EXPANDING GENERABLE SETS WITH ONE-STEP GENERATORS

Anonymous authors
Paper under double-blind review

ABSTRACT

Generative scientific discovery aims to produce valid designs beyond the support of a static pretrained model. The goal is to expand its generable set over the valid design space using only black-box verifier feedback. Existing methods steer multi-step samplers toward unexplored regions, but their sampler implicitly determines how far queries travel and how strongly they remain near the initial model scope. We introduce EXPANSE, which grows the generable set with a 1-step generator by searching directly in design space with explicit control over exploration and travel. Rejected designs push the generator away from invalid regions, while inference settings trade coverage against validity. We also derive a reachability bound giving a necessary condition for a disconnected valid region to be found, in terms of generator support and query travel. With the same verifier budget and one network function evaluation (NFE) per sample EXPANSE covers 90.7% of the connected valid region against 78.8% for the best multi-step baseline at the same gap, with higher coverage and validity overall.

1 INTRODUCTION

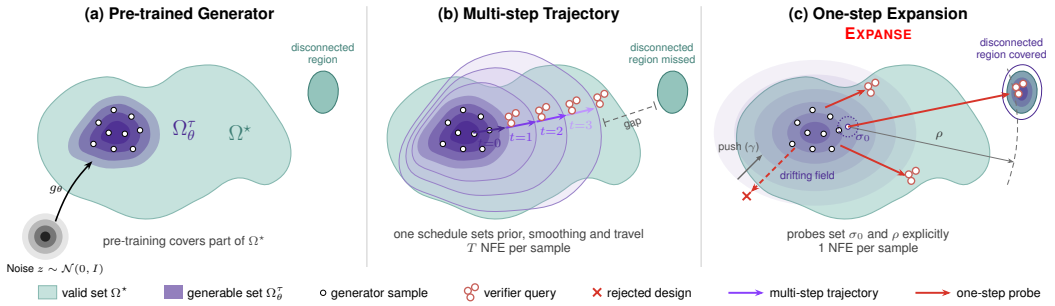

Figure 1: **EXPANSE overview.** (left) Pre-training covers only part of the valid set $\Omega^\star$ (green) with its generable set (purple). (middle) ActFlow guided multi-step sampler, where one trajectory sets prior, smoothing and travel, but disconnected region stays out of reach. (right) EXPANSE probes directly in design space with an explicit travel budget, and samples in one step.

Flow and diffusion models (Ho et al., 2020; Song et al., 2020; Lipman et al., 2022) are pre-trained to match a data distribution, so valid designs absent from that distribution receive negligible probability by construction (Fig. 1, left). This limits scientific biodesign, where the aim is molecules, peptides, and proteins new to nature (Hoogeboom et al., 2022; Watson et al., 2023; Tang et al., 2025). ActFlow (De Santi et al., 2026a) reframes the goal: a generative model approximates the valid design space $\Omega^\star$ of a black-box verifier, measured through its generable set Ω_θ^τ, and learning is expansion of that set rather than distribution matching. Since the verifier labels only proposed designs, expansion is recursive, and the model must issue queries that reveal valid regions it does not represent, then absorb the outcomes. Every round of such a loop does two separable things. It *covers*, placing queries where none have landed, and it *consolidates*, converting verified designs into terminal density so the generable set grows. Covering is geometric and uses no labels; consolidation is the only place labels enter. A multi-step sampler fuses the two, because its denoising trajectory is already a coarse-to-fine

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.

path that acquisition guidance can ride on while performing the transport that produces the sample. One schedule then sets prior strength, smoothing scale, and how far a query travels before evaluation, none separately tunable. The fusion is measurable: on QM9 Ramakrishnan et al. (2014), under a fixed budget of 16k network evaluations and *no verifier queries*, coarsening a frozen pre-trained model from 100 to 4 steps per sample drops validity from 97.5% to 27.9% while raising distinct packing centers among valid molecules from 38 to 174. More draws explain this (400 against 4000, of which 390 and 1116 are valid), but centers grow $4.6\times$ where valid pool grows $2.9\times$, and greedy packing counts grow sublinearly in pool size, so the gap understates how much coarsening also spreads samples. Solver resolution is hence a tunable knob, which no pre-training objective exposes.

What should each mechanism be, once separated? EXPANSE answers this on a one-step generator, which has no trajectory for the two to hide inside, and then transfers the same loop to a few-step molecular sampler where the decomposition still applies but the one-step property does not. Covering becomes a probe issued directly into design space, seeded from a generator sample perturbed at scale σ_0 and taking normalized steps within a travel budget ρ. Its direction combines a smoothed generator prior with the gradient of an occupancy field, a posterior variance on a fixed feature map that depletes wherever a query lands, pushing later probes into regions their predecessors left untouched; the smoothing scale anneals along the probe, which is the denoising schedule reparameterized. Nothing is transported, so ρ stays free. Consolidation absorbs outcomes through a drifting update matching accepted designs against a push distribution that mixes the generator with rejections at rate γ. Exact matching gives an affine target that extrapolates away from rejected designs rather than down-weighting them, and sampling γ during training while conditioning on it lets one model expose the whole coverage-validity frontier at inference. Each control is then measurable alone, and reachability reduces to probe geometry, since a probe entering a region must have seeded within ρ of it. Algorithm 1 states the loop and Fig. 1 contrasts it with a trajectory-guided sampler.

Contributions. **(1)** We decompose generable set expansion into covering and consolidation, and show the signal driving covering needs no verifier labels: acquisition uses query *locations* only, and validity enters solely through the generator update. **(2)** We make covering explicit by probing design space directly, so seed scale σ_0, travel budget ρ, smoothing scale h, and prior weight β become separately settable and separately measurable rather than consequences of a sampling schedule. **(3)** We prove a component-reachability bound and a disconnected-gap barrier from probe geometry alone, and show its predicate $\widehat{N}_C \geq 1$ is necessary in practice: every run that missed the disconnected component had $\widehat{N}_C < 1$. **(4)** We derive a consolidation rule in which rejections enter the push distribution as a mixture, so one conditioned generator traces the coverage-validity frontier at inference, and we report where this control transfers and where it does not.

2 RELATED WORK

Pre-trained diffusion and flow models (Ho et al., 2020; Song et al., 2020; Lipman et al., 2022) are steered toward a specified objective by post-training (Black et al., 2024; Liu et al., 2026) or inference-time guidance (Bansal et al., 2024; Uehara et al., 2025; Jensen et al., 2026), but these reweight a reward already defined on the current support and yield invalid samples when pushed past it (Gutjahr et al., 2025). Retraining on generated samples contracts support (Shumailov et al., 2024; Alemohammad et al., 2024); filtering first (Dong et al., 2023; Gulcehre et al., 2023), using synthetic data as negative guidance, or self-play (Yuan et al., 2024) target in-distribution quality. ActFlow (De Santi et al., 2026a), our closest comparison, instead makes coverage the objective, expanding Ω_θ^τ toward $\Omega^\star$ under black-box verification, where related approaches assume a differentiable verifier (De Santi et al., 2026c) or none at all (De Santi et al., 2026b). EXPANSE retains this objective, but replaces the trajectory-guided sampler, and, rather than scaling a gradient on rejected designs (Alberti et al., 2025), mixes rejections into the push distribution of a one-step update. Unlike one-step models obtained by distillation (Salimans & Ho, 2022; Song et al., 2023), we build on Drifting (Deng et al., 2026), which evolves the pushforward through an anti-symmetric field. Filling a space rather than optimizing within it is the premise of novelty search (Lehman & Stanley, 2011) and quality diversity (Mouret & Clune, 2015; Cully et al., 2015; Pugh et al., 2016), and EXPANSE uses these measures, with the generator as a variation operator, occupancy as novelty, and $r_{\min}$ as archive rule; GFlowNets (Bengio et al., 2021; 2023; Jain et al., 2022) instead amortize sampling proportional to a scalar reward over sequential construction. Our label-free signal echoes count-based exploration bonuses (Bellemare et al., 2016; Burda et al., 2018), and the posterior-variance acquisition of kernelized

bandits (Srinivas et al., 2009; Chowdhury & Gopalan, 2017; Pásztor et al., 2024), made persistent across rounds by random Fourier features (Rahimi & Recht, 2007) on a fixed map, as in neural-linear bandits (Riquelme et al., 2018). Safe Gaussian-process exploration (Sui et al., 2015; Turchetta et al., 2016; Berkenkamp et al., 2017; Sui et al., 2018) certifies regions reachable from past observations; instead of transfering ActFlow's uncertainty-oracle, we isolate geometric reach of a one-step probe.

3 METHOD

3.1 SETTING

Let $\mathcal{X} \subseteq \mathbb{R}^d$ be a closed convex ambient design space with projection $\Pi_{\mathcal{X}}$. A black-box verifier $\tilde{v} : \mathcal{X} \to \{0, 1\}$ defines the valid set $\Omega^\star = \{x : \tilde{v}(x) = 1\}$, not necessarily convex or connected. For a density p, the τ-level generable set is $\Omega^\tau(p)$=$\{x : p(x) \geq \tau\}$, and Ω^τ_θ : =$\Omega^\tau(q_\theta)$. Starting from valid data $\mathcal{D}_{\text{data}}$, we want $|\Omega^\tau(p) \cap \Omega^\star|$ large and validity $\mathbb{E}_p[\tilde{v}]$ high. We run T rounds of B queries each. Accepted and rejected designs go to buffers $\mathcal{D}^\pm_t$ with empirical laws $p^\pm_t$.

