rho_j$, then the closest point $p = \pi_j(\widehat{p}) \in S_j$ is unique and

$$\|\widehat{p} - p\|_2 \le \eta_p. \tag{1}$$

Proof. The value-error bound gives

$$|\phi_j(\widehat{p})| \le |\phi_j(\widehat{p}) - \widehat{\phi}_j(\widehat{p})| + |\widehat{\phi}_j(\widehat{p})| \le \eta_p.$$

Since ϕ_j is an exact SDF, $\operatorname{dist}(\widehat{p}, S_j) = |\phi_j(\widehat{p})|$. The strict reach condition gives uniqueness of the closest point. □

The reference point need not lie on the same grid edge as $\widehat{p}$. Consequently, (1) is a surface-projection bound, not a bound on the displacement between two roots restricted to one edge. The solver residual is retained even when bisection makes it very small.

Lemma 2 (One selector on a witness ball)**.** *At a minimum or maximum on the path selected at $\widehat{p}$, let s denote the observed selected child and k any other child. For a minimum, set $\widehat{d}_k = \widehat{f}_k(\widehat{p}) - \widehat{f}_s(\widehat{p})$; for a maximum, set $\widehat{d}_k = \widehat{f}_s(\widehat{p}) - \widehat{f}_k(\widehat{p})$. Suppose the approximate children have Lipschitz bounds L_s, L_k and value-error bounds $\varepsilon_s, \varepsilon_k$ on a box enclosing $B(\widehat{p}, \eta_p)$. If, for every competing child,*

$$\widehat{d}_k - (L_s + L_k)\eta_p - \varepsilon_s - \varepsilon_k > 0, \tag{2}$$

then the exact selector chooses child s throughout that ball.

Table 1: Assumptions, mathematical roles, and prototype support. These are producer obligations, not a claim that the prototype can verify an arbitrary producer's declarations.

Assumption	Role in the implication	Origin and prototype support
Two-sided value error	Bounds the selected primitive's distance at the numerical root (Lemma 1)	Global analytic value-error formula and numerical root residual
Gradient error $g < 1$	Bounds the exact normalized direction; stored rows additionally need normalization/evaluation error (Lemma 4)	Tube gradient bound plus a reliable stored-normal budget; the historical formula alone does not enclose all arithmetic defects
Positive surface reach	Makes the primitive projection unique (Lemma 1)	Primitive geometry with a justified lower bound; $\eta_p \geq \rho$ gives UNKNOWN
Strict path separation	Preserves each selected branch on the witness ball (Lemma 2 and Corollary 3)	Child value-error and Lipschitz bounds; any nonpositive margin gives UNKNOWN
Domain closure	Keeps every bound applicable on the point, cell, and witness-ball box	Mathematical requirement; current value/Lipschitz formulas are global and gradient use is tube-checked. A generic domain-declaration interface is not implemented
Matched rows and $\lambda > 0$	Defines the comparison and gives $H \succeq \lambda I$ (Theorem 6)	Same row indices, center, regularization, and final box; unit reference normals; actual stored offsets include their arithmetic defects

Proof. The approximate comparison gap can decrease between $\widehat{p}$ and a point in the ball by at most $(L_s + L_k)\eta_p$. Replacing the two approximate child values by the exact values decreases it by at most a further $\varepsilon_s + \varepsilon_k$. Equation (2) leaves a strictly positive gap in the required direction against every competitor. □

Corollary 3 (Path-separated signed hard-CSG witness)**.** *Apply Lemma 2 to every selector ancestor on the observed path of $\widehat{p}$, and apply Lemma 1 to its selected leaf. Then the exact expression equals $\sigma_j \phi_j$ on $B(\widehat{p}, \eta_p)$, and $p = \pi_j(\widehat{p})$ is on the exact Boolean boundary. The result applies to any finite signed min/max tree, including finite-arity selectors, under these path-separation conditions.*

Proof. Starting at the root, the selector lemma preserves the observed child at every selector; a unary negation only changes the accumulated sign. Induction along the finite selected path gives $\Psi = \sigma_j \phi_j$ on the ball. The projection p belongs to the ball and satisfies $\phi_j(p) = 0$. All ancestor comparisons are strict at p, so continuity preserves them in a neighborhood of p. The selected SDF changes sign across its smooth zero surface there. Thus Ψ also changes sign across p, proving that p is a boundary point rather than merely a zero of the expression. □

This is a query-local statement, not a global smooth-branch or topology theorem. Different rows in one cell may select different leaves, and an unresolved comparison makes that row UNKNOWN. The binary implementation checks every ancestor comparison on each selected path; a direct finite-arity implementation would also check every sibling at each selector.

3.2 Normal and centered Hermite-plane bounds

Lemma 4 (Unit-normal direction in the projection tube)**.** *Assume the projection and visibility premises above. Let $u = \sigma_j \nabla \phi_j(\widehat{p})$ and $v = \sigma_j \nabla \widehat{\phi}_j(\widehat{p})$, with $\|u - v\|_2 \leq g < 1$. Then the exact normalized direction $n^\circ = v/\|v\|_2$ and the signed exact normal n at $p = \pi_j(\widehat{p})$ satisfy*

$$\|n^\circ - n\|_2 \leq \eta_{\mathrm{dir}}(g) := \sqrt{2 - 2\sqrt{1 - g^2}} = \frac{\sqrt{2}\, g}{\sqrt{1 + \sqrt{1 - g^2}}}. \tag{3}$$

Proof. In a unique-projection tube, the gradient of an exact SDF is constant along each normal projection ray. Hence $n = u$ and $\|u\|_2 = 1$; no separate curvature-transport term is required for this correspondence. The ball of radius $g < 1$ about u excludes the origin and subtends maximum angle $\arcsin g$. The unit-direction chord is therefore at most $2\sin(\frac{1}{2}\arcsin g) = \sqrt{2 - 2\sqrt{1 - g^2}}$. □

The two expressions in (3) are equal in real arithmetic. The historical implementation evaluates the rationalized expression, because direct subtraction in the first expression can round to zero when g is small.

This stable evaluation removes the observed cancellation failure, but it does not by itself make the complete ordinary-floating-point decision outward rounded. In particular, a stored normal $\widehat{n}$ need not equal n° or have exact real unit length. To use this route for stored rows, let δ_{norm} bound $\|\widehat{n} - n^\circ\|_2$ and use the full row budget $\eta_n \geq \eta_{\mathrm{dir}}(g) + \delta_{\mathrm{norm}}$. A direct reliable enclosure of $\|\widehat{n} - n\|_2$ is also sufficient. The QEF theorem requires this full stored-row budget, not a particular normalization formula; numerical obligations are collected in Supplementary Section S7.

The additional implementation threshold $g \leq g_{\max} = 0.12$ is an experimental accuracy policy. The direction lemma itself only requires $g < 1$. Without the reach-tube premise, the normal-ray identity used in the proof is not available.

For cell center c, let $\widehat{h}$ be the actual stored offset and define the reference offset and the observed arithmetic defect by

$$h = n^{\mathsf{T}}(p - c), \qquad \omega_h = \left|\widehat{h} - \widehat{n}^{\mathsf{T}}(\widehat{p} - c)\right|.$$

Lemma 5 (Centered Hermite-plane row)**.** *If* $\|n\|_2 = 1$, $\|\widehat{p} - c\|_2 \leq R_C$, $\|\widehat{p} - p\|_2 \leq \eta_p$, *and* $\|\widehat{n} - n\|_2 \leq \eta_n$, *then*

$$|\widehat{h} - h| \leq \eta_p + R_C\eta_n + \omega_h =: \eta_h. \tag{4}$$

Proof. Expand around the observed point:

$$\widehat{h} - h = [\widehat{h} - \widehat{n}^{\mathsf{T}}(\widehat{p} - c)] + (\widehat{n} - n)^{\mathsf{T}}(\widehat{p} - c) + n^{\mathsf{T}}(\widehat{p} - p).$$

Cauchy–Schwarz and $\|n\|_2 = 1$ give the result. The cell radius bounds $\widehat{p}$, which lies on a cell edge; the exact projection p may lie outside the cell. No unit-length assumption is made on $\widehat{n}$. The ideal real-dot-product case has $\omega_h = 0$; a floating-point subtraction and dot product do not establish that equality. □

For an exact centered axis-aligned cell of longest edge h, $\sqrt{3}h/2$ bounds its half-diagonal. For actual stored coordinates, a reliable radius must instead enclose each $\|\widehat{p}_i - c\|_2$, unless exact centering has been established. A rounded center must not be silently identified with the exact box midpoint.

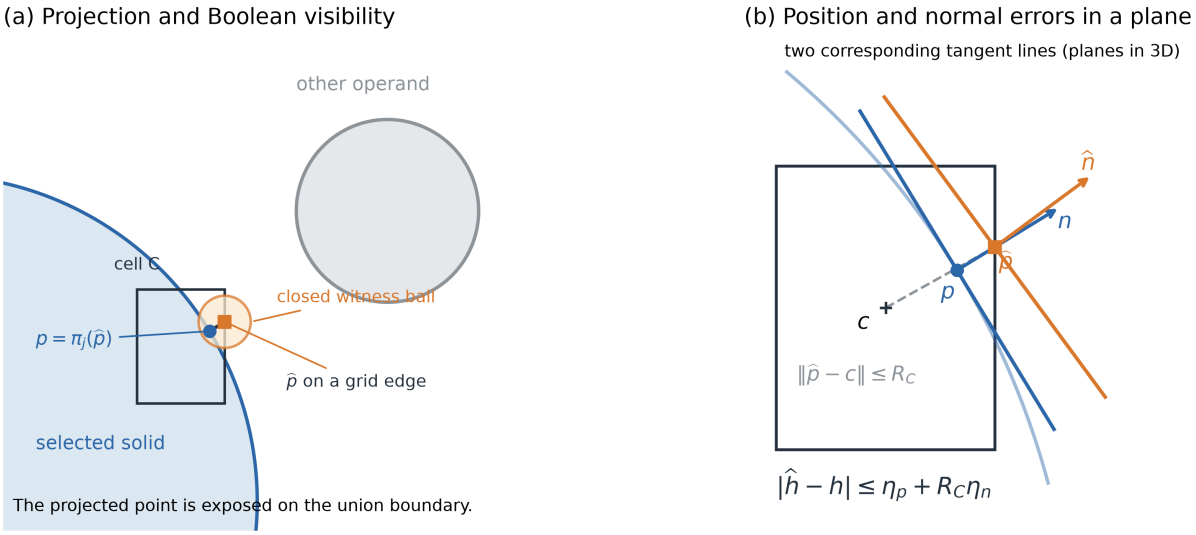


Figure 1: Objects in the local admission analysis: the point correspondence is a primitive projection, while the branch margins establish visibility in the exact hard-CSG boundary. The plane bound uses the radius of the observed point about the cell center. Panel (b) shows the ideal dot-product case $\omega_h = 0$; stored offsets require Equation (4). This two-dimensional schematic is illustrative, not an additional benchmark sample.

