

Walking into Holes: Monte Carlo Topological Derivatives for Grid-Free PDE-Constrained Design

Johannes Herter

ETH Zürich

Abstract. Creating a new interior boundary requires a different design decision from moving an existing surface. We address cavity nucleation using ideal cooled cavities in heat-generating solids. Classical topological derivatives express insertion benefit through state and adjoint fields. Topological Walk on Spheres (TWoS) estimates this coefficient with local random walks on the current geometry, without a volumetric PDE discretization or a shrinking trial hole. Cross-index products remove primal-adjoint correlation bias under explicit independence assumptions. We derive radius-independent variance and conditions for finite-insertion error control, including oracle bias, asymptotic error, and adaptive selection. A separate variational bound certifies finite Poisson-compliance benefit under stated sampling and geometry assumptions. Experiments compare actual-mesh estimators, complete insertion-refinement pipelines, and accessible channels with fixed material and heat-input budgets. Uncertainty bands and independent numerical checks retain competitive multi-grid and proximity controls. Matched-accuracy query sweeps expose strong dependence on source, objective, and field reuse; fresh-seed checks qualify apparent sparse-query gains. Conditional validation has substantial additional cost.

1 Introduction

Computational design chooses geometry to produce a desired physical response. Changing an existing boundary and deciding where a new boundary should appear are different decisions. We study the latter: selecting interior cavities in a solid to improve a PDE-defined objective. In our thermal application, prescribed cold cavity walls act as ideal internal heat sinks. Their placement changes conduction paths and the temperature of protected components. This model isolates a geometry-design problem; it does not simulate coolant flow.

A candidate could be scored by inserting a small trial hole, solving the modified PDE, and subtracting objectives. Classical topological derivatives avoid these repeated trial solves by expressing the leading change through the current state and adjoint [6, 27]. Global discretized solvers already exploit this formula. We ask how to evaluate it directly using local random walks on the boundary representation, without constructing a volumetric PDE discretization or a trial cavity.

Walk on Spheres and Walk on Stars provide the necessary point-query machinery [26, 29]. Differential WoS supplies sensitivities for existing geometry [23]. Our additional query is the capacity-normalized response to creating an interior boundary. We call its direct estimator *Topological Walk on Spheres* (TWoS). Primal and adjoint estimates can share a walk skeleton; cross-index products remove the resulting correlation bias when complete oracle pairs are independent. The estimator has no trial-radius parameter.

The problem is therefore a missing query in the grid-free design workflow: *where would a new boundary help?* Our contribution is the stochastic evaluation and deployment of classical topological sensitivity. The deterministic derivative, cross-index unbiasedness, and the Hoeffding decomposition are classical ingredients. The technical question is how their combination supports an insertion query and which additional assumptions turn that query into a finite geometric decision. We give the exact product-estimator variance, identify a shrinking-target variance barrier for specified finite-cavity couplings, and propagate sampling error, oracle bias, finite-radius error, and adaptive selection into explicit decision conditions. A separate variational argument gives a finite-benefit bound for Poisson compliance without an unknown asymptotic remainder. The




Figure 1 *Choosing where new cooled cavities should appear.* TWoS queries the current geometry to estimate insertion benefit, then adds selected cavities. Top: initial temperature and six-cavity random and voxel-adjoint designs. Bottom: the initial benefit field, an illustrative first insertion, and two TWoS allocations with similar measured cooling: six large cavities versus twelve small cavities at half the removed volume. These cavities have ideal zero-temperature walls. Temperature views share one color scale; benefit has a separate scale. The cut and inset edges are for display only, and some cavities lie outside the section. Per-design labels use independent objective evaluations; the plot summarizes three design seeds. The speed and memory callouts describe a separate 64-site query benchmark, not the full Bunny optimization. Field rendering and evaluation protocols appear in [Appendices B.6](#) and [X.5](#).

appendix contains every derivation and the additional assumptions needed by implementations.

The experiments distinguish placement quality from numerical performance. Three different solids compare complete cavity-initialization and Differential WoS refinement pipelines against tuned multigrid and geometric proximity. Matched-volume channels preserve the entire heat input. Actual-mesh estimator tests complement controlled statistical diagnostics, and a detailed Bunny study measures the crossover from sparse queries to a global solve. The evidence supports direct stochastic insertion scores; it also retains competitive deterministic and proximity controls rather than claiming universal design or runtime superiority.

2 Related Work

Grid-free PDE evaluation and differentiation. WoS estimates Dirichlet solutions using sphere exits [26]; WoSt extends the approach to mixed boundaries [29], with related estimators for Robin conditions [24]. Differential WoS computes sensitivities to problem parameters and geometry through point queries [23], and derivative WoSt improves spatial derivatives near Neumann boundaries [34]. These methods supply the PDE and existing-boundary derivative machinery on which we build. Our additional query is the coefficient of a singular cavity insertion, evaluated before its boundary exists. Fast cage cutting instead updates solutions after an explicit edit [28]. Flexible boundary representations and efficient updates are complementary to selecting where a new feature should be introduced.

Probabilistic shape optimization. Schlegel et al. [30] derive probabilistic representations of solution and objective shape derivatives for semilinear Dirichlet PDEs. Schmidt and Würschmidt [31] use a probabilistic derivative with a neural state representation to optimize shape. Our target is the capacity-normalized response to nucleating an interior cavity. Such a family is not a regular deformation of the original domain

([Proposition C.1](#)), although topological derivatives can be related to singular limits of shape derivatives on already perforated domains. We retain this established distinction without attributing topology-preservation limitations to every possible implementation of shape optimization.

Topological sensitivity and inverse design. Small-feature asymptotics and state-adjoint formulas are standard tools in topology optimization [4, 6, 11, 27]. They complement level-set boundary evolution [2, 5]; recent developments address finite-element recovery [20]. Meshfree spline methods already optimize holes using topological derivatives [19], and immersed isogeometric methods avoid remeshing using background spline spaces [33]. Neural implicit material representations offer another design parameterization [35]. In graphics, Mehta et al. [21] derive topological derivatives for inverse rendering. Our deterministic coefficient builds on this calculus; the focus is its direct random-walk evaluation and the resulting uncertainty in geometric decisions. Freiburger et al. [10] use a deterministically computed topological derivative to initialize MCMC reconstruction, a different role for Monte Carlo from estimating the derivative itself.

Sampling efficiency and operator scope. Boundary-value, harmonic, and probe caches amortize spatial queries [18, 22, 36]; neighbor reuse shares completed walks [8]. RQMC and Array-RQMC reduce sampling error [14, 15]. We examine these accelerations through an independent-randomization interface: dependent paths within one array or cache are not independent observations for cross-fitting. Learned correction introduces a separate bias-variance tradeoff [17]. Volumetric WoS, extended elliptic solvers, and transient diffusion broaden the PDE toolkit [7, 25, 37]; transferring cavity derivatives to those models requires the corresponding analysis.

Connectivity constraints. Andrade et al. [3] combine topological sensitivity with a spectral constraint on void connectivity. Their exterior-Dirichlet construction addresses whether void regions reach the exterior. Our strictly overlapping ball union guarantees an internally connected cavity, whose terminals remain inside the object. These are different feasible sets. The reproduced spectral diagnostic in [Appendix Y](#) clarifies the distinction; it is not a matched comparison of complete connectivity optimizers.

3 Problem and Topological Expansion

Design task. Given an outer solid, a heat source, boundary data, and geometric constraints, we select interior cavity locations to reduce a PDE-defined objective. The thermal experiments solve $-\Delta u = f$ with $u = 0$ on the exterior and cavity walls. Compliance $J = \int f u$ rewards cooling where heat is generated. A common fixed original-domain normalization is used when reporting its Monte Carlo objective averages, with zero contribution from removed material. For this objective, the adjoint source equals the state source; together with the self-adjoint operator, this gives $p = u$ and normalized insertion benefit $-D_T = u^2$. Self-adjointness of the operator alone does not imply $p = u$: the regional objective $J_q = \int q u / \int q$ instead uses a distinct adjoint driven by q , as evaluated in [Appendix B.2](#). The cavity walls are ideal heat sinks; cooling access and coolant transport are outside this model.

The analysis below treats the broader scalar, self-adjoint screened Poisson setting. Its regularity and oracle assumptions state the scope of the theorems; the 3D implementation specializes to constant conductivity and zero screening. Detailed proofs appear in [Appendix C](#).

3.1 State, adjoint, and reduced objective

Let $\Omega \subset \mathbb{R}^d$, $d \in \{2, 3\}$, be bounded and connected. Its fixed outer boundary is partitioned into a relatively open Dirichlet part Γ_D and a Neumann part Γ_N , with $\mathcal{H}^{d-1}(\Gamma_D) > 0$. Define

$$V := H_{\Gamma_D}^1(\Omega) := \{v \in H^1(\Omega) : \text{Tr } v = 0 \text{ on } \Gamma_D\} \quad (1)$$

and, for every measurable $D \subseteq \Omega$,

$$a_D(v, w) := \int_D \kappa \nabla v \cdot \nabla w + \sigma v w \, dx. \quad (2)$$

The state $u \in g + V$ is the unique solution of

$$a_\Omega(u, v) = \ell_\Omega(v) \quad \forall v \in V, \quad (3)$$

where g is a fixed Dirichlet lift and $\ell_\Omega \in V^*$ contains the volume source and prescribed Neumann data. In strong form this is $-\nabla \cdot (\kappa \nabla u) + \sigma u = f$.

Let $j : L^2(\Omega) \rightarrow \mathbb{R}$ be the state-dependent objective and let $\psi \in L^\infty(\Omega)$ be an optional explicit domain cost. We set

$$\mathcal{J}(\Omega) := j(u) + \int_{\Omega} \psi(x) \, dx. \quad (4)$$

The adjoint $p \in V$ is defined by

$$a_\Omega(v, p) = j'(u)[v] \quad \forall v \in V. \quad (5)$$

For example, if $j(u) = \frac{1}{2} \int_{\Omega} w(x)(u - u_\star)^2 \, dx$ with bounded $w \geq 0$, then the adjoint source is $w(u - u_\star)$.

3.2 A finite Dirichlet cavity

Let $B \subset \mathbb{R}^d$ be a fixed bounded inclusion containing the origin. For a candidate z and scale $\varepsilon > 0$, put

$$K_{\varepsilon,z} := z + \varepsilon B, \quad \Omega_{\varepsilon,z} := \Omega \setminus \overline{K_{\varepsilon,z}}. \quad (6)$$

On the new boundary we prescribe the constant value $u_{\varepsilon,z} = h_z$; in thermal design, h_z is the coolant temperature. The perforated state solves the same weak equation on $\Omega_{\varepsilon,z}$, with its test functions also vanishing on $\partial K_{\varepsilon,z}$. Extend it to all of Ω by the constant h_z inside the cavity and denote this H^1 extension by $\tilde{u}_{\varepsilon,z}$. The perturbed reduced objective is

$$\mathcal{J}(\Omega_{\varepsilon,z}) := j(\tilde{u}_{\varepsilon,z}) + \int_{\Omega_{\varepsilon,z}} \psi(x) \, dx. \quad (7)$$

The correct normalizer is the *relative operator capacity*. With $\tilde{v}$ denoting the quasi-continuous representative of $v \in V$, set

$$\rho_z(\varepsilon) := \inf \{ a_\Omega(v, v) : v \in V, \tilde{v} \geq 1 \text{ q.e. on } K_{\varepsilon,z} \}. \quad (8)$$

“Quasi-everywhere” means outside a set of zero H^1 capacity; it is the trace notion for which [Eq. \(8\)](#) is closed under weak convergence. The infimum is attained by a unique equilibrium potential $q_{\varepsilon,z}$. Its leading scale is

$$\rho_z(\varepsilon) = \begin{cases} \frac{2\pi\kappa(z)}{|\log \varepsilon|} (1 + o(1)), & d = 2, \\ \kappa(z) \text{cap}(B)\varepsilon (1 + o(1)), & d = 3, \end{cases} \quad (9)$$

uniformly on compact candidate sets. Here $\text{cap}(B)$ is the Newtonian capacity normalized by the energy $\int_{\mathbb{R}^3} |\nabla v|^2$.

Definition 3.1 (Capacity-normalized topological derivative). The Dirichlet-cavity topological derivative at z is

$$D_T(z) := \lim_{\varepsilon \downarrow 0} \frac{\mathcal{J}(\Omega_{\varepsilon,z}) - \mathcal{J}(\Omega)}{\rho_z(\varepsilon)}, \quad (10)$$

whenever the limit exists.

Assumption 3.2 (Analytic setting). Ω and B are $C^{2,\alpha}$ domains for some $\alpha \in (0, 1)$; $0 \in \text{int } B$; the mixed-boundary decomposition admits the Poincaré inequality, maximum principle, and Green kernel used in [Appendix C](#); $\kappa \in C^{1,\alpha}(\overline{\Omega})$ satisfies $0 < \kappa_- \leq \kappa \leq \kappa_+$; and $\sigma \in L^\infty(\Omega)$ is nonnegative. Candidates lie in a compact $Z \Subset \Omega$, and $u, p \in C^{1,\alpha}(U)$ on a fixed neighborhood $U \supset Z$. The values h_z are uniformly bounded on Z .

Assumption 3.3 (Objective smoothness). j is twice Fréchet differentiable on an $L^2(\Omega)$ neighborhood of u , with $\|j''(v)\|$ uniformly bounded there, and $\psi \in L^\infty(\Omega)$. The load restricted to a perforated domain is the restriction of the unperforated load, so state differences are a -harmonic away from the cavity.

Theorem 3.4 (Relative-capacity asymptotics). Under [Assumption 3.2](#), the minimizer in [Eq. \(8\)](#) exists, is unique, and satisfies [Eq. \(9\)](#). Moreover, $\rho_z(\varepsilon) > 0$ and $\varepsilon^d = o(\rho_z(\varepsilon))$ uniformly for $z \in Z$.

Theorem 3.5 (Quantitative local topological expansion). Under [Assumptions 3.2](#) and [3.3](#), there are constants $C > 0$ and $\varepsilon_0 > 0$, independent of $z \in Z$, such that for $0 < \varepsilon \leq \varepsilon_0$,

$$|\mathcal{J}(\Omega_{\varepsilon,z}) - \mathcal{J}(\Omega) - \rho_z(\varepsilon)D_T(z)| \leq C\rho_z(\varepsilon)^{1+\alpha_d}, \quad (11)$$

where $\alpha_2 = 1$ and $\alpha_3 = \frac{1}{2}$, and

$$D_T(z) = -(u(z) - h_z)p(z). \quad (12)$$

In particular, the limit in [Eq. \(10\)](#) exists uniformly on Z .

The proof in [Appendix C](#) constructs the equilibrium measure, proves its concentration on the shrinking cavity, obtains the full state corrector by an energy-minimizing lift, controls the nonlinear objective remainder in L^2 , and only then invokes the adjoint identity. Thus the sign in [Eq. \(12\)](#) follows from the convention in [Eq. \(5\)](#); changing that convention changes the displayed sign, not the algorithm.

Remark 3.6 (Why regular shape calculus is insufficient). For $d = 3$, the objective change is $O(\varepsilon)$ even though the new boundary area is $O(\varepsilon^2)$; in $d = 2$ it is $O(1/|\log \varepsilon|)$. Differentiating with respect to ε and then setting $\varepsilon = 0$ therefore does not recover a regular shape derivative. Capacity normalization is the essential singular step.

For connected outer and cavity boundaries, [Proposition C.1](#) makes this distinction explicit: a diffeomorphism preserves boundary connectedness, whereas an interior cavity creates a second boundary component. This elementary observation does not exclude recovering a topological derivative through a singular limit of shape derivatives on already perforated domains. Such connections are established prior art [\[6\]](#); the proposition makes no novelty claim about the distinction itself.

4 Topological Walk on Spheres

Equation [\(12\)](#) turns a domain-perturbation limit into two point queries on the original domain. This section constructs an estimator that remains correct under the correlated reuse patterns needed in a fast implementation.

4.1 Primal and adjoint random-walk oracles

At candidate z , let a WoS/WoSt implementation return a random pair

$$X_i(z) := (U_i(z), P_i(z)), \quad i = 1, \dots, N, \quad (13)$$

where pairs are independent and identically distributed across i , but U_i and P_i may share a seed, geometric queries, path prefixes, or a harmonic cache. We first analyze exact oracles,

$$\mathbb{E}[U_i(z)] = u(z), \quad \mathbb{E}[P_i(z)] = p(z), \quad (14)$$

and later place termination and cache error into an explicit bias budget. Screened Poisson killing, source sampling, Walk on Stars, and variance reduction can all live behind this interface.




Figure 2 A leave-one-out view of the cross-fitted estimator. (a) Independent randomizations produce pairs $X_i = (U_i, P_i)$; both components may reuse the same walk geometry. (b) For the example $i = 2$, retain $U_2 - h_z$ and average the other adjoints, excluding the potentially correlated P_2 . (c) Form $C_i = -(U_i - h_z)\bar{P}_{-i}$ and average over i . Each product has independent factors; the products themselves share samples. One adjoint sum evaluates every leave-one-out mean in $O(N)$ total work. This is exactly the estimator in Eq. (15); unbiasedness requires unbiased marginal oracles. Walks are schematic.

The independence unit is an entire randomization. For Array-RQMC or a randomly constructed cache shared across many trajectories, individual paths are not independent oracle pairs. Instead, generate $R \geq 2$ independent complete randomizations and use their field averages $X_r = (\bar{U}_r, \bar{P}_r)$ in Eq. (15), with $N = R$. Within a randomization, primal and adjoint may reuse every trajectory; neighbor queries may also share all cached paths. Different randomizations must use independent random streams and independently generated random caches. The finite-sample radius then uses $\lfloor R/2 \rfloor$, not the number of paths or cross-products. This interface preserves compatibility with Array-RQMC [14] and neighbor reuse [8] without importing an independence assumption they do not satisfy. The exact bias identities and a counterexample appear in Appendix J.

The source-integral identity used by our three-dimensional Poisson implementation is stated and proved in Proposition G.1. General objectives require a valid adjoint oracle as well as a state oracle; the product construction does not by itself solve the adjoint PDE or remove bias from a learned or approximate adjoint source. Compliance experiments exploit $p = u$; sensor-protection experiments solve the distinct adjoint driven by the observation weight.

A naive product from a same-seed pair is biased whenever U_i and P_i are correlated; its bias is $-\text{Cov}(U_i, P_i)$. Independently generating every primal and adjoint path avoids that covariance but forfeits useful reuse. We instead remove all same-index products.

4.2 Cross-fitted product estimator

Define

$$\widehat{D}_{TN}(z) := -\frac{1}{N(N-1)} \sum_{i \neq j} (U_i(z) - h_z) P_j(z) \quad (15)$$

$$= -\frac{(\sum_i (U_i - h_z))(\sum_i P_i) - \sum_i (U_i - h_z) P_i}{N(N-1)}. \quad (16)$$

The second form costs $O(N)$ rather than $O(N^2)$. It is a symmetric, order-two U-statistic [16] with kernel

$$H(X_i, X_j) := -\frac{1}{2} [(U_i - h_z) P_j + (U_j - h_z) P_i]. \quad (17)$$

The factors in every cross-index product are independent, so arbitrary dependence inside X_i is harmless. Theorems 5.1 and 5.2 formalize this observation. Equivalently, Fig. 2 keeps $U_i - h_z$, pairs it with the leave-

one-out mean $\bar{P}_{-i} = (\sum_j P_j - P_i)/(N - 1)$, and averages over i . Only the matching P_i , which may share randomness with U_i , is omitted from that term.

4.3 One set of paths, many candidate sites

The algebra above is pointwise, but the optimizer must evaluate thousands of z values. The estimator permits two complementary reuse modes; our main 3D design studies use pointwise scouting, and the separate accelerated-sampling study evaluates selected reuse mechanisms:

1. **Unbiased scout.** Boundary-value caching or stochastically truncated harmonic caching turns each completed walk into a local expansion that can be queried at multiple candidates without changing its expectation [22, 36]. We retain the independent cache index i required by Eq. (15).
2. **Approximate scout, bias-controlled validation.** A low-order deterministic cache or self-normalized probe estimator may be much faster but slightly biased [18]. It may rank sites, but it never certifies insertion. The selected sites are re-evaluated with fresh unbiased paths; alternatively, a proven cache bias b_{cache} is included in Eq. (39).

This separation is important. Reusing samples after adaptively choosing the most negative noisy location produces a winner’s-curse bias even when every pointwise estimator is unbiased. For a preselected site, fresh validation avoids reusing its selection noise. Choosing again among several freshly rescored sites requires simultaneous coverage, or another independent final check, to support a certificate.

4.4 Estimator interface

A certifying implementation must supply the tuple

$$(\widehat{D}_{TN}, r_N(\delta), b_{\text{term}}, b_{\text{cache}}), \quad (18)$$

where r_N is a statistical confidence radius and the two b terms upper-bound deterministic or algorithmic bias. The optimizer is agnostic to whether r_N comes from the bounded-sample result in Theorem 5.3, an empirical-Bernstein construction, or a robust median-of-means variant. It only requires a valid coverage statement.

5 Statistical Analysis

We now isolate the two statistical facts that make direct estimation preferable to finite-cavity differencing: its variance contains no vanishing geometric scale, and its uncertainty can be propagated into a valid topological decision. All probability statements are with respect to the random-walk and cache randomness at a fixed candidate z .

Theorem 5.1 (Unbiasedness under correlated pairs). *Suppose Eq. (14) holds, the pairs X_i are i.i.d., and the first moments exist. For every $N \geq 2$,*

$$\mathbb{E}[\widehat{D}_{TN}(z)] = D_T(z), \quad (19)$$

even when U_i and P_i are arbitrarily dependent within a pair.

Theorem 5.2 (Exact radius-independent variance). *Assume the oracle pairs are i.i.d. and $N \geq 2$. Let $A := U - h_z$, $B := P$, $\mu_A := \mathbb{E}A$, $\mu_B := \mathbb{E}B$, and $c_{AB} := \text{Cov}(A, B)$. If $\mathbb{E}A^2 < \infty$ and $\mathbb{E}B^2 < \infty$, then*

$$\text{Var}(\widehat{D}_{TN}) = \frac{4}{N}\zeta_1 + \frac{2}{N(N-1)}\zeta_2, \quad (20)$$

where

$$\zeta_1 = \frac{1}{4} (\mu_B^2 \text{Var } A + \mu_A^2 \text{Var } B + 2\mu_A \mu_B c_{AB}), \quad (21)$$

$$\zeta_2 = \frac{1}{2} (\text{Var } A \text{Var } B + c_{AB}^2). \quad (22)$$

Consequently $\text{Var}(\widehat{D}_{TN}) = O(N^{-1})$. No term depends on an insertion radius ε .

We specialize the classical Hoeffding decomposition [16]; the complete algebra appears in [Appendix D](#). Correlation inside a pair can increase or decrease the first-order variance through c_{AB} , but cannot create covariance bias because the diagonal products are absent.

Theorem 5.3 (Finite-sample confidence radius). *Assume the unbiased i.i.d. oracle-pair hypotheses of [Theorem 5.1](#), $N \geq 2$, and $|U_i - h_z| \leq M_U$ and $|P_i| \leq M_P$ almost surely, and put $M = M_U M_P > 0$. Then, for every $t > 0$,*

$$\mathbb{P}(|\widehat{D}_{TN} - D_T| \geq t) \leq 2 \exp\left(-\frac{\lfloor N/2 \rfloor t^2}{2M^2}\right). \quad (23)$$

Equivalently, for every $\delta \in (0, 1)$,

$$\mathbb{P}\left[|\widehat{D}_{TN} - D_T| \leq M \sqrt{\frac{2 \log(2/\delta)}{\lfloor N/2 \rfloor}}\right] \geq 1 - \delta. \quad (24)$$

If $M = 0$, the estimator and its target are identically zero, so coverage is deterministic. For K candidates evaluated with arbitrary shared randomness, replacing δ by δ/K gives simultaneous coverage.

The proof averages over all pairings of the observations and applies Hoeffding's lemma before a Chernoff optimization; it does not treat the $N(N-1)$ cross-products as independent. The displayed bound is deliberately conservative and assumption-transparent. Experiments may additionally report empirical-Bernstein or bootstrap intervals, but only an interval with verified coverage may drive certification.

5.1 Why finite-cavity differencing becomes noisy

Consider a coupled per-path objective difference $Y_\varepsilon = Q_\varepsilon - Q_0$ and independent replicates $Y_\varepsilon^{(1)}, \dots, Y_\varepsilon^{(N)}$. The natural finite-cavity estimator is

$$\widehat{D}_{TN,\varepsilon}^{\text{FD}} := \frac{1}{N\rho_z(\varepsilon)} \sum_{i=1}^N Y_\varepsilon^{(i)}. \quad (25)$$

Common random numbers make Y_ε zero for paths whose outcome is unaffected by the cavity. This is useful, but it exposes a rare event.

Assumption 5.4 (Rare-event coupling). There is an event A_ε and constants $c_p, C_p, c_2, C_1 > 0$, independent of sufficiently small ε , such that

$$c_p \rho_z(\varepsilon) \leq \mathbb{P}(A_\varepsilon) \leq C_p \rho_z(\varepsilon), \quad (26)$$

$$Y_\varepsilon = 0 \quad \text{on } A_\varepsilon^c, \quad (27)$$

$$\mathbb{E}[Y_\varepsilon^2 \mid A_\varepsilon] \geq c_2, \quad |\mathbb{E}Y_\varepsilon| \leq C_1 \rho_z(\varepsilon). \quad (28)$$

This captures couplings in which a path changes by $O(1)$ only when it reaches the small cavity's influence set, whose harmonic measure is $\Theta(\rho_z(\varepsilon))$. The stopped-path harmonic-correction baseline in [Appendix V](#)

explicitly uses the cavity-hit event. A generic shared-stream WoS subtraction need not vanish whenever the cavity is missed, because its walk spheres can already differ. The rare-event hypotheses must therefore be checked for the particular coupling; they do not follow from shared random numbers alone.

Theorem 5.5 (Rare-event variance barrier). *Under [Assumption 5.4](#), for every ε satisfying $\rho_z(\varepsilon) \leq c_p c_2 / (2C_1^2)$,*

$$\text{Var}[\widehat{D}_{T_{N,\varepsilon}}^{\text{FD}}] \geq \frac{c_p c_2}{2N \rho_z(\varepsilon)}. \quad (29)$$

Consequently, a variance bound independent of ε requires $N = \Omega(1/\rho_z(\varepsilon))$ paths.

This is not a universal lower bound for an importance sampler that analytically conditions on hitting the cavity. It is a precise diagnosis of the natural “insert, couple, subtract, divide” baseline.

Corollary 5.6 (Candidate-map sample complexity). *Under [Theorem 5.3](#), simultaneous error at most Δ for K candidates with probability at least $1 - \delta$ is guaranteed whenever*

$$\left\lfloor \frac{N}{2} \right\rfloor \geq \frac{2M^2}{\Delta^2} \log \frac{2K}{\delta}. \quad (30)$$

Thus $N = O(M^2 \Delta^{-2} \log(K/\delta))$, independently of ε . Under [Assumption 5.4](#), forcing the finite-difference variance below Δ^2 requires

$$N \geq \frac{c_p c_2}{2\Delta^2 \rho_z(\varepsilon)}. \quad (31)$$

5.2 Bias bookkeeping

The cross-fit removes correlation bias, but it cannot remove bias already present in either marginal PDE oracle. The next result converts marginal bounds into a topological-derivative bias budget.

Proposition 5.7 (Propagation of marginal oracle bias). *Suppose the pairs remain i.i.d. across indices but*

$$|\mathbb{E}U - u(z)| \leq \beta_U, \quad |\mathbb{E}P - p(z)| \leq \beta_P. \quad (32)$$

Then

$$|\mathbb{E}\widehat{D}_{TN} - D_T| \leq |p(z)|\beta_U + |u(z) - h_z|\beta_P + \beta_U\beta_P. \quad (33)$$

For unbiased Russian-roulette termination, the corresponding marginal bias is zero. An η -shell termination rule or a truncated cache generally has nonzero β_U, β_P . We require a conservative allocation of the right side of [Eq. \(33\)](#) into

$$|\mathbb{E}\widehat{D}_{TN} - D_T| \leq b_{\text{term}} + b_{\text{cache}}. \quad (34)$$

The bounds may come from analysis or certified a posteriori residuals. A calibration-based bound must itself control reference error and allocate confidence to the calibration; a high-accuracy reference or termination-tolerance sweep alone is insufficient. A statistical interval that omits these terms does not support the general-objective guarantee in [Appendix A](#). The specialized one-sided cooling validator instead uses the proved sign of its truncation bias ([Proposition N.4](#)).

6 Turning Scores into Geometry

Sequential design. We generate feasible candidates, scout with inexpensive walks, and rescore a shortlist with fresh samples. A selected cavity is inserted, the current domain is updated, and the process repeats. Existing interfaces can then be refined using Differential WoS. Independent rescoring reduces selection noise,




Figure 3 *From an insertion score to a finite geometric decision.* Score the current domain, freeze a center and radius, and validate with fresh samples before creating the cavity. The general-objective guard combines statistical uncertainty, oracle bias, and a finite-radius remainder bound. Its covered upper bound must be negative. A failed check leads to a valid resampling schedule or no insertion. The dashed cyan boundary shows optional subsequent shape refinement. This conditional certifying policy requires usable error bounds; the practical ranking studies and the cooling-specific finite-benefit test have separate scopes.

but selecting among several fresh estimates still requires simultaneous coverage or another final check to support a certificate. Our ranking and refinement studies evaluate a practical heuristic; they are separate from the guarded validation experiment below.

General-objective acceptance. Suppose a known bound satisfies

$$|\mathcal{J}(\Omega_{\varepsilon, z}) - \mathcal{J}(\Omega) - \rho_z(\varepsilon)D_T(z)| \leq C_z \rho_z(\varepsilon)^{1+\alpha}.$$

At a preselected site, a valid statistical radius $r_N(\delta)$ and bias budget $b = b_{\text{term}} + b_{\text{cache}}$ give the sufficient test

$$\widehat{D}_{T_N} + r_N(\delta) + b + C_z \rho_z(\varepsilon)^\alpha < 0. \quad (35)$$

The probability of accepting a non-improving insertion is then at most δ (Theorem A.2). Figure 3 shows this conditional policy. A radius chosen after validation is covered by the same event when the deterministic bound holds uniformly over allowed radii. Existence of an asymptotic constant is insufficient: the test needs a numerical upper bound. Across adaptive rounds, allocating summable conditional failure probabilities controls the familywise error (Proposition A.7). Repeatedly inspecting a fixed-sample interval without such a schedule is not covered.

A direct finite-benefit route for cooling. For $-\Delta u = f$ with bounded $0 \leq f \leq F$ with zero Dirichlet walls in a bounded Lipschitz three-dimensional domain D , let $J(D) = \int_D f u$. A contained sphere $\bar{B}_r(z) \Subset D$ has benefit

$$J(D) - J(D \setminus \bar{B}_r(z)) \geq 4\pi r (L_u - I_r^+)^2, \quad (36)$$

on an event where $L_u \leq u(z)$, provided

$$I_r^+ \geq \int_{B_r(z)} G_{B_r(z)}(z, y) f(y) dy.$$

One may use $I_r^+ = f_+ r^2/6$ for a valid local upper source bound. The difference-energy identity and the ball equilibrium potential bound the benefit by the squared average cavity trace; the center Green identity recovers that trace (Theorem A.6). A clipped, zero-terminal nonnegative-source WoS estimate enables a one-sided lower confidence bound. The complete proof, sample range, and geometric guards are in Appendix O. The guarantee remains conditional on the stated sampling laws and geometry assumptions.

Finite shape and access constraints. The sphere studies prescribe a count and radius, so all methods remove the same volume. The channel study calibrates a common radius of an overlapping ball union to a

shared material budget. Strict overlap gives a positive-width passage, while protected supports preserve the heat input. Point scoring of a channel axis remains a finite-placement heuristic, not a derived whole-channel topological derivative. Neither internal connectedness nor an exterior opening implies a coolant-flow model. Detailed geometry constraints, radius policies, and their proofs appear in the appendix.

7 Evaluation

We test the estimator, the resulting finite designs, and the additional cost of validation separately. All thermal studies use ideal zero-temperature exterior and cavity walls. The matched design studies minimize a normalized regional temperature with a distinct adjoint; the query study additionally tests distributed-heating compliance, for which $p = u$. Detailed protocols, all comparisons, and step-by-step derivations are in the appendix. The initial FEM experiment is an independent field check, not an optimized FEM design-time baseline.

Actual-geometry estimator comparison. Eight fixed feasible sites on each of three meshes are evaluated using direct TWoS, shared-stream finite-cavity subtraction, a stronger stopped-path harmonic correction, and multigrid with longest-axis grid counts 49, 65, 97, and 129. The stochastic methods use sixteen repetitions at 256 and 2,048 paths per site. A separate sixteen-batch, 8,192-path reference estimates the infinitesimal coefficient; it is not exact ground truth. Trial radii are 0.08, 0.04, 0.02, and 0.01. The stronger finite-hole baseline estimates the finite correction from a perforated-domain exit walk and a fresh original-domain state walk only when the cavity is hit. Its exact harmonic identity and proof appear in [Appendix V](#).

[Figure 4](#) separates empirical variance from error to the infinitesimal reference and from query time. At radius 0.01 and 2,048 paths, the harmonic-correction variance is $35.7\times$ direct TWoS on the housing and $13.8\times$ on the branching manifold. These are measured ratios, not confidence bounds or fitted asymptotic laws. On the annular object, all recorded harmonic corrections vanish at radii 0.01 and 0.02; the variance is unresolved at this budget. We suppress the resulting collapsed bootstrap intervals rather than interpret them as zero uncertainty. At larger radii the stronger correction can have competitive measured variance. Multigrid also remains a strong error–cost baseline. Finite-radius and discretization bias are distinct from the uncertainty shown by the sampling bands. At 2,048 paths, direct relative RMS errors are 14.9%, 25.2%, and 146%; the reference’s relative standard-error norms are 1.6%, 4.8%, and 16.2%. The annular sparse-query case is poorly resolved at these budgets, a limitation retained in the comparison.

Matched insertion and refinement. A finned housing, branching manifold, and ribbed annular spreader provide different conduction paths. Each method inserts six disjoint radius-0.065 spheres from paired candidate pools, subject to identical source, sensor, boundary, and pairwise clearance constraints. Every initialization then uses the same specialized Differential WoS center-refinement component. TWoS, multigrid at grid counts 49 and 65, and nearest-sensor placement receive eight rounds; hottest-point and random starts receive sixteen. Five paired design seeds are independently evaluated at saved checkpoints. This isolates the initialization inside a shared optimizer; it does not reproduce the original Differential WoS authors’ complete application.

Mean sensor cooling after TWoS initialization and eight rounds is 65.0%, 16.4%, and 11.4% across the three objects. The primary multigrid selector gives 65.1%, 16.3%, and 11.7%. Hottest-point starts reach 42.7%, 0.14%, and 0.01% despite twice as many rounds. However, nearest-sensor placement reaches 63.6%, 17.3%, and 12.2%. On the housing, complete mean times are 11.77 s for TWoS, 8.66 s for primary multigrid, 7.99 s for coarse multigrid, and 7.08 s for nearest-sensor placement. Thus the large benefit over hottest-point placement does not establish a quality or speed advantage over all informed controls. The cutaways show the stronger controls; the curves retain all six methods. Paired differences and their pointwise intervals are reported in [Appendix Q](#); they do not establish superiority over multigrid or equivalence to it.

Uncertainty and independent checks. The study contains 180 initial designs, 90 refinement runs, and 390 evaluated states. Curves show means, individual runs, and pointwise 95% bootstrap intervals jointly resampling five paired design seeds and four shared evaluation blocks. Tables additionally report the across-seed standard deviation. These small-sample intervals exclude numerical bias. Every saved state passes the geometric constraints, no recorded walk reaches its cap, and the maximum multigrid relative residual is below 10^{-7} . Independent grid checks use 97 and 129 nodes on the longest axis; flagged housing cases




Figure 4 Actual-geometry estimator comparisons. Each row uses eight fixed, protected query sites on one mesh. The center plots normalize summed site variances by the squared norm of the high-budget reference and use 2,048 paths per site. Bands give pointwise 95% bootstrap intervals from sixteen independent repeat vectors and sixteen independent reference batches. Direct TWoS has no trial-radius input; its horizontal band is repeated for comparison, not independently remeasured at every radius. Right: relative RMS error and wall time at both Monte Carlo budgets and all finite radii, with four deterministic grid resolutions labeled. Bars reflect repetition and reference-sampling uncertainty; deterministic grid error bars reflect only reference sampling, while their runtime bars use five repeated solves. Finite-radius, discretization, and termination biases are not enclosed. The reference is an independent, high-budget TWoS estimate, not exact ground truth. All-zero correction settings are marked unresolved and use open error-cost markers; their collapsed bootstrap intervals are suppressed.




Figure 5 *Insertion decisions affect the refined design.* Each row fixes the physical domain, protected sensor, heat input, cavity count, and exact removed volume. Views show the first predeclared seed after refinement; cyan and mint outlines mark the sensor and complete source supports. The section color shows local cooling relative to the unperforated temperature, from a common multigrid rendering field. It is not the objective’s confidence interval. Labels and curves summarize five design seeds; faint dots show individual runs. Bands are pointwise 95% bootstrap intervals for mean sensor cooling, jointly resampling seeds and independent evaluation blocks; horizontal bars resample complete runtime across seeds. Plot ranges adapt to each geometry; units and the local-cooling color scale are shared. They do not cover numerical bias. Hot and random receive up to sixteen refinement rounds; TWoS, both multigrid selectors, and nearest-sensor placement receive eight. Ghost shells show material removed only for display; some cavities lie outside the section. Independent evaluation, rendering, and mesh generation are excluded from design time.

receive grids 193 and 257. The finest housing references still differ from Monte Carlo reductions by up to 2.38 percentage points across the checked first-seed designs, so small quality differences warrant caution. The numerical audit does not certify reference accuracy.

Accessible channels at fixed material and heat input. Six open ball-union channels share a removed-volume target of 0.0168. Every heat-source support is preserved, so the exact total heat input is unchanged. Independent randomized volume audits deviate by at most 0.97% from the target; this is an observed audit error, not a certified bound. TWoS gives mean regional cooling of 37.8%, 2.29%, and 1.75%, versus 42.1%, 2.59%, and 1.93% for primary multigrid. Nearest-sensor placement gives 39.4%, 2.62%, and 2.02%. This finite-channel heuristic is useful but does not dominate the informed controls. Cutaways, volume audits, and full cost curves in the appendix report the complete comparison.

Sparse queries at matched accuracy. Four meshes, two domain states, and nested batches of 8–512 fixed sites compare TWoS against grids with 49–129 longest-axis nodes, using five paired calibration seeds. Multigrid reuses solved fields; a stronger variant retains outer masks and warm-starts after insertion. Compliance requires one solve. Common mesh setup is measured separately. Our 20% score-error target requires an eligible upper bootstrap mean-error endpoint, MG193/MG257 refinement, and independent Monte Carlo reference checks; it is not a continuum certificate. A disclosed follow-up analytically integrates the prescribed sources using classical Rao–Blackwellization ([Appendix S](#)). This reduces variance at fixed path count in ideal arithmetic, without removing termination bias. [Figure 6](#) compares objectives on the same Bunny; the full protocol, ablations, reference gates, and stream correction are in [Appendix R](#). All four calibration-selected apparent wins fail at least one confirmation rule on 20 fresh paired seeds: three compliance cases exceed the accuracy upper-endpoint threshold, while housing does not confirm a timing gain. A separately declared follow-up fixes 256 walks at eight sites in both compliance states, using 20 further fresh seeds. On the original Bunny, TWoS takes 89.9 ms versus 113.9 ms for MG49: a $1.27\times$ ratio with pointwise 95% bootstrap interval [1.05, 1.50]. Mean score error is 9.0% (upper endpoint 10.2%), versus 12.3% for MG49. After insertion, TWoS takes 83.1 ms versus 72.8 ms for warm-started MG; the speed-ratio interval [0.69, 1.11] establishes no advantage. This is a source-specific, eight-query result with geometry already loaded, not a complete optimization speedup or universal state-of-the-art claim.

