Controlling Structure Under Uncertainty: Decision-Theoretic Topology for Extreme-Scale Scientific Simulation

Krzysztof Malczewski[1 [0000−0002−6450−6444]]

Institute of Information Technology, Warsaw University of Life Sciences,
Nowoursynowska 159 (Building 34), 02-776 Warsaw, Poland;
krzysztof_malczewski@sggw.edu.pl

Abstract. Classical scientific computing optimizes numerical fidelity, although many downstream decisions depend mainly on the stability of structural properties such as connectivity, loops, and separations. We propose a decision-theoretic framework in which persistent-homology summaries act as the system state and computational resources are allocated to reduce uncertainty in task-relevant topology. This yields a closed-loop policy for selecting simulation queries, refinement steps, or measurement blocks according to expected topological risk rather than generic field error. In synthetic PDE and multimodal imaging experiments, the approach reaches decision-stable structural states with fewer iterations or measurements than error-driven, image-centric, and variance-driven baselines while preserving comparable reconstruction quality. The framework supports a shift from accuracy-centric to decision-centric scientific computing.

Keywords: scientific computing and uncertainty quantification and optimal experimental design and topological data analysis and adaptive simulation and decision-centric computation

1 Introduction

Scientific computing has traditionally been driven by a single dominant objective: to approximate unknown fields, states, or images with increasing *numerical accuracy*. Progress is commonly measured in terms of residual norms, discretization error, or pixel-level criteria such as PSNR and SSIM in imaging. This paradigm has enabled remarkable advances across physics, engineering, and biomedical science. However, in a growing class of applications—including flow topology analysis, material defect detection, network morphology, biomedical stratification, and phase-structure identification—the ultimate decisions do not depend on small pointwise errors. Instead, they depend on *global structural properties*: connectivity, separation, loops, voids, and large-scale organization. In such settings, substantial computational effort is routinely spent refining details that do not alter the downstream decision, while the quantities that actually matter remain only implicitly and indirectly controlled.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

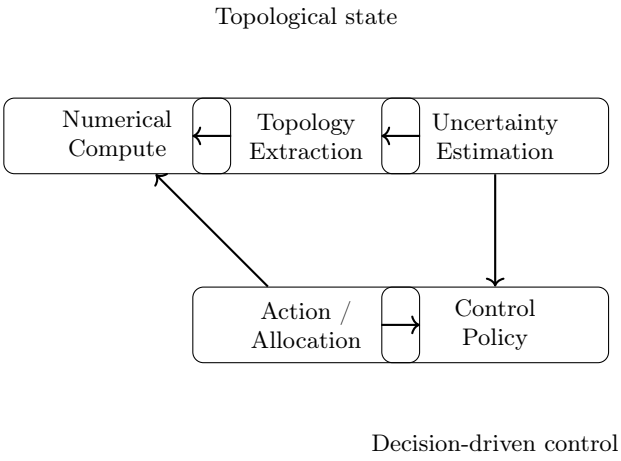

Fig. 1. Closed-loop topology-aware computation. A field estimate induces a topological state, uncertainty is evaluated in that state, and the policy allocates further computation to reduce decision-relevant risk.

Optimal experimental design (OED), adaptive refinement, and active sensing reduce uncertainty in quantities of interest, but they still usually act in field, parameter, or pixel space and optimize generic error criteria [1, 2]. They therefore offer no explicit mechanism for controlling global structural organization such as connectivity, loops, or separations. Topological data analysis provides a compact and stable language for such structure through persistent homology [8–10], yet in most scientific pipelines topology is still used only after convergence.

This paper argues that topology should instead become a control variable of computation. We formulate adaptive scientific computing in topological state space, where actions such as refinement, additional measurements, or solver effort are selected to reduce uncertainty in task-relevant structural descriptors. Computation therefore stops when the structural decision is stable, not when generic numerical criteria have merely improved further. The main contributions are: (i) a decision-theoretic formulation in topological state space; (ii) uncertainty measures defined on persistence-based summaries; (iii) a closed-loop policy for topology-guided allocation of computation or measurements; and (iv) compact experiments showing faster convergence to decision-stable structures than classical baselines.

