POSTERIOR ENTRANCES: REUSE THE INFORMATION, PRESERVE THE SAMPLING LAW

Anonymous authors
Paper under double-blind review

ABSTRACT

Recovering an image from incomplete measurements requires useful reconstructions under a finite inference budget. Yet generating posterior samples repeatedly purchases learned-prior information at related states. We ask whether this information can be reused without changing the prescribed sampling transition. Posterior Entrances (PE) couples two operations: certify a retained first-order reference under the current context, then correct its proposal with live queries. A history-conditioned rejection identity preserves the ideal block conditional and finite-horizon state law; a movement-based refresh bound and a sufficient break-even condition account for both avoided anchor work and added correction. Across two trained priors and four model–observation configurations, paired-median fresh-FORS/PE ratios at sixteen sweeps are **1.354–1.369 for total queries** and **1.135–1.296 for pipeline time**. In the representative profile, the complete query ledger shows a net saving of 80.34 queried states per endpoint and 22.85% of complete pipeline time. Site-resolved diagnostics evaluate 512-endpoint estimates. The same entrance-to-continuation design motivates E32, which carries observation-informed diffusion states into a common suffix for fixed-budget photographic reconstruction. The method makes reference lifetime a controllable inference resource while preserving the prescribed transition.

1 INTRODUCTION

An incomplete image does not specify a unique reconstruction. Missing pixels leave unobserved content, while spatial averaging removes fine detail. A posterior-oriented restoration pipeline must therefore do more than return one plausible image: it must generate useful alternatives or estimate observables within an available computation budget. We approach this problem by organizing generation around an entrance: an intermediate object that carries acquired information into the next stage of computation. PE develops this idea for retained first-order references; E32 evaluates it for observation-informed diffusion states in photographic reconstruction.

The computational question is simple: *must a sampler purchase a new first-order prior evaluation whenever the surrounding context changes?* A previous evaluation may still describe the local geometry, but treating it as current information without correction can change the transition being implemented. Conversely, acquiring an entirely fresh reference discards information that has already been paid for. We study the middle ground: reuse the evaluated information, quantify its discrepancy, and correct against the current conditional.

Two observations motivate corrected reuse. First, model evaluation accounts for 92.63% of the fresh-reference pipeline in the representative recorded workload. Second, retaining evaluated information yields a net saving of 80.34 queries per endpoint after accounting for the complete conditional draw. These observations motivate two coupled decisions: *when is a retained reference admissible, and when does correcting it cost less than refreshing it?*

Posterior Entrances answers the first question with a current-context discrepancy certificate and the second with a complete query ledger. An admitted reference shapes a Gaussian proposal; a known-term prefilter and live randomized correction produce the current block-conditional draw.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

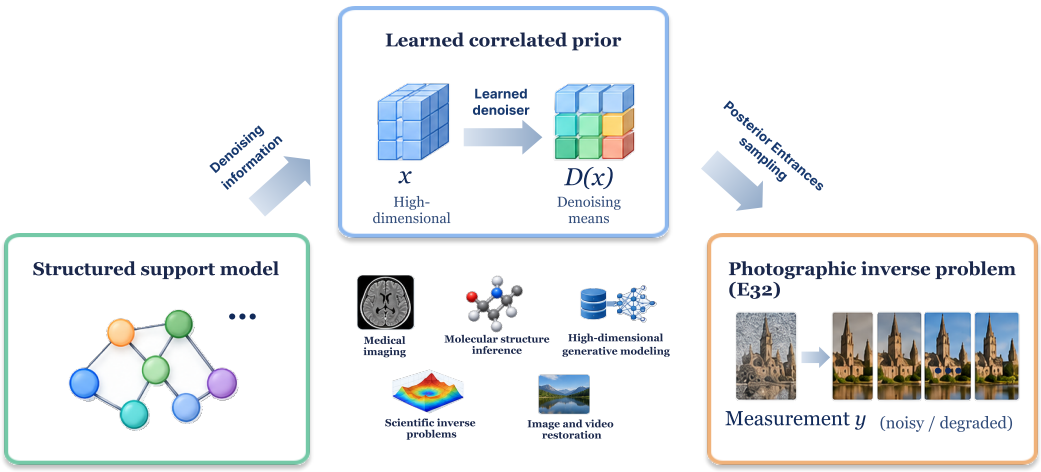

Figure 1: **Research scope and motivating applications.** The study follows entrance construction through subsequent generation: certified first-order references support PE sampling, and observation-informed diffusion states support E32 photographic reconstruction. The images and task icons are conceptual illustrations. The experimental studies cover trained digit posteriors and photographic inverse problems; the additional icons illustrate potential application domains.

The reference stays fixed during that draw and is refreshed only when its declared allowance or residual threshold is exceeded. Fresh randomness creates the output; the reusable object is the evaluated field, not a recycled sample.

Contributions and evidence. We contribute (i) a certificate-based interface for historical first-order information, together with history-conditioned preservation of the ideal conditional transition; (ii) a movement-triggered retention policy and a sufficient net-query benefit condition that explicitly charges live correction; and (iii) a matched trained-model evaluation connecting reference acquisition, pipeline time, and sitewise estimation quality. FORS (Chen et al., 2026) supplies the Gaussian/Poisson correction engine; PE contributes the historical-reference policy. We then carry the entrance-to-continuation design into photographic reconstruction with E32. The trained study measures the value of retained first-order information under a corrected transition; the photographic study measures reconstruction quality from retained diffusion states under a fixed output budget. Figure 1 presents this progression from the sampling mechanism to image reconstruction.

2 Related Work

Approximation with correction. First-order rejection sampling (FORS) supplies the Gaussian-reference and randomized-correction mechanism (Chen et al., 2026). Delayed acceptance (Christen and Fox, 2005) and scalable exact Metropolis–Hastings (Cornish et al., 2019) also exploit approximation with correction. PE's increment is a historical-reference allowance, its matched placement in both acceptance stages, and movement-based information acquisition for an evolving block conditional. Fresh FORS and a calibrated static reference therefore serve as controls of the *reference policy* under the same conditional engine.