The generator is x=$f_\theta(z, \alpha)$ with $z \sim \mathcal{N}(0, I)$ and $\alpha \geq 1$, and its law $q_{\theta,\alpha}$ is assumed to have a density. The condition scales a feature-wise affine modulation by $\log \alpha$, so α=1 is exactly the unconditioned model. Queries always use α=1, and we write q_θ : =$q_{\theta,1}$ and q^h_θ : =$q_\theta * \mathcal{N}(0, h^2 I)$. The positive buffer is initialized from $\mathcal{D}_{\text{data}}$ thinned at $r_{\min}$, which costs no verifier queries and leaves the occupancy field untouched; the released ActFlow loop starts empty, and we bound this asymmetry in Appendix C.2. ActFlow steers an iterative sampling path, which a one-step generator does not have. EXPANSE therefore builds its queries directly in design space, which exposes four controls that a denoising schedule fixes jointly: seed scale σ_0, travel budget ρ, smoothing scale h, and prior strength β.

3.2 QUERYING WITHOUT A TRAJECTORY

ActFlow picks queries by maximizing $\mathbb{E}_\pi[\sigma_t] - \beta\,\mathrm{KL}(\pi \,\|\, q_{\theta_t})$ over distributions π, where σ_t is its posterior standard deviation. With the occupancy signal u_t of Sec. 3.3 in place of σ_t, the maximizer for $\beta > 0$ is the tilted law

$$\pi^\star_t(x) \propto q_{\theta_t}(x) \exp\big(u_t(x)/\beta\big). \tag{1}$$

We use this tilt to motivate normalized ascent at successive smoothing scales h, not to sample $\pi^\star_t$. For $X \sim q_{\theta_t}$, $Z \sim \mathcal{N}(0, I)$, and the flow interpolation X_s=$sX + (1 - s)Z$,

$$\mathrm{Law}(X_s/s)$=$q^h_{\theta_t}, \qquad h$=$(1 - s)/s, \tag{2}$$

so h plays the role of ActFlow's noising level. The Gaussian mean shift at scale h satisfies $m_{t,h}(x)$=$h^2 \nabla \log q^h_{\theta_t}(x)$ (Appendix A.2). The probe direction is

$$d_{t,h}$=$\beta\, m_{t,h} + h^2 \nabla u_{t,h}. \tag{3}$$

For $\beta > 0$ it is βh^2 times the score of $q^h_{\theta_t} e^{u_{t,h}/\beta}$, and for β=0 it is pure occupancy ascent. Because Eq. equation 5 normalizes the direction, the common factor βh^2 is immaterial and only the ratio of the two terms acts, so h sets *where* the probe points and ρ=$K\delta$ sets how far it goes. This is the separation a denoising schedule does not permit.

Each round draws B generator samples x^{gen}_b and M anchors $\mathcal{A}_t$=$\{a_j\}$ from $f_{\theta_t}(\cdot, 1)$. Each probe starts from a perturbed sample and takes K normalized steps of length δ:

$$x_{t,0}$=$\Pi_{\mathcal{X}}\big(x^{\text{gen}}_b + \sigma_0 \xi\big), \quad \xi \sim \mathcal{N}(0, I), \tag{4}$$

$$x_{t,k+1}$=$\Pi_{\mathcal{X}}\big(x_{t,k} + \delta\, \hat{d}_{t,h_k}/\|\hat{d}_{t,h_k}\|\big). \tag{5}$$

The probe stays put when $\hat{d}$=0, and h_k decays geometrically from $h_{\max}$ to $h_{\min}$. In $\hat{d}$, the mean shift is estimated from anchors as $\hat{m}(x)$=$\sum_j w_j a_j - x$, with $w_j \propto e^{-\|x-a_j\|^2/2h^2}$. The occupancy term uses the closest level of a fixed ladder $\mathcal{H}$ (Sec. 3.3). Probes within a round do not interact. Diversity within a round comes from the seeds, and diversity across rounds comes from occupancy depletion.

Lemma 3.1 (Probe reach)**.** *Every probe satisfies* $\|x_{t,K} - x_{t,0}\| \leq \rho$: =$K\delta$ *(Appendix A.3).*

Together with equation 4, σ_0 controls the random seed perturbation, while ρ bounds subsequent directed travel.

3.3 WHAT TO ACQUIRE

We record where the verifier has already been queried with a posterior variance. At each level $h \in \mathcal{H}$, features ϕ_h approximate an RBF kernel with lengthscale $\ell_h = \sqrt{\ell^2 + 2h^2}$, which is the kernel smoothed at scale h (Appendix A.2). Let $\Phi_{t,h}$ stack the features of all points queried before round t. With noise σ^2,

$$v_{t,h}(x) = \sigma^2 \, \phi_h(x)^\top \big(\sigma^2 I + \Phi_{t,h}^\top \Phi_{t,h}\big)^{-1} \phi_h(x). \tag{6}$$

It is high far from past queries and low near them. It depends on where we queried, not on the labels.

Lemma 3.2 (Monotone occupancy depletion)**.** *Adding a query z never increases occupancy: $v_{t,h}^{(z)}(x) \le v_{t,h}(x)$ for all x and h (Appendix A.4).*

Querying a region therefore uses up its occupancy, whatever the label. We divide by the average over anchors,

$$u_{t,h}(x) = v_{t,h}(x) \,/\, \bar{v}_{t,h}, \qquad \bar{v}_{t,h} = \tfrac{1}{M} \textstyle\sum_j v_{t,h}(a_j). \tag{7}$$

This scaling only sets the balance against β, and has no effect when β=0. Before the first query v is nearly flat, so we set u=0. First-round queries are then the seeds themselves (β=0), or seeds pulled toward modes by mean shift ($\beta > 0$). Validity enters only through the generator update. The posterior mean of the same regression is used only in the diagnostic of Sec. 3.5.

Algorithm 1 EXPANSE.

Require: Generator f_{θ_0}, verifier $\tilde{v}$, rounds T, queries per round B.
1: $\mathcal{D}^+ \leftarrow \mathrm{Thin}_{r_{\min}}(\mathcal{D}_{\mathrm{data}})$, $\mathcal{D}^- \leftarrow \emptyset$.
2: **for** t=$0, \dots, T-1$ **do**
3: **Probe:** seed B points near f_{θ_t} and ascend $\hat{d}$ for K steps (Eq. 5).
4: **Verify:** label the probes with $\tilde{v}$.
5: **Store:** add accepted and rejected probes to $\mathcal{D}^{\pm}$ at spacing $r_{\min}$.
6: **Deplete:** update $v_{t,h}$ with the new query locations.
7: **Train:** take S drifting steps, each at a random α (Eq. 8).
8: **end for**
9: **return** f_{θ_T}.

3.4 UPDATING THE GENERATOR

We use the drifting field $\mathcal{V}_{p,q}$ of Deng et al. (2026), which pulls samples of q toward p and vanishes when q=p. The generator regresses onto its own samples, each moved one step along the field:

$$\mathcal{L}_{\mathrm{drift}}(\theta) = \mathbb{E}_z \big\| f_\theta(z,\alpha) - \mathrm{sg}\big(x + \kappa \, \hat{\mathcal{V}}(x)\big) \big\|^2, \qquad x = f_\theta(z,\alpha). \tag{8}$$

Here κ is the step size and $\hat{\mathcal{V}}$ is a finite-sample, multi-bandwidth estimate of the field (Appendix A.5). Rejections are mixed into the side that the generator is pushed away from:

$$q_\gamma^{\mathrm{push}} = (1-\gamma)\, q_{\theta,\alpha} + \gamma \, p_t^-, \qquad \alpha = 1/(1-\gamma). \tag{9}$$

The generator at condition α is pulled toward p_t^+ and pushed away from q_γ^{push}. Each step draws $\gamma \sim \mathcal{U}[0, \gamma_{\max}]$ and adds rejections to the generated batch in ratio $\gamma : (1-\gamma)$. Before any rejection exists, the push side is the generator alone. One conditional generator is trained across all sampled γ, and sampling and probing use its online weights.

Proposition 3.3 (Mixture target and its smoothed sign)**.** *Fix $\gamma \in [0,1)$ and α=$1/(1-\gamma)$. (i) The only signed measure q with $(1-\gamma)q + \gamma p_t^- = p_t^+$ is*

$$q_\alpha^\star = \alpha \, p_t^+ - (\alpha - 1)\, p_t^-. \tag{10}$$

(ii) If the verifier is deterministic, $\mathcal{D}_t^- \neq \emptyset$ and $\alpha > 1$, then $q_\alpha^\star$ has mass $-(\alpha-1)$ outside $\Omega^\star$, so it is not a distribution. (iii) For any kernel $k \ge 0$, $q_\alpha^\star * k \ge 0$ exactly where $p_t^+ * k \ge \gamma \, p_t^- * k$.

By (ii), exact mixture matching has no probability-distribution solution when $\alpha > 1$. Part (iii) describes the idealized smoothed target. The finite estimator in Eq. 17 instead uses batch-normalized affinities and is not a fixed convolution, so we use (iii) as motivation rather than as a guarantee of the implemented update. Sec. 4 measures the resulting trade-off directly. ActFlow instead scales a negative gradient during training, while here α is an input chosen at inference. Each buffer admits a design only if it is at least $r_{\min}$ from every design already there. This keeps repeatedly rediscovered regions from dominating the update. $\mathcal{D}_0^+$ is $\mathcal{D}_{\text{data}}$ thinned with no query costs.

3.5 Reachability limits

The travel budget provides a necessary condition for reaching a region C. Let $\mathcal{F}_t$ be the history before round t, and $\mathcal{B}(C, \rho)$ the points within ρ of C. Let μ_t^{seed} be the law of a round-t seed given $\mathcal{F}_t$ (Eq. 4). This law already includes the seed perturbation and projection, so ρ in Eq. equation 11 measures travel after seeding.

Theorem 3.4 (Component reachability bound)**.** *Let N_C count the queries that land in C over all rounds. Then*

$$\mathbb{E}[N_C] \le B \sum_{t=0}^{T-1} \mathbb{E}\big[\mu_t^{\text{seed}}\big(\mathcal{B}(C,\rho)\big)\big], \qquad \Pr(N_C \ge 1) \le \min\{1, \mathbb{E}[N_C]\}. \tag{11}$$

A probe can end in C only if its seed starts within ρ of C (Appendix A.6). The bound holds for any acquisition rule, and it is necessary but not sufficient.