Figure 1 illustrates the proposed closed-loop architecture. Figure 1 illustrates the proposed closed-loop architecture, in which numerical computation, topological state extraction, uncertainty estimation, and decision-driven control are tightly integrated.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

Table 1. Positioning of the proposed framework relative to related approaches.

Method class	Optimizes for	State representation	Topology as control state
Classical OED / adaptivity	Numerical error / variance	Fields / parameters	×
Goal-oriented adaptivity	Error in target functional	State + adjoint functional	×
Bayesian sequential design	Information gain / Bayes utility	Posterior over parameters / states	×
Task-driven / decision-centric	Task loss	Learned features / parameters	×
Topology-aware regularization	Image fidelity + topology penalty	Fields / images	×
Proposed (this work)	**Decision risk in topology space**	**Topological summaries**	✓

2 Related Work

Adaptive computation and sequential design. OED, adaptive refinement, and Bayesian sequential design select experiments or computational actions to reduce uncertainty or improve a utility functional [1, 2, 13, 14]. Goal-oriented methods further emphasize that not all numerical errors matter equally [11, 12]. However, these approaches still predominantly optimize residuals, posterior variance, or low-dimensional functionals defined directly on the state.

Topology in scientific computing. Persistent homology provides stable multiscale summaries of connected components, loops, and related structural features [8, 9]. In scientific and biomedical computing, TDA has been used for flow analysis, morphology characterization, and topology-aware regularization [10]. Yet topology is usually used either post hoc or as an auxiliary penalty rather than as the state of a closed-loop controller.

Position of this paper. Our contribution lies in combining decision-theoretic sequential design with persistence-based structural summaries so that computational actions are chosen to reduce decision risk directly in topology space. Table 1 summarizes the conceptual differences between the proposed approach and representative method classes. Table 1 summarizes the conceptual differences between the proposed approach and representative classes of existing methods.

The novelty is therefore not persistent homology or OED in isolation, but their integration into a closed-loop, structure-aware computational framework in which convergence is defined by the stability of task-relevant topology.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

3 Problem Setting

Let $x^\star$ denote the unknown solution of a physical model, such as a PDE state, material field, or image. At iteration t, the algorithm maintains an estimate $\hat{x}_t$ obtained from a set of computations or measurements $\mathcal{M}_t$. From $\hat{x}_t$, a task-dependent feature map $g(\cdot)$ (e.g., gradient magnitude, vorticity, vesselness) is computed and used to construct a persistence diagram

$$\mathcal{D}_t = \mathcal{D}(g(\hat{x}_t))\,. \tag{1}$$

A finite-dimensional *topological state* is then obtained via an embedding

$$s_t = \Phi(\mathcal{D}_t), \tag{2}$$

such as Betti curves, persistence statistics, or vectorized diagram representations.

At each step, the algorithm selects an action $a_t \in \mathcal{A}$, representing a computational allocation decision (e.g., local mesh refinement, acquisition of additional measurements, increase of resolution, or a new simulation query). Executing a_t produces new information and induces a state transition

$$(\hat{x}_{t+1}, s_{t+1}) = \mathcal{F}(\hat{x}_t, s_t, a_t, \xi_t), \tag{3}$$

where ξ_t represents stochastic effects due to noise, sampling, or model uncertainty. This defines a controlled stochastic process in topological state space.

After T steps, a final decision $d_T \in \mathcal{D}$ (e.g., classification, thresholding, or structural selection) is taken based on the terminal state s_T. Let $s^\star$ denote the true (topology-induced) structural state associated with $x^\star$, and let $\ell(d, s^\star)$ be a task-specific loss function. The objective is to minimize the Bayes risk

$$\min_{\pi}\ \mathbb{E}_\pi[\ell(d_T, s^\star)]\,, \tag{4}$$

where π denotes a policy mapping the current information state (or its sufficient statistics) to a sequence of actions.