3.3 Two computable regularized-QEF bounds

Stack the observed rows as $(\widehat{N}, \widehat{h})$. Row i of the reference pair (N, h) uses exactly the projection p_i, signed normal n_i, and offset established above for observed row i. Both QEFs use the same row count, correspondence, center c, and positive regularization λ. This reference QEF is not formed from independently located exact edge roots.

The following estimates are standard perturbation algebra [19] applied to this specific row correspondence. Their purpose is to make the admission test computable, not to introduce new matrix-perturbation theory. Define

$$\eta_N = \left(\sum_{i=1}^{m} \eta_{n,i}^2\right)^{1/2}, \qquad E_h = \left(\sum_{i=1}^{m} \eta_{h,i}^2\right)^{1/2}, \qquad \beta_N = \sqrt{m},$$

and

$$\eta_H = 2\beta_N \eta_N + \eta_N^2, \tag{5}$$

$$\eta_q = \beta_N E_h + \eta_N \|\widehat{h}\|_2 + \eta_N E_h. \tag{6}$$

Only the reference normals must be unit. Thus $\|N\|_2 \le \beta_N$ and $\|\widehat{N}\|_2 \le \beta_N + \eta_N$. The nonnegative term $\eta_N E_h$ in (6) is retained for comparison with historical formulas; the sharper valid choice $\eta_q' = \beta_N E_h + \eta_N \|\widehat{h}\|_2$ is also allowed.

Theorem 6 (Regularized-QEF displacement with an inexact solve)**.** *Let $N, \widehat{N} \in \mathbb{R}^{m\times 3}$, $h, \widehat{h} \in \mathbb{R}^m$, $m \ge 1$, and all data and the iterate be finite real quantities. Reference row n_i is unit, and the nonnegative bounds satisfy $\|\widehat{n}_i - n_i\|_2 \le \eta_{n,i}$ and $|\widehat{h}_i - h_i| \le \eta_{h,i}$ for the same corresponding rows. Neither observed unit length nor independence of the errors is required. Let*

$$H = N^{\mathsf{T}}N + \lambda I, \quad q = N^{\mathsf{T}}h, \qquad \widehat{H} = \widehat{N}^{\mathsf{T}}\widehat{N} + \lambda I, \quad \widehat{q} = \widehat{N}^{\mathsf{T}}\widehat{h}, \qquad \lambda > 0,$$

These products define exact real matrices from the given entries, not the rounded matrix assembled by a linear solver. Let $y = H^{-1}q$. For any computed observed iterate $\widehat{y}$, set

$$\widehat{r} = \widehat{N}\widehat{y} - \widehat{h}, \qquad \rho = \widehat{H}\widehat{y} - \widehat{q}, \quad \nu = \|\rho\|_2, \qquad E_r = \left(\sum_{i=1}^{m} \big(\eta_{n,i}\|\widehat{y}\|_2 + \eta_{h,i}\big)^2\right)^{1/2}.$$

Then both

$$B_{\text{matrix}} = \frac{\nu + \eta_H\|\widehat{y}\|_2 + \eta_q}{\lambda}, \tag{7}$$

$$B_{\text{post}} = \frac{\nu + \eta_N\|\widehat{r}\|_2 + \beta_N E_r}{\lambda} \tag{8}$$

bound $\|\widehat{y} - y\|_2$. Hence

$$\boxed{\|\widehat{y} - y\|_2 \le B_{\text{QEF}} := \min\{B_{\text{matrix}}, B_{\text{post}}\}.} \tag{9}$$

Proof. Write $\Delta N = N - \widehat{N}$ and $\Delta h = h - \widehat{h}$. Expansion about the reference rows, whose norms are controlled, gives

$$\widehat{H} - H = -N^{\mathsf{T}}\Delta N - \Delta N^{\mathsf{T}}N + \Delta N^{\mathsf{T}}\Delta N, \qquad \widehat{q} - q = -\Delta N^{\mathsf{T}}\widehat{h} - N^{\mathsf{T}}\Delta h.$$

The row-wise bounds give $\|\Delta N\|_2 \le \eta_N$, $\|\Delta h\|_2 \le E_h$, $\|H - \widehat{H}\|_2 \le \eta_H$, and $\|q - \widehat{q}\|_2 \le \eta_q$. Since $H \succeq \lambda I$, $\|H^{-1}\|_2 \le 1/\lambda$. Using the residual ρ defined above, the identity

$$H(\widehat{y} - y) = \rho + (H - \widehat{H})\widehat{y} + (\widehat{q} - q)$$

proves (7).

For the posterior estimate, put $e_r = \Delta N\widehat{y} - \Delta h$. Each component satisfies $|(e_r)_i| \le \eta_{n,i}\|\widehat{y}\|_2 + \eta_{h,i}$, hence $\|e_r\|_2 \le E_r$. A second expansion gives

$$H(\widehat{y} - y) = \rho + \Delta N^{\mathsf{T}}\widehat{r} + N^{\mathsf{T}}e_r.$$

Using $\|N\|_2 \le \beta_N$ yields (8). Both are upper bounds for the same displacement, so their minimum is valid. □

The posterior formula uses the observed row residuals instead of separately bounding terms that can cancel. In real arithmetic it is always no larger than the matrix estimate: $E_r \le \eta_N\|\widehat{y}\| + E_h$ and $\|\widehat{r}\| \le (\beta_N + \eta_N)\|\widehat{y}\| + \|\widehat{h}\|$ give

$$\lambda B_{\text{post}} \le \nu + (2\beta_N\eta_N + \eta_N^2)\|\widehat{y}\| + \eta_N\|\widehat{h}\| + \beta_N E_h \le \lambda B_{\text{matrix}}.$$

Thus the minimum is safe bookkeeping, not a choice between complementary regimes or a separate innovation. Both bounds and their ingredients are retained in the artifact; the matrix-only rule is a declared post-pilot control. The term ν allows any finite inexact iterate, and includes assembly error only when evaluated against the exact stored-row system above. A residual against a rounded assembled system alone is insufficient. The historical prototype uses ordinary floating-point evaluations; Supplementary Section S7 specifies the required reliable enclosures. Independently outward-rounded endpoints need not preserve the ordering of the exact formulas; taking their minimum is valid only if both enclose the same displacement.

Corollary 7 (Common-box projection and local admission)**.** *Let P_C denote Euclidean projection onto the same closed cell box C for the observed and reference QEFs. Under the preceding premises,*

$$\|P_C(c+\widehat{y}) - P_C(c+y)\|_2 \le B_{\mathrm{QEF}}. \tag{10}$$

Thus accepting only when $B_{\mathrm{QEF}} \le \tau$ bounds the displacement between these two box-projected, equally regularized, corresponding-row QEF vertices by τ.

Proof. Euclidean projection onto a closed convex set is non-expansive. Apply it to $c+\widehat{y}$ and $c+y$, then use Theorem 6. □

The comparison does not bound regularization bias relative to an unregularized feature intersection. Nor does it bound displacement to an independently extracted oracle Dual Contouring vertex, connectivity, manifoldness, topology, or Hausdorff distance of the completed surface.

Transfer to a stored output. For a finite returned vertex z, let $d_{\mathrm{out}} \ge \|z - P_C(c+\widehat{y})\|_2$. The triangle inequality and Corollary 7 give

$$\|z - P_C(c+y)\|_2 \le d_{\mathrm{out}} + B_{\mathrm{QEF}}. \tag{11}$$

This requires the same iterate, center and box in the certificate and in the actual output. It compares a box projection of an unconstrained reference minimizer, not the minimizer of a box-constrained QEF. A numerical admission requires an upper enclosure of the right side and an exact or lower-enclosed tolerance, including any tolerance–cell-scale product.

3.4 Optional checked spectral denominator

The scalar λ is always a lower bound for the smallest eigenvalue of the reference regularized Gram matrix. Well-spread rows can provide a larger lower bound. The following specialization uses the same perturbation estimate η_H as Theorem 6.

Corollary 8 (Spectral strengthening)**.** *Suppose $\underline{\ell} \le \lambda_{\min}(\widehat{H})$ has been verified and $\overline{\eta}_H \ge \|H - \widehat{H}\|_2$ is a valid upper bound. Define*

$$\underline{\gamma} = \underline{\ell} - \overline{\eta}_H, \qquad d = \max\{\lambda, \underline{\gamma}\}.$$

Then $H \succeq dI$. In Equations (7) *and* (8)*, replacing the denominator λ by d preserves both bounds. If the spectral check fails, or its lower endpoint does not improve λ, the original denominator remains valid.*

Proof. For any unit vector z,

$$z^{\mathsf{T}} H z = z^{\mathsf{T}} \widehat{H} z + z^{\mathsf{T}}(H - \widehat{H})z \ge \underline{\ell} - \overline{\eta}_H.$$

The regularized Gram structure separately gives $z^{\mathsf{T}} H z \ge \lambda$. Thus $\lambda_{\min}(H) \ge d > 0$ and $\|H^{-1}\|_2 \le 1/d$. Both residual identities in the proof of Theorem 6 have the same left side $H(\widehat{y} - y)$. Substituting $1/d$ for $1/\lambda$ in that proof gives the asserted bounds. Common-box projection remains non-expansive. □

The executable spectral check forms an outward-rounded enclosure of $\widehat{N}^{\mathsf{T}}\widehat{N} + \lambda I$ from the supplied binary64 rows. For a proposed shift α, it checks that outward lower endpoints of the three leading principal minors of $\widehat{H} - \alpha I$ are positive. Sylvester's criterion then proves $\alpha < \lambda_{\min}(\widehat{H})$. Bisection proposes shifts; only individually verified shifts become lower bounds. The subtraction of $\overline{\eta}_H$ is rounded downward. An ordinary eigensolver is used only for descriptive condition statistics, not to justify an accepted spectral denominator.