Denoising and distributional generation. Denoising optimization (Vincent, 2011), constrained potentials (Amos et al., 2017), and Gaussian conditional-mean identities (Dytso et al., 2021) provide the learned-field context. The trained study uses the gradient of its specified potential as its first-order interface. MCGdiff (Cardoso et al., 2024) organizes intermediate inverse-problem distributions; finite-particle SVGD (Liu and Wang, 2016; Banerjee et al., 2025) emphasizes the relationship between a finite algorithm and its evaluated law. For photographic entrance construction, PC (Song et al., 2021) and DDRM (Kawar et al., 2022) supply entrance states, while DAPS (Zhang et al., 2025), TGD (Makris et al., 2026), and Spin (Shtanchaev et al., 2026) supply recorded comparison

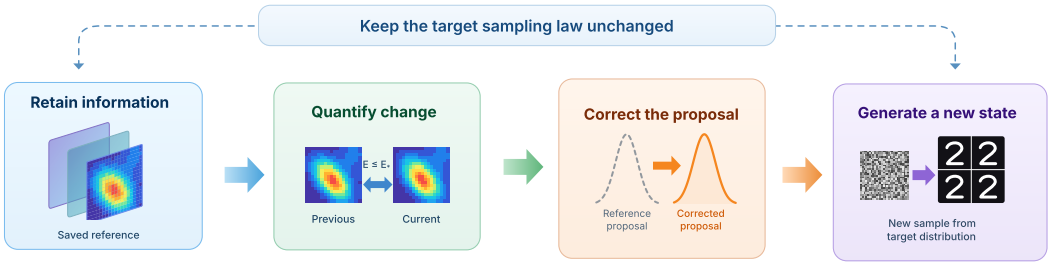

Figure 2: **The corrected-reference mechanism.** Admission or refresh precedes proposal randomness. The selected reference remains fixed throughout the conditional draw, including rejected proposals. Both reference acquisition and every live correction are charged. Icons are schematic; the target is the current block conditional.

configurations. E32 connects these entrance sources through a common continuation and evaluates the resulting finite ensemble at a fixed budget.

3 Posterior Entrances: Admit, Correct, and Refresh

3.1 Problem formulation and the reference interface

For block variable z and fixed outside context u, write

$$F_u(z) = \tfrac{1}{2} z^\top H z - h_u^\top z + V(z) - S_u(z), \qquad D_u = \nabla S_u. \tag{1}$$

The client evaluates H, h_u, and V and queries D_u. The geometry satisfies $0 \preceq \nabla^2 S_u \preceq aI$, $0 \preceq \nabla^2 V \preceq bI$, and $H - aI \succ 0$. The goal is a draw from the current block conditional, while reducing the first-order information acquired to realize that draw. A reference $(r, \widehat{D}, u_*)$ is admissible when it has a valid current allowance

$$\left\| D_u(r) - \widehat{D} \right\| \le E(u, u_*; r). \tag{2}$$

This certificate separates the conditional sampler from its model-specific bound provider. The executed policy takes $\widehat{D} = D_{u_*}(r)$ and $E = L_D \|u - u_*\|$. Define $\widehat{g} = Hr - h_u + \nabla V(r) - \widehat{D}$, $P = H - (a+\gamma)I \succ 0$, and $\theta = E^2/(2\gamma)$ for $\gamma > 0$. Propose

$$Z \sim q = \mathcal{N}(r - P^{-1}\widehat{g}, P^{-1}). \tag{3}$$

A fresh exact reference has $E = 0$ and admits $\gamma = \theta = 0$. Historical information thus changes proposal construction through a declared allowance, rather than being silently substituted for a current evaluation.

3.2 Current-context correction preserves the conditional

For $\delta = Z - r$, $s = \|\delta\|$, and $V_r(\delta) = V(r+\delta) - V(r) - \nabla V(r)^\top \delta$, the two acceptance stages use

$$\begin{aligned} U_Z &= -V_r(\delta) - \gamma s^2/2 + Es - \theta, \\ R_Z(t) &= at\, s^2 - \delta^\top [D_u(r + t\delta) - \widehat{D}] + Es, \\ B_Z(t) &= at\, s^2 + 2Es, \qquad \Lambda_Z = as^2/2 + 2Es. \end{aligned} \tag{4}$$

The geometry gives $U_Z \le 0$ and $0 \le R_Z \le B_Z$. First pass a prefilter with probability $\exp(U_Z)$. Then draw $N \sim \mathrm{Poisson}(\Lambda_Z)$ locations with density B_Z/Λ_Z, testing factors $1 - R_Z/B_Z$ sequentially. Passing every factor accepts the proposal. For $\Lambda_Z = 0$ the product equals one. Failed trials restart proposal generation without changing the admitted reference.

Theorem 1 — History-conditioned preservation. Under the stated geometry, valid allowance (2), and exact real arithmetic, the complete acceptance probability satisfies

$$q(z)\alpha(z) = c \exp[-F_u(z)], \qquad c > 0 \text{ independent of } z. \tag{5}$$

Conditional on the entire reference history, the returned draw has the current block conditional. With a common initial law and the same history-independent random-scan distribution, PE and fresh FORS have the same ideal finite-horizon state law.

The Poisson product averages to $\exp(-\int R_Z)$. Its allowance cancels with the prefilter and Gaussian proposal terms. Keeping the cache fixed throughout the draw is essential to this cancellation. This specializes the first-order correction mechanism (Chen et al., 2026) to historical references; Supplement S1 gives the complete proof. The theorem establishes preservation of the ideal transition under its stated assumptions.

Table 1: Fresh-reference implementation and PE components.

Algorithmic component	**Fresh-reference implementation**	**PE increment**
Reference	Current evaluated anchor	Retained field and its context
Allowance	$E = 0$ at acquisition	Current certificate $E(u, u_*; r)$
Correction	Gaussian / Poisson engine	Allowance in both acceptance stages
New information	Fresh-context anchor policy	Movement-triggered refresh

3.3 When is reuse worth its correction cost?

The objective is total information acquisition. Let $v_q = \mathrm{tr}(P^{-1}) + \left\|P^{-1}\widehat{g}\right\|^2$ and $A_q = \mathbb{E}_q\alpha(Z)$. The expected number of live queries in one completed conditional draw obeys

$$\mathbb{E}Q_{\mathrm{live}} \le (av_q/2 + 2E\sqrt{v_q})/A_q. \tag{6}$$

The curvature and discrepancy terms price live correction; refreshes add their anchor queries. For N visits to one block, let $V_N = \sum_{j=2}^{N} \|u_j - u_{j-1}\|$. If h_u is L_h-Lipschitz and a fresh anchor has residual at most $G_f < G_*$, the fixed-anchor policy satisfies

$$\begin{aligned} \Delta_* &= \min\{E_*/L_D, (G_* - G_f)/L_h\}, \\ R_N &\le \min\{N, 1 + V_N/\Delta_*\}. \end{aligned} \tag{7}$$

Proposition 2 — Sufficient net-query benefit. Write A_N for fresh-reference anchor work and C_N^F, C_N^P for the live-query totals. If $\mathbb{E}A_N \ge fN$, $\mathbb{E}V_N \le \nu N$, and $\mathbb{E}(C_N^P - C_N^F) \le \eta N$, then

$$\mathbb{E}Q_N^F - \mathbb{E}Q_N^P \ge N(f - \nu/\Delta_* - \eta) - 1. \tag{8}$$

Thus $g = f - \nu/\Delta_* - \eta > 0$ yields strict expected savings for $N > 1/g$.