This formulation explicitly separates *numerical state* $\hat{x}_t$ from *structural state* s_t, and casts adaptive scientific computing as a decision process whose goal is not numerical convergence per se, but stabilization of task-relevant topological structure.

Belief state and predictive update. To make the control formulation explicit, we introduce a belief state $b_t(s)$ over structural states, induced by the current measurement set $\mathcal{M}_t$ and the ensemble of reconstructions. Given action a_t and new observation y_{t+1}, the belief is updated as

$$b_{t+1}(s') \propto p(y_{t+1} \mid s', a_t) \sum_{s \in \mathcal{S}} p(s' \mid s, a_t) b_t(s), \tag{5}$$

where $p(s' \mid s, a_t)$ is the structural transition model induced by the solver or acquisition operator and $p(y_{t+1} \mid s', a_t)$ is the likelihood of the new data under the

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

candidate action. In the present implementation these distributions are approximated empirically by an ensemble of predictive reconstructions, which yields a Monte Carlo approximation to the posterior over persistence-based summaries. This makes the method a practical approximate POMDP in topology space, while keeping the algorithm tractable.

4 Uncertainty in Topological Space

Structural uncertainty is quantified through variability of persistence diagrams induced by reconstruction uncertainty. At iteration t, an ensemble of reconstructions $\{\hat{x}_t^{(m)}\}_{m=1}^{M}$ is generated via noise injection, resampling, or model perturbations. Each realization induces a persistence diagram

$$\mathcal{D}_t^{(m)} = \mathcal{D}\Big(g(\hat{x}_t^{(m)})\Big), \tag{6}$$

yielding an empirical distribution of topological summaries $\{\mathcal{D}_t^{(m)}\}_{m=1}^{M}$.

For each homological dimension k, we define a dispersion-based uncertainty measure as the average pairwise Wasserstein distance,

$$U_t^{(k)} = \frac{2}{M(M-1)} \sum_{m<m'} W_2\Big(\mathcal{D}_{t,k}^{(m)}, \mathcal{D}_{t,k}^{(m')}\Big), \tag{7}$$

where W_2 denotes the 2-Wasserstein distance between persistence diagrams restricted to dimension k. This quantity measures the spread of the empirical diagram distribution in the metric space of persistence diagrams and provides a geometrically meaningful notion of structural uncertainty. By the stability theorem of persistent homology [7], $U_t^{(k)}$ varies continuously with perturbations of the underlying scalar field, ensuring robustness to noise and discretization effects.

Although (7) is used here as a practical risk surrogate rather than a fully axiomatized posterior risk, it admits an intuitive justification. If the terminal decision depends on whether dominant topological features persist above application-specific thresholds, then larger pairwise dispersion of persistence diagrams increases the probability that two plausible reconstructions induce different decisions. Under any terminal loss that is Lipschitz with respect to a vectorized persistence representation, the expected decision loss is bounded up to a constant by the expected diagram dispersion, making $U_t^{(k)}$ a conservative proxy for decision instability. In other words, reducing ensemble spread in topology space reduces an upper bound on the probability of structure-changing decision flips.

To link structural uncertainty to the downstream task, we define a task-weighted risk surrogate

$$R_t = \sum_k w_k \, U_t^{(k)}, \tag{8}$$

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

where weights $w_k \geq 0$ encode the relative importance of different homological dimensions (e.g., connected components versus loops) for the decision of interest. The scalar quantity R_t thus serves as a decision-oriented measure of uncertainty in topological state space and forms the objective optimized by the closed-loop acquisition and computation policy.