For the supplied-row numerical comparison, all numerators and norm bounds are also evaluated outward. The λ control and the spectral variant use the same numerator and reference-row semantics, so their difference isolates the denominator. This downstream verification consumes the declared row uncertainty; it cannot verify the upstream field producer.

3.5 Using unit-row structure without an oracle

The legacy budget in (5) is valid but need not retain its quadratic term when unit-row information is available.

Lemma 9 (Unit-reference-row Hessian budget)**.** *Let $\|\Delta N\|_2 \le \eta_N$ for $\Delta N = N - \widehat{N}$. If $\|N\|_2 \le \beta$ and $\|\widehat{N}\|_2 \le \widehat{\beta}$, then*

$$\|H - \widehat{H}\|_2 \le \eta_H^{\text{unit}} := (\beta + \widehat{\beta})\eta_N.$$

For exact unit observed and reference rows, one may take $\beta = \widehat{\beta} = \sqrt{m}$, eliminating η_N^2 from (5). For binary64 observations, only the reference-unit premise is required: use $\beta = \sqrt{m}$ and an outward upper bound $\widehat{\beta} \ge \|\widehat{N}\|_F$ from the actual observed entries.

Proof. The exact identity

$$N^\mathsf{T} N - \widehat{N}^\mathsf{T}\widehat{N} = N^\mathsf{T}\Delta N + \Delta N^\mathsf{T}\widehat{N}$$

and submultiplicativity give the bound. The common regularization cancels. Unit reference rows give $\|N\|_2 \le \|N\|_F = \sqrt{m}$. □

With $\Delta h = h - \widehat{h}$, the separate identity $q - \widehat{q} = N^\mathsf{T}\Delta h + \Delta N^\mathsf{T}\widehat{h}$ also permits $\eta_q' = \beta E_h + \eta_N\|\widehat{h}\|$. To isolate one change, the fixed-data control below does *not* change either numerator: it substitutes only η_H^{unit} into the spectral denominator of Corollary 8, retaining the same verified shift. Products, sums, square roots and subtraction are enclosed in the indicated directions. Observed floating-point normals are not presumed exactly unit. The semantic reference-unit premise is supplied, not inferred or certified by these calculations. This elementary refinement is a structure-aware control, not a new matrix-perturbation theory.

3.6 When can spectral tightening change a decision?

Let $\mu = \lambda_{\min}(\widehat{N}^\mathsf{T}\widehat{N})$ and retain the verified endpoints $\underline{\ell}$ and $\overline{\eta}_H$ from Corollary 8. The following elementary characterization separates available geometric information from an actual decision change. It is an explanation of the bound, not a new eigenvalue-perturbation theorem.

Lemma 10 (Denominator and decision gains)**.** *For $\lambda > 0$, $\underline{\ell} \le \lambda_{\min}(\widehat{H})$ and $\overline{\eta}_H \ge 0$, put $d = \max\{\lambda, \underline{\ell} - \overline{\eta}_H\}$. Then*

$$d > \lambda \quad\Longleftrightarrow\quad \underline{\ell} - \overline{\eta}_H > \lambda \quad\Longrightarrow\quad \mu > \overline{\eta}_H.$$

The last implication need not reverse. For a common nonnegative finite numerator β and $\tau > 0$, set $q_{\text{dec}} = \beta/(\tau\lambda)$ and $s = d/\lambda$. Assuming all common admissibility prerequisites are satisfied and holding the input, valid numerator β, and tolerance τ fixed, changing only the denominator from λ to d yields a new acceptance exactly when

$$\frac{\beta}{d} \le \tau < \frac{\beta}{\lambda} \quad\Longleftrightarrow\quad 1 < q_{\text{dec}} \le s. \tag{12}$$

Proof. The first equivalence follows from the maximum. Since $\underline{\ell} \le \lambda + \mu$, strict denominator improvement requires $\mu > \overline{\eta}_H$. It is not sufficient for a conservative lower bound: $\widehat{N} = I$, $\lambda = 1$, $\overline{\eta}_H = 1/2$, and $\underline{\ell} = 1$ give $\mu = 1$ but $d = \lambda$. Multiplication by the positive quantities λ, d, τ gives (12). In particular, $\beta = 0$ already admits both rules, and $q_{\text{dec}} = 1$ is not a new acceptance. □

Exact row span of dimension at most two gives $\mu = 0$ and rules out a strict denominator gain for any nonnegative uncertainty. This concerns the span of the *normal rows*, not coplanarity of sample points. Near-parallel rows need not be rank deficient: if a unit n_0 satisfies $\|\widehat{n}_i - n_0\| \le \epsilon_i$, the Rayleigh quotient in any unit direction perpendicular to n_0 gives $\mu \le \sum_i \epsilon_i^2$. Thus $\sum_i \epsilon_i^2 \le \overline{\eta}_H$ is a sufficient obstruction under this additional premise. Neither all curved patches nor all multi-leaf junctions are assigned a rank by this argument.

The identities are in real arithmetic. Implementation uses downward-rounded subtraction and upward-rounded division, which may refuse a limiting case accepted by the real identities. We retain that executable decision and audit the distinction; ordinary eigenvalues appear only in explanatory plots and never replace $\underline{\ell}$.

3.7 Equal admission by tiered verification

The two nested bounds also permit a simple scheduling rule. Its purpose is to avoid unnecessary spectral work, not to establish a new perturbation result. Fix one immutable cell declaration: the observed rows and their visible reference correspondence, row count, center, candidate iterate $\widehat{y}$, common box, positive finite λ, row budgets and their validity domains, and a finite tolerance $\tau_C \geq 0$. The exact reference normals are unit vectors. The displacement is $E = \|\widehat{y} - y\|_2$ for the reference solution in Theorem 6; exact projection onto the common box also obeys the bound by Corollary 7.

Let U_m, U_p be valid nonnegative finite upper bounds on the two numerators in Theorem 6, and let $U = \min(U_m, U_p)$. Write up and down for valid outward endpoint operations. Define

$$B_0 = \min\{\mathrm{up}(U_m/\lambda), \mathrm{up}(U_p/\lambda)\}, \tag{13}$$

$$d = \max\{\lambda, \mathrm{down}(\underline{\ell} - \overline{\eta}_H)\},$$

$$B_1 = \min\{B_0, \mathrm{up}(U/d)\}. \tag{14}$$

The spectral endpoints must satisfy Corollary 8. If that check does not establish a usable improvement, take $d = \lambda$; a failed spectral computation or division retains the already valid B_0. Thus this is the `spectral_effective` rule, not a division by an unchecked or nonpositive raw spectral difference.

Proposition 11 (Equal admission for valid nested bounds)**.** *Assume the same prerequisite checks and a valid finite B_0 on every eligible declaration; if the common premises or common bound are unavailable, both policies return* Unknown*. Assume that spectral evaluation, including its fallback, supplies the same valid $B_1 \leq B_0$ to either policy. Policy A always evaluates this effective spectral rule. Policy B first accepts if $B_0 \leq \tau_C$, and otherwise evaluates the same effective spectral rule. Both use the comparison bound $\leq \tau_C$, with no audit slack. Their admitted sets are identical, and every admission implies $E \leq \tau_C$ under the declared premises.*

Proof. There are three cases. If the common prerequisites or bound are unavailable, both policies return Unknown. If $B_0 \leq \tau_C$, B accepts using $E \leq B_0$, while A accepts because $E \leq B_1 \leq B_0 \leq \tau_C$. If $B_0 > \tau_C$, both evaluate the same effective spectral bound and apply the same comparison, so both accept or both remain unresolved. In the accepting case $E \leq B_1 \leq \tau_C$. □

Equal admission does not require equal returned numbers: an early B admission may return B_0 while A returns a smaller B_1. An omitted spectral computation is recorded as omitted, never as verified spectral evidence. The proposition proves neither acceleration nor that obtaining the additional spectral admissions is worthwhile. Finite-cell regressions check its implementation; they do not replace the conditional argument.

4 Implementation and locked evaluation design

4.1 Controlled producer and verification boundary

Planes, spheres, infinite cylinders, and ring tori have surface reach bounds $\infty, r, r, \min(r, R-r)$, respectively. Bounded boxes and cylinders are formed by plane intersections; we do not assign positive smooth-surface reach to the resulting sharp box boundary. Each approximate leaf is

$$\widehat{\phi}_j(x) = \phi_j(x) + a_j \sin(k_j^{\mathsf{T}} x + \theta_j).$$

Real-arithmetic extrema of the affine phase on a box give $s_D = \max_D |\sin|$, $c_D = \max_D |\cos|$ and

$$\varepsilon_j(D) = |a_j| s_D, \quad g_j(D) = |a_j| \|k_j\|_2 c_D, \quad L_j(D) = 1 + g_j(D).$$

These bounds describe a prescribed field family, not sample maxima. The prototype evaluates them in binary64. The stable normal-chord formula in Equation (3) avoids cancellation at small g but does not provide outward rounding for the producer or its trigonometric calculations.

The downstream verifier consumes declared midpoints and radii as their exact binary64 values. It encloses arithmetic, Gram assembly, residuals, and norms outward; the observed row norms are enclosed rather than presumed exactly one. The reference normals' unit length remains an explicit semantic premise. Interval contraction uses a component-box relaxation of those declarations. Neither method establishes their field provenance, and neither changes the reference to a vertex obtained by independent exact-field extraction. Rigorous floating-point enclosures are distinct from full program formal verification [21], which is not performed.

Table 2: Guarantee objects and numerical boundaries. A downstream enclosure does not verify its upstream declaration or a completed mesh.

Layer	Object and premises	What is established here
Mathematical implication	Visible corresponding rows; valid field and row budgets; the same row count, center and $\lambda > 0$	Real-arithmetic displacement to the corresponding reference QEF solution (Theorem 6 and Corollary 8)
Historical prototype	Producer envelopes, branch/tube tests, extraction and the original bound evaluated in binary64	FLOAT-ADMITTED or UNKNOWN; not a machine-certified geometry pipeline
Outward downstream verification	Exact values of declared binary64 endpoints, valid row uncertainties, and the reference semantics; analytic controls require unit reference rows	Bound from the fixed, unclipped $\widehat{y}$ to the exact unconstrained reference regularized QEF solution; no certification of the producer
Common-box projection	Exact projection of both compared points onto the same closed box	Non-expansion (Corollary 7); actual stored output additionally requires (11). Historical outputs are not thereby recertified

For the tiered comparison, immutable input identifiers cover rows, offsets, budgets, $\widehat{y}$, correspondence, center and box, λ, tolerance, source hashes, validity-domain declarations, and unit-reference semantics. Stage-cache keys also include the loaded program version, hence its fixed spectral parameters. Source changes require a new process, not a hot update. Only common, spectral, or interval results are cached; final admission is always compared anew. No post-decision reference error is available to routing. These identities prevent accidental reuse across declarations; they do not authenticate an assumed producer or protect against arbitrary process tampering.