An exact finite-radius check. The uniformly heated ball with a concentric cooled cavity admits exact states, compliance, and a positive certificate gap ([Appendix P](#)). Rational checks at 999 radii verify the algebra, not sampler coverage.

Validation has a separate cost. A guarded six-attempt compliance diagnostic on each of the Bunny and housing uses a total failure budget of 0.01. All twelve targets pass, but validation takes 150.39 s and 52.17 s versus 2.97 s and 0.53 s for scouting and rescoring. Full bounds and independent checks are in the appendix. This additional cost is excluded from ranking-only speed claims.

8 Discussion and Limitations

The principal advantage is the direct evaluation of an insertion coefficient using point queries on the current geometry. It is distinct from the value of an adjoint: deterministic topological derivatives use the same objective information, and sensor proximity is an effective alternative on our localized tasks. The complete pipeline results do not establish a universal quality or speed advantage. A global field can be preferable for many queries, and the Bunny crossover should not be transferred to another geometry without measurement.

Choosing a solver. The practical question is whether the requested output is a small set of insertion scores or a dense reusable field. Direct queries avoid trial-cavity construction and volumetric discretization, but their work grows with the number of queries and the required sampling accuracy. A global solve can amortize its setup over many sites; caching and multigrid further change that tradeoff. Our matched-target study varies the objective on the same Bunny and retains failed fresh-seed confirmations. Localized sensor objectives can favor coarse global solves even at eight sites. Analytic source integration helps only within its stated source family; the sampling and reference criteria must be satisfied before interpreting a timing advantage. An interactive or repeatedly edited workflow remains a potential application, not a measured speed claim of this paper.

Our objects exercise detailed boundaries, narrow gaps, branching paths, and curved transport. They are




Figure 6 Same Bunny, different query workloads. Left: measured accuracy–cost fronts at eight sites; labels give paths per site or longest-axis grid nodes. Top right: fastest tested budgets passing the 20% rule on five calibration seeds. Bottom right: the separate fixed-budget confirmation on 20 fresh seeds; both methods pass the 20% error rule, but only the original-domain time ratio confirms an advantage. Bands and bars are pointwise 95% bootstrap intervals; reference bias is excluded. The analytic-source variant is restricted to the prescribed sources. Grid curves use the faster eligible reuse implementation. Calibration fronts alone do not establish a speed claim.

procedural conduction benchmarks, not a representative sample of industrial products. Zero-temperature cavities are ideal sinks, not insulating air pockets; open channels do not solve fluid flow or conjugate heat transfer. Fixed material and power budgets remove specific confounds without supplying a manufacturing model. More nested cold removal lowers compliance in this model, so costs and feature constraints are essential to a meaningful cavity-count tradeoff.

The asymptotic theory assumes smooth domains and separated cavities; triangle edges and ball-union seams do not literally meet its global smoothness hypothesis. The direct compliance bound has a separate Lipschitz-domain scope. Insulating cavities, elasticity, nonlinear PDEs, and transient problems require their own expansions and valid oracles. Cross-fitting removes correlation bias, not marginal solver bias. Bootstrap bands describe finite-sample uncertainty; they do not bound discretization, termination, finite-radius, or geometry error. The guarded validator’s sampling laws still idealize pseudorandom draws and their floating-point propagation. Its containment checks do not constitute a verified finite-precision sampler.

9 Conclusion

TWoS turns classical cavity asymptotics into grid-free insertion queries. Our analysis separates sampling variance from finite-decision error control. Matched comparisons establish a narrow sparse-query benefit and substantial limitations; cheaper validation and broader physical models remain open.

References

- [1] David R. Adams and Lars Inge Hedberg. *Function Spaces and Potential Theory*, volume 314 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, 1996. doi: 10.1007/978-3-662-03282-4.

- [2] Grégoire Allaire, François Jouve, and Anca-Maria Toader. Structural optimization using sensitivity analysis and a level-set method. *Journal of Computational Physics*, 194(1):363–393, 2004. doi: 10.1016/j.jcp.2003.09.032.
- [3] Giovanna C. Andrade, Alberto Donoso, David Ruiz, and Alex Ferrer. On a level set and topological derivative-based strategy for connectivity constraints in topology optimization. *Structural and Multidisciplinary Optimization*, 2026. doi: 10.1007/s00158-026-04327-5.
- [4] Phillip Baumann, Peter Gangl, and Kevin Sturm. A complete derivation of the topological derivative of a semilinear elliptic transmission problem in space dimensions two and three. arXiv preprint arXiv:2111.08418, 2022.
- [5] Martin Burger, Benjamin Hackl, and Wolfgang Ring. Incorporating topological derivatives into level set methods. *Journal of Computational Physics*, 194(1):344–362, 2004. doi: 10.1016/j.jcp.2003.09.033.
- [6] Jean C  a, St  phane Garreau, Philippe Guillaume, and Mohamed Masmoudi. The shape and topological optimizations connection. *Computer Methods in Applied Mechanics and Engineering*, 188(4):713–726, 2000. doi: 10.1016/S0045-7825(99)00357-6.
- [7] Iulian C  mpean, Andreea Grecu, and Arghir Zarnescu. Extended walk-on-spheres algorithm for linear and nonlinear elliptic problems of divergence-type. arXiv preprint arXiv:2608.07017, 2026.
- [8] Michael T. Czekanski, Benjamin J. Faber, Margaret E. Fairborn, Adelle M. Wright, and David S. Bindel. Walking on spheres and talking to neighbors: Variance reduction for laplace’s equation. arXiv preprint arXiv:2404.17692, 2026. Version 3, 2 July 2026.
- [9] Lawrence C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, Rhode Island, 2 edition, 2010. doi: 10.1090/gsm/019.
- [10] M. Freiburger, A. Laurain, M. Hinterm  ller, A. K  stinger, and H. Scharfetter. Using the topological derivative for initializing a markov-chain monte carlo reconstruction in fluorescence tomography. In *Diffuse Optical Imaging III*, volume 8088 of *Proceedings of SPIE-OSA Biomedical Optics*, 2011. doi: 10.1117/12.889609.
- [11] Peter Gangl and Kevin Sturm. Automated computation of topological derivatives with application to nonlinear elasticity and reaction–diffusion problems. *Computer Methods in Applied Mechanics and Engineering*, 398, 2022. doi: 10.1016/j.cma.2022.115288.
- [12] David Gilbarg and Neil S. Trudinger. *Elliptic Partial Differential Equations of Second Order*. Classics in Mathematics. Springer, Berlin, 2001. doi: 10.1007/978-3-642-61798-0.
- [13] Tom Gustafsson and G. D. McBain. scikit-fem: A python package for finite element assembly. *Journal of Open Source Software*, 5(52):2369, 2020. doi: 10.21105/joss.02369.
- [14] Val  rie N. P. Ho and Art B. Owen. Walk on spheres and Array-RQMC. arXiv preprint arXiv:2605.12844, 2026. Version 3, 8 July 2026.
- [15] Valerie N. P. Ho and Art B. Owen. Randomized quasi-monte carlo for walk on spheres. arXiv preprint arXiv:2605.08483, 2026.
- [16] Wassily Hoeffding. A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325, 1948. doi: 10.1214/aoms/1177730196.
- [17] Ethan Hsu, Ivan Ge, and Hong Meng Yam. MC²: Monte carlo correction for fast elliptic PDE solving. arXiv preprint arXiv:2605.09288, 2026.
- [18] Wanchao Huang, Yutian Zhu, Qing Fang, and Ligang Liu. Probe-based walk on spheres for efficient path reusing. *ACM Transactions on Graphics*, 2026. URL <https://t7imal.github.io/projects/2026wop/>. Proceedings of ACM SIGGRAPH 2026.

- [19] Junyoung Hur, Pilseong Kang, and Sung-Kie Youn. Topology optimization based on spline-based meshfree method using topological derivatives. *Journal of Mechanical Science and Technology*, 31:2423–2431, 2017. doi: 10.1007/s12206-017-0440-6.
- [20] Jiajie Li, Yixin Tan, and Shengfeng Zhu. Superconvergence analysis for patch-recovered finite element approximations of topological derivatives. *IMA Journal of Numerical Analysis*, 2026. doi: 10.1093/imanum/draf121.
- [21] Ishit Mehta, Manmohan Chandraker, and Ravi Ramamoorthi. A theory of topological derivatives for inverse rendering of geometry. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [22] Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. Boundary value caching for walk on spheres. *ACM Transactions on Graphics*, 42(4), 2023. doi: 10.1145/3592400.
- [23] Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. Differential walk on spheres. *ACM Transactions on Graphics*, 43(6), 2024. doi: 10.1145/3687913.
- [24] Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. Walkin’ robin: Walk on stars with robin boundary conditions. *ACM Transactions on Graphics*, 43(4):1–18, 2024. doi: 10.1145/3658153. Proceedings of ACM SIGGRAPH 2024.
- [25] Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. Solving partial differential equations in participating media. *ACM Transactions on Graphics*, 44(4), 2025. doi: 10.1145/3731152.
- [26] Mervin E. Muller. Some continuous monte carlo methods for the dirichlet problem. *The Annals of Mathematical Statistics*, 27(3):569–589, 1956. doi: 10.1214/aoms/1177728169.
- [27] Antônio André Novotny and Jan Sokołowski. *Topological Derivatives in Shape Optimization*. Springer, 2013. doi: 10.1007/978-3-642-35245-4.
- [28] Hossam Saeed, Michael Lambiri, Teseo Schneider, Derek Nowrouzezahrai, and Paul G. Kry. Fast cutting of cage based deformation with walk on spheres. In *Proceedings of the SIGGRAPH 2026 Conference Papers*, pages 1–10, New York, NY, USA, 2026. Association for Computing Machinery. doi: 10.1145/3799902.3811051.
- [29] Rohan Sawhney, Bailey Miller, Ioannis Gkioulekas, and Keenan Crane. Walk on stars: A grid-free monte carlo method for PDEs with neumann boundary conditions. *ACM Transactions on Graphics*, 42(4):1–20, 2023. doi: 10.1145/3592398. Proceedings of ACM SIGGRAPH 2023.
- [30] Luka Schlegel, Volker Schulz, Frank T. Seifried, and Maximilian Würschmidt. A probabilistic approach to shape derivatives. *Numerische Mathematik*, 2026. doi: 10.1007/s00211-026-01554-2. Published 24 August 2026.
- [31] Stephan Schmidt and Maximilian Würschmidt. A first step towards mesh-free probabilistic shape optimization. arXiv preprint arXiv:2603.01141, 2026.
- [32] Stanford University Computer Graphics Laboratory. The stanford bunny. Stanford 3D Scanning Repository, 1994. URL <https://graphics.stanford.edu/data/3Dscanrep/>. Zippered reconstruction; accessed September 2026.
- [33] Guilherme Henrique Teixeira, Nepomuk Krenn, Peter Gangl, and Benjamin Marussig. Isogeometric topology optimization based on topological derivatives. *Computer Assisted Methods in Engineering and Science*, 2026. doi: 10.24423/comes.2026.1874. URL <https://comes.ippt.pan.pl/index.php/comes/article/view/1874>. Online first.
- [34] Zihan Yu, Rohan Sawhney, Bailey Miller, Lifan Wu, and Shuang Zhao. Robust derivative estimation with walk on stars. *ACM Transactions on Graphics*, 44(6), 2025. doi: 10.1145/3763333.

- [35] Jonas Zehnder, Yue Li, Stelian Coros, and Bernhard Thomaszewski. NTopo: Mesh-free topology optimization using implicit neural representations. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [36] Zihong Zhou, Eugene d'Eon, Rohan Sawhney, and Wojciech Jarosz. Harmonic caching for walk on spheres. *ACM Transactions on Graphics*, 44(6):1–15, 2025. doi: 10.1145/3763322.
- [37] Zihong Zhou, Rohan Sawhney, Eugene d'Eon, and Wojciech Jarosz. Grid-free monte carlo for time-dependent diffusion. arXiv preprint arXiv:2609.12306, 2026. Submitted 11 September 2026.

A From Scores to Finite Geometry

For a general objective, an infinitesimal derivative does not by itself guarantee that a finite cavity improves the objective. We first combine a uniform finite-size remainder with statistical and algorithmic uncertainty, then give a direct quantitative certificate for the cooling model.

Assumption A.1 (Usable finite-radius remainder). For every admissible candidate z and $0 < \varepsilon \leq \varepsilon_0$, a known or conservatively bounded $C_z > 0$ and $\alpha > 0$ satisfy

$$|\mathcal{J}(\Omega_{\varepsilon,z}) - \mathcal{J}(\Omega) - \rho_z(\varepsilon)D_T(z)| \leq C_z \rho_z(\varepsilon)^{1+\alpha}. \quad (37)$$

For the model analyzed in [Section 3](#), [Theorem 3.5](#) establishes the existence of such a remainder bound with $\alpha = 1$ in two dimensions and $\alpha = 1/2$ in three dimensions. Certification additionally requires a numerical upper bound on C_z , rather than only its existence.

Theorem A.2 (Certified finite insertion). Suppose [Assumption A.1](#) and [Eq. \(34\)](#) hold and, with probability at least $1 - \delta$,

$$|\widehat{D}_{TN} - \mathbb{E}\widehat{D}_{TN}| \leq r_N(\delta). \quad (38)$$

If

$$\boxed{\widehat{D}_{TN}(z) + r_N(\delta) + b_{\text{term}} + b_{\text{cache}} + C_z \rho_z(\varepsilon)^\alpha < 0}, \quad (39)$$

the probability of accepting a non-improving insertion is at most δ . Equivalently, on an event of probability at least $1 - \delta$, the test implies $\mathcal{J}(\Omega_{\varepsilon,z}) < \mathcal{J}(\Omega)$. This is an error-control statement, not a posterior probability conditional on acceptance.

Corollary A.3 (Data-dependent radius selection). Assume [Eq. \(37\)](#) holds simultaneously for all $\varepsilon \in (0, \varepsilon_0]$. On the basis of the validation batch, let $\widehat{\varepsilon} \in (0, \varepsilon_0]$ be any measurable radius satisfying [Eq. \(39\)](#). Then, on an event of probability at least $1 - \delta$, acceptance implies $\mathcal{J}(\Omega_{\widehat{\varepsilon},z}) < \mathcal{J}(\Omega)$. Equivalently, the probability of accepting a non-improving radius is at most δ . No union bound over radii is needed.

The certificate separates sampling error, termination bias, cache bias, and the asymptotic remainder. If it fails, one can sample more, improve the PDE oracle, reduce ε , or reject the site.

A.1 Certifying a useful amount of improvement

For ideal zero-Dirichlet cooling, adding a feasible cavity already lowers compliance by monotonicity ([Proposition M.1](#)). The useful certification target is therefore a *minimum finite benefit*, optionally large enough to pay an explicit insertion cost.

Theorem A.4 (Minimum benefit from a topological score). *Under the hypotheses of [Theorem A.2](#), set $b = b_{\text{term}} + b_{\text{cache}}$. With probability at least $1 - \delta$, the actual benefit $B = \mathcal{J}(\Omega) - \mathcal{J}(\Omega_{\varepsilon,z})$ satisfies*

$$\begin{aligned} B &\geq \underline{B}_{\text{TD}}(z, \varepsilon), \\ \underline{B}_{\text{TD}} &= -\rho_z(\varepsilon) [\widehat{D}_{TN} + r_N(\delta) + b] - C_z \rho_z(\varepsilon)^{1+\alpha}. \end{aligned} \quad (40)$$

The event is simultaneous over radii covered by the uniform remainder. Accepting only when $\underline{B}_{\text{TD}} \geq \tau$ controls the probability of accepting an insertion with benefit below τ by δ .

This route still requires numerical bounds on capacity, bias, and remainder; [Appendix N](#) handles certified capacity intervals explicitly. For three-dimensional Poisson compliance, a variational argument supplies a second route with explicit finite-radius constants.

Theorem A.5 (Direct finite-hole cooling benefit). *Let $D \subset \mathbb{R}^3$ be bounded and Lipschitz, let $0 \leq f \leq F$ be bounded, and let $-\Delta u = f$ on D with zero Dirichlet data. Define $J(D) = \int_D f u$. Suppose $\overline{B}_R(z) \Subset D$ and known local bounds $0 \leq f_- \leq f \leq f_+$ hold on $B_R(z)$. If $\mathbb{P}\{L_u \leq u(z)\} \geq 1 - \delta$, then simultaneously for all $0 < r < R$, with probability at least $1 - \delta$,*

$$\begin{aligned} J(D) - J(D \setminus \overline{B}_r(z)) &\geq \underline{B}_{\text{cool}}(z, r) := 4\pi r m^2, \\ m &= \frac{f_-(R^2 - r^2)}{6} \\ &\quad + \frac{1 - r/R}{(1 + r/R)^2} \left(L_u - \frac{f_+ R^2}{6} \right)_+. \end{aligned} \quad (41)$$

Here the new cavity also has zero Dirichlet data. Accepting only if this lower bound exceeds a chosen target gives the same false-acceptance control as [Theorem A.4](#).

The proof combines the exact compliance difference-energy identity with ball capacity and a Poisson-kernel comparison; it needs no unknown asymptotic remainder. A finite, zero-terminal WoS walk with nonnegative source samples has expectation at most $u(z)$, enabling a one-sided confidence bound with a fully explicit sample range ([Proposition N.4](#)). Setting $L_u = 0$ instead yields a deterministic source-only bound. These conservative bounds can fail to reach the requested target even when an insertion is beneficial. [Appendix N](#) proves all steps, explicit cost and relative-benefit rules, and adaptive selection control. The manufactured check in [Appendix N.5](#) tests the formulas; the earlier Bunny and housing runs evaluate the ranking implementation. A separate application validation study is reported in [Appendix B.9](#).

Theorem A.6 (Average-trace finite-hole certificate). *Under the cooling hypotheses of [Theorem A.5](#), suppose a known I_r^+ satisfies*

$$I_r^+ \geq \int_{B_r(z)} G_{B_r(z)}(z, y) f(y) dy.$$

On the same state-coverage event, simultaneously for contained radii,

$$B(z, r) \geq \underline{B}_{\text{avg}}(z, r) := 4\pi r (L_u - I_r^+)_+^2. \quad (42)$$

One may use $I_r^+ = f_+ r^2/6$ with an upper source bound on $B_r(z)$. For affine f , the exact source integral is $f(z)r^2/6$.

The proof pairs the state-difference energy with the ball's equilibrium potential, bounding benefit by the squared *average* cavity trace; the center Green identity recovers that average. The application validator uses

the maximum of this bound and the preceding temperature-floor bound. [Appendix O](#) derives this identity, the clipped-walk validator, geometric guards, and the separate application experiment.

A.2 Scout–rank–validate–refine

One outer iteration is:

1. **Scout.** Generate candidates in free space subject to boundary clearance. Evaluate a cheap TWoS map using shared unbiased caches or a deliberately approximate ranking cache.
2. **Rank.** Keep the m most promising spatially separated candidates. Non-maximum suppression avoids repeatedly validating one basin.
3. **Validate.** Fix one finalist using only the scouting and ranking data, then draw a separate validation batch. Alternatively, validate several survivors and allocate confidence across that entire shortlist. Use a valid bias bound for the general-objective route or the one-sided construction for cooling. A sequential implementation must use either an anytime-valid confidence sequence or a predeclared stage schedule $N_1 < N_2 < \dots$ with error levels whose sum is at most the candidate’s allocated failure probability; a fixed- N interval may not be repeatedly inspected under optional stopping.
4. **Insert.** Choose the site with the largest covered benefit bound meeting the target (and any explicit cost), and instantiate $K_{\varepsilon,z}$. If no site is certified, perform only shape refinement or terminate.
5. **Refine.** Move existing interfaces using Differential WoS or derivative WoSt [\[23, 34\]](#), with a line search on independently estimated objective values.

Proposition A.7 (Adaptive selection across candidates and rounds). *Let $(\mathcal{F}_t)_{t \geq 0}$ be the history of the optimizer. At validation round t , let the scout choose (Z_t, ε_t) using only $\mathcal{F}_{t-1}$, and let a fresh validation procedure produce an event E_t such that*

$$\mathbb{P}(E_t^c \mid \mathcal{F}_{t-1}) \leq \delta_t \quad \text{almost surely.} \quad (43)$$

If the implementation accepts only when the observable test [Eq. \(39\)](#) holds, and E_t is its valid coverage event, then

$$\mathbb{P}(\text{some accepted insertion is non-improving}) \leq \sum_{t \geq 1} \delta_t. \quad (44)$$

Here t indexes individual preselected candidate tests. If a round selects among several validation estimates, coverage must hold simultaneously for that shortlist, for example by allocating δ_t/m to each of m candidates. Shared randomness across candidates is allowed; fresh data alone does not remove selection error within the validation shortlist.

A.3 Practical placement and connected refinement

The original box, Bunny, and housing ranking experiments prescribe the number and radius of cavities and use fresh shortlist estimates to rank finite insertions. Their heuristic uses lower confidence scores but supplies neither a numerical remainder bound nor full coverage after selection. These experiments therefore evaluate placement quality, not the finite-insertion certificate. The count–volume study varies prescribed radii across allocations; it optimizes centers, not individual radii.

For connected geometry, we select two interior terminals using topological scores and join them by a weighted candidate graph. The path initializes an overlapping ball union. Strict consecutive-ball overlap is enforced after each radius recalibration, so the removed interior is path connected. Differential WoS then proposes motions of its controls, and fresh paired objective estimates assess trial steps. This graph initialization is a finite geometric heuristic: the single-cavity theorem does not certify simultaneous insertion of the whole channel. Its endpoints remain internal, so connectedness does not imply exterior cooling access. [Appendix G](#) gives the parameterization, regularity assumptions, chain rule, discrete volume calibration, and conditional shape-acceptance result.

A.4 Choosing a finite radius

After estimating D_T , choose an admissible ε satisfying the minimum-benefit test, clearance constraints, and a manufacturing lower bound. For example, maximize the covered net-benefit bound over the allowed radii; the largest feasible radius need not maximize that bound. [Corollary A.3](#) shows why this data-dependent choice is valid: the same confidence event controls every radius because the deterministic remainder bound is uniform in ε . A higher-order asymptotic estimate may permit a larger step, but the leading-order certificate remains the fallback.

A.5 Multiple insertions

The single-hole expansion does not justify inserting many nearby cavities at once. The proved policy is sequential: validate and insert one cavity, update the domain, then condition on that history and validate the next one using a fresh batch. This is covered by [Proposition A.7](#). Simultaneous insertion requires a separate uniform interaction bound and is not claimed by the present theorem.

B Evaluation

We evaluate three claims: direct estimation avoids shrinking-target noise; local scores support useful cavity placement; and sparse queries can reduce computational cost relative to a global discretized solve. Matched pipelines on a finned housing, branching manifold, and annular spreader test placement and numerical cost under shared constraints. The Bunny supplies a detailed-geometry stress test and a count-radius study, while controlled experiments isolate statistical effects. Earlier thermal studies retain their original protocols and are summarized separately. [Appendices X](#) and [Y](#) retain full protocols, secondary plots, solver checks, and all measured comparisons.

Common physics and decision rule. All 3D design studies solve $-\Delta u = f$ with nonnegative internal heating and ideal zero-temperature exterior and cavity walls. The sensor-protection pipelines and accessible-channel studies minimize the regional objective $J_q = \int q u / \int q$ with a distinct adjoint; the compact sensor weight in the new matched study is normalized to integrate to one. The earlier Bunny, housing-compliance, and guarded-validation studies minimize compliance $J = \int f u$ under prescribed removal and clearance constraints. Their relative reductions use a fixed original-domain normalization and fresh objective samples. For compliance, $p = u$ and the candidate benefit is u^2 ; a cross-fitted score uses distinct walk indices even when the same samples represent both fields. Each insertion changes the domain before the next score evaluation. The ranking studies use fresh shortlist samples. [Appendix B.9](#) separately evaluates a guarded compliance selector.

Measured comparisons. The deterministic design baselines use a voxel adjoint field with either conjugate gradients or geometric multigrid; they evaluate the same topological score and share the candidate and cavity budgets. Random feasible placement tests whether the score helps beyond adding cold surface. Same-seed multiplication and finite-cavity differencing isolate estimator choices. Array-RQMC and neighbor reuse are studied separately as compatible accelerations. The external FEM experiment checks initial fields, and the probabilistic-shape and spectral-connectivity experiments check components; none is presented as a reproduced full external optimizer.

B.1 Matched design pipelines on distinct geometry

We next isolate three questions that a cavity-count comparison alone cannot answer: whether objective-aware nucleation improves subsequent boundary optimization; whether accessible-channel improvements survive fixed material and power budgets; and whether the shrinking-target comparison transfers to actual three-dimensional application meshes. The finned housing, branching manifold, and ribbed annular spreader in [Fig. 12](#) have narrow gaps, branching paths, or curved transport and a thin central bridge. Their protected sensors and complete compact heat-source supports remain unchanged under every design. [Appendix V](#) specifies the physics, normalizations, geometry constraints, and numerical checks.

Nucleation followed by Differential WoS. All methods place six disjoint spheres of radius 0.065 from the same paired candidate pools, then optimize their centers using a common specialized implementation of Differential WoS [23]. TWoS and deterministic topological derivatives use eight refinement rounds; hottest-point and random initializations also receive rounds nine through sixteen. Two complete deterministic baselines use longest-axis grid counts 49 and 65. A nearest-sensor control replaces the PDE score by Euclidean distance to the protected sensor and receives eight refinement rounds. The curves in Fig. 5 report complete design costs at independently evaluated checkpoints. This comparison isolates the value of initialization inside a shared optimizer; it does not compare against the original authors’ full application or assert that Differential WoS cannot change topology.

Across five paired seeds, TWoS followed by eight refinement rounds yields mean sensor reductions of 65.0%, 16.4%, and 11.4% on the housing, branching manifold, and annular spreader. The primary multigrid selector followed by the same refinement yields 65.1%, 16.3%, and 11.7%. Hottest-point initialization with sixteen rounds yields 42.7%, 0.14%, and 0.01%, respectively. Both TD selectors perform similarly in these comparisons; their numerical costs and the additional controls are reported in the full time–quality curves and Table 9. The nearest-sensor control yields 63.6%, 17.3%, and 12.2% across the same three objects. On the housing it requires 7.08 s, compared with 11.77 s for TWoS. Thus the large contrast with hottest-point placement does not establish a large benefit over every informed heuristic. The cutaways show this stronger geometric control alongside TWoS and multigrid; the plots retain all methods. The bands describe uncertainty in measured mean cooling, not equivalence tests or guarantees for a new geometry.

Accessible channels at matched volume and fixed power. Six parallel open ball-union channels target total removed volume 0.0168. Their common radius is recalibrated after each insertion, and independent scrambled-Sobol samples audit the final volume. All heat supports are protected, so their exact input power is invariant (Proposition V.1). Figure 13 reports these finite designs. A pointwise axis score initializes the channels; it is not a topological derivative of an entire channel, and the radius calibration is numerical rather than certified. At this matched budget, TWoS yields mean reductions of 37.8%, 2.29%, and 1.75% across the three objects, compared with 42.1%, 2.59%, and 1.93% for the primary multigrid selector. On the housing, nearest-sensor placement yields 39.4%, and the cheaper multigrid grid attains the same mean reduction as the primary grid. Thus this finite-channel heuristic does not dominate the informed controls. Table 10 reports all methods and uncertainty intervals; the independent volume and reference audits remain separate from those sampling intervals.

Variance, bias, and query cost. Figure 4 compares direct scores, common-random-number finite-cavity subtraction, a stronger stopped-path harmonic-correction baseline, and multigrid at four resolutions on eight fixed sites per object. Sixteen independent repeats resolve sampling variability. Relative variance about each method’s mean is shown separately from error to a high-budget infinitesimal reference: finite-cavity methods retain finite-radius bias, while multigrid retains discretization error. The reference itself has sampling and termination error. Bands around variance describe uncertainty in the estimated variance, not a standard-deviation envelope of field values. Appendix V gives the stronger finite-correction identity and its step-by-step proof; detailed numeric summaries accompany the source. At radius 0.01 and 2,048 paths per site, the stronger harmonic-correction baseline has $35.7\times$ and $13.8\times$ the measured direct variance on the housing and branching object. These are point estimates, not confidence bounds. On the annular object, all recorded harmonic corrections vanish at radii 0.01 and 0.02; variance is unresolved, and collapsed bootstrap intervals are suppressed. At larger radii the stronger correction can have competitive measured variance. Direct TWoS relative RMS error at this budget is 14.9%, 25.2%, and 146% across the three cases, while the reference’s relative standard-error norms are 1.6%, 4.8%, and 16.2%. The annular sparse-query configuration is poorly resolved at these budgets; it is not evidence of uniformly accurate pointwise estimation.

The maximum independent channel-volume audit deviation is 0.97% of the target. Finer housing checks change by at most 0.44 percentage points from grid 193 to 257, but MC-grid257 differences still reach 2.38 points across the checked first-seed states. Sampling intervals do not absorb this numerical reference sensitivity (Table 12).

B.2 Protecting a region with a distinct adjoint

Compliance makes the adjoint equal to the state. To test whether objective-aware scoring matters beyond finding hot spots, we keep the housing’s heat sources fixed and minimize the weighted mean temperature of a protected sensor region. The objective is $J_q = \int qu / \int q$, where q is a compact nonnegative bump in a radius-0.13 ball. The adjoint solves $-\Delta p = q$, and candidates are ranked by up . This differs from the hottest-point score u . Two sensor placements test the left and right sides of the same housing; they are two objectives, not two independent input geometries. All cavities must avoid the sensor region with an additional 0.02 gap, so no design can improve the objective by deleting the sensor. [Appendix T](#) derives the normalization, adjoint, and an exact counterexample to hottest-point placement.

Each of five paired candidate/design seeds supplies 768 candidates. Every method places six radius-0.095 cavities under the same clearance constraints. We compare sparse TWoS (12 scouting and 192 rescoring walks), a larger budget (48 and 768), WoS hottest-point placement at the sparse budget, deterministic multigrid hottest-point placement, multigrid topological derivatives, and random feasible placement. The 24-site rescore uses fresh walks. Paired state–adjoint queries share paths within each pair and use the cross-index product. Multigrid reuses one geometry hierarchy for its two right-hand sides. The primary grid is $113 \times 73 \times 57$; a supplementary sweep adds $57 \times 37 \times 29$, $85 \times 55 \times 43$, and $169 \times 109 \times 85$ without discarding unfavorable results.

Sparse TWoS reduces the left/right regional objectives by 50.87% / 50.09%, compared with 38.15% / 29.99% for deterministic hottest-point placement. The primary multigrid topological selector gives 53.30% / 52.15%. Averaged across both cases, complete sparse TWoS design takes 3.66 s, versus 8.38 s for the primary multigrid setting. The fastest tested multigrid setting (57 nodes along the first axis) takes 2.51 s and gives 53.50% / 52.53% reduction. Retaining the outer mesh in one process per complete adaptive design reduces wall time to 1.22 s for sparse TWoS, 1.41 s for coarse multigrid, and 7.13 s for primary multigrid. All 70 repeated designs have exactly the same six sequential cavity centers as their original runs, permitting reuse of the independent objective evaluations. These results demonstrate an advantage over hottest-point placement; the full budget sweep in [Fig. 14](#) and [Table 7](#) qualifies comparisons with volumetric topological derivatives.

Complete design wall times include common candidate preparation, geometry setup, source queries or PDE solves, candidate filtering, ranking, subprocess startup, and file I/O. The retained-mesh backend keeps the fixed outer boundary acceleration structure alive throughout each six-round design, for both Monte Carlo and multigrid methods. Candidate filtering, fresh rescoring, and every adaptive insertion still execute inside the measured driver; future requests are not replayed ahead of those decisions. Offline reference evaluation is measured separately. Both a fresh WoS evaluation and a separate $169 \times 109 \times 85$ multigrid solution assess every final design. These are approximate references; first-seed designs additionally receive a fourfold WoS shell refinement and a $225 \times 145 \times 113$ grid check. The designs use fixed-count ranking, not the compliance-only finite-benefit certificate. The comparison against hottest-point placement tests the adjoint’s value, which is shared with deterministic topological derivatives; the quality–cost comparison tests the choice of numerical estimator.

B.3 Accessible channels and protected components

We next replace sealed cavities with vertical through-channels in the housing. Every channel is a union of 18 radius-0.075 balls, spanning both exterior sides. Their spacing gives a guaranteed continuous passage of radius $0.075\sqrt{3}/2$; [Appendix U](#) proves this geometric statement and the component-clearance predicates. The sensor and three prescribed component regions cannot be removed. Port positions vary with the selected axes. We retain the ideal cold-wall model, so geometric access does not establish coolant delivery or realistic heat transfer.

This is an application of point scoring to a larger geometric operation: the algorithm scores candidate axes at a common interior plane and inserts the complete channel. The pointwise topological coefficient is a placement proxy here, not a derived whole-channel derivative. All five methods use the same candidate axes, six-channel count, radius, extent, and protection constraints. We repeat both sensor objectives over five paired seeds and evaluate all 50 completed designs independently. Removed volume is reported because local housing thickness varies; this is not an equal-volume experiment.

On the left and right sensors, the TWoS proxy gives reductions of 55.74% and 52.47%, versus 43.78% and 30.39% for hottest-point placement. Mean complete design times are 1.02 s for TWoS and 1.26 s for coarse

multigrid. Their mean reductions across both sensors are 54.10% and 57.55%, respectively. This preserves a useful quality–cost tradeoff under access constraints, without establishing superiority over deterministic topological scoring.

B.4 Thermal design on a finned housing

A closed 96,740-triangle procedural housing contains three localized Gaussian heat sources. Each method places six radius-0.095 cavities from 768 shared Sobol candidates. Sparse TWoS uses 12 scouting walks per feasible site and 192 fresh walks on a 24-site shortlist; a larger budget uses 96 and 768 walks. A predeclared field-error criterion selects the multigrid grid, $113 \times 73 \times 57$, from a resolution sweep. Both algorithms run in the same C++ executable. Selection time includes geometry acceleration and all six rounds of querying or global solving; it excludes file I/O, process startup, and the common independent objective evaluation.

Across five complete stochastic design seeds, sparse TWoS reduces compliance by $68.90\% \pm 0.84\%$ (standard error), compared with $70.36\% \pm 0.08\%$ for the deterministic multigrid design, whose uncertainty comes from objective evaluation. Mean selection time is 0.432 s versus 1.861 s: $4.31\times$ faster while achieving 97.9% of the measured reduction. The larger TWoS budget gives $69.39\% \pm 0.51\%$ reduction in 1.731 s, using 19.1 MiB peak RSS versus 31.4 MiB. Random placement gives $27.76\% \pm 4.57\%$. These are quality–cost tradeoffs over the tested designs; the seed count does not establish equivalence or superiority in design quality. The multigrid sweep and resource curves appear in Fig. 22. This earlier compliance comparison uses a field-error-selected grid and excludes driver costs; the regional pipelines in Appendix B.1 include those costs and both cheaper grids. Its speed ratio should not be transferred between these protocols.

An independent scikit-fem implementation on 393,616 tetrahedra checks the initial ranking. The 512-walk field has Pearson correlation 0.9943 and 95% top-20 overlap with that FEM field. Its boundary approximation and convergence are documented in Fig. 23; the FEM audit is an accuracy comparison, not an optimized FEM timing baseline.

B.5 Detailed geometry and sparse-query scaling

We repair and close the Stanford Bunny [32], obtaining 69,674 triangles, and normalize its height to one. With affine heating $f(x, y, z) = 1 + 0.65(y + 0.5)$, each method places six radius-0.05 cavities among 1,024 candidates. TWoS uses a 12-walk scout and fresh 128-walk rescoring of 20 sites. Its measured compliance reduction is 47.77%, compared with 48.06% for the $54 \times 64 \times 48$ voxel adjoint and 39.46% for random placement. This single-run test demonstrates placement on detailed geometry; the replicated allocation study below supplies a broader design comparison.

A separate matched C++ experiment isolates query cost on the original Bunny. At 64 requested sites, 512-walk TWoS takes 0.381 s and 15.4 MiB peak RSS; voxel FD/CG takes 0.528 s and 80.7 MiB. The corresponding relative score errors against an independent 16,384-walk reference are 5.89% and 6.89%. This is $1.39\times$ faster evaluation with 81% less peak memory, at similar measured error. The numerical reference itself has sampling and termination error. The methods tie near 96 sites, and at 384 sites the voxel solve is faster (0.504 s versus 2.10 s). Figure 21 reports the full sweep. These timings establish a sparse-query advantage, not a speedup for the displayed Bunny optimization.

B.6 Cavity count at fixed removed volume

Does cooling improve because placement is better, or simply because more material is removed? We separate these effects with 72 Bunny designs: two exact volume budgets, four counts $n \in \{2, 4, 6, 12\}$, three independent design seeds, and TWoS, voxel-adjoint TD, and random feasible placement. The budgets equal three or six radius-0.05 balls: 0.765% or 1.529% of the original Bunny volume. Within each allocation, every cavity has radius $r_n = 0.05(b/n)^{1/3}$, where $b \in \{3, 6\}$ is the budget in ball-volume units. Radii follow this constraint; only centers are optimized. All methods share 1,024 candidates within each seed/allocation, retain outer clearance 0.012 and pairwise surface gap 0.014, and start from the unperforated domain. TWoS retains the 12-walk scout and independent 128-walk shortlist rescoring; this ranking heuristic is not a finite-insertion confidence certificate. The compiled voxel selector uses a finer $96 \times 112 \times 84$ FD/CG grid than the original Bunny design test.

At the full volume budget, increasing the TWoS count from two to twelve raises mean cooling from 37.58% to 54.33%, while the radius decreases from 0.0721 to 0.0397. At twelve cavities, the voxel and random means are 55.02% and 38.78%, respectively. The most useful allocation comparison is across budgets: twelve small TWoS cavities at half volume yield 47.26% cooling (design range 47.03–47.67%), versus 47.45% (46.78–48.22%) for six larger cavities at full volume. The former also has 20.6% less cavity surface area, by exact sphere geometry. This is similar measured cooling with less removal, not a proved equivalence or an advantage unique to TWoS. The voxel selector shows a similar allocation trend. A fourfold thinner WoS termination shell, 1.25×10^{-4} , changes cooling by at most 0.204 percentage points for the pre-specified first-seed full-budget endpoint designs from both TD selectors. This sensitivity check does not bound the remaining bias. An additional first-seed check of the half-volume comparison, chosen after inspecting the allocation results, gives 46.88% versus 47.21% cooling with the thinner shell, respectively.

All objective comparisons use four independently scrambled 512-point interior blocks and 128 WoS paths per active point, separate from selection. Removed material contributes zero under fixed original-volume normalization; reported percentages average paired block ratios to the initial state. Design ranges and conditional block standard errors are retained separately; three seeds do not establish equivalence or a globally optimal allocation.

Why a preferred count needs a cost. Additional nested ideal cooled removal lowers compliance exactly (Proposition M.1), whereas different equal-volume allocations are not nested. At fixed removed volume V , exposed area grows as $n^{1/3}$ and the sum of isolated capacities as $n^{2/3}$ (Proposition M.2); the latter is not the interacting union capacity. A volume-only penalty is constant within this sweep. Instead, Fig. 16c illustrates $J/J_0 + \kappa n$ for an explicitly assumed count cost. This isolates the modeling choice behind a “sweet spot” without attributing it to the topological estimator or claiming an application-specific optimum.

B.7 Estimator accuracy and sampling uncertainty

The annulus diagnostic has an exact capacity and cavity-hit probability, making the rare-event comparison analytically controlled. With 256 samples per estimate and 3,000 repetitions, fitted variance slopes versus capacity are 0.00 for direct estimation, -1.10 for coupled differencing, and -1.89 for independent differencing. At the smallest tested capacity, direct variance is $8.26\times$ lower than the coupled baseline. This experiment tests the shrinking-target mechanism under its stated model; it does not rule out importance samplers conditioned on cavity hits.

Correlated-pair tests confirm the covariance bias predicted for same-index products. Cross-fitting and fully independent primal/adjoint batches both remove it; cross-fitting is not uniformly lower variance. The manufactured harmonic fields in Fig. 18 test this product algebra spatially. A separate Poisson-compliance disk test has a genuine analytic topological derivative, $D_T(z) = -(1 - |z|^2)^2/16$. There, independently randomized Array-RQMC batches reduce MSE by $19.9\times$ relative to ordinary Monte Carlo at equal path count, while taking longer in the measured Python implementation (Fig. 26). Statistical efficiency and runtime are reported separately.