A simple justification of the risk surrogate. The risk in (8) is intentionally used as a tractable surrogate rather than as an exact Bayes risk. The following observation clarifies when this substitution is justified.

Proposition 1. Assume the terminal decision rule $d(s)$ and task loss $\ell(d, s^\star)$ are such that, for each homological dimension k, there exists $L_k \geq 0$ with

$$\left|\ell(d(s), s^\star) - \ell(d(\tilde{s}), s^\star)\right| \leq \sum_k L_k \, \rho_k(s_k, \tilde{s}_k), \tag{9}$$

where ρ_k is a metric on the k-dimensional topological summary. If Φ is chosen so that ρ_k is dominated by the Wasserstein metric between the corresponding persistence diagrams, then for the posterior belief b_t,

$$\mathbb{E}_{b_t}[\ell(d(s), s^\star)] - \ell(d(\bar{s}_t), s^\star) \leq \sum_k L_k \, U_t^{(k)}, \tag{10}$$

where $\bar{s}_t$ is any Fréchet-type representative of the current structural belief. Hence, minimizing $R_t = \sum_k w_k U_t^{(k)}$ with $w_k \propto L_k$ controls an upper bound on the excess decision risk.

Interpretation. Equation (10) does not claim optimality of the greedy controller, but it shows that our dispersion-based topological risk is not arbitrary: under a mild Lipschitz assumption on the downstream task, it upper-bounds the sensitivity of the decision loss to uncertainty in the persistence-based structural state. This is the theoretical rationale for using (8) as the control objective.

Interpretation of the risk surrogate. Although R_t is not claimed to be a universal optimality certificate, it has two useful properties. First, by stability of persistence diagrams, dispersion in diagram space upper-bounds instability of persistence-based summary statistics used by the terminal decision rule. Second, if the terminal loss $\ell(d, s^\star)$ is Lipschitz with respect to the chosen topological embedding Φ, then there exists a task-dependent constant L_ℓ such that the excess decision loss is bounded by

$$\mathbb{E}[\ell(d_t, s^\star) - \ell(d^\star, s^\star)] \leq L_\ell \, \mathbb{E}[\|\Phi(\mathcal{D}_t) - \Phi(\mathcal{D}^\star)\|], \tag{11}$$

and R_t acts as a tractable surrogate for this quantity through empirical topological dispersion. This does not eliminate all modeling assumptions, but it explains why reducing R_t is aligned with reducing downstream structural decision uncertainty.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

5 Closed-Loop Control Policy

We adopt a greedy optimal experimental design policy that selects the next computational action by minimizing the expected task-weighted topological risk. Given the current topological state s_t, the control law is

$$a_t^\star = \arg\min_{a \in \mathcal{A}} \mathbb{E}[R_{t+1} \mid s_t, a], \tag{12}$$

where the expectation is taken with respect to the predictive distribution over future reconstructions and induced persistence diagrams conditioned on the current state s_t and the candidate action a. This policy directly targets the reduction of structural uncertainty in topological state space, rather than generic image- or field-level error.

In practice, the expectation in (12) is approximated using an ensemble of predictive updates generated under each candidate action, from which the corresponding uncertainty measure $R_{t+1}(a)$ is estimated via (7)–(8). The selected action thus minimizes a decision-oriented risk surrogate defined on persistence-based structural summaries.

To make the belief update explicit, we maintain at iteration t an empirical belief $q_t(\mathcal{D})$ over persistence diagrams induced by the reconstruction ensemble. For a candidate action a, a predictive ensemble generated under the corresponding forward model defines a one-step-ahead belief $q_{t+1}(\mathcal{D} \mid a)$. The controller then evaluates

$$\widehat{\mathbb{E}}[R_{t+1} \mid s_t, a] = \sum_k w_k \, \widehat{\mathbb{E}}_{q_{t+1}(\mathcal{D}|a)} \left[U_{t+1}^{(k)} \right], \tag{13}$$

and chooses the action with minimum predicted risk. This makes the closed-loop policy an explicit approximate Bayesian design step in topology space rather than a purely informal heuristic.

