A REFERENCE-FREE LOWER BOUND ON THE SYMMETRY ERROR OF 3D GENERATIVE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Scalable 3D generative models increasingly use free-form backbones even when the data or downstream task has a known symmetry, learning rotation and permutation invariance from augmentation rather than enforcing it architecturally. Augmentation teaches symmetry only approximately, so deployed models carry a residual symmetry-breaking error that is hard to measure: the reference distribution is unavailable after training, and existing equivariance diagnostics quantify the network's mismatch with the group action rather than the distance between the generated distribution and the target. We define the symmetry defect as the Kullback–Leibler divergence between a model and its own group-orbit average, computable from the model likelihood alone. A classical information-geometric Pythagorean identity makes this model-only quantity a lower bound on the error to any symmetric target, and the same identity is constructive: applying one uniformly random group element to each sample removes exactly the certified component without retraining. We separate the population quantity that carries the guarantee from its finite Monte-Carlo estimate, and exact enumeration of a finite subgroup, itself a valid lower bound, from sampled continuous-group averaging with its Jensen bias. Across known-density targets, a free-form diffusion model with a probability-flow likelihood, real ShapeNet point clouds, and pretrained conformer and point-cloud generators, finite-sample estimates respect the population bound, non-equivariant backbones read a measurable defect while equivariant counterparts read near zero, symmetrization markedly improves likelihoods, and the score ranks candidate models in strong agreement with hidden reference orderings. Because the defect is a lower bound, a near-zero reading certifies near-invariance rather than closeness to the target: the certificate falsifies rather than validates.

1 INTRODUCTION

Generative models of three-dimensional geometry have become central tools across the sciences and design, synthesizing point clouds, shapes, molecules, and protein structures. A defining property of such 3D data is symmetry: a shape is unchanged when it is rotated or translated, and a point cloud is unchanged when its points are re-ordered, so a faithful generator must place equal probability on all such variants. Honoring this symmetry is widely regarded as essential to sample quality, data efficiency, and physical validity, and a substantial body of work builds it into the model by construction through equivariant architectures (Thomas et al., 2018; Deng et al., 2021; Satorras et al., 2021; Köhler et al., 2020; Klein et al., 2023; Hoogeboom et al., 2022).

The most scalable and highest-quality 3D generators, however, increasingly move in the opposite direction, trading exact equivariance for the flexibility and scale of free-form backbones. Transformer- and point-voxel-based point-cloud and shape models, and beyond point clouds augmentation-trained structure predictors and "scalable non-equivariant" molecule generators, acquire symmetry only approximately, by learning it from rotation and permutation augmentation rather than enforcing it architecturally (Nichol et al., 2022; Mo et al., 2023; Zeng et al., 2022; Abramson et al., 2024; Ding & Hofmann, 2025). This buys scalability at the price of a residual symmetry-breaking error that pushes the generated distribution away from the symmetric target and degrades downstream use, an error that is paid silently. Quantifying it at deployment is hard: the true discrepancy to the target is unknown without exactly the expensive reference one hopes to avoid, and existing equivariance diagnostics quantify the network's mismatch with the group action at the function level, not the distance between

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.

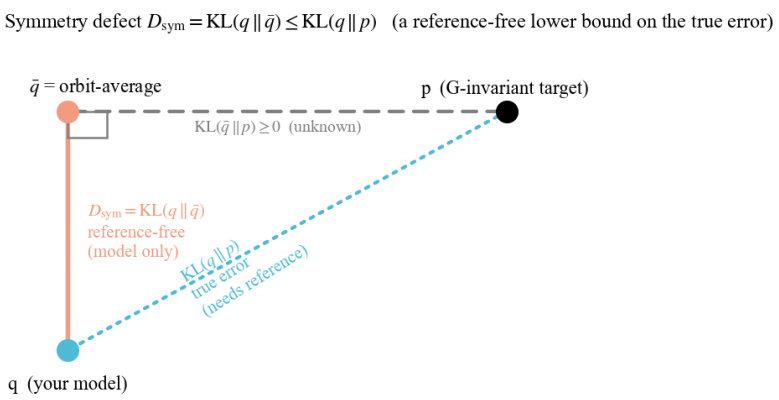

Figure 1: The symmetry defect is a reference-free lower bound. For invariant p, information projection gives $\text{KL}(q \,\|\, p) = \text{KL}(q \,\|\, \bar{q}) + \text{KL}(\bar{q} \,\|\, p)$; the model-only term $D_{\text{sym}} = \text{KL}(q \,\|\, \bar{q})$ therefore lower-bounds the true error.

the generated distribution and the target, so they guarantee nothing about the latter (Gruver et al., 2023).