This sufficient condition puts the reference lifetime and correction premium in the same inequality. Supplement S2 derives the bounds and states their expectation premises; the experiments measure anchor work, live work, and total time separately to test the corresponding operational trade-off.

4 Experiments on a theory-matched trained system

The evaluation asks three questions: does certified reuse reduce cost under a matched transition, does the gain survive live correction and pipeline overhead, and does it produce useful sitewise estimates sooner? A constrained trained prior makes the certificate and proposal geometry explicit.

4.1 Experimental setting and matched controls

The tied conservative network on 64 image coordinates is

$$\begin{aligned} S_\vartheta(x) &= \sum_{i=1}^{64}\sum_{h=1}^{m} c_h \log\cosh(w_{0h}x_i + w_{1h}x_{i+1} + b_h), \\ D_\vartheta &= \nabla S_\vartheta, \qquad p_\vartheta \propto \exp(-\|x\|^2/2 + S_\vartheta). \end{aligned} \tag{9}$$

Periodic indexing and positive constrained coefficients give $a = \sum c_h(|w_{0h}| + |w_{1h}|)^2$ and $L_D = \sum c_h|w_{0h}w_{1h}|$. Denoising optimization (Vincent, 2011) and constrained potentials (Amos et al., 2017) provide a trained field and its operational geometry. The 1,797-image optical-digit collection (Alpaydin and Kaynak, 1998) is divided into 1,200 training, 300 validation, and 297 test images. Both models use 300 Adam updates (Kingma and Ba, 2015), eight images per batch, learning rate 0.01, and unit-standard-deviation Gaussian training noise.

The 1,024- and 8,192-channel models contain 4,097 and 32,769 parameters. Validation MSE decreases from 0.232457 to 0.223229 and from 0.232698 to 0.223237, respectively. Inference fixes these checkpoints. Noisy readings and binary observations define

$$\begin{aligned}\Phi(x) = \tfrac{1}{8}\left\|x - y\right\|^2 + \tfrac{1}{2}J\sum_i (x_i - x_{i+1})^2 \\ + \sum_i \{\log(1 + e^{x_i}) - o_i x_i\}, \qquad J = 1/8.\end{aligned} \tag{10}$$

Four-coordinate blocks use $H_{jj} = 1.5$, $H_{j,j+1} = -0.125$, and logistic V. The public anchor is $r_0(u) = H^{-1}(h_u - \frac{1}{2}\mathbf{1})$. The executed settings are $E_* = 0.08$, retained-reference $\gamma = 1/16$, and $G_* = G_f + 0.5$. A refresh occurs at initialization or when the current allowance or residual exceeds its threshold.

Table 2: **Geometry for the executed model.** These checkpoint-derived values connect the reference certificate, refresh rule, and Gaussian proposal.

Operational quantity	**1,024 channels**	**8,192 channels**
Curvature a	0.335127	0.335316
Context constant L_D	0.048745	0.048803
Fresh residual G_f	3.505120	3.512566
Proposal minimum eigenvalue	0.900119	0.899930
Movement allowance Δ_*	1.641205	1.639234

Controls, budgets, and timing. The wide model uses held-out images 1020, 1598, and 36; the 1,024-channel model uses image 1020. Each method–target cell contains four production seeds and 128 endpoints per seed. A sweep is sixteen random-scan block visits. Fresh FORS uses the shared conditional engine; a calibrated static reference acquires one all-zero field and uses $E = a\left\|r_0(u)\right\| + L_D\left\|u\right\|$. Its one shared query is charged. Dedicated calibration seeds choose static γ from $\{0.25, 0.5, 0.75\}$. Timers cover geometry, preprocessing, initialization, scan generation, and sampling. Supplement S3 specifies training, inference, and statistical units.

The fresh-FORS comparison isolates PE's reference-policy contribution under a shared conditional engine. The companion `COMPARISON_RECORD.md` supplies the complete diagnostic comparisons and recorded trajectories.

4.2 Same update count, less information acquisition

At sixteen sweeps, all four configurations favor PE in paired-median total-query and complete-pipeline comparisons (Table 3). Static-reference controls also favor PE. The two model widths expose different projected-evaluation costs: 5,120 and 40,960 nonlinear activations, while both state dimensions remain 64. Together, they provide four model–observation configurations across two widths of the trained model family.

The four paired workloads define each configuration median. The complete paired ranges and record-level diagnostics accompany the paper in `COMPARISON_RECORD.md`.

4.3 Net information and pipeline savings

On the representative wide target, total acquisition decreases from 304.20 to 223.86 queried states per endpoint. The arithmetic-mean profile therefore saves **80.34 queries per endpoint**, a 26.41%

Table 3: **Four complete configuration summaries.** Each entry is a median of four paired baseline/PE ratios for 128-output workloads. Values greater than one express lower PE cost.

Channels / image	Fresh / PE queries	Fresh / PE time	Static / PE queries	Static / PE time
8,192 / 1020	1.365	1.296	3.392	3.529
8,192 / 1598	1.364	1.218	2.941	2.803
8,192 / 36	1.369	1.265	3.793	3.527
1,024 / 1020	1.354	1.135	3.363	3.685

reduction. Both anchor acquisition and live correction contribute to these totals; the companion record preserves their full decomposition.

Model evaluation accounts for 92.63% of the fresh-reference pipeline in this profile. Its time falls from 10.237 to 7.580 seconds per 128-output workload. Complete pipeline time falls from 11.052 to 8.527 seconds, a source-reported **22.85% reduction**. These profile values are arithmetic means; Table 3 reports medians of paired ratios.

4.4 Efficient sitewise estimation

Independent transfer quadrature supplies each site's mean and second moment. For $\psi_{i,1}(x) = x_i$ and $\psi_{i,2}(x) = x_i^2$, define

$$E_{\mathrm{site}}^2 = (1/128) \sum_{i=1}^{64} \sum_{j=1}^{2} \left[(\bar{\psi}_{i,j} - \mathbb{E}_\pi \psi_{i,j}) / \sqrt{\mathrm{Var}_\pi(\psi_{i,j})} \right]^2 . \tag{11}$$

Standardization and squaring precede spatial averaging. Whole-endpoint bootstrap resampling (Efron, 1979) produces the pointwise upper-percentile diagnostic. Every quality point uses 512 endpoints and the sum of their four workload times. The source analysis contains 384 method–target–horizon summaries and 49,152 paired checkpoint records, retaining target, seed, endpoint, and horizon as a hierarchy.