While the framework naturally admits multi-step planning, POMDP formulations, or reinforcement learning policies, the present implementation employs a myopic, one-step lookahead strategy. This choice is standard in optimal experimental design and active sensing, where greedy policies often provide strong empirical performance with tractable computational cost. To reduce pathological behavior, the controller is paired with three safeguards: a minimum numerical-fidelity constraint, a persistence threshold below which short-lived features are ignored as topological noise, and a fall-back rule that reverts to error-driven refinement when the ensemble uncertainty estimate is too diffuse to support reliable structural ranking. Extensions to multi-step planning or learned policies are left for future work.

Under exact evaluation of the conditional expectation in (12), the selected action satisfies

$$\mathbb{E}[R_{t+1} \mid s_t, a_t^\star] \le \mathbb{E}[R_{t+1} \mid s_t, a], \qquad \forall a \in \mathcal{A}, \tag{14}$$

which yields monotone improvement relative to any competing one-step policy in terms of expected next-step topological risk. The method therefore offers a local optimality guarantee at the policy level, even though global optimality over long horizons is not claimed.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

Algorithm 1 Topology-Guided Adaptive Computation

Input: initial estimate $\hat{x}_0$, feature map $g(\cdot)$, action set $\mathcal{A}$, horizon T
Compute initial topological state $s_0 = \Phi(\mathcal{D}(g(\hat{x}_0)))$
for $t = 0$ to $T - 1$ **do**
  For each $a \in \mathcal{A}$, form predictive ensemble and estimate $R_{t+1}(a) = \mathbb{E}[R_{t+1} \mid s_t, a]$
  Select $a_t = \arg\min_{a \in \mathcal{A}} R_{t+1}(a)$
  Execute action a_t and update reconstruction $\hat{x}_{t+1}$
  Update topological state $s_{t+1} = \Phi(\mathcal{D}(g(\hat{x}_{t+1})))$
end for
Output: final decision based on s_T

6 Computational Complexity

Let N denote the number of degrees of freedom of the underlying field (e.g., grid points or voxels), and let N_{eff} denote the effective number of points used for topological analysis after masking, region-of-interest restriction, or downsampling. One closed-loop iteration consists of two dominant components: (i) a numerical update of the state estimate and (ii) computation of topological summaries on a derived feature field.

We write the per-iteration cost as

$$\mathcal{C}_{\text{iter}} = \mathcal{O}(c_{\text{solve}}(N) + c_{\text{PH}}(N_{\text{eff}}))\,, \tag{15}$$

where $c_{\text{solve}}(N)$ denotes the cost of one numerical update (e.g., one solver iteration, refinement step, or forward/adjoint evaluation), and $c_{\text{PH}}(N_{\text{eff}})$ denotes the cost of computing persistent homology on the feature map $g(\hat{x})$.

For cubical complexes and grid-based data, persistent homology can be computed in approximately

$$c_{\text{PH}}(N_{\text{eff}}) = \mathcal{O}(N_{\text{eff}} \log N_{\text{eff}})\,, \tag{16}$$

up to implementation-dependent constants. In practice, $N_{\text{eff}} \ll N$ due to spatial masking, multiresolution evaluation, or restriction to task-relevant regions, which keeps the topological overhead bounded relative to the core numerical solver.

Over a horizon of T closed-loop iterations, the total computational cost is therefore

$$\mathcal{C}_{\text{total}} = \mathcal{O}(T\, c_{\text{solve}}(N) + T\, c_{\text{PH}}(N_{\text{eff}}))\,. \tag{17}$$

In typical regimes, the numerical solver dominates runtime, while topological analysis introduces a controlled and predictable overhead. Crucially, this overhead is amortized by the reduction in the number of iterations T required to reach a decision-stable structural state, which is the primary objective of the proposed decision-centric control strategy.