We study a cheap, reference-free population certificate of the symmetry error, applicable to any generator with a tractable likelihood. The idea is elementary. We compare the model q to its own group-orbit-average $\bar{q}(x) = \mathbb{E}_{g\sim G}[q(g\cdot x)]$, the closest symmetric distribution to it, and measure their divergence, the symmetry defect $D_{\text{sym}} = \text{KL}(q \,\|\, \bar{q})$. Because the target is symmetric, a classical information-geometric Pythagorean identity guarantees that the population D_{sym} is a lower bound on the model's population KL error (Figure 1). Numerical reporting still requires an outer expectation over model samples: exact enumeration of a finite subgroup makes the inner orbit average exact, but the reported sample mean remains a Monte-Carlo estimate unless an interval is given. As a population lower bound, the quantity isolates symmetry breaking but, by design, falsifies rather than certifies goodness. The same identity is constructive: deploying the symmetrized model, which draws a sample and then applies a uniformly random group element such as a random rotation or permutation, removes exactly the population defect. It is thus also a one-line population corrector that fixes the symmetry component it detects, with no retraining and no reference, while the complementary non-symmetry error is left untouched and uncertified.

The orbit average, its KL projection property, and even the ratio $q/\bar{q}$ have direct precedents in nearest symmetric distributions, symmetrized variational inference, relative-entropy asymmetry, and group-invariance e-values (Pavlichin, 2015; Moore, 2016; Gour et al., 2009; Koning, 2023). We build on this classical projection; our contribution is to turn it into a deployable audit for approximately-equivariant 3D generators: a population certificate/corrector interpretation, a subgroup-versus-MC estimation protocol, and empirical stress tests on modern likelihood models. This is distinct from equivariance diagnostics, which probe the map's failure to commute with the group and place no bound on the output distribution (Gruver et al., 2023), and from symmetrization methods, which apply the orbit average to enforce equivariance (Zhang et al., 2025; Puny et al., 2022; Kaba et al., 2023) rather than to audit a deployed 3D density. We also evaluate the score as an empirical model-selection rule (Section 6.6). Our contributions are threefold:

- **A 3D deployment formulation.** We specialize the classical orbit-average KL decomposition to tractable-likelihood 3D generators and state the population lower-bound and corrector consequences with their finite/compact-group and integrability assumptions.
- **An estimation audit.** We separate the population subgroup quantity (which has the guarantee) from its finite outer-MC estimate, and separate exact inner subgroup enumeration from sampled continuous-group averaging and its Jensen bias.
- **Controlled, real-data, and pretrained-model diagnostics.** Known-density flows and diffusion models check the population predictions through finite-MC estimates; a free-form diffusion model on real ShapeNet points and pretrained DiTMC supply real-data likelihood

cases; PointFlow supplies a conditional-decoder diagnosis; and a twelve-model pool tests the score as a selection rule.

2 Related Work

Exact and approximate equivariance metrics. Equivariant architectures bake symmetry into the model (Thomas et al., 2018; Deng et al., 2021; Satorras et al., 2021; Köhler et al., 2020; Klein et al., 2023; Hoogeboom et al., 2022), giving exact guarantees at the cost of specialized machinery, while the most scalable 3D generators use free-form backbones with augmentation (Nichol et al., 2022; Mo et al., 2023; Zeng et al., 2022; Abramson et al., 2024; Ding & Hofmann, 2025) and are only approximately equivariant at sampling time. That trade-off is our motivation: we do not propose another architecture, we certify the residual symmetry error of whatever model is deployed. The Local Equivariance Error and related finite-transformation mismatches measure how far the network function is from commuting with the group, e.g. $\|f(Rx) - \rho(R)f(x)\|$ (Gruver et al., 2023). This is a property of the map, not a bound on the generated distribution's divergence to the target, and the two can move independently. Our D_{sym} is a divergence between distributions and a provable lower bound on $\mathrm{KL}(q \,\|\, p)$.

Relation to symmetrization and invariance testing. Pavlichin (2015) characterize nearest G-invariant distributions and obtain arithmetic orbit averaging for the relevant KL direction. Moore (2016) explicitly use the same $\mathrm{KL}(q \,\|\, \bar{q})$ term and Monte-Carlo evaluation in symmetrized variational inference. Gour et al. (2009) study relative entropy to the nearest invariant quantum state, while Koning (2023) derives the log-optimal group-invariance e-value $q/\bar{q}$, whose expectation under q is precisely D_{sym}. Distributional invariance testing (Chiu & Bloem-Reddy, 2023) and dataset diagnostics (Lawrence et al., 2025) provide further neighboring uses. Our contribution is to turn these facts into an audited workflow for deployed tractable-likelihood 3D generators, explicitly separating population guarantees from finite outer-MC reports and comparing them with function-space diagnostics. Existing symmetrization methods (Zhang et al., 2025; Puny et al., 2022; Kaba et al., 2023; Birrell et al., 2022) motivate the sampler but do not by themselves supply this 3D deployment audit (App. H).