The representative PE estimate first meets the 0.10 level at three sweeps: $E_{\mathrm{site}} = 0.04688$, upper percentile 0.07239, and 8.324 seconds. The corresponding fresh-FORS checkpoint takes 9.877 seconds. This target-specific comparison links the preserved ideal transition to an observed reduction in the time needed for the recorded estimate. The companion Markdown contains the full eight-method trajectories and four-target crossing table.

4.5 Sensitivity of the retention policy

Seven recorded settings vary E_*, γ, and the residual threshold. Their paired-median fresh-FORS/PE query ratios range from 1.17 to 1.43 and time ratios from 1.07 to 1.32. Along the allowance axis, raising E_* from 0.04 to 0.16 changes the refresh fraction from 0.474 to 0.134 and total queries from 129.3 to 106.8 per endpoint. Both anchor and live work are included. Supplement S4 reports all seven settings and their joint query/time trade-off, documenting a usable tested range.

5 From reference entrances to photographic reconstruction

5.1 Carrying acquired information into generation

PE organizes sampling into entrance construction and subsequent generation. Its entrance is an evaluated first-order reference, and current-context correction turns that reference into a valid conditional draw. For photographs, E32 applies this organizing principle to a different computational object: a native diffusion state built using the observation. An aligned noise parameterization defines the handoff, and a common stochastic suffix carries each entrance to a reconstruction. PE supplies the conditional-preservation result for corrected references; E32 supplies the fixed-budget reconstruction evaluation for native-state entrances.

This leads to a concrete output contract: given a degraded photograph, return four saved reconstructions under a fixed denoiser-state budget. Random 70% missing observations test completion, while 4×4 spatial averaging tests fine-detail recovery from coarse measurements. Both use simulated operators on real photographs.

E32 constructs two PC (Song et al., 2021) and two DDRM (Kawar et al., 2022) entrances, using 224 evaluations per branch. It retains the resulting native states, aligns their noise parameterization, and continues each through the common 32-step suffix with fresh branch randomness. The pretrained checkpoint follows Dhariwal and Nichol (2021). Observation information enters through the prefixes and is carried forward in the retained states. The four outputs cost $4(224+32) = 1{,}024$ evaluated denoiser states, explicitly charging both entrance construction and continuation. This mirrors PE's accounting principle: evaluate the complete path from acquired information to the delivered output.

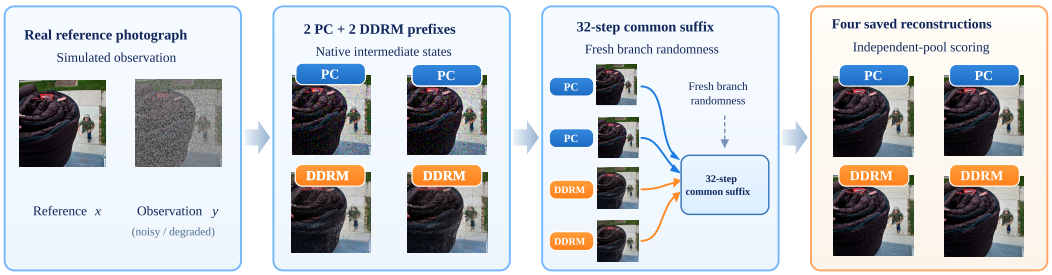

Figure 3: **The entrance-to-continuation design for photographs.** Observation-informed PC and DDRM prefixes construct four native-state entrances. Noise alignment connects them to a common 32-step suffix, yielding four reconstructions within the stated total budget. Thumbnails illustrate the computation using archived images reused across stages and methods. Figure 4 presents the saved empirical reconstructions.

5.2 Reconstruction quality from a budgeted entrance pool

With the entrance pool and continuation fixed, we evaluate the quality of the delivered reconstructions. The archive uses 100 ImageNet validation photographs (Deng et al., 2009; Russakovsky et al., 2015). For two independently initialized pools of $m = 4$ outputs and $\rho(x, z) = \|x - z\| / \sqrt{D}$, its primary score is

$$S_\times = (1/2m) \sum_{c=0,1} \sum_{i} \rho(X_i^c, x^*) - (1/2m^2) \sum_{i,j} \rho(X_i^0, X_j^1). \tag{12}$$

This energy score combines reference fit and ensemble dispersion (Gneiting and Raftery, 2007); lower is better. The output budget above is *per four-output pool*; evaluation uses two separately initialized pools, each charged its own four-output budget. Image identity is the paired statistical unit. The five released comparison cells in Table 4 satisfy the source composite improvement criterion and retain its fifteen-contrast Holm family (Holm, 1979); each includes all 100 paired images. The companion Markdown provides the broader comparison context.

Table 4: **Selected photographic advantages.** The five released cells use the same archived forward-state cap and 100-image set. Reductions are relative to the comparator score and preserve the original fifteen-contrast Holm family.

Observation	Comparator	E32	Comparator	Reduction
Random 70% missing	PC (Song et al., 2021)	0.08101	0.10238	**20.87%**
Random 70% missing	Spin (Shtanchaev et al., 2026)	0.08101	0.10207	**20.63%**
4×4 averaging	DAPS (Zhang et al., 2025)	0.08835	0.35507	**75.12%**
4×4 averaging	TGD (Makris et al., 2026)	0.08835	0.17213	**48.67%**
4×4 averaging	Spin (Shtanchaev et al., 2026)	0.08835	0.10494	**15.81%**



Figure 4: **Real reference, simulated measurements, and saved reconstructions.** The archived v45 panel uses scene `ILSVRC2012_val_00049108`, selected by the first-protocol-image rule. Every reconstruction is pool 0, draw 0; all detail strips use the same 128×64-pixel crop. Ten distinct source PNGs retain their recorded hashes. Aggregate comparison values summarize the full set of 100 paired images.

5.3 From the entrance pool to visible reconstructions

Figure 4 shows the output of this entrance-to-continuation process for one reference scene under both observation models. The sparse-pixel row shows completion under missing information; the block-average row shows image structure recovered from coarse observations. The common source image, fixed draw index, and shared crop define a consistent visual comparison. Table 4 supplies the complementary 100-image evaluation, connecting the specified entrance pool and computation budget to measured reconstruction quality.

A common computational interface. The connection between PE and E32 is operational: specify the information-bearing entrance, define the continuation that consumes it, and charge both stages. PE implements this interface with a certified reference and live correction. E32 implements the entrance-to-continuation organization with an observation-informed state and a diffusion suffix. Its fixed-quota pool turns the design into a four-reconstruction output for the tested restoration tasks.

6 Reference reuse as an inference resource

The two evaluations follow one computational question: how should acquired information enter subsequent generation? In PE, a current discrepancy certificate admits the retained reference, and live correction preserves the ideal conditional transition. The trained experiments quantify the resulting query and pipeline savings. In E32, observation-informed prefixes construct a pool of diffusion entrances, and the common suffix produces photographic reconstructions. The 100-image evaluation measures the quality of that output pool at the recorded budget. This progression connects entrance definition, continuation, and end-to-end accounting to the object each experiment delivers.