The adaptive-selection test uses null candidates with known Gaussian noise. Selecting and testing with the same samples produces almost certain false acceptance among 256 sites; fresh validation and simultaneous Bonferroni bounds produce rates near the nominal 5%. This isolates selection error. It does not measure coverage for finite 3D insertions. The general-objective certificate also requires numerical remainder and bias bounds; the specialized cooling validator is evaluated separately below.

B.8 Compatibility with connected refinement

On the housing, a topological-score-weighted candidate graph initializes a connected 22-ball cavity between two internal terminals. Every trial enforces strict overlap and the same fixed-Sobol removed volume. Six Differential WoS refinement rounds propose control-point motions; fresh paired Student- t tests with a Bonferroni allocation accept three steps. Independent evaluation reports a 2.086% objective decrease from the graph initialization. The graph initialization itself does not measurably outperform the straight connection in this example. The channel remains sealed inside the housing, and the interval coverage relies on the paired-normal model. This experiment demonstrates compatibility between cavity initialization and




Figure 7 Controlled tests of the statistical claims. (a) Direct TWoS has no insertion-radius input, so its measured variance is constant; coupled finite-cavity differencing follows the predicted ρ^{-1} rare-event law, while independent solves approach the worse ρ^{-2} law. Shading gives 95% bootstrap intervals. (b) Same-seed multiplication inherits the full covariance bias and the product of sample means retains $1/N$ of it. The leave-diagonal-out estimator remains unbiased up to Monte Carlo error; independent batches are the unbiased no-reuse reference. (c) Reusing the most negative scout loses nominal coverage as the candidate count grows. Fresh validation and a simultaneous Bonferroni scout both meet the 5% target; fresh validation controls the preselected-site test in this known-noise diagnostic.


existing-boundary refinement; it does not establish accessible coolant routing or a certified global topology optimum. Figures 24 and 25 report the geometry, derivative checks, and finite-offset sensitivity.

B.9 Minimum-benefit validation on application geometry


We separately deploy the average-trace certificate in a six-attempt design on each object. A scout-and-rescore pass fixes one finalist; fresh clipped walks then bound its finite benefit before insertion. The validator uses conservative triangle-box distance bounds, explicit source enclosures, and an error budget of 0.01 across both objects and all predeclared sampling stages. Its probabilistic guarantee remains conditional on ideal sampling laws and the stated input-geometry assumptions.

All twelve insertions pass the fixed absolute targets, yielding cumulative benefit lower bounds of 5.75×10^{-4} for the Bunny and 6.21×10^{-3} for the housing. Independent WoS evaluations estimate 8.45×10^{-4} and 1.020×10^{-2} , respectively; a separate multigrid reference agrees on total benefits within 0.7%. The average-trace bound is substantially tighter than our temperature-floor bound on the same data (Fig. 8). Eleven targets pass at 1,024 walks and one at 4,096 walks, but the full four-stage diagnostic costs 150.39s and 52.17s of validation time, versus 2.97s and 0.53s for scouting and rescoring. This is a conservative validation experiment, not an extension of the earlier ranking-only speed claim. Appendix O gives every proof, target, timing boundary, and numerical qualification.


(a) Bunny: finite insertion benefit



(b) Housing: finite insertion benefit



(c) Tightening the bound with more samples



(d) Added cost of conservative validation



Figure 8 A finite benefit bound that can be evaluated before insertion. (a,b) The average-trace certificate is tighter than the temperature-floor certificate and source-only bound on the same proposed cavities. Separate WoS and multigrid evaluations estimate the actual benefit; error bars show one spatial-block standard error for WoS. Percent axes use fixed scales 0.002 (Bunny) and 0.015 (housing), not certified baseline objectives. The dashed horizontal line marks the primary 0.1%-of-scale target. (c) Larger validation batches tighten the first candidate’s bound. (d) Complete six-attempt kernel costs, with recorded first-passing prefixes shown separately from the full diagnostic actually executed. The bounds assume the stated conditional sampling laws and valid input geometry; numerical guard audits are reported separately.

C Deterministic PDE and Capacity Proofs

This appendix proves [Theorems 3.4](#) and [3.5](#) without importing a small-hole expansion. Constants denoted by C may change from line to line, but are independent of $z \in Z$ and sufficiently small ε . We write $K = K_{\varepsilon,z}$, $\Omega_\varepsilon = \Omega_{\varepsilon,z}$, $q_\varepsilon = q_{\varepsilon,z}$, and $\rho_\varepsilon = \rho_z(\varepsilon)$ when no confusion can arise. The background facts used below—trace and compact embedding, Lax–Milgram, quasi-continuous representatives, positive distributions, and local Green-kernel bounds—are standard results of Sobolev and elliptic potential theory [[1](#), [9](#), [12](#)]; all problem-specific consequences are derived explicitly.

C.1 Why regular shape calculus cannot nucleate the cavity

Proposition C.1 (Cavity nucleation is not a regular shape path). *Assume additionally that $\partial\Omega$ and ∂B are connected. Let $V \in C_c^1(\mathbb{R}^d; \mathbb{R}^d)$ and $T_t = \text{Id} + tV$. For all sufficiently small $|t|$, T_t is a C^1 diffeomorphism, $T_t(\Omega)$ is diffeomorphic to Ω , and $\partial T_t(\Omega) = T_t(\partial\Omega)$ is connected. Hence no such regular perturbation satisfies*

$$T_t(\Omega) = \Omega \setminus \overline{(z + \varepsilon B)} \quad (45)$$

for any $\varepsilon > 0$ with $z + \varepsilon \bar{B} \in \Omega$, because the right-hand side has the two disjoint boundary components $\partial\Omega$ and $z + \varepsilon \partial B$. Consequently Eq. (10) is not a directional shape derivative along a regular deformation of the unperforated domain.

Proof of Proposition C.1. We separate the analytic and topological parts.

Step 1: T_t is a local diffeomorphism for small t . Because $V \in C_c^1(\mathbb{R}^d; \mathbb{R}^d)$, the quantity $L := \|DV\|_{L^\infty(\mathbb{R}^d)}$ is finite. Choose $|t|L < 1$. Then

$$DT_t(x) = I + tDV(x) \quad (46)$$

is invertible at every x : indeed, $\|tDV(x)\| < 1$, so the Neumann series $\sum_{k=0}^{\infty} (-tDV(x))^k$ converges to $DT_t(x)^{-1}$. The inverse-function theorem therefore makes T_t a local C^1 diffeomorphism.

Step 2: T_t is globally injective and proper. For any $x, y \in \mathbb{R}^d$, the mean-value estimate gives

$$\|T_t(x) - T_t(y)\| \geq \|x - y\| - |t| \|V(x) - V(y)\| \quad (47)$$

$$\geq (1 - |t|L) \|x - y\|. \quad (48)$$

Thus $T_t(x) = T_t(y)$ implies $x = y$. Since V has compact support, $T_t(x) = x$ outside a compact set; hence T_t is proper. A proper injective local diffeomorphism has open and closed image in connected $\mathbb{R}^d$, so its image is all of $\mathbb{R}^d$. Therefore T_t is a global C^1 diffeomorphism.

Step 3: a diffeomorphism preserves the boundary. A homeomorphism preserves closure and interior. Hence

$$\partial T_t(\Omega) = \overline{T_t(\Omega)} \setminus \text{int } T_t(\Omega) \quad (49)$$

$$= T_t(\overline{\Omega}) \setminus T_t(\text{int } \Omega) = T_t(\partial\Omega). \quad (50)$$

Because $\partial\Omega$ is connected and T_t is continuous, $T_t(\partial\Omega)$ is connected.

Step 4: an interior cavity creates a second boundary component. Put $K_{\varepsilon,z} = z + \varepsilon \bar{B} \in \Omega$ and $\Omega_{\varepsilon,z} = \Omega \setminus K_{\varepsilon,z}$. Positive separation of $K_{\varepsilon,z}$ and $\partial\Omega$ implies

$$\partial\Omega_{\varepsilon,z} = \partial\Omega \dot{\cup} (z + \varepsilon \partial B). \quad (51)$$

Both sets are nonempty, closed in the relative boundary, and separated by a positive distance. Since each is connected by assumption, they are exactly two distinct connected components. Therefore $\partial\Omega_{\varepsilon,z}$ is disconnected.

Step 5: contradiction. If Eq. (45) held, Eq. (50) would make the boundary of its left-hand side connected, whereas Eq. (51) makes the boundary of its right-hand side disconnected. This is impossible. Regular Eulerian shape calculus differentiates $J(T_t(\Omega))$ at $t = 0$; the perforated family does not lie on any such path. Its capacity-normalized limit is therefore a genuinely singular topological derivative rather than a directional derivative at the original unperforated domain. The argument does not rule out additional asymptotic analysis of shape derivatives on domains with pre-existing positive-radius cavities. $\square$

C.2 Functional-analytic foundation

Equip $V = H_{\Gamma_D}^1(\Omega)$ with

$$\|v\|_a^2 := a_\Omega(v, v). \quad (52)$$

Proposition C.2 (Boundedness and coercivity). *There are constants $c_a, C_a > 0$ such that for all $v, w \in V$,*

$$|a_\Omega(v, w)| \leq C_a \|v\|_{H^1} \|w\|_{H^1}, \quad a_\Omega(v, v) \geq c_a \|v\|_{H^1}^2. \quad (53)$$

The same coercivity constant applies to every perforated test space after zero extension across the cavity.

Proof. **Step 1: boundedness.** Hölder's inequality and the coefficient bounds give

$$|a_\Omega(v, w)| \leq \kappa_+ \|\nabla v\|_2 \|\nabla w\|_2 + \|\sigma\|_\infty \|v\|_2 \|w\|_2 \quad (54)$$

$$\leq C_a \|v\|_{H^1} \|w\|_{H^1}. \quad (55)$$

Step 2: a Poincaré inequality. Because the trace vanishes on a boundary subset of positive $(d-1)$ -dimensional measure, there is $C_P < \infty$ such that

$$\|v\|_{L^2(\Omega)} \leq C_P \|\nabla v\|_{L^2(\Omega)} \quad (v \in V). \quad (56)$$

Indeed, if this failed, a normalized sequence with unit L^2 norm and vanishing gradients would converge in L^2 to a nonzero constant; continuity of the trace would force that constant to vanish on Γ_D , a contradiction.

Step 3: coercivity. Since $\sigma \geq 0$ and $\kappa \geq \kappa_-$,

$$a_\Omega(v, v) \geq \kappa_- \|\nabla v\|_2^2 \geq \frac{\kappa_-}{1 + C_P^2} \|v\|_{H^1}^2. \quad (57)$$

If v is a perforated-domain test function, its zero extension across K belongs to V because its trace on ∂K is zero. Applying the same inequality to that extension proves the final assertion. $\square$

Proposition C.3 (Well-posed state, adjoint, and perforated state). *The weak problems in Eqs. (3) and (5) have unique solutions. For every sufficiently small ε , the perforated state with trace g on Γ_D and h_z on ∂K also exists uniquely. Its extension by h_z inside K belongs to $H^1(\Omega)$.*

Proof. **Step 1: the state.** Choose the fixed lift $g \in H^1(\Omega)$ and write $u = g + u_0$ with $u_0 \in V$. Then

$$a_\Omega(u_0, v) = \ell_\Omega(v) - a_\Omega(g, v) \quad (v \in V). \quad (58)$$

The right side is a bounded linear functional. Boundedness and coercivity from Proposition C.2 allow the Lax–Milgram theorem to give a unique u_0 , hence a unique u .

Step 2: the adjoint. Because $j'(u)$ is bounded on $L^2(\Omega)$, its restriction to V is bounded by the continuous embedding $V \hookrightarrow L^2$. Another application of Lax–Milgram gives a unique $p \in V$ satisfying Eq. (5).

Step 3: the perforated state. Define

$$V_\varepsilon := \{v \in H^1(\Omega_\varepsilon) : \text{Tr } v = 0 \text{ on } \Gamma_D \cup \partial K\}. \quad (59)$$

Since $K \Subset \Omega$, a smooth cutoff supplies an $H^1(\Omega_\varepsilon)$ lift with outer trace g and inner trace h_z . Subtracting that lift reduces the problem to V_ε . The zero-extension argument in Proposition C.2 supplies uniform coercivity, so Lax–Milgram again gives existence and uniqueness.

Step 4: extension across the cavity. The trace of $u_{\varepsilon,z}$ on ∂K equals the trace of the constant function h_z inside K . The Sobolev gluing lemma therefore shows that the piecewise definition $\tilde{u}_{\varepsilon,z} = u_{\varepsilon,z}$ on Ω_ε and $\tilde{u}_{\varepsilon,z} = h_z$ on K belongs to $H^1(\Omega)$. $\square$

C.3 Equilibrium potential and capacitary measure

Let

$$\mathcal{A}_{\varepsilon,z} := \{v \in V : \tilde{v} \geq 1 \text{ quasi-everywhere on } K_{\varepsilon,z}\}. \quad (60)$$

Lemma C.4 (Equilibrium potential). *The class $\mathcal{A}_{\varepsilon,z}$ is nonempty, closed, and convex. It has a unique minimum-energy element q_ε . This minimizer satisfies*

$$\begin{aligned} 0 \leq q_\varepsilon \leq 1 & \quad \text{quasi-everywhere,} \\ q_\varepsilon = 1 & \quad \text{quasi-everywhere on } K, \end{aligned} \tag{61}$$

$$a_\Omega(q_\varepsilon, \varphi) = 0 \quad \text{if } \varphi \in V \text{ and } \tilde{\varphi} = 0 \text{ q.e. on } K. \tag{62}$$

Proof. Step 1: nonemptiness and convexity. Because $K \Subset \Omega$, a compactly supported smooth cutoff equal to one on a neighborhood of K belongs to $\mathcal{A}_{\varepsilon,z}$. Convexity is immediate from the pointwise inequality.

Step 2: weak closedness. Let $v_n \in \mathcal{A}_{\varepsilon,z}$ converge weakly in V to v . Mazur's lemma provides convex combinations of tails of (v_n) that converge strongly to v in V . A subsequence of a strongly H^1 -convergent sequence converges quasi-everywhere. Every convex combination remains at least one quasi-everywhere on K , hence so does v . Thus $v \in \mathcal{A}_{\varepsilon,z}$.

Step 3: existence and uniqueness. A minimizing sequence is bounded by coercivity. Weak compactness and the preceding closure give a weak limit in $\mathcal{A}_{\varepsilon,z}$; weak lower semicontinuity of the squared Hilbert norm gives a minimizer. Strict convexity of $a_\Omega(v, v)$ on V gives uniqueness.

Step 4: truncation. The map $T(s) = \min(1, \max(0, s))$ is one-Lipschitz, satisfies $|T(s)| \leq |s|$, and preserves the constraint on K . The Sobolev chain rule therefore gives

$$a_\Omega(T(q_\varepsilon), T(q_\varepsilon)) \leq a_\Omega(q_\varepsilon, q_\varepsilon). \tag{63}$$

Uniqueness forces $T(q_\varepsilon) = q_\varepsilon$, proving [Eq. \(61\)](#).

Step 5: harmonicity away from the obstacle. If $\tilde{\varphi} = 0$ quasi-everywhere on K , then $q_\varepsilon + t\varphi$ is admissible for every $t \in \mathbb{R}$. Differentiating its quadratic energy at $t = 0$ in both signs gives $a_\Omega(q_\varepsilon, \varphi) = 0$. $\square$

Lemma C.5 (Capacitary measure and localization). *There is a finite positive Radon measure μ_ε supported on $\overline{K}$ such that*

$$a_\Omega(q_\varepsilon, v) = \int_{\overline{K}} \tilde{v} d\mu_\varepsilon \quad (v \in V), \quad \mu_\varepsilon(\overline{K}) = \rho_\varepsilon. \tag{64}$$

Consequently, for every $v \in V \cap C^1(U)$,

$$|a_\Omega(q_\varepsilon, v) - \rho_\varepsilon v(z)| \leq \varepsilon R_B \rho_\varepsilon \|\nabla v\|_{L^\infty(U)}, \tag{65}$$

where $B \subset B_{R_B}(0)$.

Proof. Step 1: positivity. For every nonnegative $\varphi \in C_c^\infty(\Omega)$ and $t > 0$, $q_\varepsilon + t\varphi$ is admissible. Minimality gives

$$0 \leq \frac{a(q_\varepsilon + t\varphi, q_\varepsilon + t\varphi) - a(q_\varepsilon, q_\varepsilon)}{2t} = a(q_\varepsilon, \varphi) + \frac{t}{2}a(\varphi, \varphi). \tag{66}$$

Letting $t \downarrow 0$ yields $a(q_\varepsilon, \varphi) \geq 0$. Thus $\varphi \mapsto a(q_\varepsilon, \varphi)$ is a positive distribution and hence a positive Radon measure μ_ε .

Step 2: support. By [Eq. \(62\)](#), this distribution vanishes on every test function supported in $\Omega \setminus \overline{K}$. Therefore $\text{supp } \mu_\varepsilon \subseteq \overline{K}$.

Step 3: obtain the mass without extending the measure identity. Choose $\chi \in C_c^\infty(U)$ with $\chi = 1$ on a neighborhood of K . The function $q_\varepsilon - \chi$ belongs to V and vanishes quasi-everywhere on K . Thus [Eq. \(62\)](#), applied directly in V , gives

$$\rho_\varepsilon = a(q_\varepsilon, q_\varepsilon) = a(q_\varepsilon, \chi) = \int \chi d\mu_\varepsilon = \mu_\varepsilon(K).$$

The integral identity here is used only for the smooth compactly supported function χ , for which Steps 1–2 already prove it.

Step 4: the measure does not charge polar sets. For a compact capacity-zero set $F \subset K$, the smooth-cutoff characterization of Sobolev capacity gives nonnegative $\eta_n \in C_c^\infty(\Omega)$ that are at least one near F and satisfy $\|\eta_n\|_a \rightarrow 0$. This interior characterization agrees with the relative capacity: multiply admissible approximations by a fixed interior cutoff, use norm equivalence, and the usual smooth approximation characterization of Sobolev capacity. Only compact subsets of $K \Subset \Omega$ are needed. Positivity and the already established smooth-test identity imply

$$\mu_\varepsilon(F) \leq \int \eta_n d\mu_\varepsilon = a(q_\varepsilon, \eta_n) \leq \rho_\varepsilon^{1/2} \|\eta_n\|_a \rightarrow 0.$$

Inner regularity of the finite Radon measure then proves the same assertion for every Borel capacity-zero subset of K .

Step 5: extend locally by an energy bound. For $\varphi \in C_c^\infty(\Omega)$, use the smooth nonnegative functions $T_\delta(\varphi) = \sqrt{\varphi^2 + \delta^2} - \delta$. They have the same compact support, converge to $|\varphi|$, and satisfy $\|T_\delta(\varphi)\|_a \leq \|\varphi\|_a$ by the chain rule and $\sigma \geq 0$. Consequently, dominated convergence on the finite measure space gives

$$\int |\varphi| d\mu_\varepsilon = \lim_{\delta \downarrow 0} a(q_\varepsilon, T_\delta(\varphi)) \leq \rho_\varepsilon^{1/2} \|\varphi\|_a.$$

If smooth compactly supported φ_n converge to $w \in H_0^1(\Omega)$, this inequality makes them Cauchy in $L^1(\mu_\varepsilon)$. A subsequence converges quasi-everywhere to $\tilde{w}$ and, by Step 4, also μ_ε -almost everywhere. Its $L^1(\mu_\varepsilon)$ limit is therefore $\tilde{w}$. Passing to the limit in the smooth-test identity proves $a(q_\varepsilon, w) = \int \tilde{w} d\mu_\varepsilon$ for $w \in H_0^1(\Omega)$.

Step 6: extend to the mixed-boundary space by localization. For an arbitrary $v \in V$, $\chi v \in H_0^1(\Omega)$, while $v - \chi v$ vanishes on K . Orthogonality in Eq. (62) and Step 5 give

$$a(q_\varepsilon, v) = a(q_\varepsilon, \chi v) = \int_K \chi \tilde{v} d\mu_\varepsilon = \int_K \tilde{v} d\mu_\varepsilon.$$

This avoids assuming that $C_c^\infty(\Omega)$ is dense in V , which would be false when nonzero traces are allowed on Γ_N .

Step 7: localization of a smooth observable. For $v \in V \cap C^1(U)$, every $x \in K$ obeys $|x - z| \leq \varepsilon R_B$. Using the mass from Step 3,

$$|a(q_\varepsilon, v) - \rho_\varepsilon v(z)| = \left| \int_K (v(x) - v(z)) d\mu_\varepsilon(x) \right| \quad (67)$$

$$\leq \varepsilon R_B \|\nabla v\|_{L^\infty(U)} \rho_\varepsilon. \quad (68)$$

This proves Eq. (65) and completes the extension argument. $\square$

C.4 Capacity scale

Proof of Theorem 3.4. Existence and uniqueness follow from Lemma C.4; coercivity and the fact that K contains a ball imply $\rho_\varepsilon > 0$.

The three-dimensional upper bound. Let P_B be the Newtonian equilibrium potential of B : $P_B = 1$ quasi-everywhere on B , $P_B \rightarrow 0$ at infinity, and

$$\int_{\mathbb{R}^3} |\nabla P_B|^2 dy = \text{cap}(B). \quad (69)$$

Choose a fixed cutoff $\chi \in C_c^\infty(\Omega)$ that equals one on a neighborhood of Z . The function

$$v_{\varepsilon, z}(x) := \chi(x) P_B((x - z)/\varepsilon) \quad (70)$$

is admissible for all sufficiently small ε . After the change of variables $x = z + \varepsilon y$, local uniform convergence $\kappa(z + \varepsilon y) \rightarrow \kappa(z)$ and the finite energy of P_B yield

$$\int_\Omega \kappa |\nabla v_{\varepsilon, z}|^2 dx = \varepsilon \kappa(z) \text{cap}(B) + o(\varepsilon). \quad (71)$$

The terms in which the gradient hits χ occur a fixed positive distance from z , where $P_B((x - z)/\varepsilon) = O(\varepsilon)$; they are $O(\varepsilon^2)$. Likewise $\int \sigma v_{\varepsilon,z}^2 = O(\varepsilon^2) = o(\varepsilon)$. Minimality of q_ε gives

$$\limsup_{\varepsilon \downarrow 0} \frac{\rho_\varepsilon}{\varepsilon} \leq \kappa(z) \operatorname{cap}(B). \quad (72)$$

The three-dimensional lower bound. Define $Q_\varepsilon(y) = q_\varepsilon(z + \varepsilon y)$ on $(\Omega - z)/\varepsilon$. The upper bound makes $\int |\nabla Q_\varepsilon|^2$ uniformly bounded. Moreover, Sobolev embedding and coercivity on the fixed domain give $\|q_\varepsilon\|_{L^6(\Omega)} \leq C\rho_\varepsilon^{1/2} = O(\varepsilon^{1/2})$. Rescaling yields a uniform L^6 bound for Q_ε . Local weak compactness and a diagonal subsequence therefore give a limit Q with $Q \in L^6(\mathbb{R}^3)$ and $\nabla Q \in L^2(\mathbb{R}^3)$, hence $Q \in D^{1,2}(\mathbb{R}^3)$. The cavity constraint passes to the limit by weak closedness on any fixed ball containing B , so $Q \geq 1$ quasi-everywhere on B . For every fixed R , uniform convergence of the coefficient and weak lower semicontinuity give

$$\liminf_{\varepsilon \downarrow 0} \frac{\rho_\varepsilon}{\varepsilon} \geq \liminf_{\varepsilon \downarrow 0} \int_{B_R} \kappa(z + \varepsilon y) |\nabla Q_\varepsilon|^2 dy \quad (73)$$

$$\geq \kappa(z) \int_{B_R} |\nabla Q|^2 dy. \quad (74)$$

Letting $R \rightarrow \infty$ and then using the variational definition in Eq. (69) yields

$$\liminf_{\varepsilon \downarrow 0} \frac{\rho_\varepsilon}{\varepsilon} \geq \kappa(z) \operatorname{cap}(B). \quad (75)$$

Together, Eqs. (72) and (75) prove the $d = 3$ formula.

The two-dimensional upper bound. By compactness, choose $r_0 > 0$ so that

$$B_{2r_0}(z) \Subset \Omega \quad \text{for every } z \in Z. \quad (76)$$

Let $B \subset B_{R_B}(0)$. Put $L_\varepsilon = \log(r_0/(\varepsilon R_B))$ and use the radial competitor

$$v(r) = \begin{cases} 1, & r \leq \varepsilon R_B, \\ \log(r_0/r)/L_\varepsilon, & \varepsilon R_B < r < r_0, \\ 0, & r \geq r_0. \end{cases} \quad (77)$$

A direct polar-coordinate calculation gives

$$\int_{\varepsilon R_B}^{r_0} |v'(r)|^2 2\pi r dr = \frac{2\pi}{L_\varepsilon}. \quad (78)$$

Uniform continuity of κ and then $r_0 \downarrow 0$ show

$$\limsup_{\varepsilon \downarrow 0} |\log \varepsilon| \rho_\varepsilon \leq 2\pi \kappa(z). \quad (79)$$

The zero-order energy is $O(L_\varepsilon^{-2}) + O(\varepsilon^2)$ and therefore does not alter this limit.

The two-dimensional lower bound. Choose $r_B > 0$ with $B_{r_B}(0) \subset B$. Fix r_0 as above and let $\delta > 0$ be such that $\kappa(x) \geq \kappa(z) - \delta$ on $B_{r_0}(z)$. For an admissible v , let c be its circular mean on $\partial B_{r_0}(z)$. Applying Cauchy–Schwarz to the radial derivative of the circular mean gives

$$\int_{B_{r_0}(z) \setminus B_{\varepsilon r_B}(z)} |\nabla v|^2 dx \geq \frac{2\pi(1-c)_+^2}{\log(r_0/(\varepsilon r_B))}. \quad (80)$$

On the fixed exterior domain $\Omega \setminus B_{r_0}(z)$, the trace theorem, the Poincaré inequality on functions vanishing at Γ_D , and Eq. (53) give

$$a_{\Omega \setminus B_{r_0}(z)}(v, v) \geq C_0 c^2 \quad (81)$$

with $C_0 > 0$ uniform on Z . Minimizing the sum of $A(1-c)_+^2 + C_0 c^2$, where $A = 2\pi(\kappa(z) - \delta)/\log(r_0/(\varepsilon r_B))$, gives

$$\rho_\varepsilon \geq \frac{AC_0}{A + C_0}. \quad (82)$$

First let $\varepsilon \downarrow 0$ and then $\delta \downarrow 0$ to obtain

$$\liminf_{\varepsilon \downarrow 0} |\log \varepsilon| \rho_\varepsilon \geq 2\pi\kappa(z). \quad (83)$$

This and Eq. (79) prove the $d = 2$ formula.

Uniformity and comparison with volume. All cutoffs, radii, trace constants, and coefficient moduli can be chosen uniformly because Z is compactly contained in Ω . A contradiction argument applied to any sequence $z_n \in Z$, $\varepsilon_n \downarrow 0$ upgrades the pointwise limits to uniform ones. Finally, $\varepsilon^3 = o(\varepsilon)$ in $d = 3$ and $\varepsilon^2 = o(1/|\log \varepsilon|)$ in $d = 2$, proving $\varepsilon^d = o(\rho_\varepsilon)$. $\square$

C.5 Potential and corrector estimates

Lemma C.6 (L^2 size of the equilibrium potential). For $d \in \{2, 3\}$,

$$\|q_\varepsilon\|_{L^2(\Omega)} \leq C\rho_\varepsilon. \quad (84)$$

Proof. The coercive mixed-boundary operator has a positive Green kernel $G(x, y)$. For y in the fixed interior neighborhood U , its local singularity and regular part satisfy

$$0 \leq G(x, y) \leq C \begin{cases} 1 + |\log |x - y||, & d = 2, \\ |x - y|^{-1} + 1, & d = 3. \end{cases} \quad (85)$$

These bounds follow by freezing the uniformly elliptic principal coefficient, using the Laplace fundamental solution as a local parametrix, and applying the maximum principle plus interior elliptic estimates to the regular remainder. Both right sides are square-integrable in x on a bounded domain, uniformly for $y \in U$; hence $\sup_{y \in U} \|G(\cdot, y)\|_2 \leq C_G$.

By Eq. (64), q_ε is the unique weak solution of $Lq_\varepsilon = \mu_\varepsilon$ with homogeneous outer conditions, so

$$q_\varepsilon(x) = \int_K G(x, y) d\mu_\varepsilon(y). \quad (86)$$

Minkowski's integral inequality and $\mu_\varepsilon(K) = \rho_\varepsilon$ yield

$$\|q_\varepsilon\|_2 \leq \int_K \|G(\cdot, y)\|_2 d\mu_\varepsilon(y) \leq C_G \rho_\varepsilon. \quad (87)$$

$\square$

Lemma C.7 (State corrector). Let $w_\varepsilon := \tilde{u}_{\varepsilon, z} - u$ and define

$$r_\varepsilon := w_\varepsilon + (u(z) - h_z)q_\varepsilon. \quad (88)$$

Then

$$\|r_\varepsilon\|_a \leq C\varepsilon^{d/2}, \quad (89)$$

$$|a_\Omega(r_\varepsilon, p)| \leq C\varepsilon^{d/2}, \quad (90)$$

$$\|w_\varepsilon\|_{L^2(\Omega)} \leq C\rho_\varepsilon. \quad (91)$$

In particular, $a_\Omega(r_\varepsilon, p) = o(\rho_\varepsilon)$ uniformly on Z .

Proof. Step 1: the corrector equation and trace. The state equations on Ω_ε imply that w_ε is a -harmonic against every test function vanishing on both the outer Dirichlet boundary and ∂K . The same is true of q_ε by Eq. (62). Therefore r_ε is a -harmonic on Ω_ε . On the cavity boundary,

$$r_\varepsilon = h_z - u + (u(z) - h_z) = u(z) - u. \quad (92)$$

Step 2: a localized lift. Choose $R > R_B$ and $\chi \in C_c^\infty(B_{2R}(0))$ with $\chi = 1$ on $B_R(0)$. For sufficiently small ε , define

$$\eta_\varepsilon(x) := \chi((x - z)/\varepsilon)(u(z) - u(x)). \quad (93)$$

It has the trace in Eq. (92). Local C^1 regularity gives $|u(z) - u(x)| \leq C\varepsilon$ on its support and

$$\begin{aligned} |\nabla \eta_\varepsilon| &\leq \varepsilon^{-1} |\nabla \chi((x - z)/\varepsilon)| \cdot |u(z) - u(x)| \\ &\quad + |\chi((x - z)/\varepsilon)| \cdot |\nabla u(x)| \\ &\leq C. \end{aligned} \quad (94)$$

Because the support has volume $O(\varepsilon^d)$,

$$a_\Omega(\eta_\varepsilon, \eta_\varepsilon) \leq C\varepsilon^d. \quad (95)$$

Step 3: energy minimality. On Ω_ε , both r_ε and η_ε have the same boundary trace. Thus $r_\varepsilon - \eta_\varepsilon$ is an admissible homogeneous test function, and $a_{\Omega_\varepsilon}(r_\varepsilon, r_\varepsilon - \eta_\varepsilon) = 0$. Consequently

$$a_{\Omega_\varepsilon}(\eta_\varepsilon, \eta_\varepsilon) = a_{\Omega_\varepsilon}(r_\varepsilon, r_\varepsilon) + a_{\Omega_\varepsilon}(\eta_\varepsilon - r_\varepsilon, \eta_\varepsilon - r_\varepsilon). \quad (96)$$

Inside K , the extension of r_ε equals $u(z) - u(x)$ and has energy $O(\varepsilon^d)$. Combining this fact with Eq. (95) proves Eq. (89).

Step 4: adjoint pairing. Cauchy–Schwarz in the energy inner product gives

$$|a(r_\varepsilon, p)| \leq \|r_\varepsilon\|_a \|p\|_a \leq C\varepsilon^{d/2}, \quad (97)$$

which is Eq. (90).

Step 5: L^2 size of the state perturbation. From Eq. (88), Poincaré’s inequality, and Lemma C.6,

$$\|w_\varepsilon\|_2 \leq |u(z) - h_z| \|q_\varepsilon\|_2 + \|r_\varepsilon\|_2 \quad (98)$$

$$\leq C\rho_\varepsilon + C\varepsilon^{d/2}. \quad (99)$$

The capacity scales show $\varepsilon^{d/2} = o(\rho_\varepsilon)$ for $d = 2, 3$, proving Eq. (91). The same comparison applied to Eq. (90) proves the final statement. $\square$

C.6 Objective expansion

Proof of Theorem 3.5. Fix $z \in Z$ and abbreviate $a_z = u(z) - h_z$.

Step 1: Taylor expansion of the state objective. By Eq. (91), $w_\varepsilon \rightarrow 0$ in L^2 . The integral form of the second-order Fréchet remainder is

$$j(u + w_\varepsilon) - j(u) - j'(u)[w_\varepsilon] = \int_0^1 (1 - t) j''(u + tw_\varepsilon)[w_\varepsilon, w_\varepsilon] dt. \quad (100)$$

The uniform bound on j'' and Eq. (91) imply

$$|j(u + w_\varepsilon) - j(u) - j'(u)[w_\varepsilon]| \leq C\rho_\varepsilon^2. \quad (101)$$

Step 2: explicit domain cost. The only explicit domain change is

$$\left| \int_{\Omega_\varepsilon} \psi - \int_\Omega \psi \right| = \left| \int_K \psi \right| \leq \|\psi\|_\infty |B| \varepsilon^d. \quad (102)$$

Step 3: adjoint substitution. The extension w_ε has zero trace on Γ_D , hence belongs to V . By Eqs. (5) and (88),

$$j'(u)[w_\varepsilon] = a_\Omega(w_\varepsilon, p) = -a_z a_\Omega(q_\varepsilon, p) + a_\Omega(r_\varepsilon, p). \quad (103)$$

Step 4: localization of the capacitary charge. Applying Eq. (65) to the locally C^1 adjoint gives

$$a_\Omega(q_\varepsilon, p) = \rho_\varepsilon p(z) + e_{\varepsilon, z}, \quad |e_{\varepsilon, z}| \leq C\varepsilon\rho_\varepsilon. \quad (104)$$

Combining Eqs. (90), (103) and (104) yields

$$j'(u)[w_\varepsilon] = -\rho_\varepsilon(u(z) - h_z)p(z) + O(\varepsilon\rho_\varepsilon) + O(\varepsilon^{d/2}). \quad (105)$$

Step 5: collect and express the remainder in capacity units. Equations Eqs. (101), (102) and (105) give

$$\begin{aligned} & |\mathcal{J}(\Omega_{\varepsilon, z}) - \mathcal{J}(\Omega) + \rho_\varepsilon(u(z) - h_z)p(z)| \\ & \leq C \left(\rho_\varepsilon^2 + \varepsilon\rho_\varepsilon + \varepsilon^{d/2} + \varepsilon^d \right). \end{aligned} \quad (106)$$

For $d = 3$, $\rho_\varepsilon \asymp \varepsilon$, so every term on the right is bounded by $C\rho_\varepsilon^{3/2}$. For $d = 2$, $\rho_\varepsilon \asymp 1/|\log \varepsilon|$, and the elementary limit $\varepsilon|\log \varepsilon|^m \rightarrow 0$ for every fixed m shows that every term is bounded by $C\rho_\varepsilon^2$. This proves Eq. (11) with $\alpha_3 = 1/2$ and $\alpha_2 = 1$.

Step 6: identify the derivative. Divide Eq. (106) by $\rho_\varepsilon > 0$ and let $\varepsilon \downarrow 0$. The normalized remainder vanishes, so the limit in Eq. (10) exists and equals $-(u(z) - h_z)p(z)$. Every constant above is uniform on Z , proving uniform convergence and Eq. (12). $\square$

D Cross-Fitted Estimator Proofs

Fix a candidate z . Write

$$A_i := U_i - h_z, \quad B_i := P_i, \quad \mu_A := \mathbb{E}A_i, \quad \mu_B := \mathbb{E}B_i, \quad (107)$$

and $X_i = (A_i, B_i)$. Under the exact-oracle assumption, $\theta := -\mu_A\mu_B = D_T(z)$.

D.1 Algebra and unbiasedness

Lemma D.1 (Linear-time identity). *For every $N \geq 2$,*

$$\sum_{i \neq j} A_i B_j = \left(\sum_{i=1}^N A_i \right) \left(\sum_{j=1}^N B_j \right) - \sum_{i=1}^N A_i B_i. \quad (108)$$

Proof. Distributivity gives $(\sum_i A_i)(\sum_j B_j) = \sum_{i,j} A_i B_j$. The index set of this double sum is the disjoint union of the diagonal $\{i = j\}$ and the off-diagonal $\{i \neq j\}$. Subtracting the diagonal proves the identity. Thus the estimator needs three scalar sums and $O(N)$ arithmetic. $\square$

Proof of Theorem 5.1. Step 1: factor only across different sample pairs. For $i \neq j$, X_i and X_j are independent. Therefore A_i and B_j are independent even though A_i and B_i need not be, and

$$\mathbb{E}[-A_i B_j] = -(\mathbb{E}A_i)(\mathbb{E}B_j) = -\mu_A \mu_B = \theta. \quad (109)$$

Step 2: average the ordered terms. There are exactly $N(N-1)$ ordered pairs (i, j) with $i \neq j$. Linearity of expectation and Eq. (109) give

$$\mathbb{E}\widehat{D}_{TN} = \frac{1}{N(N-1)} \sum_{i \neq j} \theta = \theta. \quad (110)$$

Finally, Eqs. (12) and (14) imply $\theta = -(u(z) - h_z)p(z) = D_T(z)$. At no point was $\mathbb{E}[A_i B_i]$ factored, so within-pair dependence is unrestricted. $\square$

D.2 Exact variance decomposition

Let $A_i^\circ = A_i - \mu_A$ and $B_i^\circ = B_i - \mu_B$. Recall the symmetric kernel

$$H(X_i, X_j) = -\frac{1}{2}(A_i B_j + A_j B_i), \quad (111)$$

for which

$$\widehat{D}_{TN} = \binom{N}{2}^{-1} \sum_{1 \leq i < j \leq N} H(X_i, X_j). \quad (112)$$

Lemma D.2 (Hoeffding projections). *The first-order and degenerate projections are*

$$h_1(X_i) := -\frac{1}{2}(\mu_B A_i^\circ + \mu_A B_i^\circ), \quad (113)$$

$$h_2(X_i, X_j) := -\frac{1}{2}(A_i^\circ B_j^\circ + A_j^\circ B_i^\circ). \quad (114)$$

They obey

$$\begin{aligned} \mathbb{E}h_1(X_i) &= 0, \\ \mathbb{E}[h_2(X_i, X_j) \mid X_i] &= 0, \\ \mathbb{E}[h_2(X_i, X_j) \mid X_j] &= 0. \end{aligned} \quad (115)$$

Proof. Because X_j is independent of X_i ,

$$\mathbb{E}[H(X_i, X_j) \mid X_i] = -\frac{1}{2}(A_i \mu_B + B_i \mu_A). \quad (116)$$

Subtracting $\theta = -\mu_A \mu_B$ and collecting centered terms gives Eq. (113). Next expand

$$\begin{aligned} H(X_i, X_j) &= -\mu_A \mu_B - \frac{1}{2}\{\mu_B(A_i^\circ + A_j^\circ) + \mu_A(B_i^\circ + B_j^\circ)\} \\ &\quad - \frac{1}{2}(A_i^\circ B_j^\circ + A_j^\circ B_i^\circ). \end{aligned} \quad (117)$$

Subtracting $\theta + h_1(X_i) + h_1(X_j)$ proves Eq. (114). Centering and independence across $i \neq j$ give all three identities in Eq. (115). $\square$

Proof of Theorem 5.2. Step 1: decompose the U-statistic. Substitute $H = \theta + h_1(X_i) + h_1(X_j) + h_2(X_i, X_j)$ into Eq. (112). Every $h_1(X_i)$ appears in $N - 1$ unordered pairs, and $\binom{N}{2} = N(N - 1)/2$, hence

$$\widehat{D}_{TN} - \theta = \frac{2}{N} \sum_{i=1}^N h_1(X_i) + \binom{N}{2}^{-1} \sum_{i < j} h_2(X_i, X_j). \quad (118)$$

Step 2: prove orthogonality. For any k and pair $i < j$, the covariance $\mathbb{E}[h_1(X_k)h_2(X_i, X_j)]$ is zero. If k is distinct from both indices, independence and centering suffice. If $k = i$, condition on X_i and use $\mathbb{E}[h_2 \mid X_i] = 0$; the case $k = j$ is identical.