3 Preliminaries

Symmetric 3D targets and model likelihoods. We consider a finite or compact group G with normalized Haar probability measure acting on distributions over N points in $\mathbb{R}^3$, $X = (x_1, \ldots, x_N) \in \mathbb{R}^{3N}$. A group G acts either by a shared rotation of every point, $g{\cdot}X = (Rx_1, \ldots, Rx_N)$ for $R \in \mathrm{SO}(3)$, or by permuting points, $g{\cdot}X = (x_{\pi(1)}, \ldots, x_{\pi(N)})$ for $\pi \in S_n$; both are measure-preserving linear maps with $|\det g| = 1$. The data distribution p is G-invariant, $p(g{\cdot}X) = p(X)$ for all g. A generative model provides a tractable density q with respect to a common base measure, such as the change-of-variables likelihood for normalizing and continuous flows or the probability-flow ODE likelihood for diffusion models (Song et al., 2021). The certificate requires only the ability to evaluate q at chosen points. We assume the displayed log ratios are integrable and the KL terms used in each identity are finite. Appendix A collects the notation.

Orbit average. For a finite group $\bar{q}(X) = \frac{1}{|G|} \sum_{g \in G} q(g{\cdot}X)$, and for a compact group with normalized Haar probability measure such as $\mathrm{SO}(3)$, $\bar{q}(X) = \mathbb{E}_{g \sim G}[q(g \cdot X)]$. Because each g preserves volume, $\bar{q}$ is a valid normalized density, G-invariant by construction, and the closest invariant distribution to q in KL: the m-projection onto the invariant family.

4 Method: the Symmetry-Defect Certificate

Problem. Given a deployed model q whose target p is G-invariant but whose true error $\mathrm{KL}(q \,\|\, p)$ is unknown for lack of a reference, we want a quantity, computable from q alone, that provably bounds that error and hence flags symmetry breaking.

Definition 1 (Symmetry defect)**.** For a model q with orbit-average $\bar{q}$ under G, the symmetry defect is $D_{\mathrm{sym}}(q) := \mathrm{KL}(q \,\|\, \bar{q}) = \mathbb{E}_{X \sim q}[\log q(X) - \log \bar{q}(X)]$.

D_{sym} uses only the model: draw $X \sim q$, evaluate q at X and at its orbit $\{g \cdot X\}$. It needs no target density p, no data, and no reference. Section 5 shows it lower-bounds the true error; here we describe estimation.

Population quantity versus numerical estimate. With M source draws $X_i \sim q$ and an inner orbit set $\mathcal{G}$, we report

$$\widehat{D}_{\mathrm{sym}} = \frac{1}{M} \sum_{i=1}^{M} \Big[\log q(X_i) - \mathrm{logmeanexp}_{g \in \mathcal{G}} \log q(g \cdot X_i) \Big]. \tag{1}$$

For point permutations S_n we enumerate the inner group for small n or sample it for large n; for $\mathrm{SO}(3)$ we average over m uniformly sampled rotations. Even when $\mathcal{G} = H$ is enumerated exactly, finite M makes equation 1 an outer-MC estimate of the exact population quantity $D_{\mathrm{sym}}^{H} = \mathbb{E}_q[\log q - \log \bar{q}_H]$. Exact enumeration removes inner-orbit error, not outer sampling uncertainty. The finite octahedral rotation subgroup ($|H|$=24, the rotations of a cube) can be enumerated exactly, and by the nesting of information projections $\mathrm{KL}(q \,\|\, \bar{q}_H) \le \mathrm{KL}(q \,\|\, \bar{q}_G) \le \mathrm{KL}(q \,\|\, p)$, so a subgroup population quantity is itself a valid, if looser, lower bound and provides an inner-exact anchor. The Monte-Carlo estimate of $\log \bar{q}$ is downward biased by Jensen's inequality, so equation 1 over-estimates D_{sym} and decreases toward the true value as m grows. We separate sound numbers from estimates throughout, and give error bars plus a delta-method debiasing in App. D. The certificate costs $|\mathcal{G}|$ likelihood evaluations per source draw, a small multiple of one density call, with no training, no data, and no expensive reference solve.

From detection to provable correction. The certificate does more than detect: the same Pythagorean identity that lower-bounds the error also says exactly how to remove the part it certifies. Read equation 3 from the other side.

Proposition 1 (The symmetrized sampler removes exactly the certified error)**.** Let G be finite or compact with normalized Haar probability measure, act by measure-preserving maps, and let p be G-invariant. For densities $p, q, \bar{q}$ with finite displayed KL terms,

$$\mathrm{KL}(\bar{q} \,\|\, p) = \mathrm{KL}(q \,\|\, p) - D_{\mathrm{sym}}(q) \;\le\; \mathrm{KL}(q \,\|\, p). \tag{2}$$

Hence $\bar{q}$ has smaller error than q whenever $D_{\mathrm{sym}} > 0$, improved by exactly the certified defect D_{sym}. Moreover $\bar{q}$ is the m-projection of q onto the invariant family, $\bar{q} = \arg\min_{r:\, G\text{-invariant}} \mathrm{KL}(q \,\|\, r)$, so it is the optimal fixed correction: no invariant distribution is closer to q.