7 CONCLUSION

Posterior Entrances treats evaluated first-order information as a reusable sampling resource. A certificate decides whether to retain it, current-context correction preserves the prescribed ideal transition, and the complete ledger determines whether reuse pays. All four trained configurations show favorable paired-median query and pipeline-time comparisons; the representative 80.34-query net saving becomes a 22.85% pipeline-time reduction. E32 carries the entrance-to-continuation design into photographic reconstruction by handing observation-informed native states to a common stochastic suffix. Together, the studies connect an explicit entrance, a defined generation procedure, and a complete computation budget to measurable sampling and reconstruction outcomes.

REPRODUCIBILITY STATEMENT

Supplement S1–S2 contains the conditional-preservation and work-bound proofs. S3–S6 specifies training, inference, timing, statistical units, sensitivity, and the photographic implementation. S7–S9 maps the sources and record hierarchy. The companion `COMPARISON_RECORD.md` provides the complete comparison and diagnostic record; the data files preserve recorded values and coordinates.

AI USE STATEMENT

This research was led and primarily carried out by the authors. Generative AI assistants were used for language polishing and auxiliary manuscript preparation, including improving the organization and clarity of author-provided material, LaTeX formatting, figure integration, and reference and consistency checks. This assistance concerned the presentation of the authors' research rather than the generation of new scientific findings or experimental observations.

REFERENCES

Ethem Alpaydin and Cenk Kaynak. Optical Recognition of Handwritten Digits. *UCI Machine Learning Repository*, 1998. doi: 10.24432/C50P49.

Brandon Amos, Lei Xu, and J. Zico Kolter. Input Convex Neural Networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 146–155, PMLR, 2017.

Sayan Banerjee, Krishnakumar Balasubramanian, and Promit Ghosal. Improved Finite-Particle Convergence Rates for Stein Variational Gradient Descent. In *International Conference on Learning Representations*, 2025.

Gabriel Cardoso, Yazid Janati El Idrissi, Sylvain Le Corff, and Eric Moulines. Monte Carlo Guided Denoising Diffusion Models for Bayesian Linear Inverse Problems. In *International Conference on Learning Representations*, 2024.

Fan Chen, Sinho Chewi, Constantinos Daskalakis, and Alexander Rakhlin. High-accuracy Sampling for Diffusion Models and Log-concave Distributions. *arXiv preprint arXiv:2602.01338v2*, 2026.

J. Andrés Christen and Colin Fox. Markov Chain Monte Carlo Using an Approximation. *Journal of Computational and Graphical Statistics*, 14(4):795–810, 2005. doi: 10.1198/106186005X76983.

Rob Cornish, Paul Vanetti, Alexandre Bouchard-Côté, George Deligiannidis, and Arnaud Doucet. Scalable Metropolis–Hastings for Exact Bayesian Inference with Large Datasets. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 1351–1360, PMLR, 2019.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-scale Hierarchical Image Database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.

Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Alex Dytso, H. Vincent Poor, and Shlomo Shamai (Shitz). A General Derivative Identity for the Conditional Mean Estimator in Gaussian Noise and Some Applications. *arXiv preprint arXiv:2104.01883*, 2021.

Bradley Efron. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 7(1):1–26, 1979. doi: 10.1214/aos/1176344552.

Tilmann Gneiting and Adrian E. Raftery. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437.

Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array Programming with NumPy. *Nature*, 585:357–362, 2020. doi: 10.1038/s41586-020-2649-2.

W. K. Hastings. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*, 57(1):97–109, 1970. doi: 10.1093/biomet/57.1.97.

Sture Holm. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 6(2):65–70, 1979.

John D. Hunter. Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.

Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising Diffusion Restoration Models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 23593–23606, 2022.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015.

Günter Last and Mathew Penrose. *Lectures on the Poisson Process*. Cambridge University Press, 2017. doi: 10.1017/9781316104477.

Qiang Liu and Dilin Wang. Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm. In *Advances in Neural Information Processing Systems*, volume 29, 2016.

Andreas Makris, Paul Fearnhead, and Chris Nemeth. Tempered Guided Diffusion. *arXiv preprint arXiv:2605.03712*, 2026.

Wes McKinney. Data Structures for Statistical Computing in Python. In *Proceedings of the 9th Python in Science Conference*, 2010. doi: 10.25080/Majora-92bf1922-00a.

K. L. Mengersen and R. L. Tweedie. Rates of Convergence of the Hastings and Metropolis Algorithms. *The Annals of Statistics*, 24(1):101–121, 1996. doi: 10.1214/aos/1033066201.

Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953. doi: 10.1063/1.1699114.

Radford M. Neal. MCMC Using Hamiltonian Dynamics. In *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, 2011. doi: 10.1201/b10905-6.

Gareth O. Roberts and Richard L. Tweedie. Exponential Convergence of Langevin Distributions and Their Discrete Approximations. *Bernoulli*, 2(4):341–363, 1996.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.

Abduragim Shtanchaev, Albina Ilina, Yazid Janati, Arip Asadulaev, Martin Takac, and Eric Moulines. Sparse Scheduled Diffusion Guidance for Inverse Problems. *arXiv preprint arXiv:2603.07860*, 2026.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*, 2021.

Pascal Vincent. A Connection between Score Matching and Denoising Autoencoders. *Neural Computation*, 23(7):1661–1674, 2011. doi: 10.1162/NECO_a_00142.

Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving Diffusion Inverse Problem Solving with Decoupled Noise Annealing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20895–20905, 2025.

POSTERIOR ENTRANCES SUPPLEMENTARY MATERIAL

Reuse the Information, Preserve the Sampling Law

S1 A COMPLETE CONDITIONAL-PRESERVATION ARGUMENT

This proof specializes the first-order rejection mechanism of Chen et al. (2026) to the history-conditioned reference in the main text. Condition on the entire history before a block visit, including its current outside state u, retained anchor r, reference $\widehat{D}$, and allowance E. These quantities are fixed during the rejection loop. The live field is $D_u = \nabla S_u$, with $0 \preceq \nabla^2 S_u \preceq aI$. For $\delta = z - r$ and $s = \|\delta\|$,

$$\begin{aligned} 0 \le \delta^\top [D_u(r+t\delta) - D_u(r)] &\le a\,t\,s^2, \\ |\delta^\top [D_u(r) - \widehat{D}]| &\le Es. \end{aligned} \tag{S1}$$

Thus the main-text residual obeys $0 \le R_z(t) \le B_z(t)$. Convexity gives $V_r(\delta) \ge 0$, and the remaining prefilter terms satisfy

$$-\gamma s^2/2 + Es - E^2/(2\gamma) = -(\gamma/2)(s - E/\gamma)^2 \le 0. \tag{S2}$$

The prefilter probability is therefore valid. Conditional on $N \sim \mathrm{Poisson}(\Lambda)$, independent locations with density B/Λ give

$$\mathbb{E}\prod_{j=1}^{N}[1 - R(T_j)/B(T_j)] = \exp\left[-\int_0^1 R(t)\,dt\right]. \tag{S3}$$

Conditional on N, the product expectation is $(1 - \int R/\Lambda)^N$; averaging the Poisson generating function proves (S3), an instance of the Poisson Laplace-functional identity (Last and Penrose, 2017). The $\Lambda = 0$ case has an empty product equal to one. Early factor evaluation computes the same event while allowing the sampler to stop after the first unsuccessful factor.