The interface separates *finite valid bound above tolerance*, *no finite valid bound*, and *prerequisites not established*. All can produce UNKNOWN, but none proves that the geometry is wrong. A spectral failure is an additional diagnostic, not a loss of a valid cheap bound. Tolerance equality is accepted; audit slack and approximate equality tests are not added to the threshold. The finite-B_0 premise of Proposition 11 is explicit: the proposition does not assert behavior for every arithmetic-overflow input of the general adapter.

Historical extraction brackets sign-changing edges, retains numerical root residuals, and records the selected signed leaf. Corner-sign tests, strict path margins, the projection tube, finite data, $\lambda > 0$, and every row's $g < 1$ are mandatory in the comparisons. An optional threshold $\max_i g_i \le \kappa$ is a separate accuracy policy. Here g_i is a dimensionless Euclidean gradient-error budget relative to the exact unit SDF gradient; it is not a normal angle or a normalized-normal chord. Removing that policy never removes $g_i < 1$. These historical floating-point tests are not independent enclosures of every premise. In particular, zero perturbation amplitude does not imply zero field-evaluation, normalization or offset error. The numerical requirements in Supplementary Section S7 are not a new evaluation population.

The separate reliable implementation distinguishes completed enclosure computation, finite usable endpoints and final admission. It checks endpoint/status consistency, invalidates prior evidence before reevaluation, and clears partial evidence on failure. Supplementary Section S7 documents unavailable-route handling and bounded regressions; they do not recertify historical experiments.

These recovery guarantees apply to the reliable producer–admission path, not to every saved-record importer retained in the artifact. Legacy replay reconstructs the historically normalized assembly rows; reliable records instead carry the actual stored rows and additional rounding budgets. The two formats are not interchangeable. Supplementary Section S7.6 maps the entry points and documents a separate version-gated import repair. Its numerical replay remains conditional on the supplied geometric budgets; it does not authenticate their source or recertify the historical producer.

Historical extraction reruns a task at its fixed finer grid if its base grid has unresolved cells. This is one whole-grid refinement, not an adaptive octree. For completion only, rejected cells receive an ordinary $\lambda = 10^{-5}$ vertex and retain UNKNOWN status. Candidates use $\lambda = 0.5$. Decisions use declarations and computed bounds only; reference audits follow them. No oracle label enters a policy.

4.2 Research questions and distinct populations

The questions are: when does spectral information change admission (RQ1); does a QEF gate add selective benefit after an optional normal policy (RQ2); how do matched uncertainty enclosures compare in tightness and cost (RQ3); what changes in a common completed mesh (RQ4); and, for the same declarations,

correspondence, tolerance and admission requirement, can tiered evaluation reduce full-spectral verification cost, and when does it fail to do so (RQ5)?

We keep the following populations separate.

- **Historical CSG:** 13 designed union, intersection, difference, and mixed-tree tasks, 21,156 active cells. Geometric and mathematical prerequisites retain 14,021; $\kappa = 0.12$ leaves the primary 11,977-cell mechanism cohort. Parameters, source hashes, meshes, and per-task records are retained. Reclassification is exploratory, not independent confirmation.
- **Historical algebraic stress:** the existing 50,000 fixed-seed row systems are replayed unchanged. Normal families, row counts, uncertainty and regularization vary. These systems are not extra CAD inputs. Their canonical scientific digest must match the frozen release.
- **Predeclared confirmation:** exactly nine tasks in three designed families, fully serialized and hashed before extraction: concentric spherical differences with inner radii $0.81, 0.85, 0.89$ and outer radius 0.96; near-tangent unions of radius-0.77 spheres with half-separations $0.65, 0.70, 0.75$; and three rotated six-plane boxes. The common box is $[-1.35, 1.35]^3$, resolution 30, $h = 0.09$, $a = 0.006$, and base frequency 22 with existing deterministic per-leaf factors. Some near-tangent spheres extend beyond this box; only the fixed local query domain is studied. No resolution, perturbation, or geometry is adjusted after outcomes are known.

The local protocol was locked at 04:30:09 UTC on 5 September 2026, before confirmation extraction; this is a local hash/time record, not public preregistration. Its contrasts are $\kappa = 0.12$ and 0.24 at primary $\tau_C = 0.5h_C$. The full 96-configuration policy grid is retained in Supplementary Section S1. These outcomes had already been observed when the later verification-cost protocol was fixed: they are *known stress tasks* for RQ5, not new blind confirmation.

Independent interception means an additional true over-tolerance rejection after the normal policy. Selective benefit additionally requires a nonempty admitted set; practical usefulness requires a defensible cost/coverage tradeoff, for which no application-specific threshold is invented here. The protocol stops similar synthetic expansion after these nine tasks if useful independent benefit remains unestablished.

4.3 Matched enclosures, measurements, and costs

The analytic controls use the same outward-rounded numerators and differ only in the denominator. Observed binary64 normal norms are enclosed rather than presumed exactly one; exact reference normals are unit by the SDF assumption. The interval control forms component boxes for normals and offsets, encloses $H = N^{\mathsf{T}}N + \lambda I$ and $q = N^{\mathsf{T}}h$, and checks an outward-rounded preconditioned contraction. A point inverse is a candidate preconditioner, not a trusted exact inverse. A failed contraction is UNKNOWN, not an unsafe solution.

These are the same source declarations, candidate and regularization, but not identical feasible uncertainty sets: component boxes relax unit-normal and cross-component dependencies; analytic norm inequalities lose other information. We report each method's finite-bound set and their common finite-bound set. Bound availability is a construction diagnostic, not admission: a finite bound above tolerance still leaves the cell UNKNOWN. Cases without a finite bound remain in admission denominators. The new common-box control in Section 5.6 removes this feasible-set mismatch for its own three methods by using one explicit component-box superset. It does not replace the historical interval runs or reproduce a general verified least-squares package (Section 1.1).

Risk uses boxed candidate–reference displacement exceeding τ. Tightness uses *unboxed* displacement divided by its bound. Point, normal, and visibility failures are separate labels. Zero denominators, nonpositive diagnostic eigenvalues, and empty sets are explicit, not epsilon-regularized. Bound audits use the disclosed numerical slack $5\,10^{-10}\max(1, \text{shift}, \text{bound})$; row audits retain $2\,10^{-8}$. Neither slack is added to an admission threshold. Cells are nested in tasks, so pooled counts do not provide independent statistical replications or a population-level zero-failure claim.

The cost study uses the same 11,977 historical cells, one full warmup and five complete timed repeats, cyclic method order and one-thread settings. We separate the common numerator and spectral increment. A separate diagnostic times selected-leaf source-cell bounds and the cached normal threshold. Neither includes extraction, enlarged visibility-box production, or an end-to-end verified producer.

Table 3: Distinct analysis populations. Eligible retains the mathematical and geometric prerequisites; the optional normal policy is a separate filter.

Use	Tasks	Cells	Interpretation
Mechanism	13	11,977	Eligible plus $\kappa = 0.12$ policy
Historical cost	13	14,021	Common eligible, optional policy removed
Known-stress cost	9	11,831	Previously observed confirmation, common eligible
Merged cost	22	25,852	Common eligible of 43,207 active cells

4.4 RQ5: fixed-service tiered verification protocol

RQ5 uses a later, separately locked verification workload. It reuses the historical and confirmation tasks without increasing the number of independent shapes. The common filter retains $g < 1$, path and tube checks, valid budgets, unchanged correspondence and the same $\lambda = 0.5$ and original absolute $\tau_C = 0.5h_C$. Removing κ does not remove a mathematical premise. The D01 cell-scale correction and superseded/aborted records remain separately indexed. Only the corrected, locked run supplies the new tables.

The four strategies are A, effective spectral on every eligible cell; B, cheap then effective spectral when needed; C, cheap only; and D, interval contraction. A and B use the same valid fallback and immutable declarations (Proposition 11). This is a fixed-grid computation: no row reselection, regularization tuning, oracle routing or adaptive refinement is allowed. The saved per-cell admission-equivalence audit (H1) confirms a finite cheap bound for every one of the 25,852 eligible cells, within the proposition's stated domain.

Cold cache means an empty *verification-stage cache* for each task and strategy, not a flushed operating-system or disk cache. Hot cache means that the same strategy has already followed its own legal fill path. Common, spectral and interval entries may be retained; an omitted B spectral stage cannot borrow an A entry. A final acceptance is recomputed from the bound and the unchanged tolerance. Cache-fill time is recorded, not treated as free.

The fixed schedule uses one warmup and five formal repeats, alternating strategy order and single-thread settings. Across four strategies, 22 tasks and two cache regimes this gives 176 warmup and 880 formal batches; these are repeated services on 22 known tasks, not 880 independent models. For a group and repeat we sum its task verification walls first, then take the median over the five group totals. We also retain all five paired A/B differences. Median ratios and the protocol's finite-range/noise rule are engineering descriptions, not confidence intervals or significance tests.

Measured wall time includes verification calls, dispatch, cache access, numerical bounds, admission decisions and result-object construction within the returned batch. It excludes input preparation, hot-cache filling, post-return destruction/garbage collection, H1 cross-checks, summary construction and output writing. Preparation and fill costs are reported separately in Supplementary Section S3. Field creation, Hermite extraction, mesh completion and any repair are outside these timings; RQ5 is not an end-to-end application benchmark.

5 Results

Unless explicitly marked as the matched-input control, populations, admission labels, tables and timings retain their historical execution semantics. Historical eligibility and audit-safe labels refer to the stated numerical tests, not new machine certifications from the corrected stored-data interface. New exact downstream checks do not certify that the historical upstream field satisfies its declared row budgets.