Proof. Rearranging equation 3 gives equation 2, and $D_{\mathrm{sym}} \ge 0$ gives the inequality; for optimality, any G-invariant r has $\mathbb{E}_q[\log r] = \mathbb{E}_{\bar{q}}[\log r]$ (Theorem 2, Step 1), so $\mathrm{KL}(q \,\|\, r) = \mathrm{KL}(q \,\|\, \bar{q}) + \mathrm{KL}(\bar{q} \,\|\, r) \ge \mathrm{KL}(q \,\|\, \bar{q})$ with equality iff $r = \bar{q}$. □

The corrector is a one-line sampler. Deploying $\bar{q}$ needs no retraining, data, or reference: draw $X \sim q$, apply one uniformly random group element ($R \sim \mathrm{SO}(3)$ or $\pi \sim S_n$), and output $g \cdot X$. The output is distributed as $\bar{q}$, is exactly G-invariant, and by equation 2 is provably at least as close to any invariant target as q, with an exact improvement of D_{sym}.

The correction has a precise population interpretation. Symmetrization removes the defect D_{sym} only; the residual $\mathrm{KL}(\bar{q} \,\|\, p)$ is unchanged, so it makes the model exactly invariant, not exactly correct. It also requires a G-invariant target. For a pose-specific target the population identity does not apply, and symmetrization may hurt (Proposition 9, App. C).

5 Theory

Theorem 2 (Pythagorean lower bound)**.** Let G be finite or compact with normalized Haar probability measure and act by measure-preserving maps. Let p be G-invariant, and let q have orbit-average $\bar{q}$. If the displayed log ratios are integrable and the KL terms are finite, then

$$\mathrm{KL}(q \,\|\, p) = \mathrm{KL}(q \,\|\, \bar{q}) + \mathrm{KL}(\bar{q} \,\|\, p), \tag{3}$$

and consequently $D_{\mathrm{sym}}(q) = \mathrm{KL}(q \,\|\, \bar{q}) \le \mathrm{KL}(q \,\|\, p)$, with equality iff $\bar{q} = p$. Moreover $D_{\mathrm{sym}}(q) = H(\bar{q}) - H(q) \ge 0$, the entropy gained by symmetrization.

Proof sketch. Write $\mathrm{KL}(q \,\|\, p) - \mathrm{KL}(q \,\|\, \bar{q}) = \mathbb{E}_q[\log \bar{q} - \log p]$. Both $\log p$ and $\log \bar{q}$ are G-invariant: p by assumption and $\bar{q}$ by construction. For any G-invariant f we have $\mathbb{E}_q[f] = \mathbb{E}_{\bar{q}}[f]$ because $\bar{q}$ is

the orbit-average of q. Hence $\mathbb{E}_q[\log \bar{q} - \log p] = \mathrm{KL}(\bar{q} \,\|\, p) \geq 0$, giving equation 3 and the bound. The identity $\mathbb{E}_q[\log \bar{q}] = \mathbb{E}_{\bar{q}}[\log \bar{q}]$ gives the entropy form. App. B gives the full proof, verified symbolically and numerically to residual $\sim 10^{-16}$ by a symbolic and Monte-Carlo checker. □

Remark 1 (Scope and limitations of the bound). D_{sym} isolates the symmetry-breaking component of the model's error, exactly the part of $\mathrm{KL}(q \,\|\, p)$ removed by projecting onto the invariant family. A large D_{sym} certifies a large true error; a small D_{sym} certifies only that q is nearly invariant, not that it is close to p, since the remaining leg $\mathrm{KL}(\bar{q} \,\|\, p)$ can be large. The certificate falsifies rather than validates.

Further theory. App. C collects additional properties of D_{sym}, including the identity $D_{\mathrm{sym}} = I(g; Y)$, subgroup monotonicity, product-group separation, and the corrector's deployment scope. In addition, App. D analyzes the orbit-sample estimator: upward bias, delta-method debiasing, a concentration bound, and a bias-cancelling split-sample estimator, all verified numerically.

6 Experiments

We test three claims in sequence: finite-MC estimates agree with the population lower bound when the target density is known; the diagnostic remains informative on free-form diffusion and real 3D densities; and the corrector identity and an empirical model-selection rule can inform deployment. The main text retains the evidence carrying those claims; complete sweeps, estimators, ablations, and implementation details are in App. F.

6.1 Known-density 3D targets

Protocol and result. On a closed-form $\mathrm{SO}(3)$- and S_n-invariant Gaussian-mixture target, anisotropic and per-point scalings inject controlled symmetry breaking. In every finite-MC cell, the outer estimate using an exactly enumerated octahedral/S_n inner average lies below the corresponding known-density KL estimate and is near zero at $\beta{=}0$. Across the nontrivial range, the ratio of estimates is 9.5–52%; sampled-$\mathrm{SO}(3)$ values are reported only as biased inner-orbit estimates anchored by the population-guaranteed subgroup quantity (Fig. 7, top).

6.2 A free-form diffusion generator, certified through its probability-flow ODE likelihood

We train a non-equivariant score model on the same targets and evaluate its probability-flow ODE likelihood (Song et al., 2021). The outer-MC estimate with exact inner subgroup enumeration is 0.037 nats on the symmetric model, rises with the injected bias, and separates rotation from permutation breaking. At $\beta{=}1$, the known-density MC reduction is 0.99 nats versus a sampled-$\mathrm{SO}(3)$ defect estimate of 1.01 (Fig. 7, top).