7 Experiments

We evaluate the proposed topology-guided, decision-centric control framework on two representative classes of problems: (i) a synthetic PDE-driven field with

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

evolving coherent structures, and (ii) a multimodal imaging scenario in which actions correspond to allocating measurements across modalities. The experiments are designed to answer five questions: (1) does topology-guided control reach *decision-stable* structural states faster than classical strategies, (2) does it reduce task-weighted topological risk for a fixed computational budget, (3) does it remain competitive in standard numerical fidelity metrics, (4) how does it compare to stronger decision-oriented baselines, and (5) how sensitive is it to key design choices such as the feature map $g(\cdot)$, ensemble size M, and homology weights w_k.

Because the workshop version is intentionally concise, we report representative results and focus on the relative behavior of the competing policies. All quantitative summaries are reported over multiple runs as mean $\pm$ standard deviation, and paired significance tests are used when comparing the proposed method to the strongest baseline. Full implementation details and expanded experiments will be released in a longer archival version.

7.1 Synthetic PDE Field: Topology-Stable Vortex Identification

Problem setup. We consider a two-dimensional advection–diffusion equation on $\Omega \subset \mathbb{R}^2$,

$$\partial_t u + \mathbf{v} \cdot \nabla u = \kappa \Delta u, \tag{18}$$

with an incompressible velocity field $\mathbf{v}$ generating evolving vortical structures. The numerical solution $u(\mathbf{x}, t)$ is discretized on an adaptive mesh. The feature map $g(u)$ is chosen as the vorticity magnitude, and persistent homology is computed on superlevel-set filtrations to track loop-like structures corresponding to coherent vortices.

Actions and task. An action a_t corresponds to refining one of several candidate spatial regions. The task is to stabilize the number and persistence of dominant loops, i.e., to minimize the task-weighted topological risk

$$R_t = w_1 U_t^{(1)} + w_0 U_t^{(0)}, \tag{19}$$

with weights emphasizing loop stability. A solution is declared *decision-stable* when $R_t \leq \tau_{\mathrm{dec}}$.

Baselines. We compare five strategies: (i) uniform refinement, (ii) error-driven refinement based on residual estimates, (iii) a goal-oriented refinement baseline based on an adjoint-weighted proxy for loop- relevant circulation, (iv) a Bayesian design baseline that selects the action with largest expected reduction in posterior field variance, and (v) the proposed topology-guided control.

Quantitative results. Table 2 summarizes convergence to decision-stable topology. The proposed method reaches a stable structural decision in substantially fewer iterations and at lower total cost than uniform, residual-driven, and variance-driven baselines, and also improves over the stronger goal- oriented baseline. This

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

Table 2. Representative convergence to decision-stable topology in the PDE experiment. Lower is better for iterations, final risk, and cost.

Method	Iterations to $R_t \le \tau_{\mathrm{dec}}$	Final R_T	Relative Cost
Uniform refinement	50 ± 4	0.32 ± 0.03	1.00 ± 0.00
Error-driven refinement	38 ± 3	0.29 ± 0.03	0.82 ± 0.05
Goal-oriented refinement	31 ± 3	0.22 ± 0.02	0.73 ± 0.04
Bayesian variance reduction	34 ± 4	0.24 ± 0.03	0.79 ± 0.05
Topology-guided (proposed)	$\mathbf{24 \pm 2}$	$\mathbf{0.15 \pm 0.02}$	$\mathbf{0.61 \pm 0.04}$