5.1 Spectral room and tolerance crossings on the fixed historical cohort

The primary mechanism population is the historical 11,977-cell shared cohort, not all active cells or a newly selected favorable subset. Of these cells, 27 have an improved verified denominator. At $0.5h$, 25 are already admitted by the matched λ bound, one remains rejected, and only one crosses the threshold. At $0.25h$ the same 27 denominators yield 14 new admissions and 13 remaining rejections; at $0.1h$ none crosses. Every case agrees with Equation (12), including the executable outward decisions.

The 27 gains occur in D01 (13 cells), M02 (2), and M03 (12). Structure-only labels place 25 at multi-leaf junctions and two in the pre-labelled thin-wall task M02; none is in the M05 shell containing the 17 historical QEF failures. The sole $0.5h$ increment is M02 cell $(16, 14, 18)$: $q_{\mathrm{dec}} = 1.03144$ and $s = 1.04002$. Its bound

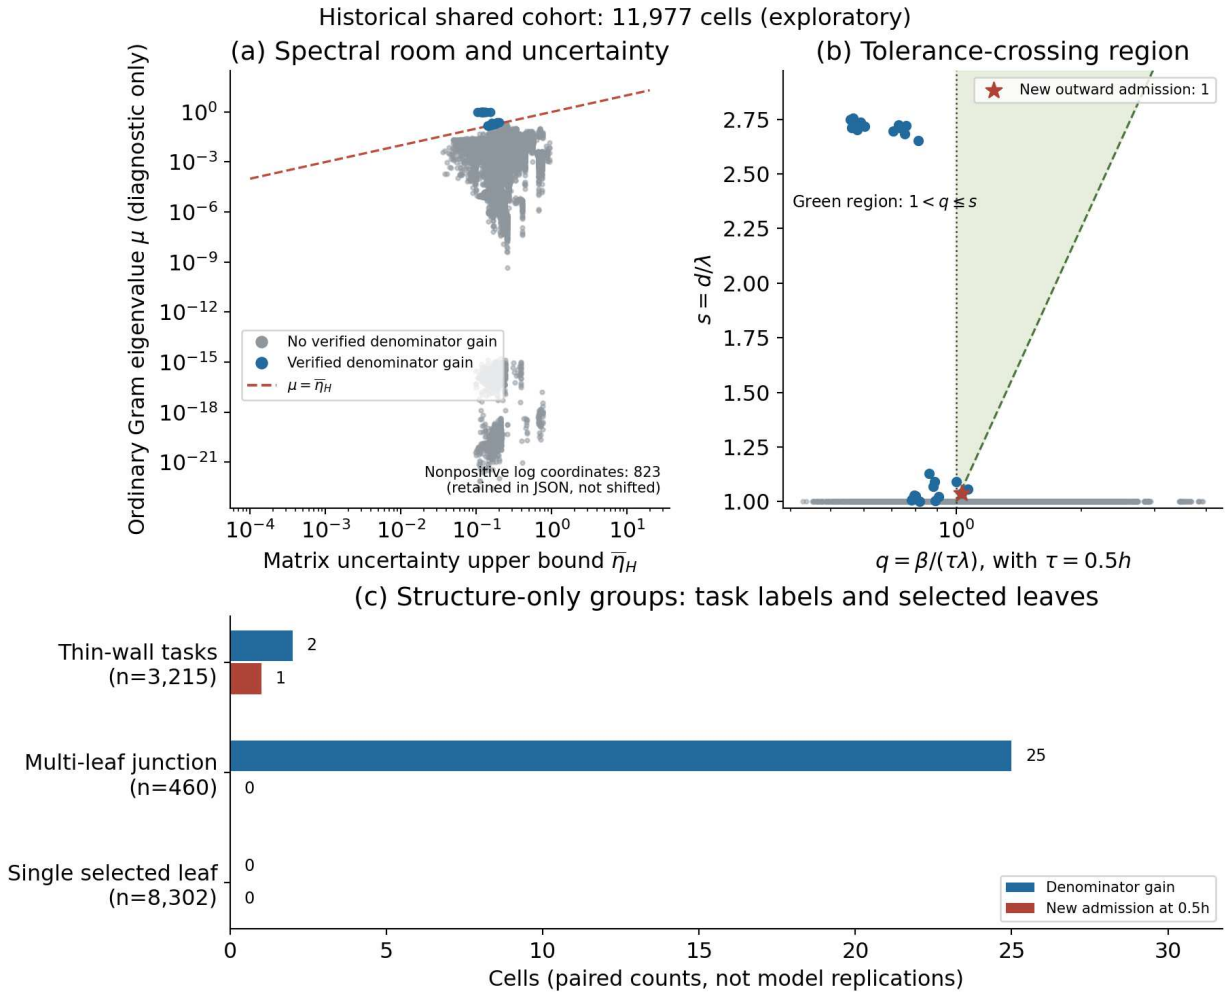


Figure 2: Historical shared-cohort mechanism diagnostics ($n = 11{,}977$). (a) Ordinary Gram eigenvalues illustrate spectral room but never certify it; 823 nonpositive diagnostic values are retained in data and omitted from log coordinates without epsilon shifts. (b) Only a tolerance crossing creates an additional admission. (c) Geometry labels use task tags and the number of selected leaves, not the observed rank or result. Cell counts are descriptive, not independent model replications.

changes from 0.0497300 to 0.0478166 across the 0.0482143 tolerance; its boxed reference displacement is 0.00467695.

Supplementary analysis of all 21,156 active cells finds 35 denominator gains; its non-shared cells are not merged into primary admission rates. An exhaustive record-level ledger is included in the artifact. The existing 50,000 algebraic systems were replayed with an unchanged canonical digest, not expanded. Their 23,577 improved denominators produce 13,934 new admissions at tolerance 0.5: 7,850 in the prescribed full-rank family and 6,084 in the mixed-random family. Neither the near-parallel nor edge-pair family contributes an increment at that tolerance. The families also differ from the CAD cells in uncertainty and regularization distributions, so this contrast supports the row-system mechanism without isolating geometry as its sole cause or establishing a CAD-level benefit.

5.2 RQ2: policy interaction and confirmation

The historical primary rule admits 3,944 of 21,156 active cells with no observed point, normal, visibility, or QEF violation. The older value-only ablation's 114 failures comprise 97 invisible projections and 17 QEF displacements, not 114 visibility failures. All 17 remaining after pointwise branch separation occur in M05. They pass the witness-ball and tube checks but fail both the optional normal threshold and QEF gate.

At $\kappa = 0.12$, removing only the QEF gate admits 11,977 cells without observed failure; removing only the normal policy keeps 3,944. Removing both admits 14,021 with those 17 failures. Thus single-factor ablations hide overlapping protections. Yet after the cheap normal policy, the minimum QEF gate rejects

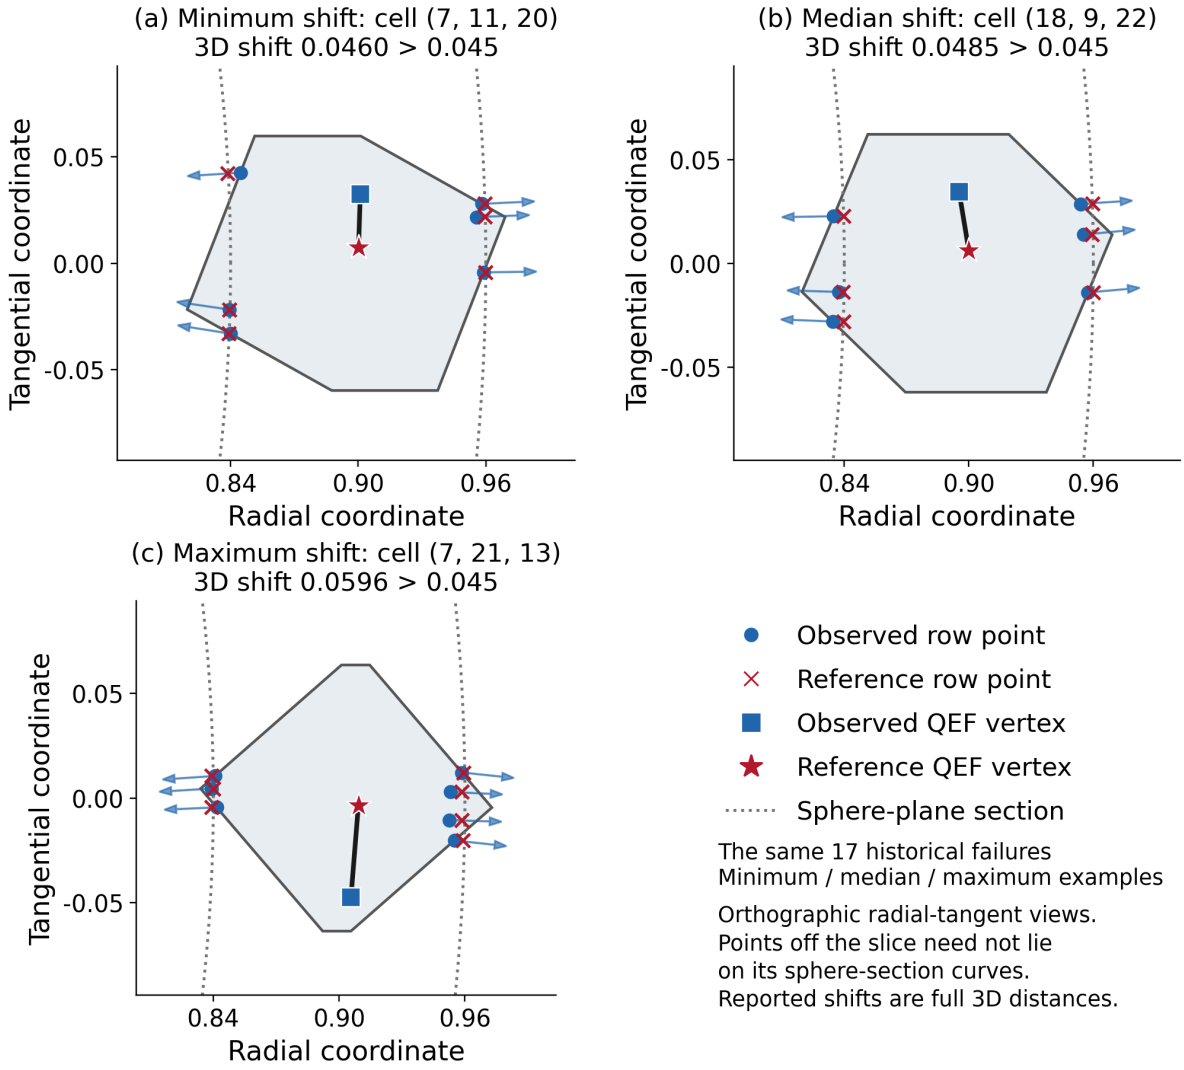


Figure 3: Three representative historical M05 failures. Points, normals, and vertices are projected onto a radial–tangential plane; displacements are measured in 3D. Both the normal policy and QEF gate reject these cells. These are historical M05 cases, not the V9W02 confirmation failures.