The first-order line integral is $S_u(z) - S_u(r)$. Combining the prefilter and Poisson terms gives

$$\begin{aligned} \log \alpha(z) = &-V_r(\delta) - (a+\gamma)s^2/2 \\ &+ S_u(z) - S_u(r) - \delta^\top \widehat{D} - \theta. \end{aligned} \tag{S4}$$

The Gaussian density is

$$\begin{aligned} q(r+\delta) = &\det(P)^{1/2}(2\pi)^{-k/2} \\ &\times \exp[-\delta^\top P\delta/2 - \widehat{g}^\top\delta - \widehat{g}^\top P^{-1}\widehat{g}/2]. \end{aligned} \tag{S5}$$

Substitute $P = H - (a+\gamma)I$ and $\widehat{g} = Hr - h_u + \nabla V(r) - \widehat{D}$. Multiplication yields

$$\begin{aligned} q(z)\alpha(z) &= c\exp[-F_u(z)], \\ c &= \det(P)^{1/2}(2\pi)^{-k/2}\exp[F_u(r) - \theta - \widehat{g}^\top P^{-1}\widehat{g}/2]. \end{aligned} \tag{S6}$$

The coefficient is positive and independent of z, proving Theorem 1 conditional on the entire admissible history. Averaging over that history preserves the conclusion.

FINITE-HORIZON STATE LAW

Fix a common initial law and random-scan distribution independent of the reference history. Both ideal algorithms use the same current block conditional at every selected visit. Induction over the finite update sequence gives the same endpoint-state law. Reference selection changes the work used to realize the transition.

S2 ACCEPTANCE, REFRESHES, AND A SUFFICIENT NET BENEFIT

POSITIVE ACCEPTANCE FROM OPERATIONAL GEOMETRY

Let $C = H + bI$ and $g = \nabla F_u(r)$. The Hessian upper bound gives a quadratic upper approximation. Gaussian integration implies

$$e^{F_u(r)} \int e^{-F_u(z)}\, dz \geq (2\pi)^{k/2} \det(C)^{-1/2} \exp(g^\top C^{-1} g/2). \tag{S7}$$

Using (S6), proposal acceptance is bounded below by

$$\begin{aligned} A_q \geq\ & \sqrt{\det(P)/\det(C)} \\ & \times \exp[-\theta - \widehat{g}^\top P^{-1} \widehat{g}/2 + g^\top C^{-1} g/2]. \end{aligned} \tag{S8}$$

The last exponential factor involving g is at least one. Dropping it provides a directly computable lower bound from the proposal and the certified residual. The proposal displacement satisfies $\mathbb{E}_q s^2 = v_q = \operatorname{tr}(P^{-1}) + \left\|P^{-1}\widehat{g}\right\|^2$ and $\mathbb{E}_q s \leq \sqrt{v_q}$.

CHARGE THE COMPLETE CONDITIONAL DRAW

The expected number of candidate Poisson locations in a proposal is at most $av_q/2 + 2E\sqrt{v_q}$. Prefiltering and sequential factor evaluation preserve this upper bound for actual live queries. Repeating independent proposals gives a renewal equation: expected work per accepted draw equals expected work in one trial divided by its acceptance probability. Hence

$$\mathbb{E}Q_{\text{live}} \leq (av_q/2 + 2E\sqrt{v_q})/A_q. \tag{S9}$$

This argument allows trial cost and trial success to be dependent. An anchor refresh is separately charged before the loop.

PATHWISE REFRESH BUDGET

For a fixed retained anchor, suppose $\|D_u(r) - D_{u_*}(r)\| \leq L_D \|u - u_*\|$ and $\|h_u - h_{u_*}\| \leq L_h \|u - u_*\|$. A fresh residual is at most G_f. While the boundary displacement is below $\Delta_* = \min\{E_*/L_D, (G_* - G_f)/L_h\}$, both admission tests hold.

Each later refresh therefore consumes at least Δ_* of cumulative context variation. The refresh intervals are disjoint, proving $R_N \leq \min\{N, 1 + V_N/\Delta_*\}$. Zero Lipschitz constants contribute infinite corresponding displacement allowances.

PROOF OF PROPOSITION 2

Use $Q_N^F = A_N + C_N^F$ and $Q_N^P = R_N + C_N^P$. Under the proposition's expectation bounds,

$$\begin{aligned} \mathbb{E}(Q_N^F - Q_N^P) & \\ &= \mathbb{E}A_N - \mathbb{E}R_N - \mathbb{E}(C_N^P - C_N^F) \\ &\geq fN - 1 - \nu N/\Delta_* - \eta N = Ng - 1. \end{aligned} \tag{S10}$$

If $g > 0$ and $N > 1/g$, this is positive. The statement is a sufficient expected-query comparison. The four-workload arithmetic profile in the main text is a separately measured decomposition. It supplies the complete observed query total under its recorded run distribution.

S3 TRAINING, CONDITIONAL GEOMETRY, AND EVALUATION UNITS

TRAINING SPECIFICATION

The source loads the 1,797-image 8×8 digit collection (Alpaydin and Kaynak, 1998). Intensities are divided by sixteen. A NumPy generator with seed 910001 permutes the images into 1,200 training, 300 validation, and 297 test indices. Per-coordinate training means define the center, and the model

input is twice the centered normalized image. The loss is the average squared error of $D_\vartheta(x+\varepsilon)$ against x, with $\varepsilon \sim \mathcal{N}(0, I)$, using a denoising objective (Vincent, 2011).

The positive coefficients use $c_h = \mathrm{softmax}(\mathrm{mass})_h \cdot [0.45\,\mathrm{sigmoid}(\mathrm{amplitude})]/(|w_{0h}| + |w_{1h}|)^2$. Adam (Kingma and Ba, 2015) uses learning rate 0.01, batch size eight, and 300 updates; gradients are clipped at norm two. Parameters are projected to $w_{0h} \in [0.15, 2]$, $w_{1h} \in [-0.08, 0.08]$, $b_h \in [-2, 2]$, and $\mathrm{mass}_h \in [-10, 10]$. The final checkpoint supplies the inference model. Validation and test corruptions use their fixed generator.