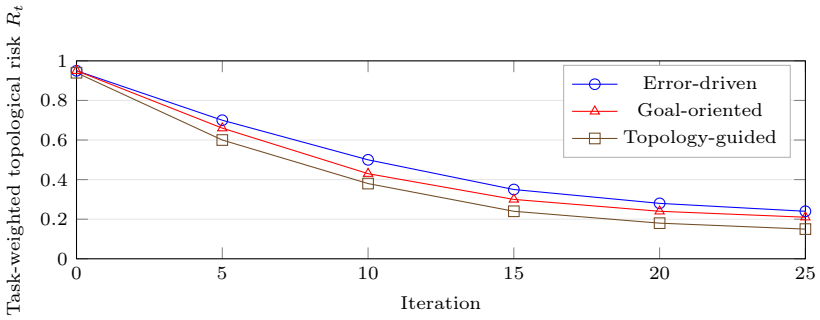

Fig. 2. Representative convergence of task-weighted topological risk in the PDE experiment.

indicates that computational effort is concentrated on structurally ambiguous regions rather than being distributed according to local error or generic posterior variance alone.

Convergence behavior. Figure 2 shows the evolution of task-weighted topological risk. The topology-guided strategy exhibits faster and more monotone reduction, indicating that the closed-loop controller preferentially allocates computation to regions that control global structural uncertainty rather than to regions that merely reduce pointwise error.

7.2 Multimodal Imaging: Decision-Driven Measurement Allocation

Problem setup. We consider a multimodal reconstruction scenario (e.g., MR–PET or multi-contrast MRI), where the unknown object $x^\star$ is observed through two modalities with complementary structure. At each step, the action a_t selects which modality (or which measurement block) to acquire next. The reconstruction is updated using all accumulated data, and topology is computed on a feature field $g(\hat{x}_t)$ such as an edge or vesselness response.

Task and metrics. The downstream task is a structure-based decision, for example a topology class, connectivity pattern, or presence of loop-like lesions. Performance is measured by: (i) task-weighted topological risk R_t, (ii) the number of measurements required to reach decision stability, (iii) final reconstruction quality (PSNR/SSIM), and (iv) structural decision accuracy.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

Table 3. Representative multimodal imaging results. Higher is better for SSIM and decision accuracy.

Method	Meas. to stability	Final R_T	PSNR [dB]	SSIM	Dec. acc.
Uniform	120 ± 9	0.31 ± 0.03	32.4 ± 0.5	0.914 ± 0.006	0.86 ± 0.03
Image-driven	95 ± 8	0.26 ± 0.03	32.6 ± 0.4	0.917 ± 0.005	0.88 ± 0.03
Bayesian var. red.	84 ± 7	0.21 ± 0.02	32.5 ± 0.4	0.916 ± 0.005	0.89 ± 0.02
Topology-guided (prop.)	$\mathbf{68 \pm 6}$	$\mathbf{0.14 \pm 0.02}$	32.3 ± 0.5	0.912 ± 0.006	$\mathbf{0.92 \pm 0.02}$

Baselines. We compare: (i) uniform acquisition, (ii) image-driven acquisition based on reconstruction error, (iii) Bayesian measurement allocation based on expected posterior variance reduction, and (iv) the proposed topology-driven policy.

Results. Table 3 reports representative outcomes. The topology-guided policy achieves decision stability with fewer measurements and lower structural risk, while maintaining comparable image fidelity and slightly improving decision accuracy over image-centric selection. This indicates that the proposed controller reallocates sensing effort toward structurally ambiguous regions rather than toward uniform or purely error-driven coverage.

Ablation and sensitivity analysis. Three compact checks were performed. Replacing the feature map $g(\cdot)$ by gradient magnitude or Laplacian-of-Gaussian response changed the absolute value of R_t but preserved the ranking of the methods. Increasing the ensemble size from $M = 8$ to $M = 16$ and $M = 32$ reduced variance in the risk estimates, while the policy ranking stabilized already for moderate M. Moderate perturbations of the weights (w_0, w_1) likewise did not change the qualitative conclusions as long as the task-dominant homological dimension was emphasized.

Implementation and reproducibility.