8,033 additional audit-safe cells with no observed incremental interception at the primary tolerance.

Relaxing to the predeclared $\kappa = 0.24$ exposes the 17 historical failures to the normal-only policy, while the minimum and spectral tests retain 3,944 and 3,945 safe candidates. This is exploratory evidence of independent interception with nonempty admission, accompanied by 10,060 and 10,059 extra safe rejections. It must not be erased by a blanket claim that QEF gates have no independent value. The question is whether this is a useful, reproducible tradeoff beyond an already observed suite.

Supplementary Section S1 retains the full policy frontiers, including empty and Unknown outcomes.

The nine confirmation tasks contain 22,051 active cells; 11,831 were classified as mathematically eligible by the historical numerical tests. At $\kappa = 0.12$, every method has empty admission. At $\kappa = 0.24$ and $\tau = 0.5h$, normal-only admits all 11,831, with 17 over-tolerance displacements in V9W02 alone. All three QEF tests admit zero: they intercept these failures but also reject 11,814 audit-safe cells. The interval test cannot construct a finite bound on 11,828 eligible cells; its three available bounds still exceed the tolerance, so all interval admission outcomes are Unknown. No numerical bound violation is found. No nonempty QEF admission appears at the smaller tolerances or looser optional policies.

Supplementary Section S1 visualizes all nine confirmation outcomes.

There is a specific design limitation: every selected leaf's source-cell phase width $h\|k_j\|_1$ is between 3.2573 and 3.5258, larger than the period π of $|\cos|$. Hence its gradient envelope reaches full amplitude, and the per-cell maximum g is at least 0.132. The $\kappa = 0.12$ rejection is analytically predictable; thresholds at or above 0.24 give the same eligible population. This confirmation cannot resolve fine differences between those thresholds. The relaxed contrast was nevertheless fixed before extraction and supplies a nonempty

Table 4: Primary-tolerance policies, $\tau = 0.5h$. Historical cells are exploratory; confirmation tasks were fixed before extraction. Extra safe rejections are relative to the normal-only policy at the same κ, using numerical audit labels. A dash denotes undefined risk for an empty admitted set.

Population / κ	Method	Admitted	QEF $> \tau$	Extra safe rej.	Risk (%)
Historical / 0.12	Normal only	11,977	0	0	0.000
Historical / 0.12	Minimum / λ	3,944	0	8,033	0.000
Historical / 0.12	Spectral	3,945	0	8,032	0.000
Historical / 0.12	Interval	2,235	0	9,742	0.000
Historical / 0.24	Normal only	14,021	17	0	0.121
Historical / 0.24	Minimum / λ	3,944	0	10,060	0.000
Historical / 0.24	Spectral	3,945	0	10,059	0.000
Historical / 0.24	Interval	2,235	0	11,769	0.000
Confirmation / 0.12	Normal only	0	0	0	–
Confirmation / 0.12	Minimum / λ	0	0	0	–
Confirmation / 0.12	Spectral	0	0	0	–
Confirmation / 0.12	Interval	0	0	0	–
Confirmation / 0.24	Normal only	11,831	17	0	0.144
Confirmation / 0.24	Minimum / λ	0	0	11,814	–
Confirmation / 0.24	Spectral	0	0	11,814	–
Confirmation / 0.24	Interval	0	0	11,814	–

Table 5: Matched 11,977-cell historical cohort. Tightness r is unboxed measured displacement divided by the bound; lower means more conservative. Available/unavailable describes construction of a finite bound, not admission at a tolerance. Own and common finite-bound medians have different denominators. Cost is the median of five complete construction batches after one warmup, not single-query latency.

Method	Available	Unavailable	Own r_{50}	Paired r_{50}	Paired r_{99}	Batch (s)
Minimum / λ	11,977	0	0.0446	0.0505	0.0936	3.212
Spectral	11,977	0	0.0446	0.0505	0.0945	26.739
Interval	7,894	4,083	0.0166	0.0166	0.3178	7.131

normal-policy comparator. The result fails to confirm *selective* QEF benefit under this design, not to prove universal uselessness. We apply the stopping rule and add no further similar synthetic cases.

5.3 RQ3: paired tightness and construction cost

The common finite-bound set contains 7,894 cells. On it the minimum bound is smaller than the interval bound in 5,760 cases; the reverse holds in 2,134. Neither method uniformly dominates. The analytic and spectral median tightness rises from 0.0446 on their own 11,977 finite-bound cells to 0.0505 on the paired set; the interval paired median is 0.0166. Reporting only each method's own quantiles would confound bound quality with Unknown selection.

The spectral implementation costs 3.75 times the interval construction in median complete-batch time on this host. Its common-numerator work takes 3.28 seconds, with a further 23.32 seconds for the checked spectral increment. This does not justify the extra cost for the single added primary-tolerance admission in the historical cohort. It is an implementation measurement, not a lower bound on spectral-verification cost or a ranking of all verified numerical methods. The separate producer diagnostic evaluates 12,672 distinct selected-leaf boxes for 11,977 cells: median costs are 4.24 microseconds per cell for source-cell bounds and 0.31 for a cached normal threshold. These partial costs are not end-to-end latency or a validated-producer benchmark.

5.4 RQ4: common completion and external diagnostics

We replay the thirteen historical tasks with the locked $\kappa = .12$ policy and $\tau = .5h$. All methods use the same candidate rows, center, $\lambda = .5$ candidate, connectivity, and $\lambda = 10^{-5}$ fallback for rejected cells. The all-small-λ mesh is a contextual control, not an admission method. Selection reads declarations only; reference-row audits are consulted after the mask is fixed. Each complete surface is measured using the

Table 6: Locked historical completion. Selected/Unknown percentages use all 21,156 active cells. Bad counts selected candidates with either a reference-row failure or boxed QEF shift above τ, not all mesh defects. Median, P95 and sample maximum are task medians of bidirectional sampled distances, in percent of each task bounding-box diagonal. NM is summed edge-nonmanifold count. Context has no admission/Unknown statistic.

Completion	Sel.%	Unk.%	Bad	Median	P95	Max	NM
Standard, $\lambda = .5$	100.00	0.00	257	0.0778	0.2083	0.5774	49
Normal policy	56.61	43.39	0	0.0738	0.1930	0.5174	49
Minimum/λ	18.64	81.36	0	0.0673	0.1651	0.5422	49
Spectral	18.65	81.35	0	0.0673	0.1651	0.5422	49
Interval	10.56	89.44	0	0.0643	0.1440	0.5330	49
Context, $\lambda = 10^{-5}$	–	–	–	0.0637	0.1332	0.5330	49

Table 7: Fixed-input validation cost: medians of five group-total wall times (seconds). A always evaluates the effective spectral bound; B first tests the cheap bound. Workloads retain the same A/B acceptance set.

Workload	A (s)	B (s)	Reduction	Range-based reading
Cold: historical (13)	31.3336	23.6530	24.51%	Finite cold-cache saving
Cold: known stress (9)	26.4842	26.4666	0.07%	No reliable gain established
Cold: combined (22)	57.7863	50.0893	13.32%	Finite cold-cache saving
Hot: combined (22)	0.5205	0.5184	0.39%	No reliable gain established

Each repeat first sums all task batch times in the stated group; the table takes the median of those five sums, not a sum of task medians. Reduction is $100(1 - \operatorname{med} T_B / \operatorname{med} T_A)$, not a paired-difference statistic. Five-run ranges support only the prespecified engineering check, not confidence intervals. Timing includes actual dispatch, cache access, numerical work, decisions and batch append; it excludes upstream preparation, meshing and post-batch checking/aggregation/output. Preparation and strategy-specific cache fill are reported in Supplementary Table S2; all five pairs are in Supplementary Table S1.

inherited 128-resolution numerical reference, 20,000 area-weighted samples in each direction, and seed 20260905.

The spectral and minimum-bound completions are geometrically identical in twelve tasks. Only M02 changes, by selecting the single additional candidate identified above. Their task-median error statistics are therefore identical in the reported table. All six methods retain 49 edge-nonmanifold incidences. The standard same-λ mesh has no degenerate triangles, whereas all four gated completions and the small-λ context have 24. Opposed-nearest-reference-face counts are orientation diagnostics, not proofs of flipped faces; edge-closed components do not establish embeddedness or vertex manifoldness.

For the spectral completion, 81.35% of cells use unverified fallback vertices. The task-median symmetric mean distance is 0.0733% of the box diagonal, versus 0.0916% for the same-λ standard and 0.0645% for the all-small-λ context. Changing the regularization of rejected cells is part of this completion, not an extension of the local theorem. We do not infer generally better meshes from these diagnostics, and no additional reference-sensitivity experiment is used to promote the tiny M02 difference into a superiority claim. The sampled distances use point-to-triangle queries; the sample maximum is not a continuous Hausdorff bound. No completion result changes the theorem object or establishes generally improved geometry. Candidate-reference displacement is distinct from surface error, and fallback regularization can dominate the complete mesh. An adaptive refinement benefit is not tested. Because we make no mesh-quality superiority claim, these fixed sampling measurements are descriptive rather than a new sampling-sensitivity argument.

5.5 RQ5: equal-admission cost and limited marginal value

Using the fixed populations, strategies and timing boundaries of Section 4.4, we ask whether the same effective spectral acceptance set can be obtained at lower verification cost. The saved environment was single-threaded Python 3.14.5 on macOS arm64, using a monotonic wall clock. The task-offset cyclic/reversed schedule uses eight method orders; five repeats do not form a complete Latin square.