Concrete conditional

For a block $z = (x_s, \ldots, x_{s+3})$ with neighboring context $u = (x_{s-1}, x_{s+4})$, the first and last coordinates of h_u receive J times their respective neighbors, with base $h = y/4 + o$. Thus $L_h = J = 1/8$. The anchor $r_0(u) = H^{-1}(h_u - \frac{1}{2}\mathbf{1})$ cancels the known quadratic-linear gradient and the logistic midpoint. If each component of D is bounded by M, the fresh residual is at most $\sqrt{k}(M + 1/2)$. At $k = 4$ this gives $G_f = 2(M + 1/2)$, matching the main-text geometry table.

Inference and timing specification

The recorded production environment is Linux x86-64, Python 3.13.5, NumPy 2.3.5 (Harris et al., 2020), and single-threaded BLAS/OpenMP settings. Model parameters are converted to double precision for NumPy inference. The source timer is `time.perf_counter`. Pipeline time covers geometry, fixed-reference preprocessing, initialization, scan construction, and generation. Offline training, calibration, checkpoint loading, numerical reference construction, checks, and file writing have separate records.

Production seeds are 930001, 930013, 930031, and 930059. Each yields 128 endpoint trajectories per method–target cell. Saved horizons are 0.25, 0.5, 0.75, 1, 1.5, 2, 3, 4, 6, 8, 12, and 16 effective sweeps. One sweep contains sixteen random-scan block selections. Dedicated static-reference calibration seeds are 920003 and 920029, using eight outputs at two sweeps and slack candidates 0.25, 0.5, and 0.75. Selected static slacks are 0.5, 0.25, 0.5, and 0.5 in the order of main-text Table 3.

Table S1: Source objects, counts, and statistical units.

Source object	**Count**	**Index / statistical unit**
Frozen trained models	2	Checkpoint identity
Production workloads	128	4 targets × 8 methods × 4 seeds
Outputs per workload	128	Endpoint trajectory
Quality aggregate points	384	Target × method × horizon
Workload checkpoints	1,536	Target × method × seed × horizon
Paired query records	49,152	Target × comparison × seed × endpoint × horizon

The checkpoint hierarchy retains shared trajectories. Paired ratios use the workload as unit; observable summaries use whole endpoints; bootstrapping (Efron, 1979) resamples each endpoint's complete feature vector. The saved 97.5th-percentile statistic is a pointwise diagnostic. Reported site-mean refinement changes under the independent quadrature rules are at most 1.96×10^{-8}.

Baseline families and source-specific variants

The companion diagnostic archive covers independence Metropolis–Hastings (IMH) (Mengersen and Tweedie, 1996), random-walk Metropolis (RWM) (Metropolis et al., 1953; Hastings, 1970), Metropolis-adjusted Langevin (MALA) (Roberts and Tweedie, 1996), and Hamiltonian Monte Carlo (HMC) (Neal, 2011). Persistence, preconditioning, the four-leapfrog HMC schedule, and RWM-E's additional potential access follow the supplied experimental protocol. RWM-E denotes the source-defined energy-access configuration.

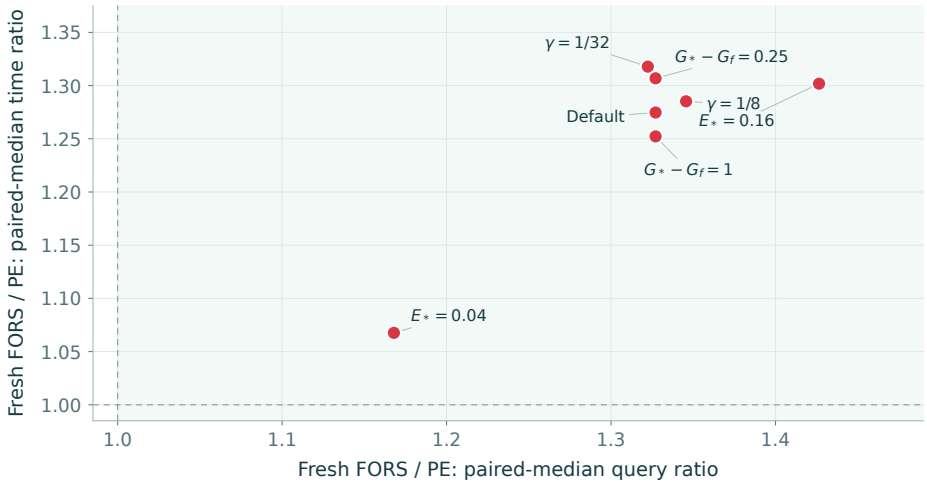

Figure S1: **Joint query and time sensitivity.** All seven recorded settings from the S4 table. Each point uses four paired fresh-FORS/PE workload ratios at eight sweeps and sixteen outputs per workload. Ratios above one favor PE. These are seven configurations, with no added or interpolated settings.

S4 Complete reference-policy sensitivity

The seven settings below use wide model / image 1020, four paired seeds, sixteen outputs per workload, and eight sweeps. Each query or time ratio is the median of four paired fresh-FORS/PE ratios. All total queries include live correction.

Table S2: Complete reference-policy sensitivity.

Policy	**Queries / endpoint**	**Query ratio**	**Time ratio**	**Refresh fraction**
Default	113.781	1.327	1.275	0.284
$E_* = 0.04$	129.281	1.168	1.068	0.474
$E_* = 0.16$	106.844	1.426	1.302	0.134
$\gamma = 1/32$	114.188	1.322	1.318	0.290
$\gamma = 1/8$	112.688	1.346	1.285	0.278
$G_* - G_f = 0.25$	113.781	1.327	1.307	0.284
$G_* - G_f = 1$	113.781	1.327	1.252	0.284

At the representative sixteen-sweep target, the fresh-FORS and PE reference refresh fractions are 0.6739 and 0.2407. The complete query ledger gives a net reduction of 80.34 queried states per endpoint. The companion Markdown preserves the component counts and proposal-acceptance measurements.

The reference-policy summary is checked in two ways: source ratios are reaggregated from the per-workload checkpoints, and the earlier positive-package summaries are compared under their stated rounding. The profile's total reduction is separately calculated from its arithmetic-mean counts. The resulting largest numerical difference in the reaggregated ratio cells is 4.45×10^{-16}.