Two different degenerate terms are also uncorrelated. Disjoint pairs are independent. If (i, j) and (i, k) share one index, then conditioning on X_i and using independence of X_j, X_k gives zero by degeneracy. Thus only identical unordered pairs contribute to the variance of the second sum.

Step 3: take variances. With $\zeta_1 = \text{Var}(h_1(X_1))$ and $\zeta_2 = \text{Var}(h_2(X_1, X_2))$, orthogonality gives

$$\text{Var}(\widehat{D}_{TN}) = \frac{4}{N^2} N \zeta_1 + \binom{N}{2}^{-2} \binom{N}{2} \zeta_2 \quad (119)$$

$$= \frac{4}{N} \zeta_1 + \frac{2}{N(N-1)} \zeta_2. \quad (120)$$

Step 4: compute ζ_1 . Squaring [Eq. \(113\)](#) and taking expectations gives

$$\zeta_1 = \frac{1}{4} (\mu_B^2 \text{Var } A + \mu_A^2 \text{Var } B + 2\mu_A\mu_B \text{Cov}(A, B)), \quad (121)$$

which is [Eq. \(21\)](#).

Step 5: compute ζ_2 . Let $c_{AB} = \mathbb{E}[A_1^\circ B_1^\circ]$. From [Eq. \(114\)](#), independence across sample pairs yields

$$\mathbb{E}h_2^2 = \frac{1}{4} \mathbb{E}[(A_1^\circ B_2^\circ)^2 + (A_2^\circ B_1^\circ)^2 + 2A_1^\circ B_1^\circ A_2^\circ B_2^\circ] \quad (122)$$

$$= \frac{1}{4} (2 \text{Var } A \text{Var } B + 2c_{AB}^2), \quad (123)$$

which is [Eq. \(22\)](#). The second-moment assumptions make all quantities finite. Since A, B , and H are evaluated on the unperforated domain, no ε occurs anywhere in [Eq. \(120\)](#). $\square$

D.3 Finite-sample concentration

Proof of [Theorem 5.3](#). Let $m = \lfloor N/2 \rfloor$ and let Π be a uniformly random permutation of $\{1, \dots, N\}$, independent of the data. Define the disjoint-pair average

$$V_\Pi := \frac{1}{m} \sum_{k=1}^m H(X_{\Pi(2k-1)}, X_{\Pi(2k)}). \quad (124)$$

Step 1: recover the U-statistic by averaging pairings. A fixed unordered pair $\{i, j\}$ is selected in one of the m slots with probability $2m/[N(N-1)]$. Its coefficient in $\mathbb{E}_\Pi[V_\Pi \mid X_1, \dots, X_N]$ is therefore

$$\frac{1}{m} \frac{2m}{N(N-1)} = \binom{N}{2}^{-1}. \quad (125)$$

Using [Eq. \(112\)](#),

$$\mathbb{E}_\Pi[V_\Pi \mid X_1, \dots, X_N] = \widehat{D}_{TN}. \quad (126)$$

Step 2: bound the moment-generating function. For every $\lambda > 0$, convexity of the exponential and [Eq. \(126\)](#) give

$$\mathbb{E} \exp\{\lambda(\widehat{D}_{TN} - \theta)\} \leq \mathbb{E} \mathbb{E}_\Pi \exp\{\lambda(V_\Pi - \theta)\}. \quad (127)$$

Conditional on a fixed permutation, the m kernel values in [Eq. \(124\)](#) use disjoint independent sample pairs. Each has mean θ and lies in $[-M, M]$ because $|A_i| \leq M_U$ and $|B_i| \leq M_P$. Hoeffding's lemma (derived explicitly for a bounded interval in Step 3 of the proof of [Proposition N.4](#)) therefore gives

$$\mathbb{E} \exp\{\lambda(V_\Pi - \theta)\} \leq \exp\left(\frac{\lambda^2 M^2}{2m}\right). \quad (128)$$

Step 3: Chernoff optimization. Markov's inequality and [Eq. \(128\)](#) imply

$$\mathbb{P}(\widehat{D}_{TN} - \theta \geq t) \leq \exp\left(-\lambda t + \frac{\lambda^2 M^2}{2m}\right). \quad (129)$$

The minimizing value is $\lambda = mt/M^2$. Substitution produces

$$\exp\left(-\frac{mt^2}{2M^2}\right). \quad (130)$$

Applying the same argument to $-\widehat{D}_{TN}$ and adding the two tail probabilities proves [Eq. \(23\)](#).

Step 4: invert the tail and cover multiple candidates. Setting the right side of [Eq. \(23\)](#) equal to δ gives [Eq. \(24\)](#). For candidate events $F_1, \dots, F_K$, each with failure probability at most δ/K , Boole's inequality gives $\mathbb{P}(\cup_k F_k) \leq \sum_k \mathbb{P}(F_k) \leq \delta$. This argument never assumes independence across candidate locations. $\square$

D.4 Marginal bias propagation

Proof of Proposition 5.7. Put $\bar{u} = \mathbb{E}U = u + e_U$ and $\bar{p} = \mathbb{E}P = p + e_P$, where $|e_U| \leq \beta_U$ and $|e_P| \leq \beta_P$. The same cross-index factorization as in Eq. (109) gives

$$\mathbb{E}\widehat{D}_{TN} = -(\bar{u} - h_z)\bar{p}. \quad (131)$$

Subtracting $D_T = -(u - h_z)p$ and expanding yields the exact identity

$$\mathbb{E}\widehat{D}_{TN} - D_T = -pe_U - (u - h_z)e_P - e_Ue_P. \quad (132)$$

The triangle inequality and the marginal bounds prove Eq. (33). $\square$

E Finite-Difference Lower Bound and Complexity

Proof of Theorem 5.5. Step 1: lower-bound the per-path second moment. By Eq. (27),

$$\mathbb{E}Y_\varepsilon^2 = \mathbb{E}[Y_\varepsilon^2 \mathbf{1}_{A_\varepsilon}] = \mathbb{P}(A_\varepsilon)\mathbb{E}[Y_\varepsilon^2 \mid A_\varepsilon] \quad (133)$$

$$\geq c_p c_2 \rho_z(\varepsilon). \quad (134)$$

Step 2: control the squared mean. By Eq. (28),

$$(\mathbb{E}Y_\varepsilon)^2 \leq C_1^2 \rho_z(\varepsilon)^2. \quad (135)$$

If $\rho_z(\varepsilon) \leq c_p c_2 / (2C_1^2)$, this is at most $\frac{1}{2}c_p c_2 \rho_z(\varepsilon)$. Subtracting it from Eq. (134) gives

$$\text{Var}(Y_\varepsilon) \geq \frac{1}{2}c_p c_2 \rho_z(\varepsilon). \quad (136)$$

Step 3: scale the independent sample mean. Independence of the N replicates gives

$$\text{Var}(\widehat{D}_{TN,\varepsilon}^{\text{FD}}) = \frac{1}{N\rho_z(\varepsilon)^2} \text{Var}(Y_\varepsilon) \quad (137)$$

$$\geq \frac{c_p c_2}{2N\rho_z(\varepsilon)}, \quad (138)$$

which is Eq. (29). $\square$

Proof of Corollary 5.6. Apply Eq. (24) to each candidate with failure probability δ/K . The common radius is at most Δ precisely when

$$M\sqrt{\frac{2\log(2K/\delta)}{\lfloor N/2 \rfloor}} \leq \Delta, \quad (139)$$

which rearranges to Eq. (30). The union bound in Theorem 5.3 then gives simultaneous error at most Δ . For finite differencing, Eq. (29) and the requirement $\text{Var}(\widehat{D}_{TN,\varepsilon}^{\text{FD}}) \leq \Delta^2$ imply Eq. (31) by direct rearrangement. $\square$

F Certification and Adaptive-Selection Proofs

Proof of Theorem A.2. Let

$$E := \{|\widehat{D}_{TN} - \mathbb{E}\widehat{D}_{TN}| \leq r_N(\delta)\}. \quad (140)$$

By hypothesis, $\mathbb{P}(E) \geq 1 - \delta$.

Step 1: upper-bound the true derivative on E . The confidence event and Eq. (34) give

$$D_T(z) \leq \mathbb{E}\widehat{D}_{TN}(z) + b_{\text{term}} + b_{\text{cache}} \quad (141)$$

$$\leq \widehat{D}_{TN}(z) + r_N(\delta) + b_{\text{term}} + b_{\text{cache}}. \quad (142)$$

Step 2: insert the deterministic remainder. Since $\rho_z(\varepsilon) > 0$, the upper side of Eq. (37) and Eq. (142) imply

$$\mathcal{J}(\Omega_{\varepsilon,z}) - \mathcal{J}(\Omega) \leq \rho_z(\varepsilon) \{D_T(z) + C_z \rho_z(\varepsilon)^\alpha\} \quad (143)$$

$$\leq \rho_z(\varepsilon) \{\widehat{D}_{TN} + r_N + b_{\text{term}} + b_{\text{cache}} + C_z \rho_z(\varepsilon)^\alpha\}. \quad (144)$$

The braced term is negative by Eq. (39); hence the objective change is negative on E . The event has probability at least $1 - \delta$. $\square$

Proof of Corollary A.3. The event E in the preceding proof is independent of which radius is chosen. On E , Eq. (144) holds simultaneously for every $\varepsilon \in (0, \varepsilon_0]$ because the remainder assumption is uniform over that interval. Therefore it also holds at the measurable random value $\widehat{\varepsilon}$. If Eq. (39) holds there, the right side is negative. Thus failure can occur only on E^c , whose probability is at most δ . Since one event covers the continuum of radii, no discretization or union bound is required. $\square$

Proof of Proposition A.7. Let B_t be the event that round t accepts a non-improving insertion. The candidate and radius are $\mathcal{F}_{t-1}$ -measurable. Conditional on that history, the proof of Theorem A.2 shows that an accepted insertion is improving whenever E_t occurs. Hence

$$B_t \subseteq E_t^c \quad \text{and} \quad \mathbb{P}(B_t \mid \mathcal{F}_{t-1}) \leq \mathbb{P}(E_t^c \mid \mathcal{F}_{t-1}) \leq \delta_t. \quad (145)$$

Taking expectations and using the tower property gives $\mathbb{P}(B_t) \leq \delta_t$. Finally, countable subadditivity yields

$$\mathbb{P}\left(\bigcup_{t \geq 1} B_t\right) \leq \sum_{t \geq 1} \mathbb{P}(B_t) \leq \sum_{t \geq 1} \delta_t, \quad (146)$$

which is Eq. (44). No independence between validation batches is used; only the conditional guarantee is required. $\square$

G Source Integration and Connected Refinement

G.1 Poisson source integration

For completeness, the inhomogeneous three-dimensional Poisson implementation used in our engineering benchmark samples the Green source term without a volume grid. The following identity also fixes the normalization used in our implementation.

Proposition G.1 (Unbiased source sample on a three-dimensional WoS ball). *Let $B_R(x) \subset \Omega$, let Θ be uniform on $\mathbb{S}^2$, and let T be independent of Θ with density $p_T(t) = 6t(1-t)$ on $(0,1)$. For every measurable source f with $\int_{B_R(x)} G_{B_R(x)}(x,y)|f(y)| dy < \infty$,*

$$\int_{B_R(x)} G_{B_R(x)}(x,y)f(y) dy = \frac{R^2}{6} \mathbb{E}[f(x + RT\Theta)], \quad (147)$$

where $G_{B_R(x)}(x,y) = (4\pi)^{-1}(\|y-x\|^{-1} - R^{-1})$ is the Dirichlet Green kernel evaluated at the ball center.

Proof. Denote the left-hand side of Eq. (147) by $I_R(x)$ and write $y = x + r\theta$. Spherical coordinates, followed by $r = Rt$, give

$$\begin{aligned} I_R(x) &= \frac{1}{4\pi} \int_{\mathbb{S}^2} \int_0^R \left(r - \frac{r^2}{R}\right) f(x + r\theta) dr d\theta \\ &= \frac{R^2}{4\pi} \int_{\mathbb{S}^2} \int_0^1 t(1-t) f(x + Rt\theta) dt d\theta. \end{aligned}$$

Uniform Θ has surface density $(4\pi)^{-1}$ and T has density $6t(1-t)$; substituting both densities into the expectation on the right of Eq. (147) gives exactly the last integral. Hence one draw of (T, Θ) yields an unbiased contribution $R^2 f(x + RT\Theta)/6$. $\square$

G.2 Connected-channel parameterization and refinement

Topology creation need not end with disconnected cavities. Let $q = (q_0, \dots, q_m) \in (\mathbb{R}^3)^{m+1}$ be control points and let a fixed subdivision matrix $W \in \mathbb{R}^{L \times (m+1)}$ have nonnegative entries and row sums one. Dense centers, the removed channel, and the retained domain are

$$\begin{aligned} s_i(q) &= \sum_{j=0}^m W_{ij} q_j, \\ H(q, r) &= \bigcup_{i=1}^L \overline{B}_r(s_i(q)), \quad \Omega(q, r) = \Omega \setminus H(q, r). \end{aligned} \tag{148}$$

Every accepted trial must satisfy $\|s_{i+1} - s_i\| < 2r$ for all consecutive centers, checked again after radius recalibration. Thus the interiors of successive balls overlap, making both the removed union and its open interior path connected. The topology stage selects two separated terminals and connects them through a topological-score-weighted candidate graph; Differential WoS then moves the interior controls while the end-points remain fixed. The current housing experiment keeps every ball strictly inside the housing; its terminals are interior design points. It therefore demonstrates an internally connected cavity, not an inlet-to-outlet conduit or a guarantee of powder evacuation. In particular, it would remain a void island under the exterior-reachability definition of Andrade et al. [3]. The two constraints must not be treated as interchangeable baselines.

Let Σ be the union of sphere-intersection seams on ∂H . For an exposed patch Γ_i and offset $\ell > 0$, define the admissible part of that patch by

$$\Gamma_{i,\ell}^{\text{safe}} = \{x \in \Gamma_i : \text{dist}(x, \Sigma) > \ell, (x, x - \ell n(x)) \subset \Omega(q, r)\}. \tag{149}$$

The Euclidean seam distance is used here; replacing the unit strip width by any fixed positive multiple changes only constants.

Proposition G.2 (Differential WoS center derivative). *Fix $r > 0$ and a regular channel configuration for which the exposed spherical patches $\Gamma_i \subset \partial H(q, r)$ are C^2 away from Σ . Assume the piecewise-smooth solution map is shape differentiable under independent center translations. Moreover, for every compact probe set $K \Subset \Omega(q, r)$, assume there are $C_K, \ell_K > 0$ such that, for $0 < \ell \leq \ell_K$ and $\xi = X_\tau$,*

$$\sup_{y \in K} \mathbb{E}_y \left[\mathbf{1}_{\{I(\xi)=i, \xi \notin \Gamma_{i,\ell}^{\text{safe}}\}} |\partial_n u(\xi)| \right] \leq C_K \ell, \tag{150}$$

$$\sup_{x \in \Gamma_{i,\ell}^{\text{safe}}} \left| \frac{u(x - \ell n(x))}{\ell} + \partial_n u(x) \right| \leq C_K \ell. \tag{151}$$

Let u solve $-\Delta u = f$ in $\Omega(q, r)$ with zero Dirichlet data. For $y \in \Omega(q, r)$, let X_τ be Brownian exit, $I(X_\tau)$ its exposed-ball index, and n the outward normal of the retained domain. If U_x is an unbiased independent WoS estimate of $u(x)$, then

$$G_{i,\ell}(y) = \mathbf{1}_{\{I(X_\tau)=i, X_\tau \in \Gamma_{i,\ell}^{\text{safe}}\}} n(X_\tau) \frac{U_{X_\tau - \ell n(X_\tau)}}{\ell} \tag{152}$$

satisfies, uniformly on compact probe sets,

$$\mathbb{E} G_{i,\ell}(y) = \nabla_{s_i} u(y) + O(\ell) \quad \text{as } \ell \downarrow 0. \tag{153}$$

For any fixed finite set of interior probes, the weighted estimate has $O(\ell)$ bias for the corresponding weighted derivative sum. To assert this rate for compliance $J(q, r) = \int_{\Omega(q,r)} f u \, dx$, additionally assume

that differentiation under the domain integral is valid and that, for all sufficiently small ℓ , the pointwise estimator error is bounded by $\ell C(y)$ with $\int_{\Omega(q,r)} |f(y)|C(y) dy < \infty$. A finite probe rule then adds its separate quadrature error.

Proposition G.3 (Control chain rule and length tangency). Suppose J is differentiable with respect to the dense centers at a regular configuration, and write $G_i = \nabla_{s_i} J$. Then

$$\nabla_{q_j} J = \sum_{i=1}^L W_{ij} G_i. \quad (154)$$

For the centerline length $\mathcal{L}(q) = \sum_{j=0}^{m-1} \|q_{j+1} - q_j\|$, let c be its gradient with the endpoint blocks set to zero. Given any endpoint-fixed proposal a , define

$$d = a - c \frac{\langle c, a \rangle}{\|c\|^2} \quad (c \neq 0), \quad (155)$$

and $d = a$ if $c = 0$. Then the endpoints do not move and $D\mathcal{L}(q)[d] = 0$. In particular, let g be the objective gradient with its endpoint blocks set to zero. Choosing $a = -g$ gives $DJ(q)[d] = -\|P_{c^\perp} g\|^2 \leq 0$. Our implementation first smooths a and therefore relies on fresh line-search validation, rather than this last inequality, for accepted descent.

The radius of every trial is recalibrated on one fixed Sobol set. If $Y_1, \dots, Y_M$ fill a box of volume V_B , define

$$\hat{V}_M(q, r) = \frac{V_B}{M} \sum_{k=1}^M \mathbf{1}_{\{Y_k \in \Omega, \min_i \|Y_k - s_i(q)\| \leq r\}}. \quad (156)$$

Proposition G.4 (Discrete equal-volume calibration). Let k be the nearest integer to $MV_\star/V_B$. If the k th and $(k+1)$ st eligible center distances are distinct, choosing r strictly between them gives exactly k removed samples and

$$|\hat{V}_M(q, r) - V_\star| \leq \frac{V_B}{2M}. \quad (157)$$

Thus all compared connected channels have the same discrete volume, up to one predeclared rounding quantum; this is a statement about the fixed QMC measure, not an exact analytic union-volume formula.

Proposition G.5 (Fresh paired shape acceptance). At shape round t , condition on the complete history $\mathcal{F}_{t-1}$, including the Differential WoS direction and finitely many trial scales. Let $\Delta_{t,k} = J(q_{t,k}, r_{t,k}) - J(q_t, r_t)$, where every $r_{t,k}$ obeys Eq. (156). Suppose simultaneous upper bounds satisfy

$$\mathbb{P}(\Delta_{t,k} \leq U_{t,k} \text{ for every } k \mid \mathcal{F}_{t-1}) \geq 1 - \delta_t. \quad (158)$$

If a trial is accepted only when its $U_{t,k} < 0$, then

$$\mathbb{P}(\text{some accepted shape step is non-improving}) \leq \sum_t \delta_t. \quad (159)$$

The conclusion is independent of bias in the proposed direction: coverage of the fresh objective difference, not gradient accuracy, certifies acceptance.

H Connected-Channel Refinement Proofs

Proof of Proposition G.2. We prove the vector identity by testing it against an arbitrary translation $v \in \mathbb{R}^3$ of ball i .

Step 1: derive the boundary condition for the shape derivative. On the interior of the exposed patch Γ_i , translation of center i has velocity $V = v$; every other open spherical patch and the outer boundary have $V = 0$. The proposition's shape-differentiability assumption permits patchwise trace differentiation; the seam itself is excluded. Write u_t for the perturbed solution and $\dot{u}$ for its Eulerian shape derivative at $t = 0$. Differentiating the identity $u_t(T_t(x)) = 0$ for $x \in \Gamma_i$ gives

$$\dot{u}(x) + \nabla u(x) \cdot V(x) = 0. \quad (160)$$

The trace of u is constant on Γ_i , so its tangential gradient vanishes there. Hence $\nabla u = (\partial_n u)n$ and

$$\dot{u} = -(n \cdot v)\partial_n u \quad \text{on } \Gamma_i. \quad (161)$$

The source f and the outer boundary are fixed. Differentiating the weak Poisson equation with test functions compactly supported in the current domain therefore gives

$$-\Delta \dot{u} = 0 \quad \text{in } \Omega(q, r), \quad (162)$$

with zero trace on all stationary Dirichlet patches. Equations Eqs. (161) and (162) are the standard differential boundary-value problem, here obtained directly from the moving trace.

Step 2: approximate the normal derivative on the safe patch. For $x \in \Gamma_{i,\ell}^{\text{safe}}$, the segment in Eq. (149) lies in the retained domain. Whenever a uniform C^2 extension is available, Taylor's theorem along this segment and $u(x) = 0$ give

$$u(x - \ell n) = u(x) - \ell \partial_n u(x) + \frac{\ell^2}{2} n^T \nabla^2 u(x - \theta \ell n) n \quad (163)$$

$$= -\ell \partial_n u(x) + O(\ell^2), \quad (164)$$

for some $\theta \in (0, 1)$. Uniform C^2 bounds on the regular exposed patches imply

$$-\partial_n u(x) = \frac{u(x - \ell n(x))}{\ell} + O(\ell) \quad (165)$$

there. In the proposition we use the explicitly stated uniform version Eq. (151); this separates the needed analytic estimate from the merely qualitative fact that Σ has zero harmonic measure. Combining Eqs. (161) and (165) yields

$$\dot{u}(x) = (n(x) \cdot v) \frac{u(x - \ell n(x))}{\ell} + O(\ell \|v\|). \quad (166)$$

Step 3: truncate the harmonic exit representation. Let X_τ denote Brownian exit from $\Omega(q, r)$, and abbreviate $\xi = X_\tau$ and $n_\tau = n(X_\tau)$. The harmonic representation of Eq. (162) first gives the exact identity

$$D_{s_i} u(y)[v] = \mathbb{E}_y [\mathbf{1}_{\{I(\xi)=i\}} (-(n_\tau \cdot v) \partial_n u(\xi))] . \quad (167)$$

Write $S_{i,\ell} = \{I(\xi) = i, \xi \in \Gamma_{i,\ell}^{\text{safe}}\}$. On $S_{i,\ell}$, Eq. (151) bounds the error made by replacing $-\partial_n u(\xi)$ with $u(\xi - \ell n_\tau)/\ell$ by $C_K \ell$. Because harmonic measure has total mass at most one, its contribution is at most $C_K \ell \|v\|$. On $\{I(\xi) = i\} \setminus S_{i,\ell}$, Eq. (150) bounds the discarded part of Eq. (167) by the same order. Therefore

$$D_{s_i} u(y)[v] = \mathbb{E}_y \left[\mathbf{1}_{S_{i,\ell}} (n_\tau \cdot v) \frac{u(\xi - \ell n_\tau)}{\ell} \right] + O_K(\ell \|v\|). \quad (168)$$

This is the precise point at which the quantitative strip assumption, rather than only measure zero of the limiting seam, is required.

Step 4: insert the nested WoS estimate. Condition on X_τ on the safe event $S_{i,\ell}$. By hypothesis,

$$\mathbb{E}[U_{X_\tau - \ell n(X_\tau)} \mid X_\tau] = u(X_\tau - \ell n(X_\tau)). \quad (169)$$

Taking iterated expectations in Eq. (152) and comparing with Eq. (168) shows that, for every v ,

$$v \cdot \mathbb{E}G_{i,\ell}(y) = D_{s_i} u(y)[v] + O(\ell \|v\|). \quad (170)$$

Duality in the finite-dimensional space $\mathbb{R}^3$ proves Eq. (153).

Step 5: differentiate compliance. Reynolds' transport formula gives

$$D_{s_i} J[v] = \int_{\Omega(q,r)} f D_{s_i} u[v] dx + \int_{\partial\Omega(q,r)} f u(V \cdot n) dS. \quad (171)$$

The boundary integral vanishes because $u = 0$ on every moving and fixed Dirichlet patch. For a fixed finite interior probe set, compact-set uniformity in Eq. (153) directly bounds the weighted sum of estimator errors by $O(\ell)$. For the continuum objective, the additional integrable-envelope hypothesis gives

$$\begin{aligned} & \left\| \int_{\Omega(q,r)} f(y) (\mathbb{E}G_{i,\ell}(y) - \nabla_{s_i} u(y)) dy \right\| \\ & \leq \ell \int_{\Omega(q,r)} |f(y)| C(y) dy. \end{aligned}$$

This proves the continuum rate under the added hypothesis. Compact-set bounds alone would not control probes approaching the boundary. Replacing the integral by a fixed weighted rule adds the separate quadrature error. $\square$

Proof of Proposition G.3. Step 1: differentiate the dense centers. For a control perturbation $\dot{q} = (\dot{q}_0, \dots, \dot{q}_m)$, linearity of Eq. (148) gives

$$\dot{s}_i = \sum_{j=0}^m W_{ij} \dot{q}_j. \quad (172)$$

Step 2: apply the ordinary finite-dimensional chain rule. At the assumed differentiability point,

$$DJ(q)[\dot{q}] = \sum_{i=1}^L G_i \cdot \dot{s}_i = \sum_{i=1}^L G_i \cdot \sum_{j=0}^m W_{ij} \dot{q}_j \quad (173)$$

$$= \sum_{j=0}^m \left(\sum_{i=1}^L W_{ij} G_i \right) \cdot \dot{q}_j. \quad (174)$$

Since this holds for every $\dot{q}$, comparison with the definition of the Euclidean gradient proves Eq. (154).

Step 3: compute the length gradient. Put $e_j = (q_{j+1} - q_j) / \|q_{j+1} - q_j\|$ and assume no consecutive controls coincide. Direct differentiation yields

$$\begin{aligned} \nabla_{q_0} \mathcal{L} &= -e_0, \\ \nabla_{q_j} \mathcal{L} &= e_{j-1} - e_j \quad (1 \leq j < m), \\ \nabla_{q_m} \mathcal{L} &= e_{m-1}. \end{aligned} \quad (175)$$

For endpoint-fixed directions only the interior blocks contribute, so replacing the two endpoint blocks by zero gives the vector c in the proposition and $D\mathcal{L}(q)[d] = \langle c, d \rangle$.

Step 4: verify tangency. When $c \neq 0$, substitution of Eq. (155) gives

$$\langle c, d \rangle = \langle c, a \rangle - \|c\|^2 \frac{\langle c, a \rangle}{\|c\|^2} = 0. \quad (176)$$

When $c = 0$, the same conclusion is immediate. Both a and c have zero endpoint blocks, hence so does d .

Step 5: verify descent for the unsmoothed projected gradient. Let g denote $\nabla_q J$ with its endpoint blocks set to zero and let $P_{c^\perp}$ be orthogonal projection within that endpoint-fixed subspace. Choosing $a = -g$ makes $d = -P_{c^\perp} g$. Therefore

$$DJ(q)[d] = \langle g, d \rangle = -\langle g, P_{c^\perp} g \rangle = -\|P_{c^\perp} g\|^2 \leq 0. \quad (177)$$

The final equality uses symmetry and idempotence of an orthogonal projection. No descent assertion follows after an arbitrary smoothing operator; this is why the implemented smoothed proposal is subjected to fresh evaluation. $\square$

Proof of Proposition G.4. For each sample define

$$d_j = \begin{cases} \min_i \|Y_j - s_i(q)\|, & Y_j \in \Omega, \\ +\infty, & Y_j \notin \Omega. \end{cases} \quad (178)$$

Let $d_{(1)} \leq \dots \leq d_{(M)}$ be the order statistics. If $d_{(k)} < r < d_{(k+1)}$, exactly the first k indicators in Eq. (156) equal one. Hence

$$\hat{V}_M(q, r) = \frac{kV_B}{M}. \quad (179)$$

Since k is a nearest integer to $MV_\star/V_B$,

$$\left| k - \frac{MV_\star}{V_B} \right| \leq \frac{1}{2}. \quad (180)$$

Multiplication by V_B/M proves Eq. (157). The distinctness assumption prevents ambiguity at the threshold. If ties occur, a deterministic tie rule retains a one-quantum bound but need not attain exactly k samples. $\square$

Proof of Proposition G.5. Let E_t be the simultaneous-coverage event inside Eq. (158), and let B_t be the event that round t accepts a non-improving trial. On E_t , every accepted index k satisfies

$$\Delta_{t,k} \leq U_{t,k} < 0. \quad (181)$$

Consequently $B_t \subseteq E_t^c$. Conditional coverage and the tower property imply

$$\mathbb{P}(B_t) = \mathbb{E}[\mathbb{P}(B_t \mid \mathcal{F}_{t-1})] \quad (182)$$

$$\leq \mathbb{E}[\mathbb{P}(E_t^c \mid \mathcal{F}_{t-1})] \leq \delta_t. \quad (183)$$

Finally, countable subadditivity gives

$$\mathbb{P}\left(\bigcup_t B_t\right) \leq \sum_t \mathbb{P}(B_t) \leq \sum_t \delta_t, \quad (184)$$

which is Eq. (159). The direction, tested scales, common random numbers within each paired difference, and deterministic radius recalibration are all part of the conditioning history. No correctness assumption about the Differential WoS proposal is used after that conditioning. $\square$

I Reproducibility and Interpretation

Detailed protocols appear in Appendices X and Y; the count-radius study retains all 72 designs and separates design-seed ranges from conditional evaluation uncertainty. Rendering cuts and scalar smoothing affect only visualization, not the physical domain or reported objectives.

The deterministic capacity expansion uses the analytic hypotheses stated in the main text. Cross-fitted unbiasedness assumes exact marginal oracle means; implemented termination error remains separate. The finite-insertion and shape-acceptance results control erroneous acceptances under their coverage premises, rather than giving posterior probabilities conditional on an accepted design. No application result is presented as a proof that its finite-size, oracle-bias, or interval assumptions hold.

J Independent Randomizations for Accelerated Oracles

Proposition J.1 (Cross-fitting independent accelerated batches). *Fix a candidate, a deterministic geometry, and $R \geq 2$. For $r = 1, \dots, R$, independently generate a full randomized estimator pair (A_r, B_r) with the same law, means (a, b) , and finite second moments. Paths within a batch can have arbitrary dependence. Define*

$$T_R = -\frac{1}{R(R-1)} \sum_{r \neq s} A_r B_s. \quad (185)$$

Then $\mathbb{E}T_R = -ab$, and [Theorems 5.2](#) and [5.3](#) apply with batch count R and the corresponding batch moments and almost-sure bounds.

Proof. For $r \neq s$, the sigma-algebras of the two complete randomizations are independent. Thus A_r and B_s are independent and $\mathbb{E}[A_r B_s] = (\mathbb{E}A_r)(\mathbb{E}B_s) = ab$. There are $R(R-1)$ ordered pairs; linearity of expectation gives $\mathbb{E}T_R = -ab$. Let $H_{rs} = -(A_r B_s + A_s B_r)/2$. Then $T_R = \binom{R}{2}^{-1} \sum_{r < s} H_{rs}$, an order-two U-statistic of the independent batch pairs. Its first and degenerate second projections are exactly those in [Appendix D](#), with the law of (A, B) now the batch law. Substitution proves the variance identity. If $|A_r| \leq M_A$ and $|B_r| \leq M_B$, then $|H_{rs}| \leq M_A M_B$. Each disjoint pairing has $\lfloor R/2 \rfloor$ independent kernels; the existing pairing/MGF proof therefore gives the stated concentration bound. No independence within a batch is used at any step. $\square$

What fails if paths are treated as independent. For arbitrary dependent pairs with fixed marginal means a, b and integrable products, direct expansion instead gives

$$\mathbb{E}T_N = -\frac{1}{N(N-1)} \sum_{i \neq j} (ab + \text{Cov}(A_i, B_j)) \quad (186)$$

$$= -ab - \frac{1}{N(N-1)} \sum_{i \neq j} \text{Cov}(A_i, B_j). \quad (187)$$

Deleting the diagonal removes none of these cross-index covariances. An exact example uses the unit disk, $u(z) = p(z) = \text{Re } z$, and $z = 0$. For a random uniform angle Θ , take $A_j = B_j = \cos(\Theta + 2\pi j/N)$, $j = 0, \dots, N-1$, $N \geq 3$. Every sample has the correct harmonic marginal, but

$$\sum_j A_j = \text{Re} \left(e^{i\Theta} \sum_j e^{2\pi i j/N} \right) = 0, \quad (188)$$

$$\sum_j A_j^2 = \frac{N}{2} + \frac{1}{2} \text{Re} \left(e^{2i\Theta} \sum_j e^{4\pi i j/N} \right) = \frac{N}{2}, \quad (189)$$

$$T_N = -\frac{(\sum_j A_j)^2 - \sum_j A_j^2}{N(N-1)} \quad (190)$$

$$= \frac{1}{2(N-1)} \neq 0 = -u(0)p(0). \quad (191)$$

Both geometric sums vanish since their ratios are nontrivial N th roots of unity. Cross-fitting independently randomized batch averages instead returns zero exactly here. This counterexample establishes a failure of the iid formula; it is not a lower bound for general Array-RQMC performance.

A common random cache is a separate independence hazard. Suppose the pairs are iid only conditional on a cache C , with conditional means $a(C), b(C)$. Conditional independence first gives $\mathbb{E}[T_N | C] = -a(C)b(C)$. Taking expectation again yields

$$\mathbb{E}T_N = -\mathbb{E}[a(C)]\mathbb{E}[b(C)] - \text{Cov}(a(C), b(C)). \quad (192)$$

Thus marginally unbiased field estimators alone do not suffice. For example, $A_i = a + C + \xi_i$ and $B_i = b + C + \eta_i$, where C is a symmetric random sign and all zero-mean noises are independent of it and each other, produce bias -1 for every N . Valid alternatives are independent caches per batch, conditionally exact fixed means, or a rigorously bounded residual cache bias.

Termination error remains a bias budget. The batch result is unbiased for the means of the implemented oracles. An ε -shell WoS implementation does not thereby become unbiased for the exact PDE. On our harmonic disk/square tests, u, p are explicit smooth polynomials. At a stopped point x and its boundary projection y , the mean-value inequality gives $|u(y) - u(x)| \leq L_u \varepsilon$ and $|p(y) - p(x)| \leq L_p \varepsilon$, where L_u, L_p bound their gradients. Optional stopping for bounded harmonic martingales then bounds the marginal oracle biases by the same quantities. The product bias is at most $|a|L_p \varepsilon + |b|L_u \varepsilon + L_u L_p \varepsilon^2$. If a path cap is exhausted, our benchmark raises an error instead of silently projecting an interior path. None of these diagnostics supplies a finite insertion certificate without the separate small-hole remainder bound.

For the Poisson-compliance disk test, the source is constant and $u(x) = (1 - |x|^2)/4$. The ball mean-value identity is $u(x) = r^2/4 + \mathbb{E}[u(Y) \mid x]$, where Y is uniform on the sphere of radius r centered at x . Iterating and stopping at the shell gives

$$u(x_0) = \mathbb{E} \left[\sum_{k < \tau} \frac{r_k^2}{4} \right] + \mathbb{E}[u(X_\tau)]. \quad (193)$$

This follows first at $\tau \wedge n$, then by monotone convergence for the nonnegative source sum and bounded convergence for the residual state. At termination, $1 - |X_\tau| \leq \varepsilon$, so $0 \leq u(X_\tau) = (1 - |X_\tau|)(1 + |X_\tau|)/4 \leq \varepsilon/2$. The implemented source estimator omits this residual and consequently has marginal bias of magnitude at most $\varepsilon/2$. Since $0 \leq u \leq 1/4$ and $p = u$, the general product-bias inequality gives the conservative bound $\varepsilon/4 + \varepsilon^2/4$. Exact constant-source integration therefore removes source quadrature error, but not shell truncation error.

K Oracle-State Shape Benchmark Derivation

The controlled shape comparison restricts the admissible domains to ellipses $\Omega_{a,b} = \{(x, y) : x^2/a^2 + y^2/b^2 < 1\}$ with $a, b > 0$. Set $c = [2(a^{-2} + b^{-2})]^{-1}$ and

$$u_{a,b}(x, y) = c(1 - x^2/a^2 - y^2/b^2). \quad (194)$$

Direct differentiation gives $-\Delta u_{a,b} = 2c(a^{-2} + b^{-2}) = 1$, while substitution into the ellipse equation gives zero boundary data. Uniqueness of the Dirichlet problem therefore identifies this polynomial as the state. The fixed target is $z_d = \max(u_{1,0.65}, 0)$, and $J(a, b) = \frac{1}{2} \int_{\Omega_{a,b}} (u_{a,b} - z_d)^2$. This exact-state setting isolates derivative estimation; it does not reproduce the neural-state training or boundary metric of Schmidt and Würschmidt [31].

Probabilistic pushforward derivative. For a smooth admissible velocity V , differentiating the Poisson problem at fixed spatial coordinates gives $-\Delta u' = 0$ in Ω and $u' = -\nabla u \cdot V$ on $\partial\Omega$. Feynman-Kac and the transport formula give, respectively,

$$u'(x) = -\mathbb{E}_x[\nabla u(Y) \cdot V(Y)], \quad (195)$$

$$J'[V] = \int_{\Omega} (u - z_d) u' dx + \frac{1}{2} \int_{\partial\Omega} z_d^2 V \cdot n ds. \quad (196)$$

All functions are bounded on the compact ellipse closure, so Fubini applies. If X is uniform on Ω and Y is a Brownian exit conditional on X ,

$$J'[V] = -\pi ab \mathbb{E}[(u(X) - z_d(X)) \nabla u(Y) \cdot V(Y)] \quad (197)$$

$$+ \frac{1}{2} \int_{\partial\Omega} z_d^2 V \cdot n ds. \quad (198)$$

Splitting the volume integral into the positive and negative parts of $u - z_d$ recovers the signed-start representation in that work, after its stated pushforward sign conversion. Uniform weighted starts estimate the same integral with a different proposal distribution. For $V_a = (x/a, 0)$ and $V_b = (0, y/b)$, parameterize the boundary by $y(\theta) = (a \cos \theta, b \sin \theta)$. Its normal times arclength is $(b \cos \theta, a \sin \theta) d\theta$, whence $V_a \cdot n ds = b \cos^2 \theta d\theta$ and $V_b \cdot n ds = a \sin^2 \theta d\theta$. These are the two surface terms integrated by the benchmark.

Independent deterministic check. Write $x = a\xi$, $y = b\eta$ on the fixed unit disk. With $q = 1 - \xi^2 - \eta^2$, we have $u = cq$ and $dx dy = ab d\xi d\eta$. Since $\partial_a c = 4c^2/a^3$, differentiation at fixed (ξ, η) gives $\partial_a u = 4cu/a^3$. Inside the target ellipse, $\partial_x z_d = -2c_*x$, where $c_* = [2(1 + 0.65^{-2})]^{-1}$; outside it the derivative is zero almost everywhere. The target is Lipschitz and its interface has area measure zero, justifying dominated differentiation of this finite-dimensional pullback. Therefore

$$\partial_a J = \int_{\Omega_{a,b}} \left[(u - z_d) \left(\frac{4c}{a^3} u - \partial_x z_d \frac{x}{a} \right) + \frac{(u - z_d)^2}{2a} \right] dx dy. \quad (199)$$

The formula for $\partial_b J$ follows by exchanging the coordinates and including the target factor 0.65^{-2} in $\partial_y z_d$. We evaluate these integrals by Gauss–Legendre radial and periodic angular quadrature and independently check central differences of J .

L Ball-Chain Connectivity and Its Scope

Proposition L.1 (Strict overlap after recalibration). *Let $r > 0$ and $s_1, \dots, s_L \in \mathbb{R}^d$ satisfy $\|s_{i+1} - s_i\| < 2r$ for every $i < L$. Then $\bigcup_i B_r(s_i)$ is path connected. If every closed ball has positive clearance from $\partial\Omega$ and is contained in Ω , their union does not meet the exterior boundary, irrespective of this connectedness.*

Proof. The midpoint $(s_i + s_{i+1})/2$ is at distance $\|s_{i+1} - s_i\|/2 < r$ from each center, so consecutive open balls intersect. Every ball is convex and hence path connected. For any two points in the union, choose balls containing them; join each point to its center by a segment, and join the intervening centers through consecutive intersection midpoints. The finite concatenation lies inside the union and is a continuous path. If each of finitely many closed balls has positive clearance from the exterior, the minimum of those clearances is positive. Their union is thus separated from the exterior. In particular, strict overlap alone cannot imply the absence of a sealed void island. Tangency would only connect closed balls at a point and is deliberately rejected for the open-channel guarantee. $\square$

M Cavity Count, Volume, and Thermal Compliance

The following facts explain the budget experiment; they are properties of the classical ideal Dirichlet cooling model, not new topological-derivative claims.