From synthetic to real data. We then train the same free-form per-point diffusion on real ShapeNet airplane points (Table 6, Fig. 4). The canonical model fires only for rotation, and its corrector improves held-out invariant-target NLL $5.03 \to -2.20$ nats/point. Rotation augmentation reduces the residual defect by $\sim 5900\times$ and makes the same corrector a near-no-op.

D_{sym} versus function-space equivariance error. A natural alternative is a function-space equivariance error (FEE) (Gruver et al., 2023). Figure 5 and Table 8 compare each model under its matching group, using the saved $\mathrm{FEE}^{\mathrm{perm}}$ for S_n-broken models. FEE responds to both types of violation, but it has no population KL lower-bound interpretation or KL calibration; the population D_{sym} does.

6.3 Point-cloud flows: controlled shapes and a real-data density

We next test tractable-likelihood point-cloud flows on controlled shapes and real ShapeNet points.

Controlled synthetic shapes. Clouds of $K{=}6$ points, of dimension $d{=}18$, are formed by randomly rotating an asymmetric template. As orientation augmentation is removed, $\widehat{D}_{\mathrm{sym}}$ rises $1.36 \to 330.60$ nats with the sliced-Wasserstein proxy. A PCA-canonicalized orbit-invariant log-score sanity check remains $\leq 10^{-3}$; because canonicalization is many-to-one, it is not a normalized density and is not reported as D_{sym} or KL (Fig. 2d).

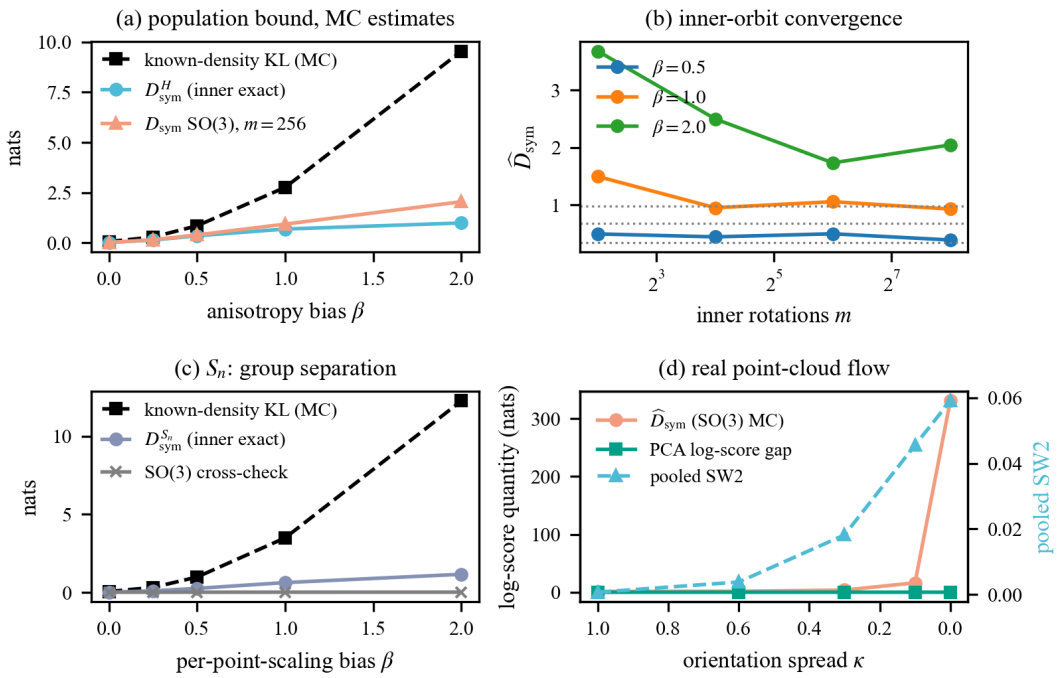

Figure 2: Controlled and real 3D validation. Finite outer-MC estimates with exact inner subgroup enumeration agree with the population bound; sampled inner orbits converge; and a real point-cloud-flow estimate co-varies with a quality proxy. The PCA-canonicalized curve is an orbit-invariant log-score sanity check.

Real ShapeNet airplane point clouds. We next fit an exact-likelihood, permutation-invariant density to real airplanes. It reads $9.2\times10^{-5}\approx 0$ on the genuine permutation symmetry and gives a sampled full-$\mathrm{SO}(3)$ estimate of 1.12 nats/point for broken rotation symmetry. A separate bounded-angle sweep is only an augmentation-spread proxy, not a Haar/group quantity or guarantee (Table 10).

A symmetry the data genuinely has. The reading above is a distance to a rotation-augmented target, so we also retrain the airplane flow under full-$\mathrm{SO}(3)$ augmentation, which makes that symmetry genuinely present in the data. The residual octahedral estimate then falls to 0.77 from the canonical model's saturated 3.18 (Table 13).