Corollary 3.5 (Disconnected-gap barrier)**.** *Let C be a valid component at distance Δ from the rest $\mathcal{M}{=}\Omega^\star \setminus C$ of the valid set. For any $r \ge 0$,*

$$\Pr(x_{t,K} \in C \mid \mathcal{F}_t) \le q_{\theta_t}\big(\operatorname{dist}(x, \mathcal{M}) > r\big) + \exp\Big[-\tfrac{1}{2}\Big(\tfrac{\Delta - r - \rho}{\sigma_0} - \sqrt{d}\Big)_+^2\Big]. \tag{12}$$

When $\sigma_0{=}0$, the second term becomes $\mathbf{1}[r + \rho \ge \Delta]$.

So a gap is crossed only by generator mass already near it, by seed noise, or by directed travel (Appendix A.6). The smoothing scale h changes where a probe goes, not how far it goes. We bound the tail mass of q_{θ_t} instead of assuming a support radius, since a generator's support is typically unbounded. We estimate the bound as $\widehat{N}_C$ from saved checkpoints, holding each checkpoint's seed law until the next one. Distances are measured to a bounding box of C, which enlarges the neighborhood, so each checkpoint term can only overestimate.

ActFlow certifies expansion with an operator $\mathcal{R}_\varepsilon$ that grows a set wherever a Lipschitz validity score stays above a threshold (Appendix A.7). Its guarantee needs a calibrated verifier, a model of the generative update, and an uncertainty-maximizing sampler. EXPANSE assumes none of these, so we use $\mathcal{R}_\varepsilon$ only as a diagnostic, with the posterior mean as the validity score. The Lipschitz constant is estimated on a grid, so the resulting set is optimistic rather than certified.

4 Experiments

We use the two-dimensional setting of De Santi et al. (2026a) with a connected three-region valid set and a disconnected component at distance Δ. The released verifier omits the component so we add it and rerun the released ActFlow code at its shipped guidance weight $w{=}100$ and at the $w{=}13$ its paper reports. EXPANSE keeps the same lengthscale and noise, but differs in its features, probe, replay and generator update, so the comparison is end-to-end rather than a swap of one component.

The grid has cell side $1/16$ and grows with Δ from 112^2 to 176^2 cells, and at every gap $|\Omega^\star|{=}3072$ cells of which 2560 are connected and 512 form the component. Here Ω^τ collects the cells whose density from 200k samples is at least τ, total coverage is $|\Omega^\tau \cap \Omega^\star|/|\Omega^\star|$, and validity is the fraction of samples in $\Omega^\star$. Connected and component coverage are the same ratio restricted to the 2560 connected cells and to the 512 component cells respectively, so total coverage is their $5:1$ weighted mean. ActFlow's level $\tau{=}0.01$ is saturated because every method except ActFlow (SDE) has a seed mean above 93% at every gap. We therefore report $\tau{=}0.05$, about 39 samples per cell, which is where the methods separate under 32k queries and 3 seeds. EXPANSE uses a residual MLP

generator of 663k parameters and runs T=500 rounds of B=64 probes with no unguided samples. Probes use M=512 anchors with σ_0=1.0 and K=40 steps of length δ=0.05, giving ρ=2.0 at β=0 with bandwidths annealed from $h_{\max}$=1.0 to $h_{\min}$=0.1. Each round takes S=100 drifting steps at $\gamma_{\max}$=0.5 with buffers at $r_{\min}$=0.05, and the cap on N_- never constrains so the realized mixture rate equals the sampled γ. The initial buffer holds the 57 of 512 pre-training points that pass separation at seed 0, which cost no queries and deplete no occupancy (Appendices C.1 and C.2).

	Δ=0	Δ=0.5	Δ=1	Δ=2
Pre-trained	2.4 / 0.0	2.4 / 0.0	2.4 / 0.0	2.4 / 0.0
ActFlow (SDE)	25.0 / 45.3	26.2 / 51.4	26.4 / 48.6	26.6 / 47.3
ActFlow (ODE, w=100)	73.9 / 80.1	72.1 / 74.5	62.4 / 69.5	69.8 / 76.0
ActFlow (ODE, w=13)	61.3 / 89.4	67.7 / 78.8	71.2 / 74.6	77.8 / 72.7
EXPANSE (small batch)	67.9 / 92.1	65.9 / 52.9	67.4 / 41.1	65.5 / 4.6
EXPANSE	**92.5 / 95.7**	**90.0 / 85.6**	**90.8 / 90.2**	**89.5 / 83.9**

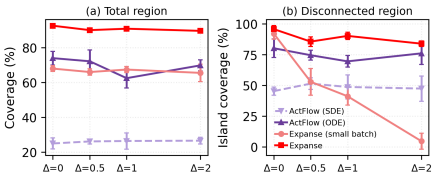

Figure 2: Total / disconnected-region coverage (%) at τ=0.05 with 32k queries and 3 seeds at α=1. Small batch is N=128 in drifting step. Right panel plots both metrics against Δ.

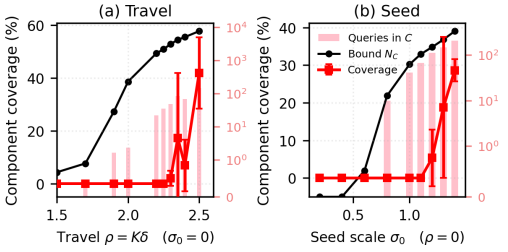

Figure 3: Travel and seed-scale sweeps at Δ=2, other mechanism off. Line on left axis is coverage, over 3 seeds. Bars and black line on right axis are retained component queries and bound $\widehat{N}_C$.

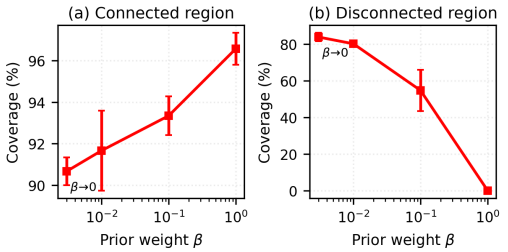

Figure 4: Connected and disconnected-region coverage against β at Δ=2 and τ=0.05 over 3 seeds with reach fixed, where β=0 is drawn at 3×10^{-3} on the log axis.

EXPANSE leads ActFlow variants on total and component coverage at every gap (Fig. 2), gaining 12 to 20 points in total and 6 to 16 on the component against the better variant at each gap, and its lead is flat in Δ where the baselines are not. At Δ=2 its connected coverage is 90.7% against 78.8% at w=13 and 68.5% at w=100, filling holes ActFlow leaves (Fig. 5c,e). Its component spread across seeds is 2.2 against 8.1 and 9.0 for ActFlow (ODE), and our 6.3-point component margin at Δ=0 sits inside its spread of 10.3. Conditioning on α traces a coverage-validity curve from one trained model (Fig. 5i–j), where raising α recovers validity at a predictable cost in coverage. At α=1.25 EXPANSE reaches 80.8 to 83.3% coverage at 97.5 to 99.2% validity, above both ActFlow guidance weights on both metrics at every gap. At similar size it samples 1.12×10^7 designs/sec at 1 NFE, which is 350× ActFlow at 100 NFE (matched batch size).

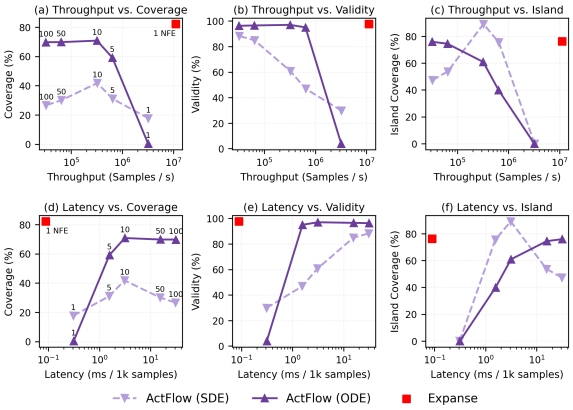

Figure 6: Throughput & latency vs coverage, validity, island at Δ=2 and τ=0.05. ActFlow at 1-100 NFE. EXPANSE at 1 NFE and α=1.25.

Making travel explicit turns reachability into a sweep (Fig. 3). Each knob is swept with the other set to zero, so the thresholds below are those of a single mechanism acting alone and are not comparable to the main configuration, which reaches the component at ρ=2.0 and σ_0=1.0 acting together. Component coverage stays at zero through ρ=2.25 at σ_0=0 and through σ_0=1.1 at ρ=0, and turns nonzero one step later in each case. It reaches $41.9 \pm 13.5\%$ at ρ=2.5 and $28.7 \pm 3.0\%$ at