Proposition M.1 (Monotonicity under additional cooled material removal). *Let $\Omega \subset \mathbb{R}^3$ be bounded and Lipschitz, let $f \in L^2(\Omega)$, and let $D_2 \subset D_1 \subset \Omega$ be nonempty Lipschitz material domains. For $i = 1, 2$, let $u_i \in H_0^1(D_i)$ solve $-\Delta u_i = f|_{D_i}$ weakly, with zero trace on every boundary component. Extend u_i by zero to Ω and define $J_i = \int_{\Omega} f u_i$. Then*

$$J_1 - J_2 = \int_{\Omega} |\nabla(u_1 - u_2)|^2 \geq 0. \quad (200)$$

Equality holds if and only if the extended states agree. If $D_1 \setminus D_2$ contains a nonempty open set on which $f > 0$ almost everywhere, the inequality is strict. Division by the fixed volume $|\Omega|$ preserves these conclusions.

Proof. Step 1: common function space. By the definition of H_0^1 , each u_i is an H^1 limit of smooth functions compactly supported in D_i . Extending these approximations by zero proves that their extensions belong to $H_0^1(\Omega)$. The same argument gives the inclusion of extended spaces $H_0^1(D_2) \subset H_0^1(D_1)$. Their gradients vanish almost everywhere outside their respective material domains.

Step 2: existence and energy identities. The bilinear form $a(v, w) = \int_{\Omega} \nabla v \cdot \nabla w$ is coercive on each extended space by Poincaré's inequality on Ω ; the right-hand side $\ell(v) = \int_{\Omega} f v$ is bounded by Cauchy–Schwarz and Poincaré. Lax–Milgram gives a unique weak solution in each space. Testing its weak equation with u_i gives $a(u_i, u_i) = \ell(u_i) = J_i$.

Step 3: the cross term. The space inclusion makes u_2 an admissible test function in the equation for u_1 . Hence $a(u_1, u_2) = \ell(u_2) = J_2$. Expanding the squared gradient norm therefore gives $a(u_1 - u_2, u_1 - u_2) = J_1 - 2J_2 + J_2$, proving Eq. (200). A zero norm forces $u_1 = u_2$ by Poincaré; the converse is immediate.

Step 4: strictness under positive interior heating. If equality held and $U \subset D_1 \setminus D_2$ were the stated open set, then $u_1 = u_2 = 0$ almost everywhere on U , so $\nabla u_1 = 0$ there. Choose a nonnegative, nonzero $\phi \in C_c^\infty(U)$. The first weak equation would give $0 = \int_U f \phi > 0$, a contradiction. Finally, multiplication by the positive constant $|\Omega|^{-1}$ proves the normalized version. $\square$

This proposition applies to nested removals, such as increasing the radius of fixed centers or adding a new cavity while retaining old ones. Replacing two large cavities by twelve small ones at the same removed volume does not produce nested material domains; monotonicity alone cannot rank those designs. Insulating, Robin, or fluid-coupled cavity walls require a different analysis.

Proposition M.2 (Exact geometry at fixed removed volume). *Let n disjoint spherical cavities of equal radius r_n have total volume $V > 0$. Their radius, total exposed area A_n , and sum S_n of isolated Newtonian capacities (energy convention) are*

$$r_n = \left(\frac{3V}{4\pi n} \right)^{1/3}, \quad (201)$$

$$A_n = (36\pi)^{1/3} V^{2/3} n^{1/3}, \quad (202)$$

$$S_n = 4\pi \left(\frac{3V}{4\pi} \right)^{1/3} n^{2/3}. \quad (203)$$

Here S_n is not the capacity of the union and does not give its exact cooling benefit.

Proof. Disjointness makes the removed volume the sum of individual ball volumes: $V = n(4\pi/3)r_n^3$. Solving gives the first formula. Each sphere has area $4\pi r_n^2$, and separation leaves that entire surface exposed; substitution into $n4\pi r_n^2$ gives the second formula. For completeness, the equilibrium potential of a radius- r ball in $\mathbb{R}^3$ is 1 inside and $r/|x|$ outside, after translation to the origin. Its energy is

$$\int_{\mathbb{R}^3} |\nabla P|^2 = 4\pi \int_r^\infty \frac{r^2}{s^4} s^2 ds = 4\pi r.$$

To check the minimizing property, write any smooth exterior competitor with the same trace and decay as $P + w$, where w has zero trace at $|x| = r$. Integration by parts gives $\int \nabla P \cdot \nabla w = 0$, since P is harmonic outside the ball and the boundary terms vanish. Thus its energy is $\int |\nabla P|^2 + \int |\nabla w|^2 \geq \int |\nabla P|^2$; density extends the identity to finite-energy, zero-trace perturbations. For the equivalent capacity definition requiring at least one on the ball, truncating a competitor at one cannot increase energy and reduces to this boundary-value problem. Summing these individual energies and substituting r_n proves the third formula. This sum treats every ball in isolation; their joint equilibrium potential satisfies a different exterior boundary-value problem, so no union capacity identity follows. $\square$

Thus equal removed volume does not equalize cooling area. Classical small-cavity scaling suggests why fragmentation can help, but the single-cavity expansion is not asserted uniformly as $n \rightarrow \infty$ with shrinking separations. Field depletion, interactions, and geometric feasibility must be evaluated at the actual finite configurations. In particular, a penalty λV is constant across a fixed-volume count sweep and cannot resolve its count preference. A cavity-count cost, surface-area cost, minimum radius, wall constraint, or explicit cooling-system model may do so; this experiment does not identify physical values for those costs.

A finite-design cost envelope, not a continuous optimum. For the finite tested designs d , let $q_d = J_d/J_0$ and let n_d be the cavity count. Given a stipulated normalized per-cavity cost $\kappa \geq 0$, minimize $q_d + \kappa n_d$ over this finite set. Designs a, b with $n_a \neq n_b$ tie precisely at $\kappa = (q_a - q_b)/(n_b - n_a)$, obtained by rearranging the two affine costs. The lower envelope can change minimizer only at such pairwise intersections. Indeed, away from intersections all pairwise differences keep their signs, so their ordering is constant on each open interval. This construction describes sensitivity to an assumed cost; neither sampling uncertainty nor the untested geometry space disappears when an envelope is plotted.

N Minimum-Benefit Certification: Complete Derivations

This appendix separates two routes. The general-objective route uses a numerically bounded asymptotic remainder. The second route uses the variational structure of three-dimensional, zero-Dirichlet Poisson compliance to bound a finite insertion directly. The latter needs no unknown remainder constant and permits Lipschitz outer geometry. It does not certify the previously reported ranking runs retrospectively; [Appendix O](#) evaluates a separate guarded implementation. All certificates below concern the stated continuum PDE under their oracle assumptions, not a fluid-cooling model or unbounded floating-point geometry error.

N.1 What is being guaranteed?

Write the actual benefit of an insertion as

$$B(z, r) = J(D) - J(D \setminus \bar{B}_r(z)).$$

A lower bound $\underline{B}(z, r)$ is useful when it exceeds a specified target $\tau \geq 0$. The statistical statement is

$$\mathbb{P}\{\text{accept and } B(z, r) < \tau\} \leq \delta, \quad (204)$$

including any covered data-dependent choice of radius or target. It is not $\mathbb{P}(B \geq \tau \mid \text{accept}) \geq 1 - \delta$: conditioning on a rare acceptance can change that probability. In ideal Dirichlet cooling, [Proposition M.1](#) already supplies nonnegative benefit. Minimum-benefit certification asks the additional quantitative question.

Proof of [Theorem A.4](#). Step 1: one event bounds the coefficient. Let E be the event in [Eq. \(38\)](#). Set $b = b_{\text{term}} + b_{\text{cache}}$ and $s = \widehat{D}_{TN} + r_N(\delta) + b$. The bias assumption gives $D_T \leq \mathbb{E}\widehat{D}_{TN} + b$, and on E the confidence bound gives $\mathbb{E}\widehat{D}_{TN} \leq \widehat{D}_{TN} + r_N$. Thus $D_T \leq s$ on E .

Step 2: translate derivative uncertainty into objective units. The upper side of the deterministic remainder inequality gives

$$J(D_r) - J(D) \leq \rho(r)D_T + C_z\rho(r)^{1+\alpha}.$$

Negating both sides and substituting $D_T \leq s$ proves, on E ,

$$B(z, r) \geq -\rho(r)s - C_z\rho(r)^{1+\alpha} = \underline{B}_{\text{TD}}(z, r).$$

Notice the factor ρ : a confidence bound on a derivative alone does not have the units of an actual objective improvement.

Step 3: radii and acceptance. For fixed z , the same E controls every radius for which the deterministic remainder bound holds uniformly. Evaluating the displayed inequality at a measurable selected radius preserves it. On E , accepting only when $\underline{B}_{\text{TD}} \geq \tau$ implies $B \geq \tau$, even for a measurable data-dependent τ . Therefore every false acceptance lies in E^c , and its probability is at most δ . This reasoning allows rejection on all outcomes and does not presume that a feasible certified radius exists. $\square$

Capacity must also be bounded. The preceding expression uses the actual relative operator capacity, not merely its leading asymptotic equivalent. If only certified numerical bounds $0 < \rho_- \leq \rho \leq \rho_+$ are available, a conservative replacement is

$$\underline{B}_{\text{int}} = \begin{cases} -\rho_-s - C_z\rho_+^{1+\alpha}, & s < 0, \\ -\rho_+s - C_z\rho_+^{1+\alpha}, & s \geq 0. \end{cases} \quad (205)$$

Indeed, $-\rho s$ is increasing in ρ when $s < 0$ and decreasing when $s \geq 0$, whereas $-C_z \rho^{1+\alpha} \geq -C_z \rho_+^{1+\alpha}$ in both cases. Adding these two lower bounds proves the claim. Replacing ρ by $4\pi r$ without an error allowance is generally invalid in this route.

N.2 A direct finite-hole bound for thermal compliance

Throughout the rest of this appendix, $D \subset \mathbb{R}^3$ is bounded and Lipschitz, $f \in L^\infty(D)$ satisfies $0 \leq f \leq F$, and $u \in H_0^1(D)$ solves $-\Delta u = f$. For $\overline{B}_R(z) \Subset D$, let known constants f_-, f_+ satisfy $0 \leq f_- \leq f(x) \leq f_+$ almost everywhere on $B_R(z)$. Set $K = \overline{B}_r(z)$ with $0 < r < R$, and let v solve the same PDE on $D \setminus K$ with zero trace on every boundary component. Extend v by zero on K . The objective is the unnormalized compliance $J(D) = \int_D f u$. Constant diffusivity has been scaled to one. Different diffusivities require the corresponding factors in both the PDE and energy.

Lemma N.1 (Positivity, boundedness, and an interior representative). *The state u is non-negative. If $\overline{D} \subset B_L(c)$, then*

$$0 \leq u(x) \leq F(L^2 - |x - c|^2)/6 \leq FL^2/6 \quad (206)$$

The inequalities hold almost everywhere in D . The state has a continuous representative on every compact subset of D .

Proof. Step 1: existence and the sign. Poincare and Cauchy–Schwarz make $\int_D \nabla u \cdot \nabla \phi = \int_D f \phi$ a coercive weak problem; Lax–Milgram gives its unique solution. The negative part $u^- = \max(-u, 0)$ belongs to $H_0^1(D)$. Testing with it yields $-\int_D |\nabla u^-|^2 = \int_D f u^- \geq 0$. Both sides must vanish; Poincare implies $u^- = 0$.

Step 2: an explicit supersolution. Put $Q(x) = F(L^2 - |x - c|^2)/6$. Since $\Delta|x - c|^2 = 6$ in three dimensions, $-\Delta Q = F \geq f$. The trace of $u - Q$ is nonpositive. Consequently $(u - Q)^+ \in H_0^1(D)$ is an admissible test in the difference of the two weak equations and gives

$$\int_D |\nabla(u - Q)^+|^2 = \int_D (f - F)(u - Q)^+ \leq 0.$$

Poincare implies $u \leq Q$. This proves [Eq. \(206\)](#).

Step 3: point values. Interior elliptic regularity for the Laplacian and bounded f gives $u \in W_{\text{loc}}^{2,p}(D)$ for finite p ; take $p > 3$ and apply the local Sobolev embedding to obtain a continuous (in fact locally $C^{1,\gamma}$ for some $\gamma > 0$) representative. These standard local regularity and trace results are the same background facts used in [Appendix C](#) [9, 12]. The almost-everywhere bounds hold everywhere for this representative in the interior. $\square$

Lemma N.2 (From a cavity temperature floor to a finite benefit). *If $u \geq m \geq 0$ on K , then*

$$J(D) - J(D \setminus K) \geq 4\pi r m^2. \quad (207)$$

Proof. Step 1: the exact finite-domain identity. The zero extension of v belongs to $H_0^1(D)$ and its gradient vanishes almost everywhere on K . Testing the state equations gives $\int_D |\nabla u|^2 = J(D)$, $\int_D |\nabla v|^2 = J(D \setminus K)$, and $\int_D \nabla u \cdot \nabla v = J(D \setminus K)$. Expanding the square therefore gives

$$J(D) - J(D \setminus K) = \int_D |\nabla w|^2, \quad w = u - v.$$

This is [Eq. \(200\)](#); it is exact for finite r .

Step 2: turn w into a capacity competitor. If $m = 0$, the energy identity already proves the result. Otherwise extend w/m by zero from D to $\mathbb{R}^3$. It belongs to the homogeneous energy space $\mathcal{D}^{1,2}(\mathbb{R}^3)$, defined as the completion of $C_c^\infty(\mathbb{R}^3)$ in gradient norm. It is at least one on K in the quasi-everywhere Sobolev

sense: $w = u$ there, the zero extension of v vanishes quasi-everywhere there, and the continuous u is at least m . By the definition of Newtonian capacity,

$$m^{-2} \int_D |\nabla w|^2 \geq \text{cap}_{\mathbb{R}^3}(K).$$

Step 3: fix the capacity normalization. The equilibrium potential is $P(x) = \min(1, r/|x - z|)$. Its exterior energy is $4\pi \int_r^\infty r^2 s^{-2} ds = 4\pi r$, and its interior energy is zero. To verify that no smaller energy is possible, truncate any admissible competitor to the interval $[0, 1]$; this cannot increase gradient energy and sets it to one on the ball. For a smooth compactly supported such competitor a , its spherical average $A(s)$ satisfies $A(r) = 1$, $A(\infty) = 0$. Jensen and Cauchy–Schwarz give

$$\begin{aligned} \int_{|x-z|>r} |\nabla a|^2 &\geq 4\pi \int_r^\infty s^2 |A'(s)|^2 ds, \\ 1 &= \left| \int_r^\infty A'(s) ds \right|^2 \leq \left(\int_r^\infty s^2 |A'(s)|^2 ds \right) \left(\int_r^\infty s^{-2} ds \right). \end{aligned}$$

The last factor is $1/r$, so the energy is at least $4\pi r$. Approximation in the exterior energy space with fixed trace extends this inequality to the admissible truncated Sobolev competitors. Equivalently, subtract P and use the zero-trace energy-density argument of [Proposition M.2](#). Thus the infimum equals $4\pi r$. Substituting into Step 2 proves the claimed finite-hole bound. $\square$

Lemma N.3 (Explicit temperature floor from one interior value). Define $t = r/R$ and

$$\begin{aligned} c(t) &= (1-t)/(1+t)^2, \\ m(L_u; r, R) &= \frac{f_-(R^2 - r^2)}{6} \\ &\quad + c(t) \left(L_u - \frac{f_+ R^2}{6} \right)_+, \end{aligned} \tag{208}$$

where $(a)_+ = \max(a, 0)$. For $L_u \leq u(z)$, this gives the floor $u(x) \geq m(L_u; r, R)$ throughout $\overline{B}_r(z)$.

Proof. Step 1: separate the source and harmonic parts inside the ball. Let $q \in H_0^1(B_R(z))$ solve $-\Delta q = f$ there. Define $h = u - q$ on that ball. Subtracting the weak equations gives $\Delta h = 0$. Its trace is $u \geq 0$, so the negative-part test proves $h \geq 0$. The polynomial functions $q_\pm(x) = f_\pm(R^2 - |x - z|^2)/6$ vanish on the sphere and solve $-\Delta q_\pm = f_\pm$. The positive-part comparison argument in [Lemma N.1](#) applied to $q - q_+$ and $q_- - q$ gives $q_- \leq q \leq q_+$. In particular,

$$h(z) = u(z) - q(z) \geq u(z) - f_+ R^2/6, \quad h(z) \geq 0.$$

Hence $h(z) \geq (L_u - f_+ R^2/6)_+$.

Step 2: derive the harmonic comparison constant. Translate z to zero. For $|x| = s < R$, the Poisson kernel representation is

$$h(x) = \int_{|y|=R} \frac{R^2 - s^2}{4\pi R |x - y|^3} h(y) dS_y.$$

At the center the kernel is $1/(4\pi R^2)$, so its pointwise ratio to the central kernel obeys

$$\frac{R(R^2 - s^2)}{|x - y|^3} \geq \frac{R(R^2 - s^2)}{(R + s)^3} = \frac{1 - s/R}{(1 + s/R)^2}.$$

Multiplication by $h(y) \geq 0$ and integration therefore imply $h(x) \geq c(s/R)h(0)$. Finally $c'(a) = (a - 3)/(1 + a)^3 < 0$ for $0 \leq a < 1$; thus $s \leq r$ gives $h(x) \geq c(r/R)h(z)$. The Poisson formula is valid with the continuous

boundary values here because $\bar{B}_R(z)$ lies strictly inside D ; the zero-Dirichlet bounded-source solution q is continuous up to that sphere.

Step 3: add the two lower bounds. On $\bar{B}_r(z)$, $q(x) \geq f_-(R^2 - r^2)/6$. Adding the lower bound on h from Steps 1–2 gives exactly Eq. (208). Both summands are nonnegative. This proof uses a finite interior ball, not a small- r Taylor approximation. $\square$

Proof of Theorem A.5. Work on the event $E_u = \{L_u \leq u(z)\}$. Then Lemma N.3 supplies a nonnegative temperature floor on the entire inserted ball. Substitution into Lemma N.2 proves $B(z, r) \geq 4\pi r m(L_u; r, R)^2 = \underline{B}_{\text{cool}}(z, r)$. Both lemmas hold for every admissible $0 < r < R$ simultaneously on the same E_u . They also hold for every admissible R if its source bounds and geometric containment are valid. Thus choosing a radius after observing L_u does not require a union bound over radii. The event $\{\text{accept}, B < \tau\}$ is contained in E_u^c , proving its probability bound. In particular, $L_u = 0$ is deterministically valid and yields

$$\underline{B}_{\text{source}}(z, r) = \frac{\pi r f_-^2 (R^2 - r^2)^2}{9}. \quad (209)$$

This source-only bound may be small or zero; its validity does not assert that it can meet the requested target. A better certified L_u can improve the bound because m is nondecreasing in L_u . $\square$

Relation to the topological score. For an exact value $L_u = u(z)$ and constant $f = f_0$ on $B_R(z)$, $u(z) \geq f_0 R^2/6$ by comparison. Sending $r \downarrow 0$ in Eq. (208) gives $m \rightarrow u(z)$ and hence $\underline{B}_{\text{cool}}/(4\pi r) \rightarrow u(z)^2 = -D_T(z)$. For merely bounded nonconstant f , one can instead choose interior radii $R(r) \rightarrow 0$ with $r/R(r) \rightarrow 0$ and uniformly bounded f_+ ; then $c \rightarrow 1$ and both source terms vanish, giving the same limit. These statements follow by direct substitution and continuity of $(\cdot)_+$, using $u(z) \geq 0$. A conservative or noisy L_u need not attain this limiting coefficient. The direct finite-radius bound builds on classical energy, capacity, and harmonic comparison; we do not claim these constituent results as new.

N.3 A one-sided WoS validator with an explicit range bound

An unbiased two-sided state estimator is not necessary for a lower temperature bound. Nonnegative source contributions permit a deliberately truncated estimator whose omitted terminal value has a known sign. This section specifies a separate validator, not an assumption about the existing application ranking code.

Proposition N.4 (Finite-walk lower confidence bound). *Fix integers $T \geq 1$, $N \geq 1$, a radius cap $s_* > 0$, and $\delta \in (0, 1)$ before observing the validation data. Start $X_0 = z$. At each active step $k < T$, choose a history-measurable $0 < s_k \leq s_*$ with $\bar{B}_{s_k}(X_k) \subseteq D$. Draw a uniform spherical direction Θ_k and set $X_{k+1} = X_k + s_k \Theta_k$. Also draw Y_k in that ball with density*

$$\frac{6}{s_k^2} G_{B_{s_k}(X_k)}(X_k, y), \quad S_k = \frac{s_k^2}{6} f(Y_k).$$

All fresh draws have their stated conditional distributions given the past; the source and exit draws may be independent. Stop at any history-measurable rule by step T . After stopping, set $S_k = 0$ and keep X_k fixed. Define $W = \sum_{k=0}^{T-1} S_k$. For independent identically distributed complete walks $W_1, \dots, W_N$, set

$$A = \frac{TFs_*^2}{6}, \quad L_u = \max \left(0, \bar{W} - A \sqrt{\frac{\log(1/\delta)}{2N}} \right). \quad (210)$$

Then $\mathbb{P}\{L_u \leq u(z)\} \geq 1 - \delta$. For $F = 0$, the solution and estimator are identically zero and the formula is read with $A = 0$.

Proof. Step 1: the one-ball identity and normalization. On a strictly contained ball, decompose u into the zero-boundary source solution and the harmonic extension of its boundary values. At the center,

the Poisson kernel is uniform on the sphere, so

$$u(X_k) = \mathbb{E}[u(X_{k+1}) \mid \mathcal{G}_k] + \int_{B_{s_k}(X_k)} G(X_k, y) f(y) dy,$$

where $\mathcal{G}_k$ is the history before the step. The kernel is $G(X_k, y) = (4\pi)^{-1}(|y - X_k|^{-1} - s_k^{-1})$. Spherical integration gives

$$\int G(X_k, y) dy = \int_0^{s_k} (a - a^2/s_k) da = s_k^2/6.$$

Consequently the stated density integrates to one and $\mathbb{E}[S_k \mid \mathcal{G}_k] = \int Gf$. An implementable sampler is $Y_k = X_k + s_k Q_k \Phi_k$, where Φ_k is uniform on $\mathbb{S}^2$ and Q_k has density $6q(1-q)$ on $(0, 1)$, as derived in [Proposition G.1](#). For constant f , $S_k = fs_k^2/6$ is exact without drawing Y_k .

Step 2: telescope at a bounded stopping time. At an active step the preceding identity is $\mathbb{E}[S_k + u(X_{k+1}) \mid \mathcal{G}_k] = u(X_k)$. At a stopped step it remains true because $S_k = 0$ and $X_{k+1} = X_k$. Each term is integrable: [Lemma N.1](#) bounds u and $0 \leq S_k \leq Fs_k^2/6$. Taking expectations and summing all T finite identities cancels consecutive state values and gives

$$u(z) = \mathbb{E}W + \mathbb{E}u(X_T), \quad 0 \leq \mathbb{E}W \leq u(z). \quad (211)$$

Thus shell termination and a finite path cap create downward bias, whose sign is sufficient for this one-sided use. No assertion of zero bias is made.

Step 3: prove the needed bounded-variable concentration. For any $Z \in [0, A]$, let $g(\lambda) = \log \mathbb{E} \exp(\lambda(Z - \mathbb{E}Z))$. Boundedness permits differentiating under the expectation. Under the exponentially tilted distribution, $g''(\lambda)$ is the variance of Z . Every distribution supported on $[0, A]$ has variance at most $A^2/4$: $\text{Var } Z = \min_c \mathbb{E}(Z - c)^2 \leq \mathbb{E}(Z - A/2)^2 \leq A^2/4$. Since $g(0) = g'(0) = 0$, integrating the second-derivative bound twice gives $g(\lambda) \leq \lambda^2 A^2/8$ for $\lambda \geq 0$. Independence of the N walks, Markov's inequality, and this bound yield

$$\mathbb{P}\{\bar{W} - \mathbb{E}W \geq a\} \leq \exp(-\lambda Na + N\lambda^2 A^2/8).$$

For $A > 0$ and $a > 0$, the choice $\lambda = 4a/A^2$ gives the bound $\exp(-2Na^2/A^2)$. With a equal to the sampling allowance in [Eq. \(210\)](#), this is δ . This supplies the one-sided Hoeffding inequality used here with all constants.

Step 4: convert it into a state bound. On the complementary event, $\bar{W} - a \leq \mathbb{E}W \leq u(z)$ by Step 2. Since $u(z) \geq 0$, replacing $\bar{W} - a$ by its maximum with zero preserves the inequality. This is precisely [Eq. \(210\)](#). $\square$

What bias has and has not been handled. In [Eq. \(211\)](#), the exact truncation bias is $u(z) - \mathbb{E}W = \mathbb{E}u(X_T) \in [0, FL^2/6]$. Its sign makes a separate magnitude bound unnecessary for the proposed lower certificate. This does not cover an approximate source integral that overestimates f , an unsafe distance query that steps outside D , approximate boundary data, or a correlated collection treated as independent walks. A rigorously known lower source $\tilde{f} \in [0, f]$ may be substituted: the one-ball inequality becomes $\mathbb{E}[S_k + u(X_{k+1}) \mid \mathcal{G}_k] \leq u(X_k)$ and telescoping still gives $\mathbb{E}W \leq u(z)$. Alternatively, if a numerical walk output $\tilde{W}$ has a certified overshoot bound $\mathbb{E}\tilde{W} \leq u(z) + \beta$ and lies in a known interval of width A , subtracting β as well as the sampling radius proves the same result. An unquantified numerical overshoot is not covered. Increasing T need not tighten the certificate because the range bound A grows with T ; this is a conservative construction, not an efficiency claim.

N.4 Costs, relative targets, selection, and repeated insertions

Proposition N.5 (Explicit costs and fixed normalization). *Suppose $B \geq \underline{B}$ on a coverage event. If the added design cost obeys $\Delta C \leq c_+$, the net improvement in $J + C$ obeys*

$$(J + C)_{\text{old}} - (J + C)_{\text{new}} \geq \underline{B} - c_+.$$

Thus $\underline{B} \geq \tau + c_+$ certifies net benefit at least τ on that event. For the objective $J/V_{\text{ref}} + C$ with fixed known $V_{\text{ref}} > 0$, the corresponding bound is $\underline{B}/V_{\text{ref}} - c_+$. For a relative target $B/J_0 \geq \eta \geq 0$, suppose

$J_0 > 0$ and a certified upper bound $J_0 \leq U_0$ is available on a second event. Then $\underline{B} \geq \eta U_0$ suffices on the intersection of both events.

Proof. Step 1: subtract costs in the correct units. By definition, the net objective reduction is $B - \Delta C$. The inequalities $B \geq \underline{B}$ and $-\Delta C \geq -c_+$ give the first claim by addition. Dividing only the compliance term by fixed $V_{\text{ref}} > 0$ preserves its order and gives the normalized claim.

Step 2: justify an explicit geometric cost example. For one disjoint interior sphere, a stipulated cost $C = \lambda V_{\text{removed}} + \gamma A_{\text{cavity}} + \kappa n$ with nonnegative coefficients increases by

$$c_+ = \lambda \frac{4\pi r^3}{3} + \gamma 4\pi r^2 + \kappa.$$

The ball volume, newly exposed sphere area, and one-unit increase of count give these terms exactly. The formula is not asserted for overlapping balls or for an unspecified manufacturing cost model.

Step 3: a relative threshold uses an upper denominator bound. On the intersection of the two events, $B \geq \underline{B} \geq \eta U_0 \geq \eta J_0$. Division by $J_0 > 0$ proves the relative claim. If the events fail with probabilities at most δ and δ_0 , their intersection fails with probability at most $\delta + \delta_0$ by subadditivity; independence is not necessary. Plugging in an unbounded noisy estimate of J_0 does not supply this implication. $\square$

Proposition N.6 (Adaptive minimum-benefit guarantees). *At round t , condition on the complete past $\mathcal{F}_{t-1}$, including all earlier geometry changes. Choose a finite shortlist of m_t candidates before collecting their validation data. Allocate nonnegative, $\mathcal{F}_{t-1}$ -measurable error levels δ_{tj} with $\sum_{t,j} \delta_{tj} \leq \delta_{\text{total}}$ almost surely. Assume every candidate certificate has conditional coverage at least $1 - \delta_{tj}$ and all its deterministic hypotheses hold on the current domain. Select at most one insertion in the round, accepting it only if its covered lower benefit (or net benefit) is at least its target τ_t . Then the probability that any accepted insertion misses its target is at most δ_{total} . On the same simultaneous event, a finite sequence of accepted insertions guarantees total benefit at least the sum of their targets, provided no other unaccounted geometry changes occur.*

Proof. Step 1: cover the shortlist before choosing its winner. Let E_{tj} be candidate j 's conditional coverage event. On $E_t = \bigcap_{j=1}^{m_t} E_{tj}$, every displayed candidate bound is true at once. Any winner selected from these data is therefore covered, as is any radius selected using its uniform radius guarantee. Conditional subadditivity gives $\mathbb{P}(E_t^c \mid \mathcal{F}_{t-1}) \leq \sum_j \delta_{tj}$. Candidates may share randomness if their marginal conditional bounds remain valid. The data used to invent the shortlist cannot silently be reused as fresh validation data.

Step 2: control all rounds. A false acceptance at round t is a subset of E_t^c . Taking expectations, then applying countable subadditivity and Tonelli to nonnegative terms gives

$$\mathbb{P}(\text{some false acceptance}) \leq \sum_t \mathbb{E} \sum_j \delta_{tj} = \mathbb{E} \sum_{t,j} \delta_{tj} \leq \delta_{\text{total}}.$$

Recomputing the PDE and containment data on the updated domain is essential: each certificate refers to that domain, not the original unperforated object.

Step 3: telescoping actual finite benefits. For accepted designs $D_0, D_1, \dots, D_k$ and their actual objective values, $J(D_{i-1}) - J(D_i) \geq \tau_i$ holds simultaneously on the covered event. Summing cancels the intermediate objectives and yields $J(D_0) - J(D_k) \geq \sum_{i=1}^k \tau_i$. The same algebra applies to $J + C$. Unvalidated shape refinement between those designs would contribute an additional objective change and cannot be omitted from this sum.

Step 4: sample-size and stopping choices. The preceding proof also covers stages indexed by (t, j, ℓ) if a predeclared sequence of fixed sample sizes is inspected with error levels whose total obeys the same budget. Each stage needs its own valid marginal coverage; the batches may be nested because subadditivity does not require independence across stages. A single fixed- N bound repeatedly checked until it passes supplies no such guarantee. Within each use of [Proposition N.4](#), independence is still required across its complete walks. Selecting one site using scouting data and then drawing a single fresh final validation batch is the special case $m_t = 1$. $\square$

N.5 An exactly solvable finite-radius check

Proposition N.7 (Concentric unit-ball reference). For $D = B_1(0)$, $f = 1$, and a concentric radius- r cavity with $0 < r < 1$, write $s = |x|$. The original and perforated states are

$$u(s) = (1 - s^2)/6, \quad (212)$$

$$v(s) = (1 - s^2)/6 - \frac{r(1+r)}{6}(1/s - 1), \quad r < s < 1. \quad (213)$$

Their exact compliance benefit is

$$B_{\text{exact}}(r) = \frac{\pi r(1-r)(1+r)^2}{9} + \frac{4\pi r^5}{45}, \quad J_0 = \frac{4\pi}{45}. \quad (214)$$

The direct capacity bound using the exact temperature floor is $B_{\text{floor}} = \pi r(1-r^2)^2/9 \leq B_{\text{exact}}$.

Proof. Step 1: solve both radial boundary-value problems. For a radial field, $\Delta a = a'' + 2a'/s$. Substitution shows $-\Delta[(1 - s^2)/6] = 1$ and $\Delta(1/s) = 0$ for $s > 0$. Regularity at the origin and zero trace at $s = 1$ give u . On the annulus write $v = u + b(1/s - 1)$; the condition $v(r) = 0$ yields $b = -r(1+r)/6$. Zero trace at one is automatic. Uniqueness of the weak Dirichlet problem verifies these candidates are the desired solutions.

Step 2: compute the exact difference energy. Inside the removed ball, $w = u$ and $|w'(s)| = s/3$, giving energy $4\pi \int_0^r (s^2/9)s^2 ds = 4\pi r^5/45$. Outside it, $w = r(1+r)(1/s - 1)/6$. Its energy is

$$\frac{4\pi r^2(1+r)^2}{36} \int_r^1 s^{-2} ds = \frac{\pi r(1-r)(1+r)^2}{9}.$$

Adding the two energies proves the benefit formula by Eq. (200). Direct radial integration gives $J_0 = 4\pi \int_0^1 (1 - s^2)s^2/6 ds = 4\pi/45$.

Step 3: independently check the lower bound. The minimum of u on $\overline{B}_r$ is $(1 - r^2)/6$. Substituting into $4\pi r m^2$ gives the stated B_{floor} . Subtracting it from the exact expression leaves

$$B_{\text{exact}} - B_{\text{floor}} = \frac{\pi r^2(1-r)(1+r)^2}{9} + \frac{4\pi r^5}{45} \geq 0.$$

This checks both the capacity normalization and the sign of the guarantee at finite radius, independently of the asymptotic theorem. $\square$

Reproducible check. A numerical check evaluates the closed-form state and compliance, checks the lower bounds over a fixed radius grid, and runs the fixed-budget validator on this analytic geometry. It records the seed, sample count, stopping policy, range bound, confidence level, target, and actual exact benefit. It is a manufactured check of the formulas and implementation, not a coverage proof or a Bunny/housing performance measurement. The confidence statement comes from the preceding derivation; numerical code remains subject to the explicitly stated geometry, source-sampling, independence, and arithmetic assumptions.

For this fixed seed and budget, the source-only rule rejects the stipulated target and the sampled rule accepts it (Table 1). The exact energy formula agrees with independent radial quadrature to 4.3×10^{-16} in the seven tested cases. The deterministic inequalities are also checked at 2,001 radii below R ; their validity for all radii comes from the proofs, not from this finite grid. The sample budget was not chosen as a solver-speed comparison, and one successful seed does not establish practical tightness across application geometries.

Table 1 Manufactured minimum-benefit check on the unit ball, with a preselected cavity radius $r = 0.1$, containing radius $R = 0.9$, and target 10% of the exact original compliance. The Monte Carlo rule allocates $\delta = 0.01$ to this one fixed test. These are numerical evaluations of the proved formulas, not an empirical estimate of coverage.

Quantity	Value
Independent walks N	65,536
Maximum steps T / radius cap s_*	32 / 0.35
Exact center state $u(0)$	0.166667
Truncated sample mean $\bar{W}$	0.162320
One-sided sampling allowance	0.003873
State lower bound L_u	0.158447
Source-only benefit bound / J_0	8.000%
Monte Carlo benefit bound / J_0	10.154%
Exact finite-hole benefit / J_0	13.614%

O Application Validation: Implementation and Proofs

The following additions deploy the minimum-benefit rule on the Bunny and housing. This is a separate experiment with new seeds, geometries selected under an acceptance rule, and retained validation samples. It does not retroactively certify the earlier ranking experiments. Three elementary constructions make the validator usable: downward clipping of nonnegative walk contributions, conservative boundary-distance queries, and explicit source enclosures. The probability guarantee retains the ideal conditional sampling assumptions of [Proposition N.4](#). Numerical containment and source safeguards alone do not prove those sampling laws for finite-precision positions and a pseudorandom generator.

O.1 A sharper certificate from the average cavity trace

The temperature-floor certificate in [Theorem A.5](#) is sufficient but can be conservative near existing cavities or a nonuniform heat source. The same energy identity admits a stronger use of a center query, without requiring a pointwise temperature floor on the new cavity.

Proof of Theorem A.6. Step 1: difference energy. Let u be the current state and v the state after removing $K = \bar{B}_r(z)$, extended by zero. Set $w = u - v$ and extend it by zero outside D . As in [Lemma N.2](#),

$$B = J(D) - J(D \setminus K) = \int_{\mathbb{R}^3} |\nabla w|^2.$$

The trace of w on ∂K is the trace of u , since v has zero trace there.

Step 2: pair with the equilibrium potential. Define $P(x) = 1$ for $|x - z| \leq r$ and $P(x) = r/|x - z|$ otherwise. On the exterior of the ball it is harmonic. Its outward normal derivative for that exterior domain, on its inner sphere, is $+1/r$. Integration by parts therefore gives

$$\int_{\mathbb{R}^3} \nabla w \cdot \nabla P = \frac{1}{r} \int_{\partial B_r(z)} u \, dS = 4\pi r \bar{u}_r, \quad (215)$$

where $\bar{u}_r = (4\pi r^2)^{-1} \int_{\partial B_r(z)} u \, dS$. To justify this calculation for H^1 functions, first use smooth compactly supported approximations to the zero extension of w . Integrate on the exterior annulus with an outer sphere beyond their support, so the outer boundary term vanishes. Pass to the limit by gradient-norm convergence and the continuous trace map on an annular neighborhood of the inner sphere. The potential has zero interior gradient and energy $4\pi r$, as computed in [Lemma N.2](#).

Step 3: apply Cauchy–Schwarz to the energy pairing. Squaring [Eq. \(215\)](#) and applying Cauchy–Schwarz gives

$$(4\pi r \bar{u}_r)^2 \leq \left(\int |\nabla w|^2 \right) \left(\int |\nabla P|^2 \right) = B 4\pi r.$$

Division by $4\pi r > 0$ proves $B \geq 4\pi r \bar{u}_r^2$.

Step 4: recover the average trace from a center value. The centered ball Green identity gives

$$u(z) = \bar{u}_r + I_r, \quad I_r = \int_{B_r(z)} G_{B_r(z)}(z, y) f(y) dy. \quad (216)$$

If $I_r \leq I_r^+$ and $L_u \leq u(z)$, then $\bar{u}_r \geq L_u - I_r^+$. Also $\bar{u}_r \geq 0$ because $u \geq 0$. Thus $\bar{u}_r \geq (L_u - I_r^+)_+$, and substitution into Step 3 proves the claimed bound. Positivity of G and $f \leq f_+$ on the inserted ball give $I_r^+ = f_+ r^2/6$, using $\int G = r^2/6$.

Step 5: affine sources and uniformity. For $f(y) = f(z) + a \cdot (y - z)$, the integral of $G(z, y)a \cdot (y - z)$ vanishes: reflection $y - z \mapsto -(y - z)$ preserves the ball and the radial kernel and reverses this integrand's sign. Hence $I_r = f(z)r^2/6$ exactly. All preceding inequalities hold simultaneously for every contained radius on the one event $\{L_u \leq u(z)\}$. They also hold together with the temperature-floor certificate, so their maximum is a valid benefit lower bound on that event. The acceptance and adaptive-selection conclusions follow from the same event-containment proof as [Proposition N.6](#); no additional confidence allocation is needed to compare deterministic certificates built from the same covered state bound. $\square$

This argument uses classical equilibrium-potential duality and the centered Green representation. Its role here is a sharper finite-insertion validator, not a claim that the underlying capacity identity is new. It avoids the harmonic minimum factor in [Eq. \(208\)](#); the local source upper bound now concerns the inserted ball itself. The earlier floor and source-only bounds remain useful when the state estimate is too weak.

0.2 Clipping gives a smaller, known sample range

Proposition O.1 (Downward-clipped validation walks). *Under the one-sided construction of [Proposition N.4](#), fix $C > 0$ before drawing a validation batch. Source samples may use a measurable lower source $0 \leq f \leq f$. Stop by a fixed step count T , possibly earlier when the accumulated nonnegative contribution reaches C , and output $W_C = \min(C, \sum_{k < \tau} S_k)$. Then*

$$0 \leq W_C \leq C, \quad \mathbb{E}W_C \leq u(z). \quad (217)$$

For N independent complete walks, the confidence rule in [Eq. \(210\)](#) holds with $A = C$.