6.4 From detection to correction: deploying the symmetrized sampler

The orbit-average corrector needs no retraining. On known-density targets, finite-MC estimates of reduction and D_{sym} agree closely (β=1: 1.03 vs. 1.05; Fig. 6). On the canonical (non-augmented) real-airplane flow, a whole-cloud shared-rotation logmeanexp changes Haar-target NLL $8.46\to3.68$ nats/point; the permutation control is a no-op.

6.5 A pretrained diffusion-transformer conformer generator

On pretrained DiTMC (Frank et al., 2025), we certify the invertible $t_{\max}$=0.97 ODE marginal with an exact 24-element inner orbit. The non-equivariant apeB backbone gives a finite-outer-sample estimate of 0.307 nats (95% CI $[0.151, 0.462]$), while the matched equivariant pe3B and permutation controls read zero (Table 1).

6.6 A real pretrained per-shape generator: diagnosing and correcting PointFlow

For sampled latent code z, PointFlow (Yang et al., 2019) exposes a tractable conditional decoder density q_z. With exact inner subgroup enumeration, its finite outer estimate saturates at 3.18 nats/cloud,

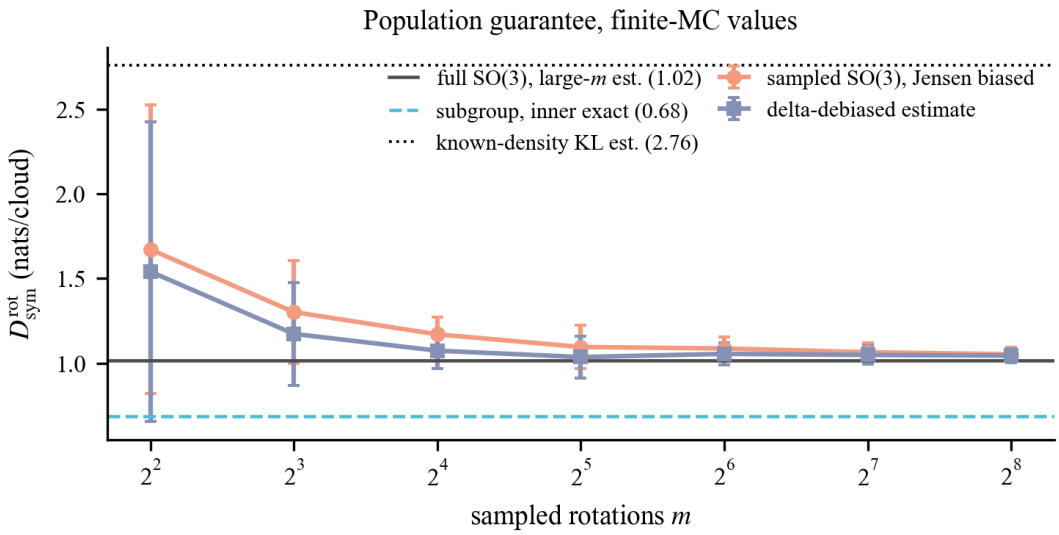

Figure 3: Inner-exact anchor versus sampled-orbit estimation. The population subgroup quantity is guaranteed; its plotted value still uses finite outer samples. The sampled-$\mathrm{SO}(3)$ inner average is Jensen-biased, while delta-method debiasing approaches the full-group estimate with shrinking uncertainty.

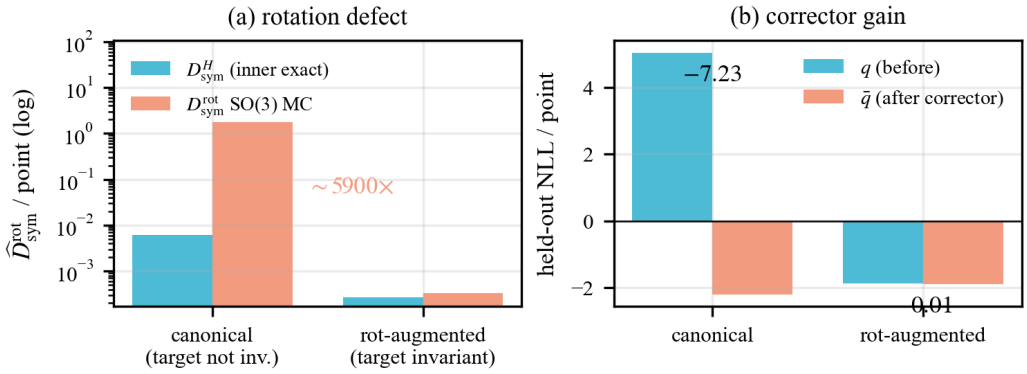

Figure 4: Real-data diffusion follows the controlled prediction. Rotation augmentation collapses the residual defect, while the corrector gives a large NLL gain only for the non-invariant canonical model.

while the permutation control is 8×10^{-4}; the same saturation and corrector response hold across the airplane, chair, and car categories (Table 14). The 1-NNA, MMD, and conditional log-score changes compare PointFlow samples with a rotation-augmented reference generated by PointFlow itself (Fig. 7, bottom). They are internal orbit-self-consistency diagnostics, not standard real-data quality improvements or evidence about the marginal latent-variable density.