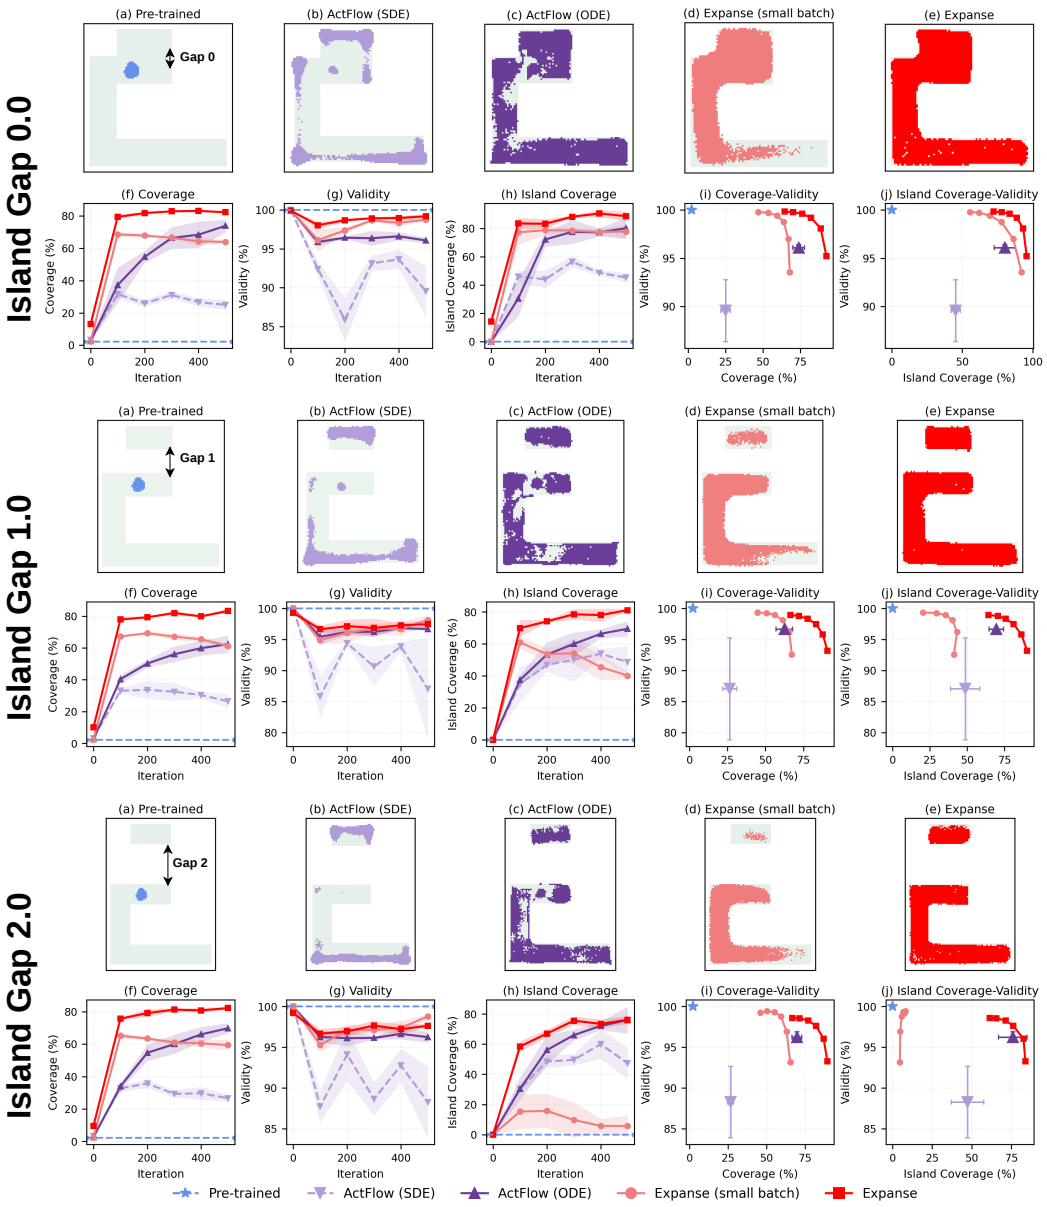

Figure 5: **Expansion at three gaps.** Rows are Δ=0, 1, 2. (a–e) generable set $\Omega^{0.05}$ at the final iteration for seed 0 at α=1 over the valid set in light green. (f–h) Coverage validity and component coverage against iteration over three seeds with EXPANSE at α=1.25 and the pre-trained level dashed. (i–j) Validity against total and component coverage where the EXPANSE curves trace $\alpha \in [1, 2]$ and the baselines are single points at 100 NFE.

σ_0=1.4, and between the thresholds its seed spread exceeds its mean. Retained queries grow with both knobs from none below the threshold to 149 ± 18 and 203 ± 12 respectively, against bounds of 7600 ± 200 and 330 ± 25. Every run with $\widehat{N}_C < 1$ missed the component and every run that found it had $\widehat{N}_C \geq 1.08$, with the bound loose on travel because a seed near C enters it only when the ascent points there. Because β is an explicit input we can sweep it (Fig. 4), and raising it from 0 to 1 lifts connected coverage from 90.7 to 96.6 while component coverage falls to zero. Total coverage is flat up to β=0.01 and then declines because the component is lost faster than the connected region fills. Every setting that spreads queries without reaching the component fills 94 to 97% of the connected region, while the prior term alone (β=1, occupancy off) gives 48.6% since mean shift pulls probes

to modes. All other results use β=0, which is pure normalized occupancy ascent with no generator prior in the probe direction. ActFlow (ODE) is no worse at 10 NFE than at 100 (Fig. 6), and even against that operating point EXPANSE is higher on coverage validity and component coverage at no less than 35× the throughput. The SDE sampler at 10 NFE beats us on the component alone at 89.2%, but buys it with 41.8% total coverage and 60.8% validity. Per round EXPANSE spends 576 generator evaluations and 2560 occupancy-gradient evaluations on acquisition, against ActFlow's 6400 guided velocity evaluations with backward passes; consolidation costs a further S=100 drifting steps at N=512, against ActFlow's guidance-weight training over the same round. As a consistency diagnostic, we iterate $\mathcal{R}_\varepsilon$ from the pre-trained generable set for 40 iterations at ε=0.05 and η=0.5, scoring validity by the finest-level posterior mean. The resulting set holds 2264 to 2346 cells at every $\Delta > 0$, and after 32k queries 94 to 98% of it lies inside the achieved generable set.

Fig. 7 reports connected and island coverage at Δ=2, where removing the probe entirely (σ_0=ρ=0) drops EXPANSE from 90.7 / 83.9 to 52.5 / 0.0 and the prior term alone gives 48.6 / 0.0. Seed noise without ascent (σ_0=1.0, ρ=0) gives 96.1 / 0.0, so random seeding fills the connected region well but never crosses the gap, while directed travel makes the crossing reliable. Fixing h=$h_{\min}$ costs about six points on the component, so annealing helps modestly. Removing separation lowers both metrics and widens the component spread across seeds to 21.9 to 64.3, so replay construction matters about as much as acquisition. Setting γ=0 costs about five points on the component and none on the connected region but collapses the curve to a point, so the frontier comes at no measured cost.

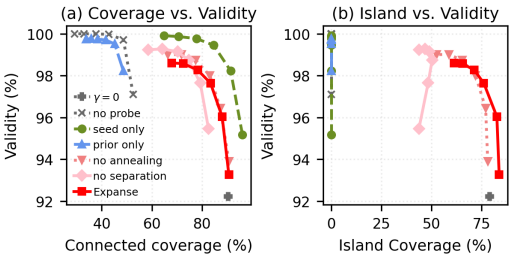

Figure 7: Ablation of region coverage vs validity at Δ=2, τ=0.05 over 3 seeds, curve sweeps $\alpha \in [1, 2]$ with α=1 at lower right.

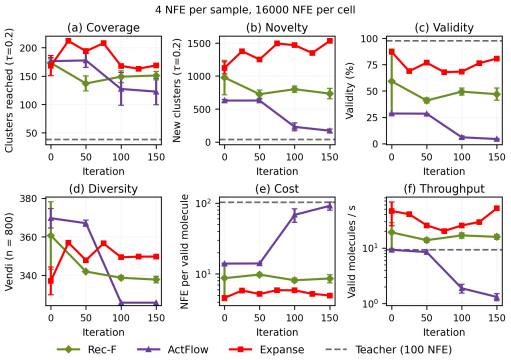

Figure 8: **QM9 over the expansion loop.** Sampling at 4 NFE under 16k NFE per cell at α=1, with panels for coverage, novelty, validity, diversity, cost, and throughput.

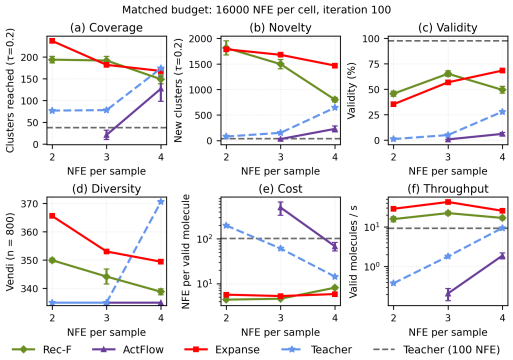

Figure 9: **QM9 at a matched budget.** The same six panels at iteration 100 under 16k NFE per cell at α=1, with the teacher shown as a few-step curve and as the dashed line at its native 100 NFE.

We repeat the comparison on QM9 Ramakrishnan et al. (2014), where the verifier is RDKit validity and the design space is the latent space of a pre-trained flow model. Discovery is a greedy packing of the valid generated molecules at Tanimoto similarity τ=0.2 on Morgan fingerprints, and novelty packs the valid molecules lying outside a 238-center reference packing of a 20k QM9 pool at the same cutoff. Each loop runs 200 rounds of 64 probes for 12.8k verifier queries plus 512 free anchor positives, while scoring separately spends 16k function evaluations per cell. A method sampled at f NFE therefore draws $16000/f$ molecules and no metric is censored, so discovery here is discovery bought at fixed compute where low validity is paid for in draws. The generator is a FiLM-conditioned fine-tune of the teacher velocity field sampled at 4 NFE and trained with the drifting loss plus a flow-matching anchor, so the one-step claim covers the two-dimensional benchmark only. Rec-F retrains the teacher by flow matching on its own filtered accepted buffer with no acquisition and no mixture (Dong et al., 2023; Gulcehre et al., 2023), and ActFlow is the released method run in the

same loop. Appendix D gives the conditional sampler, the latent probe, the bucketed replay, and the exact definition of every metric reported here, which differ from the two-dimensional definitions above. Steps per sample set a coverage-validity trade for every method and the fixed budget makes it visible (Fig. 9). The frozen teacher alone finds 38 centers at 97.5% validity with 100 steps and 174 at 27.9% with 4, which is 4.6× the count at less than a third of the validity. At 4 NFE and iteration 100 EXPANSE finds 168 centers, 1468 new centers and 68.4% validity against 149 ± 11, 800 ± 46 and 49.4 ± 3.4 for Rec-F and 127 ± 29, 229 ± 60 and 6.1 ± 1.4 for ActFlow. At 2 NFE it finds the most centers of any method while matching Rec-F on novelty, and at 3 NFE the two are close on centers while EXPANSE leads on novelty. Its cost per valid molecule stays near Rec-F's at every step count. ActFlow needs an order of magnitude more at 4 steps and produces nothing at 2. At fixed sample size we lead Rec-F on diversity, while ActFlow rarely produces enough valid molecules to be scored.