Proof. **Step 1: lower sources preserve a one-step inequality.** Positivity of the ball Green kernel gives $\int Gf \leq \int Gg$. The one-ball representation therefore implies

$$\mathbb{E}[S_k + u(X_{k+1}) \mid \mathcal{G}_k] \leq u(X_k)$$

at every active step. After stopping, pad with $S_k = 0$ and $X_{k+1} = X_k$; the inequality remains true. Reaching the cap after a step is a permissible stopping rule because the decision uses only the history available at the next step.

Step 2: telescope and then clip. Summing the T finite conditional inequalities and taking expectations yields $\mathbb{E} \sum_{k < \tau} S_k + \mathbb{E}u(X_\tau) \leq u(z)$. Since $u \geq 0$ and $0 \leq W_C \leq \sum_{k < \tau} S_k$, it follows that $\mathbb{E}W_C \leq u(z)$. Clipping also ensures $W_C \leq C$ pathwise. The bound makes no assertion that clipping preserves the mean; additional downward bias is expected.

Step 3: apply bounded-variable concentration. The moment-generating-function proof in [Proposition N.4](#), Step 3, uses only independence and the range width. Substituting C for A gives

$$\mathbb{P} \left\{ \bar{W}_C - \mathbb{E}W_C > C \sqrt{\frac{\log(1/\delta)}{2N}} \right\} \leq \delta.$$

Subtract this allowance and clip the resulting state lower bound at zero. Step 2 and $u(z) \geq 0$ prove its coverage. A smaller C improves the sampling allowance but can increase clipping bias, so it need not improve the final benefit certificate. $\square$

0.3 A conservative distance query on triangle geometry

Proposition 0.2 (Bounding-box distance and contained steps). *Let a closed embedded triangle surface $S = \bigcup_i T_i$ bound the material domain, and let $Q_i \supset T_i$ be the coordinate-aligned bounding box of triangle i . For any point x ,*

$$d_Q(x) := \min_i \text{dist}(x, Q_i) \leq \text{dist}(x, S). \quad (218)$$

If x is inside the current material domain, subtracting retained spherical cavities gives the valid distance lower bound

$$d_*(x) = \min \left(d_Q(x), \min_j (|x - c_j| - r_j) \right).$$

For $0 < s < d_*(x)$, the closed ball $\overline{B}_s(x)$ lies strictly inside the material domain. A computed endpoint x' is still interior whenever a certified upper bound on $|x' - x|$ is smaller than a certified lower bound on $d_*(x)$.

Proof. Step 1: set containment reverses distance. For each i , the infimum of $|x - y|$ over the larger set Q_i is at most its infimum over T_i . Taking minima over i proves Eq. (218). A box may reach into the material and give zero distance at an interior point; stopping there is conservative.

Step 2: evaluate the box distance. If $Q = \prod_{a=1}^3 [l_a, h_a]$, its nearest point is the coordinatewise clamping of x_a to $[l_a, h_a]$. Indeed, the squared Euclidean distance is a sum of three independent nonnegative coordinate objectives. Thus

$$\text{dist}(x, Q)^2 = \sum_{a=1}^3 \max(l_a - x_a, x_a - h_a, 0)^2.$$

Lower bounds on each nonnegative coordinate gap, their squares, their sum, and its square root give a lower bound on the distance. The implementation rounds each of these elementary operations outward using adjacent binary64 numbers, without fast-math reassociation. This statement assumes the usual correctly rounded IEEE arithmetic and square root for those operations.

Step 3: justify hierarchy pruning. Store in every internal hierarchy node a box containing all descendant triangle boxes. Its lower distance is a lower bound on the distance to each descendant triangle. If the current returned candidate lower bound is b and a node's lower distance is at least b , every triangle in that node has true distance at least b , so pruning that node cannot make b exceed the global true surface distance. Evaluating all unpruned leaves and taking their minimum lower bound preserves this invariant. Using lower bounds rather than nearest-point estimates is essential to this argument.

Step 4: include old cavities and strict containment. For x outside $\overline{B}_{r_j}(c_j)$, its distance to that sphere is $|x - c_j| - r_j$. Taking the minimum with the outer-surface lower bound gives a lower bound for the distance to every boundary of the current domain. If a point of $\overline{B}_s(x)$ left the domain, the line segment from interior x to that point would first cross a boundary at distance at most $s < d_*(x)$, a contradiction. The same segment argument with $|x' - x| < d_*(x)$ proves the computed-endpoint claim. $\square$

The initial finalist is classified by exact rational positive-axis ray parity, treating stored binary64 mesh coordinates as exact dyadic rationals. Projected bounding boxes cull triangles; the remaining projected orientation and intersection-sign tests use rational arithmetic. On an edge or coplanar hit the code retries another axis and rejects the classification if all axes are degenerate. For a closed embedded surface and a nondegenerate ray, each crossing changes interior status; the ray ends outside the bounded object, so an odd crossing count is equivalent to an interior origin. Prior cavity exclusion is checked by exact squared-distance comparison. These predicates presuppose the closed embedded input surface; manifold diagnostics are not an exhaustive proof of absence of every self-intersection.

After initialization, every source location and walk endpoint passes the upper-distance guard in [Proposition 0.2](#). A guard failure invalidates the validation record; it is not silently counted as a zero-valued sample. Steps use $s = \min(s_*, 0.95 \underline{d}_*)$, and the source integration and positive accumulation are rounded downward.

A rounded endpoint remaining inside is a geometric guarantee; it does not prove that its distribution is exactly the ideal uniform spherical measure.

O.4 Explicit local source bounds

Proposition O.3 (Affine and Gaussian source enclosures). *For $f(x) = 1 + a(x_2 + 1/2)$ with $a \geq 0$, valid extrema on $B_R(z)$ are $f_{\pm} = 1 + a(z_2 + 1/2 \pm R)$. For a nonnegative Gaussian mixture*

$$f(x) = b + \sum_j a_j \exp\left(-\sum_{k=1}^3 \frac{(x_k - c_{jk})^2}{w_{jk}^2}\right),$$

assume $b, a_j \geq 0$ and $w_{jk} > 0$. define

$$t_j^{\min} = \sum_k \frac{\max(|z_k - c_{jk}| - R, 0)^2}{w_{jk}^2},$$

$$t_j^{\max} = \sum_k \frac{(|z_k - c_{jk}| + R)^2}{w_{jk}^2}.$$

Then $b + \sum_j a_j e^{-t_j^{\max}}$ and $b + \sum_j a_j e^{-t_j^{\min}}$ are respectively lower and upper source bounds throughout $B_R(z)$. Moreover, for $t \geq 0$ and integer $n \geq 1$,

$$\max(0, 1 - t/n)^n \leq e^{-t} \leq (1 + t/n)^{-n}. \quad (219)$$

Proof. Step 1: enclose coordinates. Every point in $B_R(z)$ has $z_k - R \leq x_k \leq z_k + R$. Monotonicity of the affine source gives its two bounds. The triangle and reverse-triangle inequalities give

$$\max(|z_k - c_{jk}| - R, 0) \leq |x_k - c_{jk}| \leq |z_k - c_{jk}| + R.$$

Squaring, dividing by $w_{jk}^2 > 0$, and summing proves that the exponent argument lies in $[t_j^{\min}, t_j^{\max}]$.

Step 2: enclose each exponential. The function $t \mapsto e^{-t}$ is decreasing. Apply it to the exponent interval, multiply by $a_j \geq 0$, and sum with b to obtain the source enclosures. Optimizing each coordinate separately relaxes the ball to a box; this weakens the bounds but does not invalidate them.

Step 3: remove dependence on a rounded exponential function. For $0 \leq a < 1$, integration of $1/(1-s) \geq 1$ over $[0, a]$ gives $\log(1-a) \leq -a$. With $a = t/n < 1$, multiply by n and exponentiate to get the lower inequality in Eq. (219); for $t \geq n$ its left side is zero and the result is immediate. Similarly, integration of $1/(1+s) \leq 1$ gives $\log(1+a) \leq a$ for $a \geq 0$. Negate, multiply by n , and exponentiate to get $(1+t/n)^{-n} \geq e^{-t}$. The implementation uses $n = 1024$, computed by ten squarings. Downward rounding in the lower formula and a lower denominator followed by upward reciprocation in the upper formula preserve the enclosures. For a sampled source point, only the nonnegative lower formula is used. It supplies the lower-source condition in Proposition O.1. $\square$

O.5 Predeclared protocol and scope of its probability statement

There are six attempted insertions per object. Each attempt scouts the current retained material, rescores a shortlist, and fixes one finalist before collecting any certificate samples. The shortlist's scores are used only for selection; they supply no certificate. A separate validation stream evaluates that finalist at the four nested sample sizes $N \in \{1024, 4096, 16384, 65536\}$. We allocate $\delta_{tj\ell} = 1/4800$ to each of the $2 \cdot 6 \cdot 4 = 48$ possible tests. Their total is exactly 0.01. Proposition N.6 therefore covers the maximum of the four bounds, selection of the containing radius, and all accepted updates, conditional on the stated oracle assumptions.

For a conservative numerical concentration allowance, use

$$a_N = C\sqrt{8.5/(2N)} \geq C\sqrt{\log(4800)/(2N)}.$$

Table 2 New guarded designs, with one finalist per attempt and a total ideal-sampling error allocation of 0.01 across both objects and all 48 validation stages. Compliance is unnormalized. The two numerical references are approximations, not exact ground truth. Times cover the respective kernels over six attempts; reference evaluation, process startup, and hierarchy construction are excluded.

Quantity	Bunny	Housing
Accepted / attempted	6/6	6/6
Target per insertion	$2 \cdot 10^{-6}$	$1.5 \cdot 10^{-5}$
Cumulative lower benefit	$5.75 \cdot 10^{-4}$	$6.21 \cdot 10^{-3}$
WoS reference benefit	$8.448 \cdot 10^{-4}$	$1.020 \cdot 10^{-2}$
Multigrid reference benefit	$8.455 \cdot 10^{-4}$	$1.013 \cdot 10^{-2}$
Scout + rescore (s)	2.97	0.53
First passing prefixes (s)	2.82	1.48
Full validation sweep (s)	150.39	52.17

The constant needs no numerical logarithm: the positive exponential series gives $e^{17/2} > \sum_{k=0}^{25} (17/2)^k / k! > 4800$, where the last comparison is checked as an exact rational inequality. Monotonicity of $\log$ proves $\log 4800 < 8.5$. Sample means are recomputed by exact dyadic summation of the stored binary64 values, then rounded down. The allowance is rounded up, and the subsequent benefit expression is evaluated with lower arithmetic bounds. The π lower constant is obtained from Machin’s identity $\pi = 16 \arctan(1/5) - 4 \arctan(1/239)$ and alternating rational partial sums (12 terms for the first arctangent, 5 for the second). The alternating-series error signs give a lower first sum and upper second sum, hence a lower π . For completeness, the tangent double-angle formula gives $\tan(2 \arctan(1/5)) = 5/12$ and $\tan(4 \arctan(1/5)) = 120/119$. The tangent subtraction formula then gives

$$\tan(4 \arctan(1/5) - \arctan(1/239)) = 1.$$

This angle lies in $(0, \pi/2)$, so it equals $\pi/4$, proving the stated identity. These scalar arithmetic precautions do not bound the error in the random walk’s ideal sampling distribution.

The primary targets are fixed absolute compliance reductions: 2×10^{-6} for the Bunny and 1.5×10^{-5} for the housing. They equal 0.1% of stipulated reference scales 0.002 and 0.015, respectively. These scales are fixed design units, not confidence bounds on the true baseline compliance. Additional 0.5%, 1%, 2% thresholds are diagnostic tests on the same proposals and current domains, not separately optimized designs. A percentage of a noisy measured baseline is never used as a guaranteed relative target.

Only accepted cavities enter the next iteration’s material domain. For each proposed insertion, separate objective evaluations compare its current and trial domains with four randomized spatial blocks, 512 original-domain probes per block, and 256 WoS paths per active probe. Their differences estimate the actual finite benefit and its block standard error; these approximate references do not participate in acceptance. The original-domain volume normalization includes zero values inside removed cavities.

All four validation budgets are computed for a tightness diagnostic, even when an earlier budget already passes. We report that complete diagnostic cost. The first passing stage and its prefix time describe the smaller computation that a staged stopping policy could use; they do not erase the extra work actually executed. Selection, validation, hierarchy construction, and independent reference evaluation are kept distinct.

O.6 Measured validation results

Table 2 reports a separate six-attempt design on each object. Eleven insertions pass the fixed absolute target at the first predeclared stage of 1,024 walks; the fifth housing insertion first passes at 4,096 walks. Later samples tighten the reported benefit bounds; they are not required to make these particular accept/reject decisions. The maximum of the four covered stage bounds supplies each reported lower benefit. Its sum over accepted insertions is justified by Proposition N.6. Displayed cumulative lower bounds are rounded down, and the raw samples and full-precision bounds are retained.

The new designs reduce the independently estimated objective by 47.18% on the Bunny and 68.65% on the housing. These measured percentages use their respective new reference baselines and are not themselves

certified relative reductions. The full-precision cumulative bounds are 68.1% and 61.0% of the corresponding estimated absolute benefits. Per insertion, the average-trace bound is $2.4\text{--}16.9\times$ the temperature-floor bound on the Bunny and $59.4\text{--}365.4\times$ on the housing, evaluated on the same validation data and current domains. This compares two derived certificates, not two forward-solver runtimes or an external state-of-the-art optimization baseline.

For an additional deterministic check, geometric multigrid uses grids of $97 \times 113 \times 85$ on the Bunny and $113 \times 73 \times 57$ on the housing, with relative residual tolerance 10^{-8} . Its positive state values are recovered from the solver’s squared-state score output, and objectives use the same independent spatial probes as the WoS reference. Its total-benefit estimates differ from WoS by 0.085% and 0.645%, respectively. Every reported insertion bound lies below both approximate reference estimates. This agreement is a numerical consistency check; it does not bound either method’s continuum error or establish exact coverage.

The complete diagnostic collects 786,432 independent-walk records under the statistical model and executes 45,470,847 guarded steps with no containment failure. The audit recomputes all 48 state bounds from retained samples, checks the twelve initial classifications with exact rational predicates, compares conservative center distances with libigl, and tests 240 source evaluations using high-precision arithmetic. The geometric and source inequalities are justified by [Propositions O.2](#) and [O.3](#); the finite audit is an implementation check, not their proof.

The conservative query backend and large sample sweep are expensive. The first-passing prefix times in the table are recorded kernel times from the executed stages, illustrating what staged acceptance could save; the complete sweep actually ran. Neither those prefix times nor the earlier ranking-only speedups establish that fully validated optimization is faster than the deterministic baseline. Validation tightness, overhead, and the remaining ideal-sampling assumption are reported as separate properties.

P An Exact Three-Dimensional Finite-Cavity Check

This check separates the deterministic finite-benefit argument from Monte Carlo reference error. It is an analytic verification case, not an additional application benchmark or a validation of the floating-point random-walk law. Let $D = B_1(0) \subset \mathbb{R}^3$, let $f = 1$, and impose zero Dirichlet data on all boundaries. Remove the concentric ball $\overline{B}_a(0)$, $0 < a < 1$, and use compliance $J = \int u$. The source is restricted to the retained domain; this check therefore does not belong to the fixed-total-heat design comparison.

Proposition P.1 (Exact benefit and certificate gap). *Writing $s = |x|$, the original and perforated states are*

$$u(s) = \frac{1 - s^2}{6}, \tag{220}$$

$$v_a(s) = \frac{1 - s^2}{6} - c_a \left(\frac{1}{s} - 1 \right), \quad c_a = \frac{a(1 + a)}{6}, \quad a \leq s \leq 1. \tag{221}$$

The exact finite benefit $B_a = J(D) - J(D \setminus \overline{B}_a)$ is

$$\frac{B_a}{4\pi} = \frac{a(1 + a)^2(1 - a)}{36} + \frac{a^5}{45}. \tag{222}$$

The average-trace lower bound is $L_a = 4\pi a(1 - a^2)^2/36$, and its gap is

$$\frac{B_a - L_a}{4\pi} = \frac{a^2}{180} [5(1 - a^2) + a(5 - a^2)] > 0. \tag{223}$$

Moreover, for the relative capacity $\rho_a = 4\pi a/(1 - a)$,

$$\lim_{a \downarrow 0} (-B_a/\rho_a) = -1/36 = -u(0)^2.$$

Proof. Step 1: solve the radial equations. In three dimensions, $\Delta s^2 = 6$ and $\Delta(1/s) = 0$ for $s > 0$. Thus both displayed states satisfy $-\Delta u = -\Delta v_a = 1$ on their respective domains. The original state is regular at zero and vanishes at $s = 1$. The perforated state vanishes at $s = 1$, and

$$v_a(a) = \frac{1-a^2}{6} - \frac{a(1+a)}{6} \frac{1-a}{a} = 0.$$

Uniqueness of the Dirichlet problem identifies them as the states.

Step 2: integrate the objective directly. Polar integration gives $J(D)/(4\pi) = 1/45$. On the shell, an antiderivative of $v_a(s)s^2$ is

$$F_a(s) = \left(\frac{1}{6} + c_a\right) \frac{s^3}{3} - \frac{s^5}{30} - \frac{c_a s^2}{2}.$$

Consequently $B_a/(4\pi) = 1/45 - F_a(1) + F_a(a)$. Substitution of $c_a = a(1+a)/6$ and collection of powers of a gives

$$\frac{B_a}{4\pi} = \frac{a + a^2 - a^3 - a^4}{36} + \frac{a^5}{45},$$

which is [Eq. \(222\)](#).

Step 3: verify the difference-energy identity independently. Extend v_a by zero inside the cavity and put $w = u - v_a$. Then $w'(s) = -s/3$ for $s < a$, and $w'(s) = -c_a/s^2$ for $a < s < 1$. It follows that

$$\frac{1}{4\pi} \int_D |\nabla w|^2 = \int_0^a \frac{s^4}{9} ds + c_a^2 \int_a^1 s^{-2} ds \quad (224)$$

$$= \frac{a^5}{45} + c_a^2(1/a - 1) = \frac{B_a}{4\pi}. \quad (225)$$

The contribution inside the removed ball is necessary; dropping it would not give the exact finite benefit.

Step 4: evaluate and compare the trace certificate. The trace is constant on the sphere: $\bar{u}_a = (1 - a^2)/6$. The center Green integral is exactly $a^2/6$, so $u(0) - a^2/6 = \bar{u}_a$. Thus the exact-state version of [Theorem A.6](#) returns $L_a = 4\pi a \bar{u}_a^2$. Subtracting it from Step 2 yields [Eq. \(223\)](#). Both terms in its square brackets are positive for $0 < a < 1$.

Step 5: check the infinitesimal normalization. The relative equilibrium potential equals one for $s \leq a$ and $q_a(s) = a(1/s - 1)/(1 - a)$ for $a < s < 1$. Integrating its squared radial derivative gives

$$\rho_a = 4\pi \frac{a^2}{(1-a)^2} \int_a^1 s^{-2} ds = \frac{4\pi a}{1-a}.$$

Dividing [Eq. \(222\)](#) by this expression and taking $a \downarrow 0$ gives the claimed $-1/36$ derivative. This also distinguishes relative capacity from the free-space ball capacity $4\pi a$ at finite radius. $\square$

We check the boundary trace, the two independent expressions for benefit, and the certificate gap using exact rational arithmetic at $a = i/1000$, $i = 1, \dots, 999$. All algebraic residuals are exactly zero. The finite-benefit lower bound retains 99.00%, 95.00%, 89.99%, and 74.80% of the exact benefit at radii 0.01, 0.05, 0.10, and 0.25, respectively. These values describe this symmetric model only. They neither establish certificate tightness on general meshes nor verify sampling coverage for the implemented validator.

Q Claim Scope and Paired Comparisons

The deterministic topological derivative and the off-diagonal product construction are established ingredients. The claimed contribution is their use as a direct grid-free insertion query, the accompanying error accounting, and its evaluated integration into finite geometric decisions. The exact sphere check in [Appendix P](#) is verification evidence, not a claim that the radial Poisson solution is new.

Table 3 Paired mean cooling difference, in percentage points, with pointwise 95% joint bootstrap intervals. MG uses longest-axis grid count 65.

Geometry	TWoS minus MG	TWoS minus nearest	TWoS minus hottest
Finned housing	-0.11 [-1.24, +0.99]	+1.43 [-0.32, +3.39]	+22.25 [+12.76, +37.94]
Branching manifold	+0.14 [-0.74, +1.63]	-0.93 [-1.40, -0.55]	+16.26 [+15.22, +17.43]
Annular spreader	-0.32 [-1.11, +0.18]	-0.84 [-1.99, +0.66]	+11.35 [+10.39, +12.24]

Three different levels of guarantee. The direct-score unbiasedness and variance statements concern ideal marginal PDE oracles and independent complete pairs. Finite-tolerance runs approximate those oracles. The general-objective insertion theorem additionally requires numerically usable bounds on marginal bias and the finite-radius remainder; we have not supplied those bounds for every displayed mesh and objective. The specialized Poisson-compliance validator instead avoids an unknown remainder through a variational bound and handles truncation through a one-sided construction. Its probability guarantee still assumes the stated conditional sampling laws; the floating-point implementation is not a verified sampler. The design pipelines and channel-placement heuristic are not certified by either theorem. Confidence bands in the plots quantify sampling uncertainty, not those missing numerical error budgets.

Paired differences distinguish placement value from solver value. Table 3 uses the existing paired bootstrap outputs, without selecting new seeds or rerunning optimization. Differences compare TWoS after eight refinement rounds with multigrid/proximity after eight, and hottest-point placement after sixteen. Positive values favor TWoS. The intervals are pointwise and exclude numerical bias; they are not simultaneous comparisons or equivalence tests. An interval containing zero does not prove that the methods are equivalent. The uncertainty and reference limitations in the main text apply to every entry.

R Matched-Accuracy Sparse-Query Workloads

Question and frozen protocol. The comparison asks when evaluating a small set of topological insertion scores is cheaper than constructing a reusable volumetric primal–adjoint field. It does not presume that a grid-free method is faster. Before the measured sweep, we fixed four geometries, query counts 8, 32, 128, 512, five random seeds, TWoS path budgets 256, 1024, 4096, and longest-axis multigrid counts 49, 65, 97, 129. The objects are the existing housing, branching manifold, annular spreader, and the 69,674-triangle Bunny. All use distinct compact primal and adjoint sources. The Bunny has sensor center (0.15, −0.15, 0) and radius 0.07, a primary source centered at (−0.15, −0.15, 0) with radius 0.085 and amplitude 35, and sensor heating with radius 0.06 and amplitude 4. Thus this test differs from the earlier Bunny compliance crossover; it does not retrospectively replace its objective.

A scrambled Sobol pool with seed 17092026 fixes feasible sites without consulting scores. The last admissible pool point determines a radius-0.04 cavity. The 512 query sites avoid this prospective cavity by a center-distance margin of 0.10. Both methods evaluate the original domain and the domain with this fixed cavity. The edit is a controlled workload perturbation, not an optimized insertion or a claimed improvement in objective value. Nested batches contain 8, 24, 96, 384 new sites, giving the stated prefix counts. Five paired random seeds are 1611, 2711, 3811, 4911, 6011.

Giving the grid its legitimate reuse. The multigrid baseline uses the same masked finite-difference operator and geometric V-cycle as the previous study. Its primal and adjoint share the geometry hierarchy. The new query backend retains both solved finest-grid fields, so additional sites on an unchanged domain require only trilinear interpolation. A cavity insertion invalidates the cached fields; the common outer-mesh acceleration structure remains available to both methods. Each method–budget–seed combination has a fresh process, preventing peak memory from being inherited from another method. We measure mesh setup separately and cumulative wall time through every prefix, including field construction where incurred, request handling, and output writing. These measurements describe this implementation, not an optimized lower bound on multigrid query latency. All timing runs execute serially on four threads; method order rotates with seed. Independent reference construction is excluded from each method’s query cost. Cached repeats must reproduce all primal, adjoint, and score values within eight machine epsilons after CSV serialization and

make no geometry queries. Across 360 repeat checks the maximum absolute discrepancy is 1.36×10^{-20} .

Accuracy and reference gates. For a query prefix of size Q , let s^{257} denote the score vector from MG257 and $\hat{s}_{m,b,r}$ the tested method m , budget b , and seed r . We measure

$$e_{m,b,r}(Q) = \frac{\|\hat{s}_{m,b,r} - s^{257}\|_2}{\|s^{257}\|_2}.$$

This is reference-relative numerical error, not known continuum error. In particular, the grid reference and multigrid baseline share a discretization family. We therefore report a refinement discrepancy $\delta_Q = \|s^{193} - s^{257}\|_2 / \|s^{257}\|_2$ and an independent Monte Carlo check at 32 fixed, evenly indexed query sites. Eight independent 8,192-path batches give mean $\bar{s}$ and componentwise standard errors. Let η be the norm of that standard-error vector divided by the norm of the grid reference on those sites, and d the similarly normalized $\|\bar{s} - s^{257}\|_2$. For a target τ , the declared diagnostic gate is

$$\delta_Q \leq \tau/4, \quad \eta \leq \tau/4, \quad d \leq \tau/4 + 2\eta.$$

These gates are diagnostic tolerances, not probabilistic error certificates: the factor two is not a proved multivariate confidence radius, and the spot check does not cover every site. Grid convergence alone cannot rule out shared bias. A zero or nonfinite reference norm is unresolved. The primary target is $\tau = 0.20$; 0.10 and 0.40 are predeclared sensitivity targets. No gate is relaxed after inspecting results.

Uncertainty and selection. We resample the five paired seeds 10,000 times with fixed bootstrap seed 20260917, reporting pointwise percentile intervals for mean relative error and mean time. A budget is eligible only if its upper 97.5th percentile for mean error is at most τ and the reference gates pass. We select the fastest eligible tested budget within each method. Missing eligible budgets are reported explicitly, never converted to infinite speedups. Intervals for selected runtime ratios resample paired seeds jointly; they are descriptive after selection and are not simultaneous or selection-adjusted confidence statements. Timing and sampling variation do not enclose termination, floating-point, geometric, or discretization bias. All tested budgets, reference diagnostics, and unsuccessful targets accompany the paper.

Measured results. Table 4 reports the predeclared 20% target. A dash means an unresolved reference or no eligible tested budget; it never denotes an infinite speedup. The analytic-source and incremental-grid extensions were introduced after the first housing result and frozen before their own measurements. Their results are therefore separate diagnostic extensions, not retroactively part of the initial protocol.

Fresh-seed confirmation and its failures. The four apparent 20%-target wins selected from calibration did not pass both fresh-seed rules. The three compliance cases exceeded the mean-error upper-endpoint threshold; the housing timing interval included no advantage. We then declared a separate, fixed 256-path follow-up for both eight-site compliance states, with 20 further fresh seeds and no subsequent budget search. All results are retained below. These pointwise bootstrap checks are not simultaneous confidence guarantees.

Random-stream indexing audit and correction. The retained legacy implementation XORs hashes of a seed, site index, and walk index. When a driver seed equals a walk-index offset, this construction can cancel or swap identifiers and reuse a stream across nominal repetitions. For the present five seeds and four query batches, this affects twelve indices at the 4,096-path budget; budgets through 1,024, the high-seed reference batches, and the fresh holdout seeds do not have that structural overlap. We correct the affected calibration draws, retain the superseded outputs, and enumerate the resulting stream keys to check uniqueness. This is a disclosed correctness amendment, not a selection based on outcomes.

Let I be the fixed set of affected indices. We replay their original paths and replacement paths with unused indices, obtaining (U_i, P_i) and (U'_i, P'_i) . The corrected sufficient statistics are

$$\begin{aligned} S'_U &= N\bar{U} + \sum_{i \in I} (U'_i - U_i), & S'_P &= N\bar{P} + \sum_{i \in I} (P'_i - P_i), \\ S'_{UP} &= \sum_i U_i P_i + \sum_{i \in I} (U'_i P'_i - U_i P_i), & \hat{S}' &= \frac{S'_U S'_P - S'_{UP}}{N(N-1)}. \end{aligned}$$

Table 4 Calibration results at the 20% reference-relative accuracy target. Times include every batch through the listed query count. MG is the fastest eligible tested resolution and reuse variant. RB denotes exact-source integration in ideal arithmetic. All quantities are empirical; source restrictions and numerical-reference gates apply.

Object	State	Queries	TWoS (s)	TWoS-RB (s)	MG (s)	MG/RB	RB paths
Branch	edited	8	0.997	0.823	0.049	0.06	4096
Branch	edited	512	53.502	13.063	0.053	0.00	1024
Branch	original	8	0.982	0.877	0.090	0.10	4096
Branch	original	512	49.404	11.476	0.095	0.01	1024
Bunny	edited	8	1.635	0.409	0.122	0.30	1024
Bunny	edited	512	77.790	17.474	0.128	0.01	1024
Bunny	original	8	1.406	0.322	0.201	0.62	1024
Bunny	original	512	71.133	18.535	0.208	0.01	1024
Bunny compliance	edited	8	–	0.023	0.073	3.17	64
Bunny compliance	edited	512	–	4.372	0.079	0.02	256
Bunny compliance	original	8	–	0.028	0.116	4.19	64
Bunny compliance	original	512	–	4.430	0.120	0.03	256
Housing	edited	8	–	0.084	0.031	0.37	1024
Housing	edited	512	22.283	4.875	0.036	0.01	1024
Housing	original	8	–	0.084	0.085	1.02	1024
Housing	original	512	21.718	4.595	0.089	0.02	1024
Ring	edited	8	–	0.255	0.039	0.15	1024
Ring	edited	512	54.190	11.832	0.043	0.00	1024
Ring	original	8	–	0.239	0.102	0.43	1024
Ring	original	512	49.844	12.061	0.109	0.01	1024

Table 5 Independent reference diagnostics. Refinement discrepancies are shown at 8 and 512 sites; the 32-site Monte Carlo diagnostic is common to each state. Values are percentages. These quantities are not rigorous error bounds.

Object	State	δ_8	δ_{512}	MC SE norm	MC-grid discrepancy
Branch	edited	1.34	1.29	3.07	5.37
Branch	original	1.34	1.29	3.35	5.62
Bunny	edited	0.84	0.74	3.79	2.45
Bunny	original	0.84	0.74	3.86	2.64
Bunny compliance	edited	1.16	1.27	0.57	2.49
Bunny compliance	original	1.16	1.27	0.56	2.44
Housing	edited	0.70	3.03	2.87	6.11
Housing	original	0.70	3.03	2.87	6.11
Ring	edited	1.36	1.03	2.98	7.01
Ring	original	1.36	1.03	2.98	7.01

To verify the reconstruction, split each original sum into indices inside and outside I : subtraction removes exactly the old inside terms, and addition supplies exactly their replacements. Thus these are precisely the sums of the corrected N pairs in exact arithmetic. Selection of I depends only on stream identities, so it does not condition on observed path values. Under the ideal independent-stream model, unique keys give the intended independent pairs and the usual off-diagonal-product proof applies. Distinct PRNG keys alone are not a mathematical proof of PRNG independence. We check floating-point reconstruction against complete clean runs on each affected workload; all replay work is charged to corrected query costs.

Fresh-seed confirmation of apparent speed wins. Budgets selected from five calibration seeds are not independently validated by those same seeds. Before measuring the extensions, we therefore fix a confirmation protocol: every apparent 20%-target speed win retains its selected configurations and uses 20 fresh seeds, alternating method order. The upper bootstrap mean-error endpoint must remain below 20%, and the lower paired runtime-ratio endpoint must exceed one, to call the result a confirmed win. Failed confirmations remain reported. The same numerical reference gates apply; these pointwise checks are neither simultaneous

Table 6 *Fresh-seed query confirmations. Times exclude common geometry loading. Error upper endpoints and speed intervals use 10,000 paired bootstrap resamples of 20 seeds. Passing requires error upper endpoint $\leq 20\%$ and speed-ratio lower endpoint > 1 . Every row uses MG49.*

Workload / state	Stage	Sites	Paths	TWoS / MG (ms)	Error upper	MG/TWoS interval	Pass
Bunny compliance / edited	Selected	8	64	23.3 / 78.4	21.8%	[2.51, 4.57]	no
Bunny compliance / original	Selected	8	64	28.7 / 341.3	21.3%	[6.21, 19.36]	no
Bunny compliance / original	Selected	32	64	74.0 / 136.8	20.8%	[1.58, 2.13]	no
Housing / original	Selected	8	1024	137.4 / 131.9	13.6%	[0.58, 1.54]	no
Bunny compliance / edited	Fixed 256	8	256	83.1 / 72.8	10.3%	[0.69, 1.11]	no
Bunny compliance / original	Fixed 256	8	256	89.9 / 113.9	10.2%	[1.05, 1.50]	yes

tests nor continuum-error certificates. Offline tuning and reference construction are excluded from query costs, so this comparison is not an online accuracy controller.




Figure 9 All measured budgets on the original domain, at 8 sites (top) and 512 sites (bottom). Bars are pointwise 95% bootstrap intervals over five seeds; they exclude numerical-reference bias. Dashed horizontal lines mark the 20% target before applying the independent reference gates. Numbers label path counts or longest-axis grid counts. Original and source-integrated TWoS and both grid reuse variants are retained.




Figure 10 All measured budgets on the edited domain, at 8 sites (top) and 512 sites (bottom). Bars are pointwise 95% bootstrap intervals over five seeds; they exclude numerical-reference bias. Dashed horizontal lines mark the 20% target before applying the independent reference gates. Numbers label path counts or longest-axis grid counts. Original and source-integrated TWoS and both grid reuse variants are retained.




Figure 11 Measured query-count scaling at fixed budgets in the initially frozen sweep. Shaded bands show pointwise 95% bootstrap intervals for mean time. Original-source TWoS scales with the number of sites, whereas solved-field multigrid amortizes its setup. These fixed-budget curves do not assert matched accuracy.

S A Source-Specific Rao–Blackwellized Variant

This extension was added after the first housing comparison exposed poor accuracy–cost performance of the original source-point estimator. It is reported as a separate ablation, with the original measurements retained. Analytic source integration and Rao–Blackwellization are standard variance reduction ideas; we make no priority claim for either. The extension applies to the prescribed radial polynomial loads in these benchmarks, not arbitrary spatial source fields.

Step 1: construct a particular solution. For $A \geq 0$ and $R > 0$, consider one source $f(x) = A(1 - \|x - c\|^2/R^2)_+^2$, set $d = \|x - c\|$ and

$$\Psi(d) = \begin{cases} AR^2\left(\frac{1}{6} - \frac{d^2}{6R^2} + \frac{d^4}{10R^4} - \frac{d^6}{42R^6}\right), & d \leq R, \\ \frac{8AR^3}{105d}, & d > R. \end{cases}$$

Inside the source, direct differentiation gives

$$\Psi'(d) = -A\left(\frac{d}{3} - \frac{2d^3}{5R^2} + \frac{d^5}{7R^4}\right), \quad -(d^2\Psi'(d))' = d^2f(d).$$

Thus $-\Delta\Psi = f$ for $0 < d < R$. At the origin the derivative vanishes and the polynomial is smooth. Outside R , $1/d$ is harmonic. At $d = R$, the two values equal $8AR^2/105$ and the two derivatives equal $-8AR/105$. Consequently there is no distributional surface source at the join, and $-\Delta\Psi = f$ globally. Linearity permits summing sources and using the normalized sensor bump for the adjoint.

Step 2: integrate a walk ball exactly. Here r denotes a walk-ball radius, not the inserted-cavity radius. For the radius- r ball centered at the current walk point x , the Poisson representation gives the nonnegative source contribution

$$\begin{aligned} I(x, r) &= \int_{B_r(x)} G_{B_r(x)}(x, y) f(y) dy \\ &= \Psi(\|x - c\|) - \frac{1}{4\pi} \int_{\mathbb{S}^2} \Psi(\|x + r\omega - c\|) d\omega. \end{aligned}$$

This follows by subtracting the harmonic extension of Ψ from Ψ in the ball; their difference has zero boundary data and source f . For $d > 0$, parameterize the polar angle by $t = \cos \theta$. With $s^2 = d^2 + r^2 + 2drt$, we have $dt = s ds/(dr)$ and therefore

$$\frac{1}{4\pi} \int_{\mathbb{S}^2} \Psi(\|x + r\omega - c\|) d\omega = \frac{1}{2dr} \int_{|d-r|}^{d+r} s\Psi(s) ds.$$

An antiderivative for $s \leq R$ is

$$F(s) = A\left(\frac{R^2s^2}{12} - \frac{s^4}{24} + \frac{s^6}{60R^2} - \frac{s^8}{336R^4}\right).$$

For $s > R$, continue it by $F(R) + 8AR^3(s - R)/105$. Hence the mean is $[F(d + r) - F(|d - r|)]/(2dr)$. For $d = 0$ it is simply $\Psi(r)$. If $d \geq r + R$, the source support misses the ball and $I = 0$ exactly. If $d + R \leq r$, the sphere mean is $8AR^3/(105r)$. These cases also avoid unnecessary cancellation in the implementation.

Step 3: retain the same score target. Let $\mathcal{H}_i$ contain the entire sphere-exit path of walk i , including its termination, and let (U_i, P_i) denote the original source-sampled primal–adjoint pair. Source-point draws are independent of the sphere-exit draws in the ideal sampling model. Exact ball integration yields $\bar{U}_i = \mathbb{E}[U_i \mid \mathcal{H}_i]$ and $\bar{P}_i = \mathbb{E}[P_i \mid \mathcal{H}_i]$. Different walks remain independent. For $i \neq j$, conditional independence of their source draws gives

$$\mathbb{E}[U_i P_j \mid \mathcal{H}_1, \dots, \mathcal{H}_N] = \bar{U}_i \bar{P}_j.$$

Summing and dividing by $N(N - 1)$ proves

$$\hat{S}_{\text{RB}} = \frac{\sum_{i \neq j} \bar{U}_i \bar{P}_j}{N(N - 1)} = \mathbb{E}[\hat{S} \mid \mathcal{H}_1, \dots, \mathcal{H}_N].$$

The tower property preserves expectation. For finite second moments, total variance gives the precise guarantee

$$\text{Var}(\hat{S}) = \text{Var}(\hat{S}_{\text{RB}}) + \mathbb{E}[\text{Var}(\hat{S} \mid \mathcal{H}_1, \dots, \mathcal{H}_N)] \geq \text{Var}(\hat{S}_{\text{RB}}).$$

This compares the same number of independent exit paths under exact source integration; it does not guarantee lower wall time. It also preserves any termination bias of the underlying walk model rather than removing it. The earlier ideal-arithmetic variance upper bounds therefore remain valid for this conditional estimator under the same moment assumptions.

Numerical implementation and scope. The implementation evaluates these expressions in long double and clamps small negative roundoff to zero. Very small d/R uses the centered limit. These are numerical approximations to the ideal conditional expectation; the variance inequality is not a finite-precision certificate. An independent quadrature check accompanies the source. The measured extension uses the same sites, seeds, accuracy gates, and reference fields as the original study. Before measuring this variant, we additionally fix path budgets 16 and 64 alongside the original 256, 1,024, and 4,096, to test whether reduced variance permits cheaper queries. It is compared to a stronger multigrid extension that retains outer geometry masks and warm-starts the primal and adjoint after cavity insertion. Both extensions are explicitly distinguished from the initially frozen sweep.

Checking the earlier distributed-heating regime. A separate Bunny experiment uses the earlier PDE type $f(x) = 1 + 0.65x_y$ and compliance $J = \int_D f u$, so $p = u$. It uses the new fixed query sites and cavity, not the exact historical site pool. Here the sphere-centered Green kernel is radial. Odd moments vanish, hence

$$\int_{B_r(x)} G(x, y) f(y) dy = f(x) \int_{B_r(x)} G(x, y) dy = \frac{r^2}{6} f(x).$$

The final equality follows by integrating $\int_0^r (t - t^2/r) dt = r^2/6$. Thus affine-source integration is exact in ideal arithmetic without using the radial-bump formula. TWoS estimates u^2 with the same off-diagonal product. The 257-node reference required increasing the maximum V-cycle count from 120 to 800 while retaining the 10^{-7} residual tolerance; the same cap is available to every grid budget in this compliance workload. The multigrid baseline solves only one field and reuses it for the adjoint; it also retains its mask and warm-starts after insertion. This diagnostic was added after the three engineering-geometry results, with the same five seeds, query counts, reference gates, and targets. It uses the predeclared path budgets 16, 64, 256, 1024, 4096. Since cavity removal restricts the volumetric heat source, it is a query comparison, not a fixed-total-heat design study.