An empirical reference-free selection rule. Across twelve separately trained candidates varying in augmentation, capacity, and seed, the tested rule ranks against the hidden known-density KL estimate at Spearman 0.97 and yields 0.003-nat estimated regret (worst case 2.89). Within a fixed augmentation level, D_{sym} cannot see all fit variation, so a separate fit check is required (App. Fig. 16).

The certificate has a deliberately limited interpretation: D_{sym} is a lower bound, so a near-zero reading certifies near-invariance rather than closeness to the target, and the corrector inherits that boundary. The tractable-likelihood requirement, the estimator's Jensen bias for continuous $\mathrm{SO}(3)$, and each experiment's reach are stated in App. I.

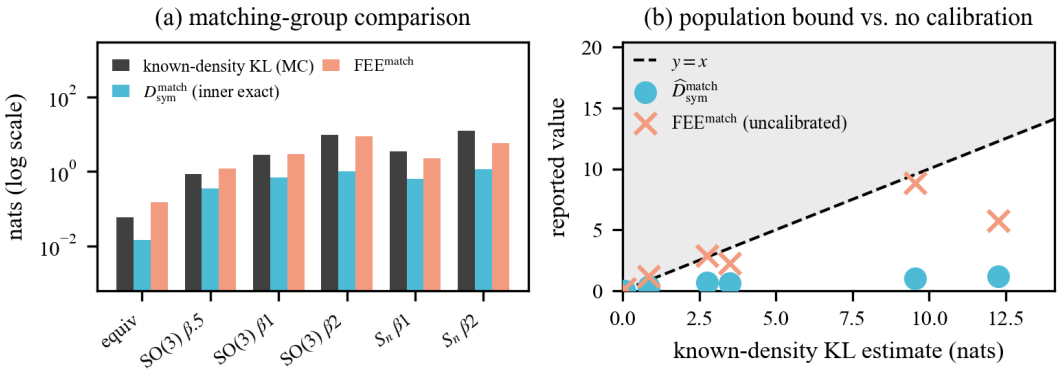

Figure 5: Matching-group distributional and function-space quantities. Each family uses rotation FEE or permutation FEE as appropriate. Population D_{sym} has a KL lower-bound interpretation (plotted values are finite outer-MC estimates); FEE has neither a KL lower bound nor KL calibration.

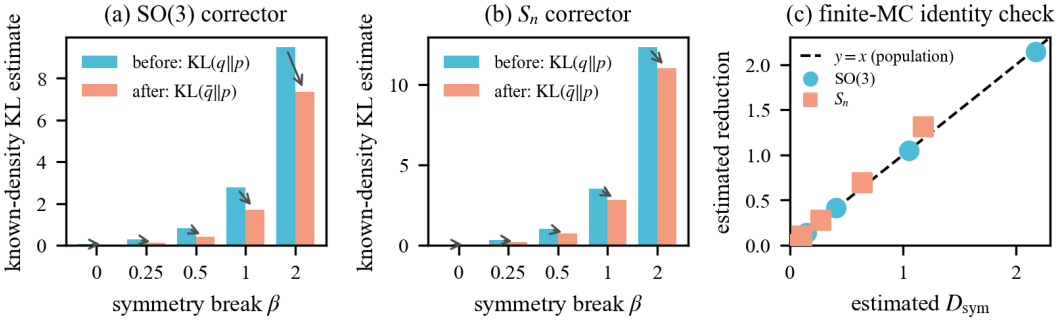

Figure 6: Finite-MC check of the population corrector identity. For both rotation and permutation, known-density KL estimates decrease in every tested cell, and estimated reduction lies near the estimated D_{sym} parity line. Exact equality is the population proposition.

7 Conclusion

As 3D generative modeling trades exact equivariance for scale, residual symmetry error warrants an explicit audit. The symmetry defect $D_{\mathrm{sym}} = \mathrm{KL}(q \,\|\, \bar{q})$ is a reference-free lower bound on the error to any invariant target, computable from the model likelihood alone, and the symmetrized sampler removes exactly that component at no training cost. On a free-form diffusion model trained on real ShapeNet points, the corrector improves held-out invariant-target NLL from 5.03 to -2.20 nats per point; on pretrained DiTMC, the audit measures a 0.307-nat defect for the non-equivariant backbone and zero for its equivariant counterpart; and on a twelve-model pool, the score ranks candidates against hidden KL at Spearman 0.97.

The practical conclusion is narrower than a generic quality guarantee. The estimation audit separates population identities from finite outer-MC estimates and from the Jensen bias of sampled continuous-group averages, while the PointFlow experiment supports a conditional-decoder diagnosis rather than a marginal-density claim. These boundaries should accompany any deployment decision based on the defect.