Over the loop the methods separate on validity rather than on reference coverage (Fig. 8), where EXPANSE holds its center count across 150 iterations while novelty rises by about a third. The added mass is therefore new rather than a redistribution over known chemistry (Fig. 10), where at a matched budget EXPANSE places valid molecules outside the grey reference packing while the baselines stay inside it, and its cost per valid molecule stays flat throughout. ActFlow moves the other way as validity and novelty both fall and cost per valid molecule rises sixfold, and by iteration 100 it produces zero valid molecules from 8000 draws at 2 NFE. Rec-F stays roughly flat, which is what filtered self-training predicts since it cannot place mass where the teacher has none. The α control does not transfer to this setting. Raising it lowers validity and coverage, so we report α=1 on QM9.

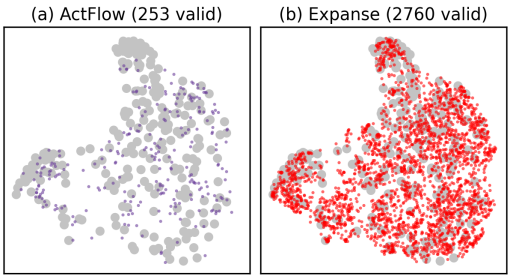

Figure 10: **Valid molecules in fingerprint space.** Joint UMAP on Jaccard distance at iteration 100 and 4 NFE under 16k NFE per cell for seed 0, where grey is the reference packing and counts are valid molecules from the same budget.

5 DISCUSSION

Trajectory methods set prior strength, smoothing and travel implicitly through the sampler. EXPANSE builds queries in design space instead, exposing σ_0, ρ, h and β as separate knobs, and each mechanism then reads off its own sweep. Seed noise spreads queries around the generator and nearly saturates the connected region, but crosses a gap only by chance; directed travel makes the crossing reliable, at a cost of about five points where seeding alone would have sufficed. Making travel explicit also turns reachability into a necessary condition computable before the budget is spent: every run with $\widehat{N}_C < 1$ missed the component and every run that found it had $\widehat{N}_C \geq 1$, though the converse fails because a seed near C still needs the ascent to point there. Acquisition needs no calibrated classifier, since occupancy depletes on query locations alone and validity enters only through the generator update. The prior term hurts here because mean shift pulls probes back toward modes, which suggests ActFlow's KL regularizer mainly keeps guided samples on the data manifold, a role a one-step generator does not need. Replay separation matters as much as acquisition: without it a few dense regions dominate the update, halving component coverage and widening its seed spread. The mixture makes rejections an inference-time input, so one model traces a coverage-validity curve and at α=1.25 beats ActFlow on both metrics at every gap, while γ=0 costs five points on the component and none on the connected region, so the control is close to free. **Limitations.** Choosing σ_0 and ρ needs some prior sense of how far valid regions lie, and our theory gives necessary conditions rather than a coverage guarantee. Consolidation is batch-limited: at N=128 connected coverage falls 13 to 29 points and the component is lost as Δ grows, and the requirement grows with the number of regions the generator must hold at once, uncharacterized here. The mixture control does not transfer, since raising α on QM9 lowers validity and coverage together, because $\mathcal{D}_t^-$ concentrates where the decoder fails generically or because atom-count buckets holding no rejections realize a rate of zero. QM9 occupancy is in the teacher's latent space, and other high-dimensional designs may need a purpose-built representation. EXPANSE starts from free positives, ActFlow starts empty.

AI USE STATEMENT

Generative AI tools were used only to improve the grammar and readability of the manuscript, which falls under the recommended disclosure category.

ETHICS STATEMENT

Not Applicable.

REPRODUCIBILITY STATEMENT

The source code is submitted as Supplementary Material in a zip folder.

REFERENCES

Silas Alberti, Kenan Hasanaliyev, Manav Shah, and Stefano Ermon. Data unlearning in diffusion models. In *International Conference on Learning Representations*, volume 2025, pp. 3084–3100, 2025.

Sina Alemohammad, Josue Casco-Rodriguez, Lorenzo Luzi, Ahmed Imtiaz Humayun, Hossein Babaei, Daniel LeJeune, Ali Siahkoohi, and Richard Baraniuk. Self-consuming generative models go mad. In *International Conference on Learning Representations*, volume 2024, pp. 53581–53608, 2024.

Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Roni Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *International Conference on Learning Representations*, volume 2024, pp. 51304–51323, 2024.

Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems*, 29, 2016.

Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Advances in neural information processing systems*, 34:27381–27394, 2021.

Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio. Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023.

Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. *Advances in neural information processing systems*, 30, 2017.

Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning. In *International Conference on Learning Representations*, volume 2024, pp. 4965–4987, 2024.

Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.

Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pp. 844–853. PMLR, 2017.

Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.

Riccardo De Santi, Bruce Lee, Cristian Perez Jensen, Kimon Protopapas, Sophia Tang, Cheng-Hao Liu, Pranam Chatterjee, Yisong Yue, and Andreas Krause. Active flow expansion for out-of-distribution discovery: from theory to molecules. *arXiv preprint arXiv:2606.08802*, 2026a.

Riccardo De Santi, Kimon Protopapas, Ya-Ping Hsieh, and Andreas Krause. Verifier-constrained flow expansion for discovery beyond the data. *arXiv preprint arXiv:2602.15984*, 2026b.

Riccardo De Santi, Marin Vlastelica, Ya-Ping Hsieh, Zebang Shen, Niao He, and Andreas Krause. Flow density control: Generative optimization beyond entropy-regularized fine-tuning. *Advances in neural information processing systems*, 38:11056–11088, 2026c.

Mingyang Deng, He Li, Tianhong Li, Yilun Du, and Kaiming He. Generative modeling via drifting. *arXiv preprint arXiv:2602.04770*, 2026.

Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. Raft: Reward ranked finetuning for generative foundation model alignment. *arXiv preprint arXiv:2304.06767*, 2023.

Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, et al. Reinforced self-training (rest) for language modeling. *arXiv preprint arXiv:2308.08998*, 2023.

Sven Gutjahr, Riccardo De Santi, Luca Schaufelberger, Kjell Jorner, and Andreas Krause. Constrained molecular generation via sequential flow model fine-tuning. In *Generative AI and Biology (GenBio)" Workshop the International Conference on Machine Learning*, 2025.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Emiel Hoogeboom, Vıctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International conference on machine learning*, pp. 8867–8887. PMLR, 2022.

Moksh Jain, Emmanuel Bengio, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Bonaventure FP Dossou, Chanakya Ajit Ekbote, Jie Fu, Tianyu Zhang, Michael Kilgour, Dinghuai Zhang, et al. Biological sequence design with gflownets. In *International conference on machine learning*, pp. 9786–9801. PMLR, 2022.

Cristian Jensen, Luca Schaufelberger, Riccardo De Santi, Kjell Jorner, and Andreas Krause. Value matching: Scalable and gradient-free reward-guided flow adaptation. In *International Conference on Learning Representations*, volume 2026, pp. 144926–144958, 2026.

Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, 2011.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Wanli Ouyang. Flow-grpo: Training flow matching models via online rl. *Advances in neural information processing systems*, 38:40783–40818, 2026.

Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.

Barna Pásztor, Parnian Kassraie, and Andreas Krause. Bandits with preference feedback: A stackelberg game perspective. *Advances in Neural Information Processing Systems*, 37:11997–12034, 2024.

Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.

Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007.

Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):140022, 2014.

Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. *arXiv preprint arXiv:1802.09127*, 2018.

Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.

Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. Ai models collapse when trained on recursively generated data. *Nature*, 631(8022):755–759, 2024.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023.

Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.

Yanan Sui, Alkis Gotovos, Joel Burdick, and Andreas Krause. Safe exploration for optimization with gaussian processes. In *International conference on machine learning*, pp. 997–1005. PMLR, 2015.

Yanan Sui, Vincent Zhuang, Joel Burdick, and Yisong Yue. Stagewise safe bayesian optimization with gaussian processes. In *International conference on machine learning*, pp. 4781–4789. PMLR, 2018.

Sophia Tang, Yinuo Zhang, and Pranam Chatterjee. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion. *ArXiv*, pp. arXiv–2412, 2025.

Matteo Turchetta, Felix Berkenkamp, and Andreas Krause. Safe exploration in finite markov decision processes with gaussian processes. *Advances in neural information processing systems*, 29, 2016.

Masatoshi Uehara, Yulai Zhao, Chenyu Wang, Xiner Li, Aviv Regev, Sergey Levine, and Tommaso Biancalani. Inference-time alignment in diffusion models with reward-guided generation: Tutorial and review. *arXiv preprint arXiv:2501.09685*, 2025.

Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.

Huizhuo Yuan, Zixiang Chen, Kaixuan Ji, and Quanquan Gu. Self-play fine-tuning of diffusion models for text-to-image generation. *Advances in Neural Information Processing Systems*, 37: 73366–73398, 2024.

A Derivations and proofs

A.1 Initialization

With fresh latents $\{z_i\}_{i=1}^n$ and a minibatch $\{x_j\}_{j=1}^n$ of $\mathcal{D}_{\text{data}}$,

$$\mathcal{L}_{\text{CD}}(\theta)=\tfrac{1}{2}\,\mathbb{E}\Big[\tfrac{1}{n}\sum_i \min_j \|f_\theta(z_i,1)-x_j\|^2+\tfrac{1}{n}\sum_j \min_i \|f_\theta(z_i,1)-x_j\|^2\Big]. \tag{13}$$

Chamfer rather than a likelihood, because the target is a generable set and every subsequent update is a pushforward step that ignores likelihood. Training stops once $\text{cov}@0.01$ reaches the pre-trained flow's 0.033, so the initial generable set is matched to the baseline's at that level rather than trained past it. We do not claim equality at τ=0.05, and Appendix C.2 bounds what the remaining difference can explain.