Predeclared extension stopping rule. Before measuring either source-integrated extension, we fix the following rule. Budgets 16, 64, and 256 are measured at every query count and both states. Budgets 1,024 and 4,096 are measured only through the largest remaining unresolved prefix in each state. A case closes when a previous budget meets the strictest reference-resolved target among 10%, 20%, and 40%, using the unchanged upper bootstrap mean-error rule. A case with no resolved reference target remains explicitly unresolved and does not trigger more sampling. The saved records include every stopping decision. Thus the extensions identify target-attaining budgets; they are not exhaustive high-budget error-cost fronts. The original four-workload sweep is unchanged.

T A Distinct Adjoint for Regional Thermal Protection

T.1 Objective, adjoint, and the insertion score

Fix a nonnegative weight q supported in a protected region $S \Subset D$. Require all admissible cavities to be disjoint from S and set $Q = \int_S q > 0$. The physical source f and the observation weight q serve different roles. For the zero-Dirichlet Poisson state define

$$-\Delta u_D = f \quad \text{in } D, \quad J_q(D) = Q^{-1} \int_S q u_D. \quad (226)$$

Let p_D solve $-\Delta p_D = q$ in D with zero boundary data. This unnormalized adjoint is convenient in the implementation; the actual adjoint for J_q is p_D/Q .

Proposition T.1 (Regional sensitivity and monotonicity). *Under the smoothness and separation assumptions of Theorem 3.5, the capacity-normalized derivative for a zero-temperature cavity at $z \notin S$ is*

$$D_T J_q(z) = -u_D(z)p_D(z)/Q. \quad (227)$$

For finite nested domains obtained by removing cavities outside S , J_q is nonincreasing. This monotonicity does not establish a minimum benefit, or certify a particular finite-radius prediction.

Proof. Step 1: differentiate the objective in the state. For every $v \in H_0^1(D)$, the linear objective has derivative $j'(u)[v] = Q^{-1} \int_D qv$. The adjoint weak equation is therefore $\int_D \nabla v \cdot \nabla(p_D/Q) = Q^{-1} \int_D qv$. The fixed support and cavity exclusion keep both q and Q unchanged by every admissible insertion.

Step 2: apply the cavity expansion with this adjoint. The state trace imposed on the inserted cavity is zero. Substitution of $h_z = 0$ and adjoint p_D/Q into the established local coefficient $-(u_D(z) - h_z)p(z)$ gives Eq. (227). Multiplication by the positive constant Q does not change candidate rankings. This specialization uses the classical coefficient; it is not a claim that the underlying topological formula is new.

Step 3: compare finite-domain states.

Write $u_D^- = \max(-u_D, 0) \in H_0^1(D)$. The weak equation gives

$$-\int_D |\nabla u_D^-|^2 = \int_D f u_D^- \geq 0.$$

Both the gradient norm and the nonnegative right-hand side must vanish. Poincaré's inequality and the zero trace imply $u_D^- = 0$, hence $u_D \geq 0$. For $D' = D \setminus K$, the difference $w = u_D - u_{D'}$ is weakly harmonic in D' , has zero outer trace, and has nonnegative trace u_D on the new cavity. Consequently $w^- = \max(-w, 0)$ has zero trace on all of $\partial D'$ and belongs to $H_0^1(D')$. The harmonic weak equation gives $-\int_{D'} |\nabla w^-|^2 = 0$. Poincaré's inequality again yields $w^- = 0$, or $w \geq 0$. Hence $J_q(D) - J_q(D') = Q^{-1} \int_S qw \geq 0$. Applying this argument at each insertion proves the finite nested-domain statement. It supplies no positive quantitative lower bound beyond zero. $\square$

T.2 An exact counterexample to hottest-point placement

Proposition T.2 (A hottest point need not maximize topological benefit). *On the unit ball in $\mathbb{R}^3$, take $f = 1$ and $q(x) = 1 + bx_1$ for $0 < |b| < 1$. For the unnormalized linear objective $\int qu$, the hottest point is the origin, but the origin does not locally maximize the insertion benefit up.*

Proof. Step 1: solve the state. Writing $r = |x|$, the function $u = (1 - r^2)/6$ vanishes on the boundary and satisfies $-\Delta u = 1$, because $\Delta r^2 = 6$. Uniqueness of the weak Dirichlet solution identifies this function as the state. It is strictly decreasing in $r > 0$, so its unique maximum is at the origin.

Step 2: solve the adjoint. The product rule gives $\Delta(x_1 r^2) = x_1 \Delta r^2 + 2\nabla x_1 \cdot \nabla r^2 = 6x_1 + 4x_1 = 10x_1$. Consequently

$$p(x) = (1 - r^2) \left(\frac{1}{6} + \frac{bx_1}{10} \right)$$

has zero trace and satisfies $-\Delta p = 1 + bx_1 = q$. Again uniqueness identifies the adjoint.

Step 3: inspect the benefit at the hottest point. Direct multiplication gives

$$u(x)p(x) = (1 - r^2)^2 \left(\frac{1}{36} + \frac{bx_1}{60} \right), \quad \partial_{x_1}(up)(0) = \frac{b}{60} \neq 0.$$

Moving a sufficiently small distance in the direction $\text{sign}(b)e_1$ therefore increases up , while it strictly decreases u . Thus hottest-point ranking and topological ranking are distinct even for the same self-adjoint Poisson operator. This analytic example establishes a mechanism, not the size of the benefit in the housing experiment. $\square$

T.3 Paired random walks and what cross-fitting removes

At a candidate, one walk accumulates both the physical-source contribution U_i and the observation-source contribution P_i . The spherical exits and Green-distributed source samples are shared within a pair. Different complete pairs use separate random streams. For $N \geq 2$ the implementation evaluates

$$\hat{s} = \frac{(\sum_i U_i)(\sum_i P_i) - \sum_i U_i P_i}{N(N-1)}. \quad (228)$$

Proposition T.3 (Expectation of the implemented product). *For independent identically distributed pairs with finite second moments, write $\mu_U = \mathbb{E}U_i$, $\mu_P = \mathbb{E}P_i$. Then $\mathbb{E}\hat{s} = \mu_U \mu_P$, whereas $\mathbb{E}[N^{-1} \sum_i U_i P_i] = \mu_U \mu_P + \text{Cov}(U_i, P_i)$.*

Proof. Expanding the product of sums and subtracting its diagonal leaves $\sum_{i \neq j} U_i P_j$. There are exactly $N(N-1)$ ordered distinct pairs. For each, independence gives $\mathbb{E}[U_i P_j] = \mathbb{E}U_i \mathbb{E}P_j = \mu_U \mu_P$. Divide the sum by its number of terms to obtain the first assertion. For the diagonal, the defining identity for covariance is $\mathbb{E}[U_i P_i] = \mu_U \mu_P + \text{Cov}(U_i, P_i)$; averaging preserves it. $\square$

Cross-fitting removes product-correlation bias. It does not remove WoS termination error. In the ideal contained-ball model with nonnegative sources, dropping the terminal values makes each expectation a lower approximation of its state. The measured code uses the existing approximate distance and finite-precision sampling implementation and does not assert a rigorous bias magnitude. The specialized compliance certificate cannot be substituted into Eq. (226): its square-energy identity is specific to compliance. These new designs use fresh ranking at a fixed count. Radius-independent variance does not imply independence from the observation region: a very small support of q can itself make source sampling difficult. The regional study fixes that support before varying the estimator budget. Its source-sampling difficulty is distinct from encountering a shrinking trial cavity.

T.4 Sensor normalization and protected geometry

The two observation regions are radius- $a = 0.13$ balls centered at $(\pm 0.55, \pm 0.18, -0.03)$, with the signs chosen together. Define $q(x) = (1 - |x - c|^2/a^2)_+^2$. Radial integration gives

$$\begin{aligned} Q &= 4\pi a^3 \int_0^1 t^2 (1 - t^2)^2 dt \\ &= 4\pi a^3 \left(\frac{1}{3} - \frac{2}{5} + \frac{1}{7} \right) = \frac{32\pi a^3}{105}. \end{aligned}$$

The sensor-ball volume is $4\pi a^3/3$, so its average weight is $8/35$. Uniform sensor-ball probes therefore estimate J_q as the sample mean of qu divided by $8/35$. Candidate centers obey $|z - c| \geq a + r + 0.02$; by the triangle inequality every point of the radius- r cavity is at least 0.02 away from the protected ball. Thus reducing the objective cannot be accomplished by deleting the monitored region.

T.5 Replicated designs and numerical reference checks

The benchmark retains all final geometries, per-round candidate fields, solver residuals, work counts, and timings. Five seeds change the shared candidate set and stochastic streams. The two target regions are evaluated separately on the same housing. Each final objective uses four independent scrambled Sobol blocks of 128 uniform sensor-ball probes, with 256 fresh WoS paths per probe. Evaluation streams are independent of selection and shared between methods for paired comparisons. The table reports variation across complete designs, not a confidence guarantee for arbitrary new inputs. The separate reference solver evaluates the state at the same probes. Across first-seed designs, the fourfold shell refinement changes measured reduction by at most 0.087 percentage points. Refining the evaluation grid from $169 \times 109 \times 85$ to $225 \times 145 \times 113$ changes it by at most 0.255 points. Across all primary and supplementary designs, WoS

Table 7 Regional design comparison. Reductions are means $\pm$ one standard error across five designs per sensor. Both wall-time columns include candidate preparation and the complete selection driver, averaged across sensor cases. Retained-mesh timings use one fresh server per design; dashes denote configurations not rerun with that backend. The extra multigrid resolutions are a supplementary sweep.

Method	Left reduction (%)	Right reduction (%)	Reload mesh (s)	Retain mesh (s)
Random	10.55 ± 1.54	10.23 ± 3.14	0.44	0.39
WoS hottest	36.47 ± 2.22	29.26 ± 2.24	3.63	1.19
MG hottest	38.15 ± 1.03	29.99 ± 1.77	6.07	4.88
TWoS sparse	50.87 ± 1.46	50.09 ± 1.34	3.66	1.22
TWoS larger budget	52.13 ± 1.08	51.84 ± 0.52	5.04	2.62
MG TD (57)	53.50 ± 0.62	52.53 ± 0.84	2.51	1.41
MG TD (85)	53.57 ± 0.66	52.45 ± 0.82	4.54	–
MG TD (113)	53.30 ± 0.57	52.15 ± 0.88	8.38	7.13
MG TD (169)	53.57 ± 0.66	52.15 ± 0.88	40.00	–

and the $169 \times 109 \times 85$ reference differ by at most 1.864 points. These sensitivity checks do not bound the remaining discretization or sampling bias.

The primary design protocol and sampling budgets were fixed before the first completed seed. That seed was inspected before the other four ran. The supplementary grid resolutions were specified after that pilot and before their measurements. All are retained. No improvement in the cooling model or guarantee of accessible cooling follows from changing the observation objective. The theoretical counterexample above and the observed improvement over hottest-point placement show why temperature alone need not characterize objective sensitivity. The adjoint supplies the missing observation dependence in the topological formula; deterministic topological derivatives exploit the same information.

U Accessible Ball-Union Channels

U.1 A geometric access guarantee

The following construction gives an explicit passage through an object. It does not assert a cooling rate, pressure drop, or manufacturing tolerance. Let e be a unit vector, let $m \geq 1$ be an integer, and put $c_j = c_0 + jhe$ for $j = 0, \dots, m$, where $0 < h < 2r$. Define the open channel $C = \bigcup_{j=0}^m B_r(c_j)$ and $\rho = \sqrt{r^2 - h^2/4} > 0$.

Proposition U.1 (A positive-width passage). *The channel C is path connected and contains the straight tube*

$$T = \{c_0 + te + v : 0 \leq t \leq mh, v \cdot e = 0, |v| < \rho\}.$$

If c_0 and c_m lie strictly on opposite sides of a slab containing the closed object, the removed set C supplies an open route between the two exterior sides. In particular, the route has positive width; it is not merely a point-contact chain.

Proof. Step 1: choose a neighboring center. For each $t \in [0, mh]$, rounding t/h to a nearest integer gives $j \in \{0, \dots, m\}$ with $|t - jh| \leq h/2$. **Step 2: bound the distance.** Orthogonality of v and e gives

$$|c_0 + te + v - c_j|^2 = (t - jh)^2 + |v|^2 < h^2/4 + \rho^2 = r^2.$$

Thus every point of T belongs to an open ball in C . **Step 3: connect all channel points.** Any point of a ball can be connected to its center by a segment in that ball. The centerline belongs to T . Concatenating two radial segments with a centerline segment connects any two points of C . **Step 4: reach the exterior.** The endpoint centers lie in the exterior by the slab assumption. Following the centerline from one endpoint to the other therefore gives a route in C , including its portions removed from the object. The strict tube

inequality establishes positive transverse clearance. Intersecting the tube with the object need not give a single component on an arbitrary shape, but every such portion communicates with the same open route. This statement does not require the remaining material to be connected. $\square$

Proposition U.2 (Protection and channel separation). *Let $a > 0$ and $g \geq 0$. For parallel channels of radius r , a transverse distance at least $r + a + g$ from a protected ball’s center to the channel axis gives a gap of at least g from that radius- a ball. An axis separation of at least $2r + g$ gives a gap of at least g between channels.*

Proof. **Step 1: enclose the channel.** Every channel ball lies in the radius- r cylinder around its axis, so the union lies in that cylinder. **Step 2: project protected points.** Orthogonal projection onto the plane perpendicular to the axis cannot increase distances. A point in the protected ball projects at distance at most a from its center’s projection. A channel point projects at distance at most r from its axis. The reverse triangle inequality therefore bounds their separation below by $(r + a + g) - a - r = g$. **Step 3: apply the same argument to two channels.** Their two projected radial distances are each at most r , giving the lower bound $(2r + g) - 2r = g$. $\square$

U.2 What the construction does not prove

For the nonnegative-source, zero-Dirichlet model and a protected observation region, domain removal still gives a nonincreasing regional objective, provided the weak solutions and trace comparison used in [Appendix T](#) apply. To see this, set $D' = D \setminus \overline{C}$ and use exactly the difference $u_D - u_{D'}$ and its negative-part test from that proof. The new trace is nonnegative and the source agrees on D' , so the weighted integral difference is nonnegative. No assertion that the channel is infinitesimal is needed for this comparison.

However, $-u(z)p(z)$ is the coefficient for a separated small ball. A channel consists of many overlapping balls and meets the exterior. The single-ball expansion therefore does not justify summing their coefficients or assigning a certified finite benefit to the whole channel. The experiment uses the point score as a placement heuristic, then independently evaluates the actual channel geometry. A stronger channel-shape derivative or finite-insertion theorem remains separate work. Nor does an open route justify ideal isothermal walls: flow, convective heat transfer, external manifolds, and pumping power are absent from the model.

U.3 Reproducible experiment

The new comparison uses two sensor positions, five paired candidate seeds, and five methods. Each design adds six vertical channels of radius 0.075. Eighteen balls span $z = -0.45$ to $z = 0.825$ with spacing 0.075, so [Proposition U.1](#) gives the exact core radius $0.075\sqrt{3}/2$. Endpoint coordinates are checked against the actual boundary mesh’s bounds. Channel axes avoid the sensor ball and three radius-0.07 component balls, with a 0.02 additional gap; channels have a 0.025 mutual gap. Components are centered at the three prescribed source peaks. The Gaussian source tails remain a volumetric field on surviving material; this is not a fixed-power electronic-component model.

Each seed gives 512 common candidate axes. Scores are queried at $z = -0.03$ in the current domain and all channel balls are added after each selection. Sparse TWoS and WoS hottest-point placement use 12 scout walks and 192 fresh rescore walks on 24 candidates. Multigrid evaluates the same pointwise topological coefficient at $57 \times 37 \times 29$ or $113 \times 73 \times 57$ nodes; random feasible placement supplies a geometric control. Both solver families retain the outer mesh through each complete six-round design. Timings include candidate preparation, geometry setup, adaptive selection, I/O, and process startup and shutdown. Reference evaluations are separate.

Four independently scrambled 128-point sensor-probe blocks and 256 fresh walks per probe evaluate every final design. A separate $113 \times 73 \times 57$ grid provides another numerical reference. First-seed designs additionally receive a fourfold shell refinement and a $169 \times 109 \times 85$ grid check. Neither reference is exact. Four independent 32,768-point scrambled blocks estimate removed material volume and the removed fraction of the prescribed source. All methods share channel count, radius, axial extent, and constraints; the nonuniform housing thickness means their removed volumes are not exactly equal. We report those volumes rather than

Table 8 Accessible-channel comparison. Regional temperature reductions are means $\pm$ one standard error across five designs per sensor. Time and removed material are averaged across both sensors. All methods use six channels with the same radius and extent; volume varies with local housing thickness. Point scores are placement proxies, not derivatives of an entire channel.

Method	Left reduction (%)	Right reduction (%)	Design time (s)	Removed material (%)
Random	17.30 ± 2.23	10.74 ± 2.48	0.26	4.28
WoS hottest	43.78 ± 1.89	30.39 ± 4.17	1.00	4.24
TWoS proxy	55.74 ± 1.35	52.47 ± 0.56	1.02	4.22
MG proxy (57)	58.14 ± 1.11	56.95 ± 0.28	1.26	4.37
MG proxy (113)	58.14 ± 1.11	57.35 ± 0.41	6.55	4.36

describing this as an equal-volume comparison. Port positions are optimized, not fixed connections to an external manifold.

The saved-data audit verifies all 50 channel geometries and 240 nonrandom selections. All walks terminate before the step cap and all recorded multigrid solves meet the 10^{-8} residual tolerance. The largest first-seed shell-refinement change is 0.084 percentage points; the grid-refinement change is 1.721 points. Across all designs, the two primary numerical references differ by at most 3.133 points. These are sensitivity checks, not bounds on the remaining numerical error.

V Matched Geometry and Pipeline Study

V.1 A fixed-power, protected-region problem

The additional study uses the same Poisson and ideal zero-Dirichlet cooling model as the regional experiment, but replaces Gaussian source tails with compact, fully protected heat components. This removes an ambiguity in the fixed-count channel study: a lower objective can no longer result from deleting a portion of a heat-generating source. All dimensions and temperatures are nondimensional. No fluid or convective boundary model is implied.

For a center c , support radius $a > 0$, and amplitude $A \geq 0$, define

$$b_{c,a}(x) = \left(1 - \frac{|x - c|^2}{a^2}\right)_+^2, \quad f(x) = \sum_{j=1}^L A_j b_{c_j, a_j}(x).$$

The sensor weight is $q = b_{s,b}/Q_b$, where $Q_b = 32\pi b^3/105$. The objective is $J(D) = \int_D q u_D dx$; the adjoint satisfies $-\Delta p_D = q$ with homogeneous Dirichlet data. Component supports and the sensor support are required to remain strictly inside the material. The final protocol includes a low-power heat component of radius 0.085 and amplitude 4 at the protected sensor. This choice was fixed before the measured campaign: an excluded remote-source-only pilot left the branching sensor almost unheated and produced poorly resolved rare-event gradients. The pilot is retained separately. Main heat sources remain stronger and spatially distinct, so f is not proportional to q .


Proposition V.1 (Exact normalization and preserved input power). Suppose $\overline{B_{a_j}(c_j)} \subset D$ for every j , and a removed set C is disjoint from these closed balls. Then

$$\int_{D \setminus \overline{C}} f dx = \int_D f dx = \sum_{j=1}^L \frac{32\pi A_j a_j^3}{105}.$$

If C also avoids $\overline{B_b(s)}$, then $\int_{D \setminus \overline{C}} q = 1$. For a point X uniformly distributed in $B_b(s)$,


$$J(D) = \mathbb{E} \left[\frac{35}{8} \left(1 - \frac{|X - s|^2}{b^2}\right)^2 u_D(X) \right].$$

Finned housing




Parallel fins, narrow gaps,
and mounting bosses
138,348 triangles; watertight input surface

Bifurcating manifold



Branching transport, thin ribs,
and an elevated return loop
57,488 triangles; watertight input surface

Ribbed annular spreader



Curved transport, a central bridge,
and closely spaced ribs
82,360 triangles; watertight input surface

Figure 12 Distinct application geometries. Three watertight procedural test objects exercise different conduction paths and small geometric features. Every method uses the same surface, source configuration, and protected sensor within a case. Triangle count describes the input, not numerical difficulty. These are stylized conduction designs, not industrial cooling validations.

Proof. Step 1: restrict to the support. The positive-part definition makes $b_{c,a}$ zero outside $B_a(c)$. Since this ball is contained in D and untouched by C , its integral over the original and the perforated domain is the same. **Step 2: evaluate one radial integral.** Translation and spherical coordinates give

$$\begin{aligned} \int b_{c,a}(x) dx &= 4\pi \int_0^a r^2 (1 - r^2/a^2)^2 dr \\ &= 4\pi a^3 \int_0^1 (t^2 - 2t^4 + t^6) dt \\ &= 4\pi a^3 \left(\frac{1}{3} - \frac{2}{5} + \frac{1}{7} \right) = \frac{32\pi a^3}{105}. \end{aligned}$$

Step 3: sum the components. Linearity of the integral gives the displayed power identity. Disjointness between heat components is unnecessary; only protection from removal is used. **Step 4: normalize the sensor.** Apply Step 2 with $(c, a) = (s, b)$ and divide by Q_b . Its support is unchanged, so the integral of q remains one. **Step 5: express the objective as an expectation.** The uniform-ball density is $1/V_b$ with $V_b = 4\pi b^3/3$. Consequently $\int qu = V_b \mathbb{E}[q(X)u(X)]$ and $V_b/Q_b = 35/8$, proving the formula. This identity concerns exact integrals; finite quadrature and Monte Carlo estimates retain their numerical errors. $\square$

Proposition V.2 (The local insertion coefficient for the new objective). *If a candidate z lies outside all protected supports and satisfies the regularity, separation, and cavity-expansion assumptions of the regional result in [Appendix T](#), then*

$$\lim_{r \downarrow 0} \frac{J(D \setminus \overline{B_r(z)}) - J(D)}{4\pi r} = -u_D(z)p_D(z).$$

Proof. Step 1: identify the state and observation data. The compact bump and its first derivative vanish at its support boundary; in particular the finite sums defining f and q are bounded and continuous. The support-protection assumptions make both data independent of the inserted cavity for all sufficiently small r . **Step 2: specialize the regional objective.** Its observation density is exactly the fixed, nonnegative q above; the adjoint is the solution of $-\Delta p = q$. The regional cavity expansion therefore has coefficient $-u_D(z)p_D(z)$. **Step 3: use the three-dimensional capacity convention.** For the Laplacian convention

used throughout the paper, $\text{cap}(B_r) = 4\pi r$. Divide the expansion by this quantity and let r decrease to zero. The normalized remainder tends to zero under the cited assumptions. This substitution provides no new remainder estimate, and does not extend the separated-ball result to an overlapping channel. $\square$

V.2 Baselines and complete design costs

Every pipeline starts from the same unperforated mesh, source configuration, sensor, and paired candidate seed. We compare TWoS, deterministic multigrid TD, deterministic hottest-point placement, and random placement of six radius-0.065 spheres. The spheres remain disjoint and fully contained, so every initialization and every subsequent translation removes exactly $6(4\pi/3)(0.065)^3$ units of volume. Source and sensor clearance and a positive inter-cavity gap are checked at every saved refinement state. Sphere centers must remain at least $r + 0.012$ from the outer boundary and at least $r + a + 0.018$ from each protected support center of radius a ; pairs of sphere centers must be separated by more than $2r + 0.025$.

Each initialization is subsequently refined using our specialized implementation of Differential WoS [23], Algorithm 2 and Eq. (13). For zero boundary data and inward offset ℓ , a cavity-boundary hit contributes $V_n \hat{u}(x - \ell n)/\ell$ to the objective derivative; for a center translation the parameter velocity has normal component $V_n = n \cdot v$. The same radius, offset, gradient estimator, constraints, and backtracking rule are used for all initializations. At each trial step, cavities are updated in a cyclic order; each displacement is halved up to twelve times until feasible. This avoids freezing all parameters when one cavity reaches a constraint. An earlier global-rejection variant is retained as an excluded implementation diagnostic; refinement and final evaluation were rerun for every method. TWoS and multigrid TD receive eight refinement rounds; the hottest-point and random initializations additionally receive rounds nine through sixteen. Independent evaluations at rounds 0, 4, 8, and, where applicable, 16 expose quality versus complete cost rather than comparing only equal iteration counts. The two round allowances are not identical wall-time budgets: distant initial cavities can produce few nonzero boundary-gradient samples and consequently cheaper iterations. We therefore report measured time explicitly and do not infer dominance at unmeasured budgets or over all possible Differential WoS sampling schedules.

This is an attributed reimplement of a paper component inside a common benchmark, not the original authors' application optimizer. We checked the terminal contribution and recursive boundary offset against the original authors' implementation. Differential WoS is not in general restricted to topology-preserving changes. Here its fixed six-cavity parameterization isolates initialization quality; this restriction is a choice of comparison protocol, not a limitation claimed for that method.

TWoS uses 384 candidates, 48 scouting paths and 1,024 fresh rescore paths on a 16-site shortlist per insertion. Multigrid solves for the primal and distinct adjoint on a shared hierarchy with at most 65 nodes along the longest bounding-box axis; other axis counts scale with the aspect ratio. It is our embedded Cartesian finite-difference solver, not a finite-element implementation of another paper. A separate complete-design baseline uses at most 49 nodes along the longest axis with otherwise identical constraints and refinement. An additional nearest-sensor control ranks candidates by Euclidean distance to the sensor center and uses eight refinement rounds. It performs no PDE solve during initialization, but shares every candidate and feasibility constraint. This control tests whether simple proximity explains a benefit in the localized observation task; it was added during the campaign as a baseline-strengthening check, not selected through a parameter sweep. Independent resolution checks use longest-axis counts 97 and 129. Convergence to a 10^{-7} relative residual is checked separately from discretization error.

Complete design times include constructing the retained Embree geometry, solver calls, request serialization, cavity selection, and all refinement or radius-calibration decisions. Geometry generation, shared input preprocessing, and independent evaluation are excluded. Backend high-water resident memory is reported separately from driver memory. Instrumented outer-boundary distance queries and parity rays are distinct counters; neither is inferred from the number of walk steps. The measured campaign uses four OpenMP threads and single-threaded BLAS on a Linux environment exposing an Intel Xeon Platinum 8573C, with an eight-CPU quota and a 20 GiB memory limit. GCC 13.3 builds the C++17 backends with `-O3`. Runs execute serially without simultaneous rendering or other experiments. This is a managed shared environment, not a dedicated timing node. The primary four methods rotate order by seed; the two additional controls run subsequently, so their timings should also be read in this execution context.

Guarded cavity offsets. A single ray-parity query can be ambiguous at a triangle edge. The following geometric implication supplies a fallback specifically for contained spherical cavities; it does not relax source protection or inter-cavity separation.

Lemma V.3 (An enclosing-ball test for a normal offset). *Let D be open, $c \in D$, and $d = \text{dist}(c, \partial D) > r + \ell$. For any unit vector e , the point $c + (r + \ell)e$ belongs to D .*

Proof. Step 1: an open distance ball lies in the domain. If some y with $|y - c| < d$ were outside D , the line segment from the interior point c to y would first meet ∂D at a point whose distance from c is at most $|y - c| < d$, a contradiction. Thus $B_d(c) \subset D$. **Step 2: apply the radius inequality.** Since $|c + (r + \ell)e - c| = r + \ell < d$, this offset point belongs to $B_d(c)$, and hence to D . $\square$

The implementation retains its usual membership and analytic sphere-segment checks. If the former rejects an offset, a positive enclosing-ball margin and a checked interior cavity center provide an additional membership test. A replay of the encountered case and an unchanged control are retained with the benchmark. This is a numerical geometry guard under the stated mesh assumptions, not a proof that floating-point distance and parity routines are exact, and not the separate conservative application certificate.

V.3 Finite-cavity estimator comparison

Query timings use a retained outer-boundary acceleration structure, with construction and an untimed warmup excluded for every estimator. Eight protected, interior query sites are fixed per geometry before measuring any estimator. A high-budget direct estimate uses sixteen independent batches of 8,192 paths per site. Direct TWoS and common-random-number finite-cavity differencing use sixteen repeats at 256 and 2,048 paths per site. For the latter, the two state walks start at the same independently sampled sensor point and use the same random stream in the original and perforated domain. Their weighted difference is divided by $4\pi r$. Trial radii are 0.08, 0.04, 0.02, 0.01. This is a specified local coupling, not a claim that all possible finite-difference couplings have been optimized.

We additionally evaluate a stronger stopped-path baseline. It estimates the finite cavity correction directly, rather than subtracting two noisy source integrals. The following identity specifies precisely which quantity it estimates.

Proposition V.4 (A finite-cavity harmonic-correction estimator). *Let D be bounded and let $D_r = D \setminus \overline{B_r(z)}$, with the cavity strictly contained and disjoint from the sensor support. Let u_D and u_{D_r} solve the same Poisson equation $-\Delta u = f$ on their respective domains with zero Dirichlet data. Assume the domains are regular for the Dirichlet problem and the state solutions are continuous up to their boundaries. Let Y be uniform in $B_b(s)$, let X be a Brownian path started at Y , and let τ_r be its first exit from D_r . Write $W(Y) = (35/8)(1 - |Y - s|^2/b^2)^2$ and $H_r = \{X_{\tau_r} \in \partial B_r(z)\}$. Suppose $U_0(x)$ is a conditionally unbiased, independent estimate of the original-domain temperature $u_D(x)$. Then*

$$Z_r = \frac{W(Y) \mathbf{1}_{H_r} U_0(X_{\tau_r})}{4\pi r} \quad \text{satisfies} \quad \mathbb{E} Z_r = \frac{J(D) - J(D_r)}{4\pi r}.$$

The statement assumes the displayed expectation is integrable. Exact paths and conditionally unbiased state estimates are mathematical assumptions here; finite-termination implementations retain their numerical errors.

Proof. Step 1: subtract the two PDEs. On D_r , both states have the same source, so $h = u_D - u_{D_r}$ satisfies $-\Delta h = 0$. On the old boundary, both states vanish and hence $h = 0$. On the cavity boundary, $u_{D_r} = 0$ and hence $h = u_D$. **Step 2: derive the harmonic exit representation.** Fix $y \in D_r$. The difference h is harmonic in D_r and continuous on its compact closure, hence bounded. Localize the Brownian motion to compact interior subdomains and a finite time t . Itô's formula for the localized process has drift $(1/2)\Delta h = 0$; its stochastic integral has zero mean because the localized derivatives are bounded. Thus the expectation of the stopped h remains $h(y)$. Exhausting the domain gives the same equality for $h(X_{t \wedge \tau_r})$ by

dominated convergence. The exit time τ_r is finite almost surely: D_r is contained in a ball, whose Brownian exit time has finite expectation, and exiting D_r occurs no later than exiting that ball. Letting $t \rightarrow \infty$ and using boundedness and boundary continuity therefore yields

$$h(y) = \mathbb{E}_y[h(X_{\tau_r})] = \mathbb{E}_y[\mathbf{1}_{H_r} u_D(X_{\tau_r})].$$

The second equality uses exactly the two boundary values from Step 1; it introduces no source-integral subtraction. **Step 3: replace the boundary temperature by a fresh estimate.** Condition on Y , the entire exit path, and its terminal point. Conditional unbiasedness of the independent nested estimate gives $\mathbb{E}[U_0(X_{\tau_r}) \mid Y, X] = u_D(X_{\tau_r})$. Applying the tower property therefore leaves the expression in Step 2 unchanged. **Step 4: integrate the protected observation.** The cavity does not remove any sensor weight. Proposition V.1 and the tower property give $\mathbb{E}[W(Y)h(Y)] = \int q(u_D - u_{D_r}) = J(D) - J(D_r)$. **Step 5: normalize.** Divide this equality by $4\pi r > 0$. The result is a finite-radius difference; identifying its limit with $u_D(z)p_D(z)$ still requires the separate cavity asymptotics. No finite-radius equality with that point product has been asserted. $\square$

The implementation uses a perforated-domain boundary walk followed by one independent original-domain source walk only on a new-cavity hit. It reports the number of such hits, distance queries, time, and estimator variance. This baseline tests the rare-event mechanism after removing unnecessary subtraction noise. The common-random-number subtraction baseline is retained as a separate implementation, rather than conflating the two couplings.

Variance about each estimator’s own mean is kept separate from RMS error to the high-budget infinitesimal reference. A finite-radius difference has a different expectation from a topological derivative; its error includes this remainder as well as sampling and termination error. The reference also has sampling and termination error and is not labeled ground truth. Aggregate errors weight high-influence sites more strongly; a reference with no sampled contribution at a remote site does not establish a zero exact coefficient. The data release reports reference standard error and the number of zero reference estimates. A degenerate bootstrap interval from zero observed hits likewise does not establish zero population variance. Multigrid query estimates use longest-axis counts 49, 65, 97, and 129. Additional finite-hole resolution checks explicitly expose grid differences that fail to resolve the smallest cavities; a small linear-solver residual cannot resolve that issue.

V.4 Matched finite channels

Each channel is the union of balls at spacing 0.04 along a vertical axis, with centers from $z = -0.48$ to 0.80 . These endpoints lie beyond every geometry’s vertical bounds. For radius $r > 0.02$, Proposition U.1 gives a positive core radius $\sqrt{r^2 - 0.02^2}$. Parallel axes protect all compact sources and the sensor by Proposition U.2. The shared candidate pool is conservatively generated with maximum radius 0.11, including the source-clearance and axis-separation tests, even though calibrated radii are usually smaller. This restriction is common to every method; the experiment does not optimize over all feasible channel locations at the final radius.

After the k th axis is selected, a shared common radius is calibrated to removed volume $0.0028k$ using a fixed scrambled Sobol sample of 2^{20} bounding-box points. Recalibration can change previous channels. Four independent 2^{20} -point scrambles audit the final volume; their dispersion is a numerical integration diagnostic, not a certified enclosure. Six channels target total volume 0.0168. Neither equal radii nor equal channel counts are used as a substitute for the volume measurement. The axis-point topological score remains a heuristic for a finite channel, followed by independent objective evaluation.

V.5 Evaluation, geometry, and visualization

Five paired seeds are used on a finned housing, a bifurcating manifold with thin ribs and an elevated return loop, and a ribbed annular spreader with a central bridge. They are procedural application test objects, not standardized industrial datasets. The latter two test branching and curved transport in addition to the housing’s narrow gaps and extended fins; surface triangle count alone is not treated as evidence of numerical difficulty.

Each saved objective is evaluated in four independent scrambled-Sobol sensor blocks of 128 points with 512 walks per point. The exact weight in [Proposition V.1](#) is used. Paired-seed summaries retain the variation between designs and separately report evaluation-block uncertainty. Additional paired intervals jointly bootstrap the five design seeds and the four evaluation blocks; block resampling is shared across methods and seeds to preserve the common-random-number dependence. These are approximate, small-sample intervals, not certified coverage guarantees. Reusing the same evaluation blocks across methods supports paired differences; those blocks never participate in candidate selection or line-search acceptance. The housing-channel comparison exposed several percentage points of disagreement between independent numerical references. We therefore added a uniform refinement rule: a first-seed case is reevaluated at longest-axis counts 193 and 257 if any method changes by more than 1 percentage point from grid 97 to 129 or differs by more than 2 points between Monte Carlo and grid 129. Every method in a flagged case is checked. This rule was introduced during reference auditing, rather than claimed as a preregistered test; the added solves are excluded from design times. Line search itself is heuristic and is not reported as certified.

Uncertainty bands and bars. Let $J_{m,k,b}$ be the independently evaluated objective for method m , design seed $k \in \{1, \dots, 5\}$, and evaluation block $b \in \{1, \dots, 4\}$, and let $J_{0,k,b}$ denote its unperforated reference. The reported mean reduction is

$$\bar{R}_m = \frac{1}{5} \sum_{k=1}^5 100 \left(1 - \frac{\sum_b J_{m,k,b}}{\sum_b J_{0,k,b}} \right).$$

For each of 4,000 bootstrap replicates, we resample five seed indices and four block indices with replacement, then recompute this expression. The same indices are used across methods and checkpoints, retaining the paired design and common-random-number structure. Shaded trajectory bands are the pointwise 2.5th and 97.5th percentiles. Isolated channel results use vertical bars instead. Horizontal bars resample measured complete-design times over the same seed indices; they include both algorithmic and machine-timing variation. Faint points show the individual design runs. Bands between checkpoints are visual interpolation, not evaluations at every intervening time or simultaneous confidence regions for an entire optimization path.

For an estimator setting, let $S_i \in \mathbb{R}^8$ be its vector of scores in repeat i , and let $T_j \in \mathbb{R}^8$ be one of sixteen high-budget reference batches. With $\bar{T} = 16^{-1} \sum_j T_j$, the displayed relative variance and RMS error are respectively

$$V = \frac{\sum_{a=1}^8 \widehat{\text{Var}}_i(S_{i,a})}{\|\bar{T}\|_2^2}, \quad E = \sqrt{\frac{\frac{1}{n} \sum_{i=1}^n \|S_i - \bar{T}\|_2^2}{\|\bar{T}\|_2^2}}.$$

We independently resample the sixteen repeat vectors and sixteen reference batches 4,000 times. Sites remain together within each vector; treating sites as independent experimental replicates would be inappropriate. Percentile bands around V measure uncertainty in the estimated variance, not the standard deviation of an individual score. RMS-error bars additionally reflect sampling uncertainty in the reference. A deterministic multigrid setting has one output: its error bar reflects only reference sampling, and no Monte Carlo field variance is inferred. Timing intervals use five additional repetitions per grid with rotated resolution order and a retained outer mesh; each call rebuilds its multigrid hierarchy. All repeated fields are checked against the original deterministic output. For the eight-site query benchmark, a scheduling audit found that the original dynamic chunk of four exposed only two work chunks to four threads. We repeat the direct-score timing with chunk size one, verify every score and component field bitwise against the retained original output, and use these repeats for query-time plots. Original timings remain in the raw records. This correction changes neither the numerical samples nor the recorded complete-design timings, which retain their original scheduling. Numerical discretization, termination, offset, and finite-radius bias are not enclosed by these bootstrap intervals. With only five design seeds and sixteen estimator repeats, the intervals remain approximate descriptive evidence, not established frequentist coverage or a finite-insertion certificate. In particular, resampling cannot reveal rare events absent from the recorded repeats; a zero observed score does not establish zero temperature or influence. Site-level raw scores and cavity-hit counts are retained for this reason. When all recorded correction vectors are zero, we report the empirical output but suppress collapsed variance and RMS-error intervals and mark variance unresolved. The earlier Bunny allocation bands in [Figs. 1](#) and [16](#) remain observed ranges over three designs; they are not relabeled as confidence intervals.

Figures use the first predetermined seed, rather than selecting the visually best design. Camera, section plane, display window, lighting, and map scale are held fixed within each geometry. The local cooling map

displays $1 - u_D(x)/u_{D_0}(x)$ on the visible material section using a longest-axis-193 multigrid field, clipped to $[0, 1]$ for display. Material is gray where the unperforated section temperature is below 10^{-4} of its maximum, avoiding misleading relative changes in essentially cold material. Quantitative objective values come from the separate evaluations above. A faint depth-tested shell shows the physically present upper structure removed only for presentation. The display cut, tessellated cavity surfaces, and lighting denoiser never enter the PDE or its objective. Section colors are applied from the interpolated numerical field after radiance denoising.




Figure 13 Accessible designs with controlled material and power budgets. Six open ball-union channels are selected under identical source and sensor protection constraints. A common radius is calibrated to total removed volume 0.0168; the displayed audit range spans the independent mean volume estimates over all methods and seeds in each geometry, not a certified interval. Heat input is exactly invariant under the protected-support model. Views use the first predeclared seed; quantitative summaries use five seeds. Vertical bars give pointwise 95% bootstrap intervals for mean cooling, resampling seeds and independent evaluation blocks together across methods. Horizontal bars describe uncertainty in mean runtime across seeds; plot ranges adapt to each geometry while retaining the same physical units. The local cooling field shares a scale across designs. No coolant flow is solved, and the point score remains a heuristic for these extended removals.