S5 SITEWISE ESTIMATION PROTOCOL

The trained evaluation compares each site's first and second moments with independent transfer-quadrature references. Standardization and squaring precede spatial averaging. Each quality checkpoint combines 512 endpoints from four equal-sized workloads; its elapsed time sums those four workloads. Whole-endpoint resampling preserves each endpoint's complete feature vector.

The saved horizons are 0.25, 0.5, 0.75, 1, 1.5, 2, 3, 4, 6, 8, 12, and 16 sweeps. The 0.10 level is applied to the saved pointwise upper-percentile diagnostic. The companion `COMPARISON_RECORD.md` contains the full trajectories for every trained target and reference policy.

S6 E32 IMPLEMENTATION AND PHOTOGRAPHIC PROVENANCE

Following the entrance-to-continuation organization in the main text, E32 uses observation-informed native states as entrances. Its recorded configuration comprises two ordinary PC prefixes (Song et al., 2021) and two DDRM prefixes (Kawar et al., 2022), 224 evaluated states per root, followed by the common 32-step suffix. The pretrained diffusion checkpoint follows Dhariwal and Nichol (2021). The handoff aligns the noise parameterization; native states continue with independent branch randomness. The geometric suffix scale grid contains 33 values from the selected entry scale to 0.01. Writing d_θ for the denoising mean, the update is

$$\begin{aligned} w_k &= (s_{k+1}/s_k)^2, \\ X_{k+1} &= w_k X_k + (1-w_k) d_\theta(X_k, s_k) \\ &\quad + s_{k+1}\sqrt{1-w_k}\,\xi_k, \qquad \xi_k \sim \mathcal{N}(0, I). \end{aligned} \tag{S11}$$

The output transformation is $X_{32}/\sqrt{1+0.01^2}$. The observed information enters the prefixes. The four-output pool uses fixed two-PC/two-DDRM source quotas. Two separately initialized pools define the cross-pool energy estimator in the main text, based on energy scoring (Gneiting and Raftery, 2007); image identity is the paired unit in the recorded comparison.

The displayed scene is `ILSVRC2012_val_00049108`, from ImageNet/ILSVRC (Deng et al., 2009; Russakovsky et al., 2015). Reference and output PNGs are the source files in v45's selected-pixel archive. Every panel uses cold pool 0, draw 0. Detail rectangles are $x = 16 \ldots 144$ and $y = 48 \ldots 112$ in the 256×256 image, giving identical 128×64 crops across all methods. Ten distinct files supply the two rows because they share the reference photograph. Each file matches its original SHA-256. The photographs are displayed using their archived pixel values.

The main-text photographic table gives five selected released comparison cells. The original test uses the source's fifteen-contrast Holm adjustment (Holm, 1979), and all five selected rows carry $p = 0.00015$ after that adjustment. Each row also satisfies its source pooled-energy and cold-pool-discrepancy criterion. Percent reductions are calculated from the unrounded released means, then displayed to two decimals.

Table S3: **Released photographic score differences.** The difference is E32 minus comparator, so negative values favor E32. Confidence intervals and Holm-adjusted p values reproduce the supplied comparison records. Each row uses the same 100 paired image identities.

Observation	**Comparator**	**Difference**	**95% interval**	**Holm p**
70% missing	PC	-0.021372	[-0.023249, -0.019573]	0.00015
70% missing	Spin	-0.021057	[-0.022508, -0.019634]	0.00015
4×4 averaging	DAPS	-0.266727	[-0.283442, -0.250195]	0.00015
4×4 averaging	TGD	-0.083777	[-0.088850, -0.078931]	0.00015
4×4 averaging	Spin	-0.016588	[-0.018086, -0.015084]	0.00015

The photographic comparator sources are PC (Song et al., 2021), DAPS (Zhang et al., 2025), TGD (Makris et al., 2026), and Spin (Shtanchaev et al., 2026); the table reports the archived configurations.

S7 SOURCE MAPPING AND REPRODUCIBLE CHECKS

Table S4: Source mapping and material used.

Input	Material used
SiteResolved package	Recorded workloads, geometry, counters, summaries, source implementation
Positive-results package	Arithmetic-mean resource profile and source-defined aggregation
Broader package	General certificate notation and conditional break-even comparison
v45 archive	Approved comparison panel, selected scores, crop definitions, genuine pixels
DenseScatter figure package	Saved counter pairs, site errors, all-method quality coordinates, sensitivity figures

The tables and retained sensitivity plot identify their source and statistical unit. The reviewed source records reaggregation of sixteen primary ratio cells, checks of 49,152 paired-record keys, ten photographic file-hash checks, and a conditional-cancellation check on 128 fixed-input numerical cases. Its largest reported algebraic residual is 2.00×10^{-15}. Mathematical identities are proved in S1–S2; the numerical records describe checks on their stated finite inputs.

The supplied diagnostic archive records six data-file hashes, 24,576 unique static-reference/PE counter-pair keys, and 24,576 unique site-estimate keys. Its eight component-comparison rows and 36 overlapping quality rows agree with the earlier supplied CSVs. The archive also records reconstruction of 96 protocol-first E_{site} values and plotted times from the saved workload summaries. The companion Markdown retains the diagnostic figures and their saved bootstrap percentiles.

The supplied counter file contains 24,576 static-reference/PE pairs. Each counter dot represents one endpoint's cumulative first-order query count. Workload totals charge the shared static setup query once. The quality CSV reports medians across workload-level mean costs; the representative profile uses arithmetic means, and the main comparison table uses medians of paired ratios. These definitions preserve the source aggregation units.

The editable sources compile the combined paper and both standalone documents. The companion data retain the original figures, photographs, coordinate CSVs, and source metadata. The accompanying Markdown collects the full comparison record, diagnostic plots, and editorial transfer log. Author–year citations resolve to the shared bibliography; the standalone supplement includes its own cited-reference list. The conceptual overview and archived experimental photographs retain their respective source identities. The archived rendering scripts use NumPy (Harris et al., 2020), pandas (McKinney, 2010), and Matplotlib (Hunter, 2007).

S8 RECORD ORGANIZATION AND STATISTICAL UNITS

The paired-record hierarchy indexes target, comparison, seed, endpoint, and saved horizon. Each nested checkpoint retains its trajectory identity. Workload-level comparisons use four paired workloads per target, and each workload contains 128 endpoints.

The site-estimate hierarchy indexes method, workload, horizon, and coordinate. A saved estimate summarizes one 128-endpoint workload. The companion Markdown supplies the complete counter-pair and signed-error figures alongside their coordinate definitions.

S9 COMPANION COMPARISON RECORD

The companion `COMPARISON_RECORD.md` contains the complete paired timing ranges, all target-specific first crossings, full diagnostic plots, and the source comparison record. It preserves the target, seed, horizon, and statistical unit attached to each result. Main-text summaries retain their original paired-median definitions, and the photographic table identifies its five selected comparison cells.