A.2 GAUSSIAN SMOOTHING, KERNEL INFLATION, AND MEAN SHIFT

For $X \sim q_{\theta_t}$, $Z \sim \mathcal{N}(0, I)$ independent and $X_s{=}sX + (1-s)Z$, dividing by s gives $X_s/s{=}X + hZ$ with $h{=}(1-s)/s$, proving Eq. equation 2: the family $\{q^h_{\theta_t}\}_h$ a solver visits along its trajectory is available at any h from the terminal generator alone. Differentiating the same convolution gives $m_{t,h}(x){=}h^2\nabla \log q^h_{\theta_t}(x)$ for the Gaussian mean shift $m_{t,h}(x){=}\big[\int (y-x)e^{-\|x-y\|^2/2h^2} q_{\theta_t}(y)dy\big]/\big[\int e^{-\|x-y\|^2/2h^2} q_{\theta_t}(y)dy\big]$.

For the RBF kernel k_ℓ and independent $\xi, \xi' \sim \mathcal{N}(0, h^2 I)$, writing $r{=}x-y$ and $Z{=}\xi-\xi' \sim \mathcal{N}(0, 2h^2 I)$,

$$\mathbb{E}[k_\ell(x+\xi, y+\xi')]{=}c_h\, k_{\ell_h}(x,y), \qquad \ell_h{=}\sqrt{\ell^2+2h^2}, \qquad c_h{=}\big(\ell^2/(\ell^2+2h^2)\big)^{d/2}. \tag{14}$$

Features are $\phi_{h,j}(x){=}\sqrt{2/F}\cos(w_j^\top x/\ell_h + b_j)$ from a shared RFF bank, $w_j \sim \mathcal{N}(0, I_d)$, $b_j \sim \mathrm{Unif}[0, 2\pi]$. The implementation keeps ℓ_h exactly and drops c_h. Since c_h scales $v_{t,h}$ by a constant within a level and Eq. equation 7 divides by the anchor mean at that level, it cancels in $u_{t,h}$ and the probe direction is exact within a level; what does not cancel is the relative weighting *between* levels as h_k anneals, so the probe traverses a normalized multi-scale approximation rather than an exact smoothed-GP posterior. Fixing $h{=}h_{\min}$ removes the annealing entirely and costs about six points on the component, which bounds what this approximation can be worth.

A.3 PROBE REACH

$\Pi_{\mathcal{X}}$ is non-expansive and fixes $\mathcal{X}$, so each step of Eq. equation 5 moves at most δ and the triangle inequality gives $\|x_{t,K} - x_{t,0}\| \le K\delta{=}\rho$, proving Lemma 3.1. Adding the seed and Gaussian norm concentration, for any $\delta_f \in (0,1)$,

$$\|x_{t,K} - x_t^{\mathrm{gen}}\| \le \rho + \sigma_0\big(\sqrt{d} + \sqrt{2\log(1/\delta_f)}\big) \tag{15}$$

with probability at least $1-\delta_f$. The two terms are not interchangeable: travel is a hard budget attained by every consistently ascending probe, whereas reaching distance R by seeding costs $\exp[-\frac{1}{2}(R/\sigma_0 - \sqrt{d})^2]$ per query. Two configurations with equal nominal reach therefore query a distant region at very different rates, which is the asymmetry Appendix B.1 reads off Fig. 3.

A.4 OCCUPANCY MONOTONICITY

With $A_{t,h}{=}\sigma^2 I + \Phi_{t,h}^\top \Phi_{t,h}$, adding a query z replaces $A_{t,h}$ with $A_{t,h} + \phi_h(z)\phi_h(z)^\top$, and Sherman–Morrison gives

$$v^{(z)}_{t,h}(x){=}v_{t,h}(x) - \sigma^2 \frac{\big(\phi_h(x)^\top A_{t,h}^{-1}\phi_h(z)\big)^2}{1+\phi_h(z)^\top A_{t,h}^{-1}\phi_h(z)}, \tag{16}$$

proving Lemma 3.2. The subtracted term scales with the squared kernel-space inner product, so a query depletes a neighborhood of radius ℓ_h rather than a point, and it contains no label, so a rejected query depletes exactly as an accepted one does.

The lemma is exact at fixed coordinates in the space where features live. On the two-dimensional benchmark that is the design space, so probe and depletion share coordinates. On QM9 the probe searches a latent while features come from a frozen ϕ on a decoded graph. Within a round θ_t is fixed and the two coordinate systems are related by a fixed map, so depletion is monotone in both; across rounds a feature location depleted at round t may be reached from a different latent later, which the probe will not treat as depleted. What the loop actually relies on survives: an already queried feature location has low occupancy whenever it is reached again, since verifier and buffers act on decoded graphs. We claim nothing about latent-space revisiting.

A.5 DRIFTING ESTIMATOR

Each step forms a generated set $X{=}\{f_\theta(z_i, \alpha)\}_{i=1}^N$, a rejection set X^- of $N_-{=}\min(\mathrm{round}(N\gamma/(1-\gamma)), N_{\max})$ draws from $\mathcal{D}_t^-$, and an acceptance set X^+ of N_+ draws from $\mathcal{D}_t^+$. The push set $P{=}X \cup X^-$ realizes Eq. equation 9 up to rounding and the cap. With λ the mean distance from X to

$X \cup X^+$, S_{ij}=$-\|x_i - y_j\|/(\lambda R)$ over $y_j \in P \cup X^+$ excluding self-pairs, and A_{ij}=$(\mathrm{softmax}_j S_{ij} \cdot \mathrm{softmax}_i S_{ij})^{1/2}$, let $a_i^{\pm}$ be the affinity mass of x_i on X^+ and P and $\bar{y}_i^{\pm}$ the affinity-weighted means. Then

$$\hat{V}_R(x_i)=a_i^P\, a_i^+ \left(\bar{y}_i^+ - \bar{y}_i^P\right), \qquad \hat{\mathcal{V}}=\sum_{R\in\mathcal{R}} \frac{\hat{V}_R}{\|\hat{V}_R\|_F/\sqrt{Nd}}, \tag{17}$$

each bandwidth normalized by its RMS over generated samples and coordinates, with field and loss computed in coordinates scaled by $\lambda/\sqrt{d}$. The double softmax makes the estimator batch-coupled where the population field is not: the affinity a generated sample receives from an accepted sample falls when other generated samples compete for it, and $\bar{y}_i^+ - \bar{y}_i^P$ averages only what is in the batch. A sparsely represented region of $\mathcal{D}_t^+$ therefore contributes a small a_i^+ to its few neighbors and nothing to the rest. This is the mechanism behind Appendix B.3.

A.6 Proofs

Proposition 3.3. (i) Solving $(1-\gamma)q + \gamma p_t^- = p_t^+$ gives $q=\alpha p_t^+ - (\alpha-1)p_t^-$, unique because $q \mapsto (1-\gamma)q$ is injective, with total mass 1. (ii) $q_\alpha^\star(\mathcal{X} \setminus \Omega^\star) = -(\alpha-1)$. (iii) Convolution is linear, so $q_\alpha^\star * k = \alpha(p_t^+ * k) - (\alpha-1)(p_t^- * k)$; divide by $\alpha > 0$ and use $(\alpha-1)/\alpha=\gamma$.

Theorem 3.4. Given $\mathcal{F}_t$, each seed has law $\mu_t^{\mathrm{seed}}=[\Pi_{\mathcal{X}}]_\#(q_{\theta_t} * \mathcal{N}(0,\sigma_0^2 I))$. If $x_{t,K}^b \in C$ then Lemma 3.1 gives $x_{t,0}^b \in \mathcal{B}(C,\rho)$, so $\mathbb{E}[\mathbf{1}[x_{t,K}^b \in C] \mid \mathcal{F}_t] \le \mu_t^{\mathrm{seed}}(\mathcal{B}(C,\rho))$. Summing over b,t and taking expectations gives Eq. equation 11; Markov gives the probability statement. The only inequality used is the containment $\{x_{t,K} \in C\} \subseteq \{x_{t,0} \in \mathcal{B}(C,\rho)\}$, so the slack is exactly the seed mass that starts within ρ of C and does not end in it.

Corollary 3.5. Non-expansiveness gives $\mathrm{dist}(x_{t,0},\mathcal{M}) \le \mathrm{dist}(x_t^{\mathrm{gen}},\mathcal{M}) + \sigma_0\|\xi_t\|$ and Lemma 3.1 adds ρ. Since every point of C is at distance at least Δ from $\mathcal{M}$, on the event $\mathrm{dist}(x_t^{\mathrm{gen}},\mathcal{M}) \le r$ reaching C requires $\sigma_0\|\xi_t\| \ge \Delta - r - \rho$. With $s=(\Delta - r - \rho)/\sigma_0 - \sqrt{d}$, the bound is trivial for $s \le 0$ and Gaussian norm concentration gives $e^{-s^2/2}$ for $s > 0$; for σ_0=0 the event is empty unless $r + \rho \ge \Delta$. The complementary event has probability $q_{\theta_t}(\mathrm{dist}(x,\mathcal{M}) > r)$. The free r splits the two ways a gap is crossed: small r charges it to seed noise and travel, large r to generator mass already placed away from $\mathcal{M}$. The bound is informative while the generator is concentrated on $\mathcal{M}$, which is the regime before the component is absorbed and not after.