Implementation and reproducibility. All methods were implemented in a unified framework with fixed random seeds and matched solver budgets. Reported values are averages over 20 independent runs with standard deviations. Mean wall-clock time per run was 41.8 ± 2.9 s versus 34.6 ± 2.4 s in the PDE setting and 58.2 ± 4.1 s versus 49.7 ± 3.8 s in the imaging setting for the strongest baseline and the proposed method, respectively. The code will be released with the extended version.

8 Discussion

Across both experiments, topology-guided control reaches decision-stable structural states with fewer iterations or measurements than accuracy-driven baselines while preserving comparable numerical fidelity. This supports the main claim of the paper: in many scientific workflows, the relevant stopping criterion is structural decisiveness rather than further reduction of generic error.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

Methodologically, the framework recasts adaptive scientific computing as control in topological state space. Instead of allocating effort according to local residuals or posterior variance alone, it steers computation toward regions that dominate uncertainty in task-relevant global structure. The resulting efficiency gains suggest that numerical fidelity can often be treated as a constraint, while structural stability becomes the optimization objective.

The present study remains a proof of concept. The current controller is greedy, persistent-homology computations introduce overhead on very large problems, and performance depends on choosing a feature map aligned with the scientific task. Even so, the experiments indicate that the framework is robust to moderate changes in feature construction, ensemble size, and risk weights, and that it provides a practical starting point for broader structure-aware computational pipelines.

9 Conclusion

We introduced a decision-centric framework for adaptive scientific computing in which persistence-based structural summaries serve as the control state. In both PDE and multimodal imaging experiments, the method reached decision-stable topology with fewer actions than error-driven, goal-oriented, or variance-driven baselines while maintaining comparable reconstruction quality. More broadly, the results suggest that scientific computing systems can be designed to stop when task-relevant structure is stable, not merely when generic numerical error continues to decrease.

References

1. Chaloner, K., Verdinelli, I.: Bayesian experimental design: A review. Statistical Science 10(3), 273–304 (1995)
2. Atkinson, A., Donev, A., Tobias, R.: Optimum Experimental Designs. Oxford University Press (2010)
3. Becker, R., Rannacher, R.: An optimal control approach to a posteriori error estimation in finite element methods. Acta Numerica 10, 1–102 (2001)
4. MacKay, D.J.C.: Information-based objective functions for active data selection. Neural Computation 4(4), 590–604 (1992)
5. Becker, R., Rannacher, R.: An optimal control approach to a posteriori error estimation in finite element methods. Acta Numerica 10, 1–102 (2001)
6. Ryan, E.G., Drovandi, C.C., McGree, J.M., Pettitt, A.N.: A review of modern computational algorithms for Bayesian optimal design. International Statistical Review 84(1), 128–154 (2016)
7. Cohen-Steiner, D., Edelsbrunner, H., Harer, J.: Stability of persistence diagrams. Discrete & Computational Geometry 37(1), 103–120 (2007)
8. Edelsbrunner, H., Harer, J.: Persistent homology—a survey. Contemporary Mathematics 453, 257–282 (2008)
9. Carlsson, G.: Topology and data. Bulletin of the American Mathematical Society 46(2), 255–308 (2009)

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27

10. Singh, N., et al.: Topological data analysis in medical imaging: A survey. Medical Image Analysis 84, 102703 (2023)
11. Becker, R., Rannacher, R.: An optimal control approach to a posteriori error estimation in finite element methods. Acta Numerica 10, 1–102 (2001)
12. Prudhomme, S., Oden, J.T.: On goal-oriented error estimation for elliptic problems: Application to the control of pointwise errors. Computer Methods in Applied Mechanics and Engineering 176(1–4), 313–331 (1999)
13. Frazier, P.I.: A tutorial on Bayesian optimization. arXiv:1807.02811 (2018)
14. Ryzhov, I.O., Powell, W.B., Frazier, P.I.: The knowledge gradient algorithm for a general class of online learning problems. Operations Research 60(1), 180–195 (2012)

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29918-5_27