Where the tiered implementation saves time. Every saved batch's admission digest agrees with the H1 identifier sets; A and B both accept exactly 3,945 cells, with no set difference. This finite regression is implementation evidence for the conditional equal-admission argument, not its proof. Table 7 shows a 24.51% reduction of median historical cold-cache validation time and a 13.32% reduction for the combined cold workload. The latter passes the prespecified engineering check: positive median paired saving exceeds the larger five-repeat time range, 11 of 13 historical tasks have positive median paired savings, and historical aggregate savings remain positive after removing each historical task in turn. The known-stress group does not show a reliable saving. Neither does the combined hot-cache workload, whose observed paired differences change sign and whose variability exceeds the median saving. Supplementary Table S1 shows every pair rather than a smoothed trend. These small-sample engineering checks are not statistical confidence intervals or evidence of speedup on new workloads.

Why not use only the cheap bound? The saving relative to A is not a demonstration that the additional spectral service is worth requesting. C accepts 3,944 cells: B gains only 1 cell, identified in the archived record as `M02_hollow_pipe_tee:16,14,18`. In the combined cold workload, C's median validation time is 7.4639 seconds, and B costs 6.71 times as much. The difference between their two total-time medians is approximately 42.63 seconds. This is neither a median paired difference nor a general price per newly accepted cell, and it does not establish an application break-even threshold. B still invokes the spectral stage for 21,908 of 25,852 eligible cells (84.74%) in every formal cold repeat.

D accepts 2,235 cells in 7.9899 seconds of median combined cold validation. Its set shares 1,779 cells with A and contains 456 cells not accepted by A. Thus D is not a subset of A, and its lower total count is not sufficient to declare it inferior. A versus B measures implementation cost under the same admission service; B versus C asks about the cost of requesting that extra service; A/B versus D compares different enclosures and partially different admission sets. Keeping these questions separate explains why the finite cold-cache result does not overturn the limited marginal-value conclusion of the mechanism study. It also does not prove that spectral information is universally useless.

Structure-aware control. On the same 25,852 eligible declarations, the untimed unit-structure control replaces only the Hessian uncertainty by Lemma 9. Thirty final bounds become smaller, but the admitted identifier set remains exactly 3,945. This known-data control removes one avoidable source of slack; it does not show that the remaining bound is optimal. Supplementary Section S5 retains the full aggregate and taskwise records, with no new timing claim.

5.6 RQ3 follow-up: a genuinely common-box numerical control

The historical comparison does not isolate solver quality on an identical feasible set. A subsequent known-data pilot therefore fixes one family for all three controls:

$$|N_{ij} - \widehat{N}_{ij}| \le \eta_{n,i}, \qquad |h_i - \widehat{h}_i| \le \eta_{h,i}.$$

This component-box family is a common *superset* of the unit-normal, Euclidean-row-ball family. In particular, every control uses $\|\Delta n_i\| \le \sqrt{3}\eta_{n,i}$ and a general-row norm bound; none takes an extra unit-normal advantage. The analytic and scalar interval controls are recomputed for this family, not copied from the historical tables.

Before evaluating the new methods, we locked eight evenly spaced eligible cell identifiers per historical task. Twelve tasks contribute 96 cells; M05 has no eligible records in that selection source and contributes zero. Selection uses previously recorded prerequisites, not the new bound or an oracle error. This is a small known-data pilot, not independent confirmation. All 200 saved small-grid cell records form a separate software-fixture stratum, reported only in Supplementary Section S8.

The stronger control retains componentwise residual information and also evaluates a centered row-perturbation form. For nonnegative enclosures $M \ge |I - RH|$ and $b \ge |R(q - H\widehat{y})|$, the rational certificate checks

$$w > 0, \quad Mw < w, \qquad u \ge 0, \quad b + Mu \le u.$$

The weighted contraction proves $|y - \widehat{y}| \le u$. The nominal preconditioner, residual and comparisons use exact rational arithmetic; norm endpoints are enclosed. Zero row radii therefore do not erase a nonzero stored-iterate residual. Supplementary Section S8 gives the argument, tests and all tolerances. This is a componentwise numerical control, not a new verification principle or a reproduction of Rump's algorithm.

Let s_C be the largest exact stored box width. At $0.5s_C$ the componentwise method admits 71 of 96 declarations, compared with 27 and 22 for the common-box analytic and scalar controls. It constructs

Table 8: New locked pilot on 96 stored historical cells from 12 tasks. All three methods cover the same component-box family. The last three columns count conditional admissions at fractions of the largest exact stored box width s_C, including output-rounding defect. No-finite-bound cases remain in every denominator. These are not the old unit-normal bound results.

Method	Finite	No finite bound	$0.1s_C$	$0.25s_C$	$0.5s_C$
Analytic common-box	96	0	0	0	27
Scalar contraction	64	32	3	12	22
Componentwise	92	4	26	62	71

92 finite bounds, retaining four failures as Unknown. Among its 92 paired finite cases with the analytic control, it is smaller in 85 and larger in seven; the median ratio is 0.276. Against the scalar control it is smaller in all 64 common-finite cases, with median ratio 0.111. Failed constructions are not dropped from the admission comparison.

The changes combine centered-form enclosures with a componentwise weighted contraction; this pilot does not isolate their individual contributions. This changes the interpretation of the earlier negative results: part of the conservatism is specific to the bound construction, not an established limitation of QEF certification. It does not establish superiority over the original smaller unit-normal uncertainty family, all 11,977 historical cells, mature verified least-squares packages, or new CAD inputs. Upstream historical budgets remain conditional declarations. No per-method timing, new extraction or surface-accuracy advantage is claimed for this pilot.

6 Discussion and limitations

A tighter bound, a different admission, a cheaper implementation, and a useful application are four distinct outcomes. The historical mechanism separates the first two; RQ5 separates the last two. The tiered strategy can reduce some cold-cache costs of an already required full-spectral service. It does not reverse the limited marginal-value conclusion: the cheap-only control gives almost the same admissions at much lower verification cost. Nor do finite negative results establish that QEF bounds are universally useless. The interval control admits a different set and cannot be ranked solely by its total count.

The producer remains a controlled analytic error family evaluated in ordinary floating point. Downstream outward bounds verify consequences of declared endpoints, not that a real field lies within them (Table 2). The stored-data corrections specify additional proof obligations; they do not validate these historical producer declarations. The separate restricted analytic realization discussed in Supplementary Section S7 supplies pointwise enclosures under a different information model, not new evidence for RQ1–RQ5 or the surface-audit study. Strictly conservative neural distance fields [17] do not automatically supply two-sided errors, normal bounds, positive reach or unique-zero information. The separate sampled-grid pilot encloses interpolation quantities but is not integrated with Hermite/QEF admission. The historical STEP diagnostic starts from final solids and reference-derived errors, not CSG histories or an independent producer. Both tracks, including their exclusions and unavailable inputs, remain in Supplementary Section S2; neither is pooled with analytic CSG or used to support the cost claim.

The experiments concern finite designed populations. Historical policy sweeps are exploratory; the nine original confirmation tasks share the perturbation model and a phase-saturated gradient envelope. They are already known in RQ5. Five timing repeats estimate variation in fixed task totals, not variation over independent models or a confidence interval for future workloads. The historical reference audit is not a formal proof. The new componentwise control compares one stronger enclosure on a shared component-box model; it does not benchmark mature verified least-squares packages or retain all unit-normal and inter-row correlations. The Python host implementation, small matrices, richer input access and explicitly excluded costs limit transfer to production kernels.

6.1 What a useful recomputation experiment must establish

A zero displacement bound is not a surface certificate. Take $C = [-1, 1]^3$, $c = 0$, the exact plane $x_1 = 9/16$ and the four crossing-edge rows $n_i = (1, 0, 0)$, $h_i = 9/16$. With $\lambda = 1/2$ and identical exact observed/reference data, both regularized QEF solutions are $(1/2, 0, 0)$. Row, residual and output errors may all be zero, so the displacement bound is zero, whereas the distance to the plane is $1/16$. This example satisfies the theorem, including common-box projection. It isolates reference-fit bias rather than exhibiting a counterexample.

An explicit, additional geometry obligation. Let $z_* = P_C(c+y)$ and let $p_i \in S$ be a proved visible reference witness. For an actual returned vertex z, write $B_{\text{stored}} = B_{\text{QEF}} + d_{\text{out}}$. Then

$$d(z,S) \le B_{\text{stored}} + \|z_* - p_i\|.$$

The second term is not supplied by the QEF displacement theorem. If only approximations $\bar{z}_*$ and $\bar{p}_i$ are available, with verified radii e_* and $e_{p,i}$, a usable upper bound is

$$B_{\text{stored}} + e_* + \overline{\|\bar{z}_* - \bar{p}_i\|} + e_{p,i}.$$

All upper bounds must be valid, and their production must be charged to the method. Solving a reference problem solely for an offline audit does not make its value free information for admission.

There is also a cheaper independent control: the already established $\|\widehat{p}_i - p_i\| \le \eta_{p,i}$ implies $d(z,S) \le \overline{\|z - \widehat{p}_i\|} + \eta_{p,i}$ without a QEF displacement bound. Any proposed practical benefit should survive comparison with this direct-witness test. These inequalities are elementary triangle inequalities, not new distance theory. Vertex tests alone do not certify whole output triangles, the reverse surface distance, topology or missing components.

The engineering protocol therefore requires a fixed geometric error target, a valid reference-fit or direct-witness producer, and an actual local recomputation operation. Uniform recomputation, the cheap witness control and QEF-guided recomputation must use the same final checker. Selection may not read free offline oracle labels. Producer, certificate, fallback, repair and final-check costs belong in the total. The source preflight in Supplementary Section S9 finds no complete existing interface satisfying these requirements for the saved engineering models. Its execution gate remains closed; no unrelated analytic surface experiment is substituted.

Finally, the guarantee controls displacement between matched regularized systems, with common-box projection justified in exact arithmetic and actual stored output requiring the defect in (11). It does not remove regularization bias, ensure that every component is sampled, prove topology or manifoldness, or bound global mesh error. These are different objectives [24, 23]. A useful downstream deployment would have to save an actual expensive operation while retaining the same output requirements and accounting for unknown-region fallback, producer, verification and repair costs. Neither whole-CAD speed, peak-memory reduction nor a CFD accuracy claim is established by these records. Supplementary Section S6 preserves the separate directed fixed-mesh audit, including its residual Unknown region and local-search cost. It is not counted as evidence for this application's missing QEF-to-surface link.