A.7 Relation to ActFlow reachability

ActFlow models $\Pr[y{=}1 \mid x]=\varsigma(g(x))$ with $\Omega^\star=\{x : \varsigma(g(x)) \ge \eta\}$, $L=L_\varsigma L_g$, and $\mathcal{R}_\varepsilon(S)=\{x : \exists x' \in S,\ \varsigma(g(x')) - L\|x - x'\| - \varepsilon \ge \eta\}$. The two results answer different questions. $\mathcal{R}_\varepsilon$ is sufficient and local: it certifies a region is valid and worth expanding into, given a calibrated verifier, a Lipschitz score, an assumption on the generative update, and an uncertainty-maximizing sampler. Theorem 3.4 is necessary and global: it bounds whether a region can be queried at all from probe geometry alone, and says nothing about validity. A method can satisfy our bound and never expand; a region can lie in $\mathcal{R}_\varepsilon(S)$ and never be queried. We prove the necessary direction because it holds for any acquisition rule and needs no calibrated verifier, which is the assumption Expanse is built to avoid.

The diagnostic substitutes the finest-level posterior mean for $\varsigma \circ g$ and its largest grid gradient norm for L. Both substitutions are optimistic, so the resulting set upper-bounds what is certifiable, which is why we report only that the achieved generable set contains 94 to 98% of it. Theorem 3.4 does not use the operator.

B Reading the mechanisms

B.1 Seed scale, travel, and why the bound is loose on one and not the other

At ρ=0 the query is the seed, so a round is i.i.d. from μ_t^{seed} and reaching distance R costs the Gaussian tail of Appendix A.3: zero component coverage through σ_0=1.1, then $28.7 \pm 3.0\%$ at σ_0=1.4. Buying that tail inflates the whole seed law, scattering queries isotropically and paying in

validity. At σ_0=0 each query is displaced by exactly ρ along the ascent direction with no tail to pay once ρ exceeds the distance and the gradient points across: zero through ρ=2.25, then $41.9 \pm 13.5\%$ at ρ=2.5, above anything the seed sweep attains. Between the thresholds the seed spread exceeds the mean, because whether the gradient points across depends on where earlier queries depleted occupancy. The combined setting (σ_0=1.0, ρ=2.0) reaches 83.9%, above either mechanism alone at any value tested; they compose rather than duplicate, which is the split Eq. equation 11 already makes between μ_t^{seed} and $\mathcal{B}(C, \rho)$. Practically, σ_0 is set by how far generator mass is expected to lie from the regions of interest and ρ by how far one query should travel once placed; the former costs validity, the latter K occupancy-gradient evaluations per probe.

The bound overestimates retained component queries by $\approx 51\times$ on the travel sweep (7600 ± 200 against 149 ± 18) and under $2\times$ on the seed sweep (330 ± 25 against 203 ± 12), and the single containment in the proof accounts for both. At ρ=0 the containment is an equality, since $\mathcal{B}(C, 0)$=C; the residual factor of 1.6 is explained by measuring distances to a bounding box of C and by the bound counting all queries where we report those retained after thinning, leaving nothing for the acquisition rule. At ρ=2.5 the gap is almost entirely the volume the neighborhood introduces: C occupies 512 cells of side $1/16$, so its bounding box has side ≈ 1.41 and dilating by ρ gives $s^2 + 4s\rho + \pi\rho^2 \approx 35.8$, a ratio near 18. The looseness therefore tracks a volume ratio the bound accepts in exchange for holding under any acquisition rule, not the rule itself.

This fixes what the predicate is evidence for. Every run with $\widehat{N}_C < 1$ missed the component and every run that found it had $\widehat{N}_C \geq 1.08$, but $\widehat{N}_C$ crosses 1 near where ρ and $\sigma_0\sqrt{d}$ together exceed Δ, so it is a calibrated restatement of that threshold rather than independent evidence for it. Its value is that the restatement is computable from checkpoints before the budget is spent, while Δ is unknown for a real design space: a practitioner can rule a configuration out in advance, not in. Two caveats on $\widehat{N}_C$ itself. μ_t^{seed} is not stationary, since once the component is absorbed the generator places mass there and later rounds reach it at a rate early ones do not predict, so holding each checkpoint's seed law until the next overestimates for a second reason beyond the bounding box. And success is determined by the first hit rather than the total, which is why we report the predicate on $\Pr(N_C \geq 1)$.

B.2 What the probe buys, what it costs, and where labels enter

Undirected seeding with no ascent reaches 96.1% of the connected region and 0% of the component, against 90.7% and 83.9% for the full method. On the connected region the probe costs about five points, and Eq. equation 6 says why: occupancy is depleted by queries rather than by coverage, so the probe leaves a densely queried region even where the generator has not yet placed mass, while undirected seeding keeps returning. Every configuration we test that spreads queries without reaching the component lands between 94 and 97% connected coverage, so that region is close to saturated by the seed law and the redirected query is one it would have used anyway. The component is the opposite case: the probe converts a rate falling exponentially in distance into one flat up to ρ. The β sweep traces the same trade from the other side, since β=1 recovers 96.6% connected and loses the component entirely, $m_{t,h}$ pointing toward modes of $q_{\theta_t}^h$ and cancelling the outward part of $\nabla u_{t,h}$. We run β=0 throughout, and read ActFlow's KL term as playing a role a one-step generator does not need, keeping guided samples on the data manifold during a trajectory where here validity is enforced by the update rather than by the query. The reading is conditional: on a connected valid set the initial generator touches, undirected seeding at a well-chosen σ_0 is stronger, and the probe is worth its cost only when valid regions are suspected outside the seed law's reach, which is the case EXPANSE is configured for.

Labels enter once. The occupancy field is a posterior variance, which for a Bayesian linear model depends on $\Phi_{t,h}$ and not the responses, as Eq. equation 16 confirms; the probe direction is built from $u_{t,h}$ and $m_{t,h}$, neither of which sees a label; validity first appears in Eq. equation 8 through the buffer partition. Acquisition therefore needs no calibrated classifier. ActFlow's uncertainty oracle is a validity model whose posterior standard deviation drives exploration, so miscalibration steers queries toward regions the model is wrong about rather than regions that are unqueried; occupancy cannot be miscalibrated in that sense because it makes no validity claim. What it gives up is any preference among unqueried regions: at β=0 a region is equally worth querying whether or not every neighboring query was rejected. That is cheap where the valid set occupies much of the reachable

space and would not be where valid designs are rare, which is the main regime in which we expect the label-free signal to be the wrong choice.

B.3 Batch size and replay separation

Reducing N from 512 to 128 costs connected coverage 13 to 29 points and collapses the component once Δ grows, from 92.1 at Δ=0 to 4.6 at Δ=2 against 95.7 and 83.9 at N=512. Appendix A.5 gives the mechanism: a newly discovered region contributes few members to both sides of the affinity, so it receives a small a_i^+ and its generated neighbors are displaced toward it weakly. The locus is consolidation rather than covering, since N is the generated-batch size of the drifting step and appears nowhere in Eqs. equation 4–equation 5; both arms probe with identical B, M, σ_0, ρ, h ladder and β. The requirement scales with the number of distinct regions the generator must hold simultaneously rather than with dimension, since what matters is the expected count of $\mathcal{D}_t^+$ members per region in a batch of size N. We do not measure the constant, and flag this as the mechanism most likely to bind on a design space with many disconnected modes; the QM9 bucketed replay sidesteps it by partitioning the batch on atom count, which is a domain-specific fix rather than a general one.

Replay separation is the same mechanism from the buffer side, which is why removing it lowers both metrics and widens the component spread from 2.2 to a range of 21.9 to 64.3. Without thinning, $\mathcal{D}_t^+$ is dominated by whichever regions the generator samples most, a batch drawn from it holds few members of a new region, and absorption is slow. Separation and occupancy both discourage reprocessing a region but act on different sets, occupancy on every query and $r_{\min}$ on what is stored, so a repeatedly queried region keeps depleting occupancy while contributing a bounded number of buffer entries.

B.4 What the mixture target does and does not say

Proposition 3.3(i) says the implied target $\alpha p_t^+ - (\alpha - 1)p_t^-$ subtracts a multiple of p_t^- outright rather than down-weighting it, which is the distinction from scaling a negative gradient, whose effect is bounded by the scale. Part (ii) says the target is not a probability distribution for $\alpha > 1$, so no generator realizes it, and (iii) identifies where its smoothing is nonnegative. The implemented update realizes none of this exactly, since Eq. equation 17 uses batch-normalized affinities over a finite sample and is not a fixed convolution; (iii) motivates the construction rather than describing it.

Whether the intended extrapolation occurs is empirical. On the two-dimensional benchmark it does: one model traces a coverage-validity curve in α, α=1.25 gives 80.8 to 83.3% coverage at 97.5 to 99.2% validity, and γ=0 costs about five points on the component and nothing on the connected region while collapsing the curve to a point. On QM9 it does not, since raising α lowers validity and coverage together. Two mechanisms could produce that and we do not separate them: the realized $\widehat{\gamma}_k$ of Appendix D.2 may be near zero for most updates, since many atom-count buckets hold no rejections and the generator is then conditioned on an α its push batch never realized; or $\mathcal{D}_t^-$ may concentrate where the decoder fails generically, so extrapolating away from it moves the generator away from decodable latents. Logging $\widehat{\gamma}_k$ per update would distinguish them and we did not. Until then the α control is a property of the two-dimensional loop and we report α=1 on QM9.

C Experimental protocol

C.1 Two-dimensional hyperparameters

Occupancy ladder $\mathcal{H}$=$\{1.0, 0.6, 0.36, 0.21, 0.13, 0.1\}$, F=1024 random features, ℓ=0.1, σ^2=$\log 2$. Drifting uses N=N_+=512, $N_{\max}$=1024, $\mathcal{R}$=$\{0.02, 0.05, 0.2\}$, κ=0.05. The generator has latent dimension 16, width 256, 4 residual blocks, AdamW at 2×10^{-4} with gradient clipping at norm 2. Phase 1 fits Eq. equation 13 with Adam at 10^{-3} and batch 256, stopping at cov@0.01=0.033. Probes use M=512 anchors, K=40 steps of length δ=0.05, bandwidths annealed geometrically from $h_{\max}$=1.0 to $h_{\min}$=0.1. Evaluation samples 200k designs from online weights.