Table 9 Complete cavity-design pipelines. Means over five paired design seeds; SD is the sample standard deviation of their evaluated cooling reductions in percentage points, including evaluation noise. Cooling intervals are pointwise 95% bootstrap intervals jointly resampling seeds and independent evaluation blocks. Backend RSS is process high-water memory, not total driver memory; a dash denotes no PDE backend call. Times include all design stages but exclude independent evaluation. Hot and random use sixteen refinement rounds; TWoS, multigrid, and nearest-sensor use eight.

Geometry	Method	Cooling (%)	SD (pp)	95% interval	Time (s)	RSS (MiB)
Finned housing	TWoS	64.99	1.41	[63.72, 66.20]	11.77	24.2
	Multigrid TD (49)	64.51	0.67	[63.67, 65.26]	7.99	24.2
	Multigrid TD (65)	65.10	0.59	[64.25, 65.94]	8.66	24.4
	Nearest sensor	63.56	2.60	[61.31, 65.60]	7.08	24.2
	Hottest point	42.74	16.50	[27.88, 51.06]	13.14	24.3
	Random	21.88	15.37	[10.87, 34.92]	10.69	24.1
Branching manifold	TWoS	16.40	0.94	[15.50, 17.38]	51.24	14.1
	Multigrid TD (49)	16.44	1.99	[14.68, 17.91]	43.21	14.0
	Multigrid TD (65)	16.27	1.92	[14.57, 17.69]	43.76	14.1
	Nearest sensor	17.33	0.80	[16.43, 18.27]	41.49	14.0
	Hottest point	0.14	0.04	[0.01, 0.35]	13.43	14.0
	Random	5.44	2.01	[3.90, 7.37]	46.65	13.9
Annular spreader	TWoS	11.36	0.88	[10.42, 12.26]	35.78	17.4
	Multigrid TD (49)	11.83	1.53	[10.50, 13.17]	30.30	17.5
	Multigrid TD (65)	11.68	1.36	[10.50, 12.88]	29.67	17.6
	Nearest sensor	12.19	2.12	[10.26, 13.63]	29.96	17.3
	Hottest point	0.01	0.00	[0.00, 0.02]	10.73	17.4
	Random	1.59	3.54	[-0.00, 4.84]	18.63	17.5

Table 10 Matched-volume accessible channels. Means over five paired design seeds; SD is the sample standard deviation of their evaluated cooling reductions in percentage points, including evaluation noise. Cooling intervals are pointwise 95% bootstrap intervals jointly resampling seeds and independent evaluation blocks. Backend RSS is process high-water memory, not total driver memory; a dash denotes no PDE backend call. Times include all design stages but exclude independent evaluation. Independent removed-volume audits are reported separately; every compact heat support is preserved.

Geometry	Method	Cooling (%)	SD (pp)	95% interval	Time (s)	RSS (MiB)
Finned housing	TWoS	37.79	5.48	[32.92, 42.35]	4.48	24.1
	Multigrid TD (49)	42.11	3.00	[39.13, 44.68]	2.37	24.2
	Multigrid TD (65)	42.11	3.00	[39.13, 44.68]	2.84	24.1
	Nearest sensor	39.37	4.77	[34.93, 43.05]	1.69	—
	Hottest point	29.48	5.64	[24.56, 33.80]	2.73	24.3
	Random	7.31	1.74	[5.57, 9.29]	2.10	—
Branching manifold	TWoS	2.29	0.14	[1.74, 2.89]	6.55	14.1
	Multigrid TD (49)	2.57	0.23	[2.00, 3.20]	1.16	13.9
	Multigrid TD (65)	2.59	0.16	[2.05, 3.19]	1.63	14.0
	Nearest sensor	2.62	0.18	[2.05, 3.23]	0.83	—
	Hottest point	1.25	1.14	[0.40, 2.18]	1.60	14.0
	Random	0.93	0.29	[0.57, 1.21]	0.93	—
Annular spreader	TWoS	1.75	0.27	[1.48, 1.99]	8.14	17.6
	Multigrid TD (49)	1.89	0.13	[1.73, 2.01]	2.91	17.4
	Multigrid TD (65)	1.93	0.16	[1.74, 2.12]	2.72	17.3
	Nearest sensor	2.02	0.11	[1.88, 2.15]	1.25	—
	Hottest point	0.00	0.00	[-0.00, 0.01]	2.87	17.3
	Random	0.20	0.25	[0.02, 0.44]	1.40	—

Table 11 Independent resolution check. First predeclared seed, before sphere refinement. Ranges span all six methods on each geometry. The grid columns give longest-axis node counts. Differences are percentage points of objective reduction. These comparisons test reference sensitivity; neither numerical method is exact ground truth.

Study	Geometry	Grid 129 minus 97 (pp)	MC minus grid 129 (pp)
Pipeline	Finned housing	[-1.30, -0.07]	[1.56, 3.18]
Pipeline	Branching manifold	[-0.01, 0.49]	[-0.15, 1.82]
Pipeline	Annular spreader	[-0.00, 0.37]	[0.00, 1.26]
Channels	Finned housing	[-2.60, -0.05]	[2.03, 4.79]
Channels	Branching manifold	[-0.01, 0.21]	[-0.01, 0.34]
Channels	Annular spreader	[-0.11, 0.03]	[-0.08, 0.04]

Table 12 Finer reference checks for flagged cases. A case is refined when any method changes by more than 1 pp from grid 97 to 129, or differs by more than 2 pp between MC and grid 129. Every method in that case is checked. Ranges span all six methods; sphere states are before refinement.

Study	Geometry	Grid 257 minus 193 (pp)	MC minus grid 257 (pp)
Pipeline	Finned housing	[-0.07, 0.44]	[0.47, 2.38]
Channels	Finned housing	[-0.17, 0.03]	[0.71, 2.12]

W Additional Figures for Earlier Thermal Studies

These figures retain the earlier studies' original protocols and data. Their budgets differ from the matched-volume, fixed-power, complete-pipeline comparison in [Appendix B.1](#); their reported speed ratios and uncertainty summaries should not be transferred between experiments.




Figure 14 Regional thermal protection with a distinct adjoint. (a–c) Initial temperature, sensor influence, and insertion scores on the same horizontal cut; each field has its own color scale. Cyan outlines the protected sensor; the light disk in (c) excludes cavity centers by the sensor radius, cavity radius, and safety gap. (d) Projected cavity spheres selected by sparse TWoS and deterministic hottest-point placement for the first right-sensor seed. (e) Regional temperature reductions, with one standard error across five complete designs. S and L denote the two TWoS sampling budgets. (f) Mean design quality and complete wall time across both sensor cases using the retained-mesh backend, showing both sampling budgets and the coarse/primary multigrid settings. The complete resolution sweep and both backend timings are in Table 7. Numerical references remain approximate; no finite-benefit certification is enabled in this experiment.




Figure 15 Accessible cooling with objective-aware placement. Top: actual saved first-seed right-sensor designs for hottest-point, TWoS, and coarse multigrid placement, with identical cameras, display cuts, lighting, and an absolute temperature scale. Magnified views show the protected sensor (cyan); cavity walls retain geometric shading. Faint, depth-tested shells indicate bars removed only for the display cut. The displayed cut does not enter the physical PDE. Section temperatures use a separate $169 \times 109 \times 85$ multigrid reference and bilinear display interpolation. Per-design labels use independent WoS evaluations. Bottom left: the original temperature and relative admissible up priority; circles mark the final TWoS axes, selected sequentially as fields change. Bottom right: mean cooling and complete design cost over both sensor positions and five seeds; error bars are one standard error across the five paired-seed averages. The full 50-design study, including random and finer-grid baselines, is in Table 8. All methods share channel count, radius, and extent; removed volume varies. The point score remains a whole-channel placement heuristic, with no channel derivative or finite-benefit certificate asserted.




Figure 16 *Cavity allocation and a cost-dependent count.* (a–b) Cooling at two fixed removed volumes; radii are listed in count order. Lines are means over three complete design seeds, and shaded ranges are their minima and maxima, not confidence intervals. Each design is evaluated using four independent randomized spatial blocks with fresh WoS paths; blocks are shared across designs. (c) The preferred tested full-budget TWoS count after adding an assumed per-cavity cost to normalized compliance; the unperforated option is included. This is a finite-design sensitivity plot, not an estimated physical cost or a global optimum. All cavities have ideal cooled walls; at fixed volume, their total surface area increases with count.




Figure 17 Measured topology design on a stylized finned conduction housing. (a) The closed 96,740-triangle input. (b) Initial temperature on the central cut, with mint crosses at the two chip-source centers. (c) Direct grid-free topological scores inside a shaded cutaway. (d) Six cavities selected by scout-rank-fresh-validate; solid spheres lie in the retained half and dashed outlines lie in the removed half. (e) Compliance for one canonical run under independent reevaluation after every insertion. (f) Mean final reduction across five complete stochastic designs, with standard error and mean total selection time; multigrid uses one deterministic design and ten fresh objective streams. In this earlier figure, “validation” denotes fresh shortlist ranking samples; it is not the finite-benefit certificate.

X Detailed Experimental Protocols and Results

X.1 Controlled statistical diagnostics

Annulus rare-event benchmark. Let the outer disk have radius one, start Brownian motion at $r_0 = 0.35$, and insert a concentric cavity of radius $\varepsilon < r_0$. Both the relative capacity $\rho(\varepsilon) = 2\pi/\log(1/\varepsilon)$ and the cavity-hit probability $q_\varepsilon = \log(1/r_0)/\log(1/\varepsilon)$ are exact. This diagnostic uses a pointwise harmonic response, outside the L^2 -smooth objective assumptions of [Theorem 3.5](#). Taking the coupled path difference to be minus the hit indicator gives the exact normalized target $D_T = -\log(1/r_0)/(2\pi) = -0.16708$. We compare direct cross-fitted harmonic-measure queries, the common-random-number hit estimator, and two independent objective samples. We use $N = 256$, 3,000 independent replications, twelve radii from 10^{-1} through 10^{-24} , percentile-bootstrap variance intervals, and fit slopes on the smaller-capacity half of the sweep.

Correlation benchmark. We draw 30,000 batches of $N = 32$ Gaussian oracle pairs with means $(0.8, 0.65)$, standard deviations $(0.5, 0.6)$, and within-pair correlation swept from -0.9 to 0.9 . The exact target is $D_T = -0.52$. Same-seed multiplication, the product of sample means, cross-fitting, and fully independent primal/adjoint batches are evaluated on identical parameter settings. Dotted curves in [Fig. 7\(b\)](#) are the exact covariance-bias formulas.

Adaptive-selection benchmark. We search K independent null sites whose estimates are Gaussian with known standard error $1/\sqrt{64}$. A one-sided 5% upper bound drives insertion. Across 40,000 replications, we compare reusing the selected scouting estimate, a Bonferroni simultaneous bound on the scout, and a fresh pointwise validation estimate. Wilson intervals quantify Monte Carlo uncertainty.

The fitted slopes in [Fig. 7\(a\)](#) are 0.00, -1.10 , and -1.89 for direct, coupled, and independent estimation, respectively. At the smallest tested capacity, direct TWoS has variance 7.06×10^{-4} , $8.26\times$ below coupled finite-cavity differencing and $80.4\times$ below independent differencing. This crossover is the intended claim: finite differencing can be competitive for large holes, but becomes progressively worse in the topological limit.

At correlation 0.9, the same-seed product has bias -0.2708 , the product of means has bias -0.00919 , cross-fitting has bias -7.50×10^{-4} , and independent batches have bias 7.50×10^{-4} . The latter two deviations are within their Monte Carlo standard errors. Cross-fitting does not dominate independent batches in variance; it preserves unbiasedness while permitting shared trajectories or caches. After selecting among $K = 256$ null sites, naive reuse falsely inserts in 100.0% of trials. Bonferroni and fresh validation yield 4.985% and 5.003%, with 95% Wilson intervals $[4.776, 5.203]\%$ and $[4.793, 5.221]\%$, respectively.

X.2 Spatial manufactured WoS field

We next test the complete spatial estimator rather than only abstract oracle pairs. In the unit disk, choose the harmonic polynomials

$$u(x, y) = 0.28 + 0.58x - 0.30y + 0.22(x^2 - y^2), \quad (229)$$

$$p(x, y) = 0.46 - 0.38x + 0.34y + 0.40xy. \quad (230)$$

Their boundary traces define two independent harmonic Dirichlet problems. The analytic target is the product $-up$. Because the second field has nonzero boundary data, it is not the homogeneous-boundary adjoint in [Eq. \(5\)](#); this is a product-estimator diagnostic, not a manufactured design objective. The Poisson-compliance disk experiment in [Appendix Y](#) tests an actual topological derivative. A vectorized WoS implementation queries 6,359 sites in the compact disk $\|z\| \leq 0.88$, using 96 paired walks per site and a 10^{-7} termination shell. Walks take 22.5 spheres on average.

X.3 Three-dimensional internal-cooling reference

We next exercise scout-rescore-insert placement in three dimensions. The solid is the box $[-1.2, 1.2] \times [-0.8, 0.8] \times [-0.45, 0.45]$ with zero-temperature Dirichlet data on its exterior and on every inserted cooling cavity. We solve $-\Delta u = f$ for the positive affine source $f(x, y, z) = 1 + 0.45x/1.2 - 0.25y/0.8$ and minimize the thermal compliance $J = \int_\Omega f u \, dx$. Since the adjoint equals the state, the normalized spherical-cavity




Figure 18 *Manufactured spatial validation with actual WoS paths. The two harmonic fields (a–b) and exact product $-up$ (c) supply analytic reference values. Multiplying same-seed primal and adjoint samples creates a coherent covariance-error field (d), with RMSE 0.108 and mean bias 0.0987. The product of sample means (e) suppresses the error to RMSE 0.0195 but retains its finite- N covariance bias. Cross-fitted TWoS (f) has RMSE 0.0196 and mean bias -1.20×10^{-4} , matching the sampling error while removing the bias exactly in expectation.*

derivative is $D_T = -u^2$. Affinity of f makes the source contribution on every three-dimensional WoS ball exact: $f(x)R^2/6$. Both the box and the union of cavities are queried analytically; the TWoS method therefore constructs no volume grid.

We choose seven cavities of radius 0.115 among 2,048 scrambled-Sobol candidate centers, subject to identical exterior and pairwise clearances for all methods. Our method uses 16 walks per scouting site, freshly rescores a 24-site shortlist with 192 walks per site, and maximizes the resulting 95% lower confidence bound. We compare against reusing the noisy scout, a $49 \times 33 \times 19$ voxel adjoint/TD solve, and random feasible insertion. Every design is evaluated on a fixed independent Sobol probe set with fresh WoS paths; the optimization samples are never reused for the reported objective.

The initial TWoS and voxel TD fields have Pearson correlation 0.977 and relative RMSE 0.104. After seven insertions, validated TWoS reduces compliance by 43.99%, compared with 43.15% for scout reuse, 46.17% for the voxel reference, and 25.47% for random insertion. Thus the direct grid-free estimator realizes 95.3% of the voxel baseline’s measured reduction and $1.73\times$ the reduction of random placement. The small gap between fresh validation and scout reuse is not statistically resolved in this single design run; the post-selection coverage experiment in Fig. 7(c), rather than this trajectory, supports the statistical selection mechanism in an idealized known-noise setting. It does not establish finite-cavity coverage for this 3D example. We likewise withhold a timing claim: the current TWoS reference is vectorized Python, whereas the sparse voxel solve runs in compiled code. The separate compiled Bunny and housing studies measure performance with matched C++ implementations.




Figure 19 *Correctness-first 3D internal-cooling benchmark. (a) The initial temperature field on the central plane. (b) The direct cross-fitted topological score, with selected cavity centers projected onto the plane. (c) The resulting grid-free 3D cavity configuration. (d) Independently evaluated compliance throughout optimization. (e) Final improvement under the same seven-cavity budget. (f) The final central-plane temperature. This analytic-geometry pilot validates the complete application loop; it is not the final equal-time mesh benchmark.*

X.4 Stanford Bunny curved-mesh stress test

The analytic box isolates the optimization mechanism but does not exercise a production-style curved boundary. We therefore repeat the complete loop inside the canonical Stanford Bunny [32]. This experiment is a geometry stress test, not an engineering claim about cooling a rabbit. The original zippered reconstruction stores 35,947 vertices and 69,451 triangles and contains five holes in its bottom. Our deterministic preprocessing removes 1,113 unreferenced vertices, consistently orients the faces, seals boundary loops of lengths 22, 42, 39, 40, and 80, and verifies a closed edge- and vertex-manifold surface with Euler characteristic two. The result has 34,839 vertices and 69,674 triangles.

After normalizing the Bunny height to one, we solve the same Poisson-compliance problem with affine source $f(x, y, z) = 1 + 0.65(y + 0.5)$. We insert six cavities of radius 0.05 among 1,024 interior candidate centers. Exact closest-triangle distances come from a libigl AABB tree; candidate containment uses fast winding numbers. TWoS uses 12 scouting walks and freshly rescores a 20-site shortlist with 128 walks per site. The voxel reference uses a $54 \times 64 \times 48$ Cartesian grid clipped by the Bunny.

The noisy 48-walk initial TWoS field has Pearson correlation 0.911 with the voxel TD field. Validated TWoS reduces compliance by 47.77%, versus 45.99% for scout reuse, 48.06% for the voxel reference, and 39.46% for random feasible placement. It therefore realizes 99.4% of the voxel reference’s reduction without a volume grid; the final-objective point-estimate gap is smaller than the reported probe-sampling standard errors. This result supports geometric applicability; we next isolate field-evaluation performance with matched compiled implementations.




Figure 20 Curved-mesh geometry stress test on the Stanford Bunny. (a) The repaired closed input mesh. (b) Interior cross-fitted topological-derivative samples viewed through the translucent surface. (c) Six selected cooling cavities. (d) Compliance measured with fresh paths after every insertion. (e) Improvement under the common cavity budget. (f) Final temperature on the central slice. Bunny data courtesy of the Stanford University Computer Graphics Laboratory.

X.5 Compiled output-sensitive scaling

We implemented both methods in the same C++17 executable. The grid-free kernel uses Embree 4 closest-point queries and OpenMP over independent paths; the voxel reference uses the same Embree scene for parity classification, a seven-point finite-difference stencil, and Eigen conjugate gradients. Both receive eight CPU threads. Timings include BVH construction and all work needed to return the field, exclude process startup and file I/O, and report the median of three runs. Peak resident set size is measured for each isolated process.

We evaluate 384 common interior Sobol sites against an independent 2,048-walk-per-site TWoS reference. We sweep 4–512 walks per TWoS site and eight voxel grids from $24 \times 28 \times 20$ to $110 \times 128 \times 96$. A separate output-sensitivity sweep fixes 512 walks per site and a $96 \times 112 \times 84$ grid, which have relative field RMSE 0.0757 and 0.0732, respectively, on all 384 sites, and varies only the requested output count.

At 64 outputs, TWoS takes 0.381 seconds at relative RMSE 0.0622; the voxel reference takes 0.528 seconds at RMSE 0.0711. Thus TWoS is $1.39\times$ faster and its 15.4 MiB peak is $5.25\times$ smaller than the voxel process’s 80.7 MiB peak. The speedup is $2.49\times$ at eight outputs and the methods are effectively tied at 96 (0.566 versus 0.561 seconds). At 384 outputs the ordering reverses: TWoS takes 2.10 seconds and the voxel solve 0.504 seconds. This crossover is the claimed performance result, not a universal speedup. In the full-resolution surface sweep, Embree BVH construction takes 9.74 ms and the scalar closest- boundary kernel sustains 1.77 million queries per second. Embree distances agree with libigl to maximum absolute error 3.04×10^{-7} ; the independent C++ and Python TWoS fields correlate at 0.991, and the C++ and SciPy voxel fields agree to relative RMSE 1.11×10^{-13} at the validation grid.




Figure 21 Matched C++ scaling on the 69,674-triangle Bunny. (a) Actual Embree-accelerated sphere-walk trajectories in a shaded cutaway shell. (b–c) Direct TWoS and voxel topological scores at identical exposed interior sites and color scale. (d) Relative RMSE against an independent 2,048-walk field versus end-to-end time. (e) The same accuracy sweep versus measured process peak memory. (f) At nearly matched field error, the grid-free method is output-sensitive: the methods tie near 96 requested sites, and the global voxel solve wins for dense output. Each point is measured rather than analytically extrapolated.

Independent high-sample check for the teaser. For the 64-site comparison in Fig. 1, we recomputed the reference using an independent seed and 16,384 walks per site. Against this reference, the retained TWoS and voxel score fields have relative RMSE 0.05888 and 0.06886, respectively. The root-sum-square reference score standard error divided by the reference score norm is 0.01160, in addition to error from the unchanged termination shell $\varepsilon = 5 \times 10^{-4}$. This is a numerical accuracy check, not exact ground truth; its timing is not included in the performance comparison. The teaser separately visualizes fresh 256-walk temperature and insertion-benefit estimates on a 160×200 section at $z = 0.09$ for the saved cavity designs. The physical domains do not contain the display opening. A one-grid-cell Gaussian filter is applied only to the displayed scalar colors; raw fields, accuracy errors, cooling objectives, and timing data are retained without it. The placement study and sparse point-query timings are distinct experiments: the latter does not establish a speedup for the displayed full design loop.

X.6 Finned-housing design benchmark

We next test the complete compiled design loop on a license-free stylized engineering geometry rather than an animal scan. An analytic union of a rounded enclosure, nine heat-sink fins, and four mounting bosses is extracted by marching cubes. The resulting surface has 48,372 vertices and 96,740 triangles, is connected, closed, edge- and vertex-manifold, and has Euler characteristic two. This procedural construction is deliberately auditable; it resembles a conduction housing but is not proprietary CAD.

Inside the housing we solve $-\Delta u = f$ with zero Dirichlet data on the outer surface and every inserted cavity. The positive source is a background plus three anisotropic Gaussian hot spots representing two




Figure 22 Tuning and resource tradeoff on the housing benchmark. (a) Five multigrid resolutions, each timed three times. The star is the smallest grid meeting the predeclared relative-RMSE target 0.09 against the independent 512-walk field. (b) Final design quality versus complete six-round selection time; marker area encodes peak memory and vertical bars show standard errors. Sparse TWoS occupies the fast regime, while the near-matched-time budget approaches the tuned multigrid solution without allocating a volume grid.

chips and a regulator; the objective is the thermal compliance $J = \int_{\Omega} f u$. On each WoS ball, we sample its Green source integral with the unbiased beta-radial estimator in Proposition G.1. Each method inserts six radius-0.095 spherical cavities from the same 768 Sobol candidates. Sparse TWoS uses 12 scouting walks per feasible candidate and 192 fresh walks on a 24-site shortlist; a near-matched-time variant uses 96 and 768 walks, respectively. Our strongest deterministic baseline is a matrix-free geometric multigrid solver for the same embedded seven-point operator, with four levels, three weighted-Jacobi pre- and post-sweeps, residual-minimizing coarse correction, and relative residual 10^{-6} . A predeclared sweep selects the smallest tested grid whose field error is at most 0.09, yielding $113 \times 73 \times 57$. We also retain the $96 \times 64 \times 54$ diagonal-preconditioned FD/CG baseline and random feasible insertion.

All trajectory objectives use 384 fixed Sobol probes and 160 independent walks per active probe. We rerun each stochastic selector end to end for five independent seeds; every final design is then evaluated with at least five fresh 384-walk streams, while the canonical and deterministic designs use ten. Error bars for stochastic selectors are standard errors across complete design runs; deterministic-baseline error bars isolate objective-evaluation noise. Selection timings include BVH construction and every scouting/validation walk for TWoS, or hierarchy construction, solve, and interpolation for multigrid and FD/CG; all exclude process startup, file I/O, and the common independent objective evaluation.

Across the five complete design runs, sparse and near-matched-time TWoS reduce compliance by $68.90\% \pm 0.84\%$ and $69.39\% \pm 0.51\%$, respectively. Tuned multigrid obtains $70.36\% \pm 0.08\%$, FD/CG obtains $69.96\% \pm 0.10\%$, and five random designs obtain $27.76\% \pm 4.57\%$. Thus sparse TWoS realizes 97.9% of the multigrid reduction in 0.432 seconds rather than 1.861 seconds, a $4.31\times$ speedup. At near-matched time, TWoS realizes 98.6% of the multigrid reduction in 1.731 seconds and uses 19.1 MiB peak RSS rather than 31.4 MiB, a $1.65\times$ memory reduction. The near-matched-time mean is about one percentage point below multigrid. The small seed count and the displayed standard errors do not establish statistical equivalence; we report the quality–cost tradeoff directly.

Relative to a separate 512-walk TWoS field, the initial 128-walk TWoS, multigrid, and FD/CG fields have relative RMSE 0.163, 0.0852, and 0.0814, respectively. The multigrid resolution sweep in Fig. 22(a) therefore strengthens rather than weakens the deterministic reference. The design result does not claim that the inexpensive scouting map is more accurate: shortlist validation recovers nearly all multigrid design quality despite noisier global ranking.




Figure 23 Independent external-solver field audit. (a) Cutaway boundary facets and high-ranking sites on the finest retained tetrahedral mesh. (b) Its central-slice u_h^2 field. (c) The independently assembled FEM density and 512-walk TWoS score at the identical 768 candidates after one fitted positive scale; color encodes x position. (d) End-to-end time and field convergence over three resolutions; marker area encodes measured peak RSS and the triangle denotes the finest reference.

X.7 Independent tetrahedral FEM audit

We next remove implementation coupling from the initial-field comparison. A separate Python program uses scikit-fem 11.0.0 [13] to assemble continuous piecewise-linear tetrahedra and SciPy’s sparse direct solver. The only shared inputs are the analytic housing SDF, the analytic load, and the CSV containing the 768 candidate coordinates; no C++ field, grid, walk, or interpolation code is reused. Background tetrahedra are kept when their centroids lie inside the SDF, all exposed facets receive homogeneous Dirichlet data, and the reported ranking density is u_h^2 , since compliance is self-adjoint for this load case.

We solve at $41 \times 27 \times 23$, $61 \times 39 \times 33$, and $81 \times 51 \times 43$ background resolutions. The finest retained mesh has 79,851 vertices and 393,616 tetrahedra; its free-DOF relative algebraic residual is 5.45×10^{-14} . End-to-end construction, assembly, solve, and bounded-memory evaluation at all sites takes 20.37 s and reaches 1.67 GiB peak RSS in the untuned Python process. The two coarser fields have relative RMSE 0.1285 and 0.0802 against the finest field, respectively.

Against the finest FEM field, 128-walk TWoS has Pearson correlation 0.9792, Spearman correlation 0.9849, scale-fitted relative RMSE 0.1514, and 80% top-20 overlap. At 512 walks these become 0.9943, 0.9957, 0.0789, and 95%, respectively; top-50 overlap is 92%. The independent check also supports the in-house reference: the $133 \times 85 \times 67$ multigrid field has Pearson correlation 0.9989, relative RMSE 0.0416, and 95% top-20 overlap against FEM. This experiment validates the field and ranking across discretization families. It is not a claim that this centroid-classified mesh is a production-quality body-fitted discretization, nor a comparison against a complete level-set FEM optimizer.

X.8 Connected cooling-channel growth and shape refinement

The preceding comparison permits six disconnected cavities. We next impose a different, practically motivated topology: one connected channel must join two separated terminals while removing the same nominal volume as the six radius-0.095 spheres. This constrained result is not compared as a quality replacement for the unconstrained cavities; the latter remain an informative unconstrained reference, not a proved optimum or upper bound.

We score 3,072 interior Sobol candidates with 48-walk TWoS, select high-score terminals on opposite




Figure 24 From a topological field to a connected, refined cooling channel. (a) A TWoS field selects two terminals and weights the grid-free connection graph. (b) The shortest weighted path becomes a 22-ball connected channel. (c) Nested Differential WoS moves seven controls; the dashed centerline is the pre-refinement path. (d) Fresh paired line-search differences with Bonferroni-adjusted Student- t whiskers under the paired-normal model; only a wholly negative interval is accepted. (e) Six independent high-signal tangent directions compare the fixed-radius Differential WoS derivative with central finite differences. (f) Independent final compliance reduction at matched nominal volume. The disconnected result is a topology-unconstrained reference, not a connected competitor.

sides of the housing, and construct a 24,365-edge k -nearest-neighbor graph. An edge’s Euclidean length is multiplied by a penalty for low endpoint topological score, and Dijkstra’s algorithm supplies the initial connection. Resampling and safe smoothing give seven controls. Fixed segment subdivisions produce a connected union of 22 overlapping balls, as in Eq. (148). A straight channel between the same terminals is the geometry-only baseline.

The target removed volume is 0.0215482. Every connected trial radius is an order-statistic calibration on the same 2^{20} -point Sobol set; all three reported channels attain QMC volume 0.0215479, within one quantization bound from Proposition G.4. At each shape round, 4,096 Differential WoS paths per active probe estimate all 66 center coordinates jointly; the executable also reports the 22 radius derivatives, hence 88 parameters. With the explicit safe-strip test, the six refinement passes take 3.62–3.82 seconds. The normal offset is $\ell = 0.004$ and the one-offset-width admissible patch is Eq. (149). We apply the exact transpose map in Eq. (154), fix the endpoints, smooth the control proposal, and project out first-order centerline-length change. Each of the three step scales uses eight fresh paired 320-walk objective streams. A one-sided Student- t upper bound uses a predeclared Bonferroni allocation of 0.05 across all 18 possible scale tests. This follows the conditional protocol of Proposition G.5 under a paired-normal interval assumption; exact coverage is not established for the sampled differences. Final values use eight separate 768-walk streams.

The straight and topology-weighted graph channels reduce compliance by $57.269\% \pm 0.080\%$ and $57.325\% \pm 0.048\%$, respectively; their paired objective difference is not resolved ($-8.60 \times 10^{-6} \pm 1.00 \times 10^{-5}$ SE). With safe-strip truncation, the familywise procedure accepts the first three Differential WoS steps and rejects the remaining tested updates. Independent evaluation of the accepted shape gives $58.215\% \pm 0.037\%$




Figure 25 Finite-offset and seam-strip audit on the frozen pre-refinement channel. (a) Central finite differences versus Differential WoS for six directions selected on a separate pilot stream, shown at the smallest, nominal, and largest offsets. (b) FD-on-Differential-WoS slope and cosine similarity across all seven normalized offsets. (c) The discarded harmonic exit mass, both over all exits and conditional on a cavity exit. The dashed line is an $O(\ell)$ reference, while $\hat{p}$ is an unconstrained log-log fit. Error bars are standard errors over three independent streams; the random streams are matched across offsets.

reduction. Relative to the unrefined graph channel, the objective decreases by 2.086%; its paired absolute change is $-1.356 \times 10^{-4} \pm 8.00 \times 10^{-6}$ SE. The unconstrained disconnected design reaches $68.964\% \pm 0.050\%$, showing a measured quality gap between these particular constructions; it does not establish the optimal cost of connectivity.

We audit the fixed-radius gradient at the pre-refinement graph channel with an independent 8,192-walk Differential WoS pass and central differences of step 0.012, each using three paired 4,096-walk replicates. To avoid a noise-dominated correlation statistic, the predeclared seeded directions are the descent direction plus five feasible tangent directions spanning the top 30% of predicted signal; this audit therefore makes no claim about arbitrarily weak directions. Across the six signed directions, Pearson correlation is 0.9982, cosine similarity is 0.9982, and the least-squares slope of finite difference on Differential WoS is 1.341. The coherent scale difference is consistent with finite-offset and finite-difference-step effects in Proposition G.2; signs and relative magnitudes, which drive the validated line search, agree closely.

To isolate the normal offset, we freeze the same seven-control, 22-ball channel and use an independent 8,192-walk pilot to select six feasible high-signal length-tangent directions. The reference central differences use three paired 4,096-walk replicates; each of seven offsets from 0.002 to 0.016 then uses three independent 4,096-walk Differential WoS streams, with seeds matched across offsets. All six signs agree at every offset. Through $\ell/r = 0.102$, cosine similarity is at least 0.9981; at the deliberately large $\ell/r = 0.203$, it falls to 0.9811 and the scale-corrected relative ℓ_2 error rises to 19.3%. The smallest offset has greater $1/\ell$ sampling noise: its FD-on-estimate slope is 1.303, compared with 1.216 at $\ell = 0.003$. Linear extrapolation from offsets through 0.008 has cosine 0.9993 with the finite differences and FD-on-extrapolate slope 1.167; the residual scale difference also includes the finite FD step.

Across 32,858,112 outer paths, the discarded all-exit mass rises from 0.003874 ± 0.000048 to 0.054123 ± 0.000064 . A log-log fit gives exponent 1.269 with $R^2 = 0.99991$; only three additional normal segments fail the retained-domain check (9.13×10^{-8} of paths). This is strong numerical evidence for the geometric exit-mass component of Eq. (150) on this frozen configuration, but it neither measures the $|\partial_n u|$ weight nor proves the asymptotic hypothesis.

The connected experiment demonstrates compatibility with existing-boundary refinement. Neither the topological graph initialization nor the measured line search supplies an unconditional certificate for the full geometry.

X.9 Result summary

Table 13 *Measured results. Controlled diagnostics validate statistical mechanisms; the final rows report matched compiled comparisons, stochastic housing designs, and the equal-volume connected-channel refinement.*

Experiment	Metric	Baseline	Ours	Outcome
Smallest annulus capacity	Variance	5.83×10^{-3}	7.06×10^{-4}	8.26× below coupled finite cavity
Correlation 0.9	Absolute bias	2.708×10^{-1}	7.50×10^{-4}	same-seed error floor removed
256 null candidates	False insertion	100.0%	5.003%	fresh validation meets 5% target
6,359-site WoS field	Spatial RMSE	0.1078	0.0196	5.50× below same-seed product
64-site C++ field	Wall time (s)	0.528	0.381	1.39× faster than voxel FD/CG
384-site, ≈ 0.075 RMSE	Peak RSS (MiB)	80.6	15.6	5.17× smaller than voxel FD/CG
Housing, sparse budget	Compliance reduction	70.36%	68.90%	97.9% of tuned multigrid reduction
Housing, sparse budget	Wall time (s)	1.861	0.432	4.31× faster than tuned multigrid
Housing, near-matched time	Compliance reduction	70.36%	69.39%	98.6% of tuned multigrid reduction
Housing, near-matched time	Peak RSS (MiB)	31.4	19.1	1.65× smaller than tuned multigrid
External FEM field audit	Pearson correlation	1.000	0.9943	512-walk TWoS on identical 768 sites
External FEM field audit	Top-20 agreement	100%	95%	cross-discretization ranking agreement
Connected housing channel	Compliance reduction	57.325%	58.215%	2.086% lower objective after safe-strip DiffWoS refinement
Channel derivative audit	Pearson correlation	—	0.9982	six high-signal tangent directions
Channel offset audit	Discarded-mass exponent	$O(\ell)$ reference	1.269	$R^2 = 0.99991$ over 32.86M paths

Y Accelerated Sampling and Baseline Components

A genuine Poisson-compliance derivative. On the unit disk, solve $-\Delta u = 1$ with homogeneous Dirichlet data and $J = \int_{\Omega} u$. Then $p = u = (1 - |z|^2)/4$ and $D_T(z) = -(1 - |z|^2)^2/16$. Each WoS ball contributes the exact constant-source Green integral $r^2/4$, with the same implementation for MC, scrambled-Sobol RQMC, and shifted-Korobov Array-RQMC [14]. We use five fixed sites, four independently randomized batches, 64, 256, or 1,024 paths per batch, and 32 complete replications at each setting. All products are cross-fitted across batches, as required by Appendix J; trajectories inside an array are never treated as independent samples. The shell tolerance is 10^{-7} , and exhaustion of the 256-step cap raises an error. The maximum marginal truncation bias is $\varepsilon/2$ here; it is not silently equated with zero.

At 1,024 paths per batch, Array-RQMC has mean TD MSE 2.51×10^{-8} versus 5.00×10^{-7} for MC, a 19.9× reduction at equal launched-path budget. The corresponding median times are 1.008 s and 0.073 s. Thus the measured Python implementation trades lower error per path for substantial sorting overhead. It establishes no uniform runtime advantage. RNG initialization and distance queries are included; the timing applies only to these implementations and scenes. The final Poisson campaign runs without other benchmark jobs; harmonic timings come from a shared runtime and are exploratory. Figure 26 reports error against paths, time, and query count; bands are percentile bootstrap intervals for mean MSE across the 32 full replications, not certified insertion intervals.




Figure 26 Measured accelerated sampling with independent-batch cross-fitting. Top: the actual Poisson-compliance topological derivative on the disk. Bottom: a manufactured harmonic-product diagnostic on the square, including neighbor reuse. All methods have equal launched-path budgets at corresponding markers. The three columns expose the difference between statistical efficiency, Python wall time, and distance-query cost. Bands describe uncertainty in empirical MSE across 32 independent full runs.

Neighbor reuse on harmonic diagnostics. On both disk and square, we reuse the manufactured harmonic polynomials from the earlier field diagnostic at 25 fixed sites in $[-0.55, 0.55]^2$. These nonzero-boundary harmonic pairs test the product-estimator algebra; they are not the state and homogeneous-boundary adjoint of the Poisson compliance problem above. We implement the deterministic variance-bound weights in Czekanski et al. [8], equations (17)–(18), and retain the first sphere exit and final boundary value of every walk. Both queries and weights are computed inside each independent batch. At the largest square budget, neighbor reuse reduces product MSE by $1.96\times$ relative to MC using the same launched paths. Our 13-bit Hilbert implementation matches the authors’ Array-RQMC update routine in four common-randomization cases to within 10^{-12} . We do not apply terminal-only neighbor reweighting to a source PDE: its first-ball source contribution would need a separate correct treatment.

Fixed-topology probabilistic shape control. We specialize a probabilistic derivative [31] to a two-parameter ellipse family, giving both optimizers the same exact Poisson state. The target has semiaxes $(1, 0.65)$; the initial semiaxes are $(0.85, 0.85)$. We compare the uniform-start, importance-weighted exit formula with a separately derived deterministic pullback gradient (Appendix K). The deterministic gradient agrees with central differences to within 10^{-8} ; 48 replications at each of three path budgets audit the stochastic error. With identical bounds and Armijo line searches, five stochastic 30-step runs reduce the tracking objective from 1.35×10^{-3} to the range $[1.41 \times 10^{-8}, 9.39 \times 10^{-8}]$. This validates the shape-derivative implementation on a fixed-topology design task. It does not measure neural-state training, boundary FEM smoothing, or the complete published optimizer, and does not compare full optimizers.




Figure 27 Baseline components and geometric feasibility. (a–c) Exact-state ellipse tracking tests the probabilistic shape derivative and descent, with five stochastic runs. (d–e) The reproduced two-phase spectral operator distinguishes sealed islands from exterior-connected voids; two open components are allowed. (f) Existing 3D ball-chain designs satisfy the newly explicit strict-overlap guard. These are component experiments, not complete PINN/FEM or level-set optimizer reproductions.

Spectral connectivity reproduction and geometry correction. We reproduce the two-phase Dirichlet eigenproblem of Andrade et al. [3], equation (4), with P_1 triangles and consistent mass. The cavity coefficients are $k_1 = m_1 = 1$; the other phase has $k_0 = m_0 \in \{10^{-2}, 10^{-4}, 10^{-6}\}$. Five geometries and three mesh resolutions produce 45 eigenproblem solves. At the finest 129×129 node resolution and contrast 10^{-6} , the sealed connected strip has $\lambda_1 \approx 2.24 \times 10^{-5}$, while one open strip and two disconnected open strips both have $\lambda_1 \approx 2.47$. The latter is close to the limiting mixed-boundary value $\pi^2/4$. Thus their constraint tests exterior reachability, which is different from connectedness of our interior channel. At nonzero contrast, even an isolated cavity has a positive eigenvalue; a threshold needs its own contrast and resolution study.

After each radius recalibration, consecutive balls must overlap strictly. All three retained designs satisfy this check, with maximum gap/diameter ratios from 0.275 to 0.434. The geometric guarantee and its exterior-access limitation are proved in Appendix L. Figure 27 separates these validation results from the complete application comparisons, which were not performed for these two external methods.