7 Conclusion

The visible-row correspondence makes a local regularized-QEF displacement claim well-defined for approximate hard CSG. Under its stated geometric and numerical premises, the bound is valid even when using it as a gate is unattractive. Spectral tightening requires verified room beyond uncertainty and adds a decision only when it crosses the tolerance. The historical and fixed confirmation studies explain why these opportunities can be sparse and why conservative rejection must not be confused with practical benefit.

Valid nested bounds permit equal admission with less spectral work, not an automatic acceleration theorem. The historical locked workload shows a 13.32% reduction in median cold-cache verification time, without reliable savings on known stress tasks or hot-cache calls. The tiered service still costs about 6.71 times the cheap-only service for one extra admission. Thus the cost result improves a specific verification workload; it does not establish general application value for additional QEF gating. We retain the narrower contribution: a conditional local guarantee, a mechanism for its conservatism, and reproducible evidence of limited marginal value. The unit-structure control removes avoidable Hessian slack but leaves the fixed-data admission set unchanged. Its untimed result strengthens the interpretation of this particular negative finding, not a universal claim about all possible QEF bounds. The new common-box pilot, in contrast, adds substantial conditional admission on its small matched population. This positive control makes the interpretation more specific: the historical negative outcome is not an intrinsic barrier to stronger numerical enclosures. A larger comparable study and genuine engineering use remain unestablished. The engineering preflight does not change the historical cohorts or establish a whole-surface advantage. Stored-data corrections strengthen the meaning of the conditional guarantee, while the explicitly missing reference-fit and recomputation interface prevents claiming end-to-end engineering closure.

Data, code and reproducibility

The reproduction material is organized into a manuscript bundle and a companion historical-evidence bundle. The first contains the LaTeX sources, figures, the common-box pilot's fixed inputs and results, and a portable entry point. The companion retains the hash-identified archives needed for the earlier QEF, cost and stored-data studies. Its guide maps each result family to an archive and separates checksum checks, table regeneration, software tests and optional scientific re-execution. Neither checksums nor successful recompilation constitute independent scientific replication.

A separate local addendum provides the hardened runtime, tests and source-protection records (Supplementary Section S7: 140 passing tests and two existing one-cell comparisons). It neither replaces the frozen archives nor adds scientific results.

A separate saved-record repair follows the retained interface audit. It adds strict Boolean guards, explicit format dispatch, recomputed same-target QEF/output budgets, and separate candidate/returned-vertex status fields. Its 31 new tests and the retained 140 tests pass. Bounded old/new replay preserves 96 historical records' numerical outputs; replay of explicitly upgraded copies of 200 existing fixtures preserves their candidate budgets without promoting the 42 uncertified outputs. The reproduction guide distinguishes this repair from frozen historical code and from independent geometric recertification (Supplementary Section S7.6).

Supplementary Section S4 records the study chronology and archive aliases; Sections S3 and S6 preserve measurement windows, excluded costs and the failed predecessor of the mesh audit. Original records are not overwritten. Some auxiliary history remains nonportable: 517 original STEP paths are not restored, and an early workspace-integrity check requires unbundled build and cache files. These limits do not turn a retained result into a freshly reproduced one. Source attribution and license notices are retained; the bundles do not assert unrestricted redistribution rights for all third-party models.

The stored-data obligations and software fixtures do not recertify historical upstream budgets. The engineering input guard still rejects an incomplete interface; no full geometry-utility run is claimed. Detailed reproduction instructions and access status accompany the bundles. A public deposit and reviewer download access have not yet been established. Evidence-index and assistant tools are optional navigation aids, not additional scientific evidence. No independent human proof review of the present revision, formal program verification or assistant-productivity evaluation is claimed.

Funding and competing interests

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. The author declares no conflict of interest.

References

[1] T. Ju, F. Losasso, S. Schaefer, and J. Warren. Dual contouring of Hermite data. *ACM Transactions on Graphics*, 21(3):339–346, 2002. doi:10.1145/566654.566586.

[2] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, 1987. doi:10.1145/37402.37422.

[3] L. P. Kobbelt, M. Botsch, U. Schwanecke, and H.-P. Seidel. Feature sensitive surface extraction from volume data. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 57–66, 2001. doi:10.1145/383259.383265.

[4] Z. Chen, A. Tagliasacchi, T. Funkhouser, and H. Zhang. Neural dual contouring. *ACM Transactions on Graphics*, 41(4), Article 104, 13 pages, 2022. doi:10.1145/3528223.3530108.

[5] R. Sundararaman, R. Klokov, and M. Ovsjanikov. Self-supervised dual contouring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4681–4691, 2024. doi:10.1109/CVPR52733.2024.00448.

[6] Z. Marschner, S. Sellán, H.-T. D. Liu, and A. Jacobson. Constructive solid geometry on neural signed distance fields. In *SIGGRAPH Asia Conference Papers*, Article 121, 2023. doi:10.1145/3610548.3618170.

[7] C. Bálint, G. Valasek, and L. Gergő. Operations on signed distance function estimates. *Computer-Aided Design and Applications*, 20(6):1154–1174, 2023. doi:10.14733/cadaps.2023.1154-1174.

[8] P. Trettner and L. Kobbelt. Fast and robust QEF minimization using probabilistic quadrics. *Computer Graphics Forum*, 39(2):325–334, 2020. doi:10.1111/cgf.13933.

[9] S. Schaefer, T. Ju, and J. Warren. Manifold dual contouring. *IEEE Transactions on Visualization and Computer Graphics*, 13(3):610–619, 2007. doi:10.1109/TVCG.2007.1012.

[10] S. Plantinga and G. Vegter. Isotopic approximation of implicit curves and surfaces. In *Symposium on Geometry Processing*, pp. 251–260, 2004. doi:10.2312/SGP/SGP04/251-260.

[11] W. Barbier, M. Sanchez, A. Paris, E. Michel, T. Lambert, T. Boubekeur, M. Paulin, and T. Thonat. Lipschitz pruning: Hierarchical simplification of primitive-based SDFs. *Computer Graphics Forum*, 44(2):e70057, 2025. doi:10.1111/cgf.70057.

[12] H.-T. D. Liu, M. Agrawala, C. Yuksel, T. Omernick, V. Misra, S. Corazza, M. McGuire, and V. Zordan. A unified differentiable Boolean operator with fuzzy logic. In *ACM SIGGRAPH Conference Papers*, Article 109, 2024. doi:10.1145/3641519.3657484.

[13] X. Carrera, N. Wang, C. Batty, O. Stein, and S. Sellán. Dual contouring of signed distance data. In *ACM SIGGRAPH Conference Papers*, Article 38, 12 pages, 2026. doi:10.1145/3799902.3811116.

[14] M. Kohlbrenner and M. Alexa. Contouring signed distance fields by approximating gradients. *Computer Graphics Forum*, 45(2), Article e70373, 2026. doi:10.1111/cgf.70373.

[15] M. Kohlbrenner and M. Alexa. Isosurface extraction for signed distance functions using power diagrams. *Computer Graphics Forum*, 44(2):e70037, 2025. doi:10.1111/cgf.70037.

[16] S. Sellán, Y. Ren, C. Batty, and O. Stein. Reach for the arcs: Reconstructing surfaces from SDFs via tangent points. In *ACM SIGGRAPH Conference Papers*, 2024. doi:10.1145/3641519.3657419.

[17] I. Ludwig and M. Campen. Strictly conservative neural distance fields. *Computer Graphics Forum*, Early View, Article e70528, 2026. doi:10.1111/cgf.70528.

[18] H. Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959. doi:10.1090/S0002-9947-1959-0110078-1.

[19] R. Bhatia. *Matrix Analysis*. Springer, New York, 1997.

[20] A. Neumaier. *Interval Methods for Systems of Equations*. Cambridge University Press, Cambridge, 1991. doi:10.1017/CBO9780511526473.

[21] S. M. Rump. Verification methods: Rigorous results using floating-point arithmetic. *Acta Numerica*, 19:287–449, 2010. doi:10.1017/S096249291000005X.

[22] M. Burr, J. D. Hauenstein, and K. Lee. Certified surface approximations using the interval Krawczyk test. In *Proceedings of the 2026 International Symposium on Symbolic and Algebraic Computation*, pp. 77–86, 2026. doi:10.1145/3815436.3815444.

[23] A. Chattopadhyay, S. Plantinga, and G. Vegter. Certified meshing of radial basis function based isosurfaces. *The Visual Computer*, 28(5):445–462, 2012. doi:10.1007/s00371-011-0627-2.

[24] S. Plantinga and G. Vegter. Isotopic meshing of implicit surfaces. *The Visual Computer*, 23(1):45–58, 2007. doi:10.1007/s00371-006-0083-6.

[25] C. Berenfeld, J. Harvey, M. Hoffmann, and K. Shankar. Estimating the reach of a manifold via its convexity defect function. *Discrete & Computational Geometry*, 67:403–438, 2022. doi:10.1007/s00454-021-00290-8.

[26] S. M. Rump. Verified bounds for least squares problems and underdetermined linear systems. *SIAM Journal on Matrix Analysis and Applications*, 33(1):130–148, 2012. doi:10.1137/110840248.

[27] S. M. Rump. Improved componentwise verified error bounds for least squares problems and underdetermined linear systems. *Numerical Algorithms*, 66(2):309–322, 2014. doi:10.1007/s11075-013-9735-6.

[28] S. M. Rump and T. Ogita. Verified error bounds for matrix decompositions. *SIAM Journal on Matrix Analysis and Applications*, 45(4):2155–2183, 2024. doi:10.1137/24M165096X.

[29] J. R. Shewchuk. Adaptive precision floating-point arithmetic and fast robust geometric predicates. *Discrete & Computational Geometry*, 18:305–363, 1997. Author's publication and software page.

[30] P. Trettner and L. Kobbelt. Sampling from quadric-based CSG surfaces. *Computer Graphics Forum*, 40(8):41–56, 2021. doi:10.1111/cgf.14380.

[31] Y. Kang, M.-H. Kyung, S.-H. Yoon, and M.-S. Kim. Fast and robust Hausdorff distance computation from triangle mesh to quad mesh in near-zero cases. *Computer Aided Geometric Design*, 62:91–103, 2018. doi:10.1016/j.cagd.2018.03.017.

[32] L. Sacht and A. Jacobson. Cascading upper bounds for triangle soup Pompeiu–Hausdorff distance. *Computer Graphics Forum*, 43(5):e15129, 2024. doi:10.1111/cgf.15129.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementTVC.pdf