C.2 Initialization, threshold, and baselines

EXPANSE begins with the 57 positives obtained by thinning 512 pre-training samples at $r_{\min}$, where the released ActFlow loop begins empty. Three facts bound what this can explain. The points cost no verifier queries, so the 32k budget is matched. They never enter $\Phi_{t,h}$, so the occupancy field starts flat for both and acquisition is unaffected. And they are drawn from $\mathcal{D}_{\text{data}}$, whose generable set is the 2.4 / 0.0 row of Fig. 2 at every gap, so they contain no point of the component and cannot contribute to the component result, which is where our margin is largest. What they provide is a pull side for the first drifting steps, a cold-start requirement specific to a pushforward update: with $\mathcal{D}_t^+=\emptyset$ the field of Eq. equation 17 has $a_i^+=0$ and the update is inert. ActFlow has no equivalent requirement because it steers a pre-trained sampler already covering that 2.4%. We do not run the crossed initializations, so this is an argument bounded by the pre-trained row rather than a measurement.

ActFlow reports $\tau=0.01$, at which every method except ActFlow (SDE) exceeds 93% seed-mean total coverage at every gap, so that level cannot separate them. We report $\tau=0.05$, about 39 samples per cell. Raising τ demands more samples per cell identically for every method, so a method leading at one level loses the lead at a higher one only if its density is more concentrated in fewer cells; the margins in Fig. 2 compress toward $\tau=0.01$ by construction and would widen above 0.05 as the weaker methods fall below the level first. The single reported level is therefore the conservative choice among those that separate.

For baselines, the released ActFlow verifier omits a disconnected component, so we add one at distance Δ and otherwise rerun the released code unmodified, at the shipped $w=100$ and the $w=13$ its paper reports, at both samplers, and across 1 to 100 NFE in Fig. 6. ActFlow (ODE) is no worse at 10 NFE than at 100 and we compare against that operating point as well as its native one. On QM9 it reaches $6.1 \pm 1.4\%$ validity at 4 NFE and no valid molecule from 8000 draws at 2 NFE by iteration 100, far enough from its published behavior that a configuration failure in our molecular port is at least as likely as a property of the method; we did not sweep its guidance weight there and rest no claim on that margin. The QM9 comparison that carries weight is against Rec-F, which shares the retraining structure and differs only in acquisition and mixture.

EXPANSE and ActFlow differ in acquisition signal, query mechanism, replay construction and generator update, so Fig. 2 compares two methods and no margin there is attributable to one difference. A component swap is not available: the probe needs terminal design coordinates to ascend in, which a guided denoising trajectory does not expose until its final step, and the mixture needs a pushforward update, which guidance-weight training does not have. Attribution is internal, through Fig. 7. The cross-method comparison supports one claim: the loop taken whole reaches more of this valid set under this budget.

D QM9 implementation and metrics

The toy benchmark probes terminal design coordinates directly; the molecular experiment probes a structured latent and evaluates acquisition only after decoding. The definitions here are authoritative for QM9.

D.1 Sampler and probe

The model reuses the pretrained QM9 vector field and adds an α-condition through FiLM on its scalar embedding, $e \mapsto e \odot (1 + s_\alpha) + b_\alpha$, where (s_α, b_α) are produced from $\log \alpha$ and multiplied by $\log \alpha$, so $\alpha=1$ is exactly the unconditioned network. Both vector-field and FiLM parameters are updated. At n NFE a prior state $x_0 \sim p_0$ is propagated with exactly n evaluations under the pretrained scheduler, giving $G^{(n)}_{\theta,\alpha}(x_0)$. This is a fine-tune of a teacher sampled at 2 to 4 NFE, not a one-step generator, and we do not report it as one: the one-step claim concerns the two-dimensional benchmark, where $f_\theta(z, \alpha)$ is trained from scratch and every sample is one forward pass. What transfers is the covering/consolidation decomposition, which depends on neither architecture nor step count; what does not is the argument that a one-step generator *forces* the separation, which here is imposed by construction.

With S_w the blockwise RMS normalization over coordinates, atom logits, charge logits and edge logits, the probe starts from $z_0=S_w(x_0) + \sigma_{\text{seed}}\epsilon$, maps a candidate back to a graph state, decodes it

with the same n_{exp}-NFE map used for verifier queries, and evaluates a frozen extractor ϕ, a copy of the pretrained model at a fixed feature timestep. The RFF posterior lives in this feature space, not in fingerprint space. With $f_i(z_i)=\phi(G_{\theta,\alpha}^{(n_{\text{exp}})}(S_w^{-1}(z_i)))$, posterior variance is differentiated through the decode,

$$z_{i,k+1}=z_{i,k}+\delta_z\,\nabla_{z_i} v_{t,h_k}(f_i(z_{i,k}))\big/\big\|\nabla_{z_i} v_{t,h_k}(f_i(z_{i,k}))\big\|, \qquad \delta_z=\texttt{probe_step}\sqrt{D_z},$$

with the batch objective divided by its detached mean, which does not change the direction after per-candidate normalization. The verifier is evaluated only after decoding. The implementation also supports a positive-anchor log-sum-exp term at $\beta > 0$ over accepted feature vectors; this is *not* the mean-shift prior of Eq. equation 3, so the population reading of that equation is not claimed for it, and at the β=0 we report it is inactive. Lemma 3.2 applies exactly at fixed coordinates in this frozen feature space, and Appendix A.4 states what that does and does not give across rounds.

D.2 Replay and drifting update

Verified graphs are stored by atom-count bucket, accepted graphs thinned in the frozen feature space before entering the positive buffer, rejected graphs stored without that rule. A drifting step picks a bucket, draws replay from it, and generates graphs of the same atom count, which keeps the field from comparing incompatible sizes. For requested γ the code takes $C_-=\text{round}(C_g\gamma/(1-\gamma))$ negatives for C_g generated samples when the bucket has them, so the empirical fraction is $\widehat{\gamma}_k=C_-/(C_g+C_-)$ up to rounding and $\widehat{\gamma}_k$=0 when it has none, while the condition remains $\alpha=1/(1-\gamma)$ from the requested γ. Whenever $\widehat{\gamma}_k \neq \gamma$ the correspondence of Prop. 3.3 fails for that update, which is the first of the two candidate explanations in Appendix B.4. Graphs are flattened and block-normalized before the multi-bandwidth field is computed, generated and rejected on the push side, accepted on the pull side. A teacher-anchor loss on cached teacher trajectories regularizes the sampler; it is separate from the mixture algebra and present in all EXPANSE QM9 runs including γ=0.

D.3 Metrics and matched-compute scoring

A terminal graph is discretized and converted to an RDKit molecule, and $\text{Val}=N_{\text{valid}}/N_{\text{drawn}}$ keeps invalid graphs in the denominator. Fingerprints are Morgan, radius 2, 2048 bits. The scorer's $\texttt{cov}\tau$ is a greedy packing count on valid generated molecules: in sample order a fingerprint becomes a center if its maximum Tanimoto similarity to existing centers is below τ. At τ=0.2 these are discovery counts, not the fraction of a reference packing hit, and the figures label them as discovery. Reference coverage is separate: $\mathcal{R}$ is built greedily from the reference pool at Tanimoto-distance radius r_{ref}, a center is hit if some valid generated fingerprint has similarity at least $1-r_{\text{ref}}$, and `ref_frac` is the hit fraction; r_{ref}, not τ, controls this packing. Among novelty quantities, `outside` is the fraction of valid molecules outside every reference center and `new_total` a second greedy packing among those, so the first counts molecules and the second centers. `new_total` is the least confounded, since a molecule must be valid, outside the reference packing and distinct from other novel molecules to contribute; $\texttt{cov}\tau$ is a joint function of draws and validity under a fixed budget; `ref_frac` is what to read for coverage of known chemistry. The QM9 claims in Sec. 4 are stated in that order of confidence.

Each score cell has budget B_{NFE}, so an arm at n NFE receives $N_{\text{draw}}=\lfloor B_{\text{NFE}}/n \rfloor$ draws. Metrics are computed on the complete valid pool from those draws and a cell with no valid molecule receives zero rather than being censored, so low validity is paid for in draws rather than hidden. Packing counts are thus discovery per fixed compute by construction, not an NFE-invariant estimate of a support. Shape statistics that depend directly on sample size use fixed-n subsamples. One vector-field call counts as one NFE and the teacher and ActFlow checkpoints are evaluated under the same per-cell budget.

The frozen teacher illustrates the budget: 38 centers at 97.5% validity with 100 steps, and 174 at 27.9% with 4, from 4000 draws rather than 400, of which 390 and 1116 are valid. Greedy packing counts grow sublinearly in pool size, so the 4.6× center ratio and 2.9× valid-pool ratio are not directly comparable. The direction of that bias is against the larger pool: sublinearity means the 4-step arm would need *more* than 2.9× the centers of the 100-step arm to be explained by pool size alone, so observing 4.6× is conservative evidence that coarsening the solver spreads samples as well as multiplying them. A pool-matched subsample would tighten the estimate and we do not report one.

D.4 SCOPE ON QM9

Theorem 3.4 is evaluated on the toy benchmark, where probed coordinates and disconnected geometry share a Euclidean space. QM9 controls travel in the block-normalized latent metric and we assume no Lipschitz relation from it through the decoder to Tanimoto distance, so the theorem is not a molecular-space reach certificate. Appendix B.1 carries over in form, since σ_0 still enters through a Gaussian norm and ρ as a hard budget, but the distances those controls bound are not the ones the metrics measure. Prop. 3.3 does not imply that raising α increases validity for a finite learned sampler, so the QM9 behavior is a failure of the two-dimensional frontier to transfer rather than a contradiction. Finally, the $350\times$ sampling advantage follows from evaluating one network per sample rather than one hundred and would hold for any correctly distilled sampler of comparable width; it is a property of the deployed artifact, not evidence for the expansion procedure, and no claim about covering or consolidation rests on it.