AI use statement

Generative AI tools were used for manuscript polishing and typesetting, organizing code and experimental results, and code standardization. The authors reviewed all AI-assisted work and take responsibility for the final manuscript.

Table 1: Pretrained-model diagnostics. Population subgroup guarantees are evaluated by finite outer Monte Carlo with exact inner octahedral enumeration; only the displayed DiTMC interval is a reported CI.

model / check	subgroup estimate	negative control	scope
DiTMC apeB (t=0.97)	0.307 [0.151, 0.462]	pe3B -0.001 ± 0.007	ODE marginal
PointFlow conditional q_z	$3.18 = \log 24$	perm. 8×10^{-4}	internal conditional diagnosis

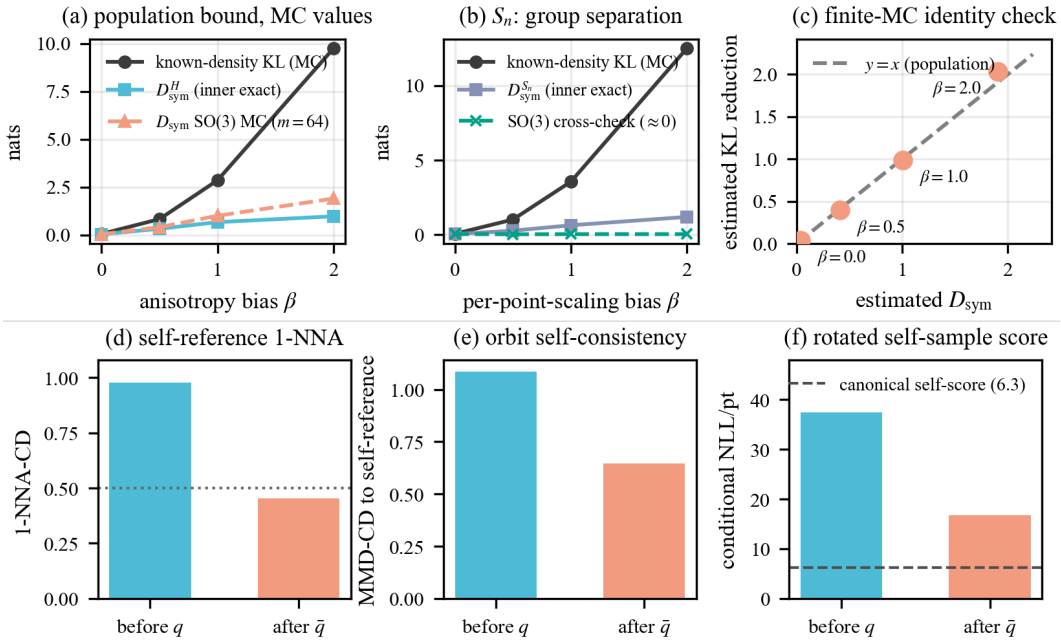

Figure 7: Known-density validation and conditional diagnosis. Top: finite-MC estimates on a known-density diffusion model agree with the population subgroup bound and corrector identity. Bottom: conditional-decoder diagnostics on pretrained PointFlow (Sec. 6.6).

Reproducibility Statement

The complete proof of Theorem 2, the properties of the symmetrized sampler, and the orbit-estimator analysis are in Apps. B, C, and D. Fixed seeds, settings, implementation details, and experiment-specific limitations are documented in Apps. G and I. DiTMC and PointFlow use public checkpoints; all other reported analyses use the saved runs and controls described in the appendix.

References

Josh Abramson et al. Accurate structure prediction of biomolecular interactions with AlphaFold3. *Nature*, 2024.

Shun-ichi Amari. *Information Geometry and Its Applications*. Springer, 2016.

Jeremiah Birrell, Markos A. Katsoulakis, Luc Rey-Bellet, and Wei Zhu. Structure-preserving GANs. *International Conference on Machine Learning (ICML)*, 2022.

Kenny Chiu and Benjamin Bloem-Reddy. Non-parametric hypothesis tests for distributional group symmetry. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Congyue Deng, Or Litany, Yueqi Duan, Adrien Poulenard, Andrea Tagliasacchi, and Leonidas J. Guibas. Vector neurons: A general framework for SO(3)-equivariant networks. In *International Conference on Computer Vision (ICCV)*, 2021.

Yuhui Ding and Thomas Hofmann. Scalable non-equivariant 3d molecule generation via rotational alignment. In *International Conference on Machine Learning (ICML)*, 2025.

J. Thorben Frank, Winfried Ripken, Gregor Lied, Klaus-Robert Müller, Oliver T. Unke, and Stefan Chmiela. Sampling 3d molecular conformers with diffusion transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

Gilad Gour, Iman Marvian, and Robert W. Spekkens. Measuring the quality of a quantum reference frame: The relative entropy of frameness. *Physical Review A*, 80(1):012307, 2009. doi: 10.1103/PhysRevA.80.012307.

Nate Gruver, Marc Finzi, Micah Goldblum, and Andrew Gordon Wilson. The lie derivative for measuring learned equivariance. In *International Conference on Learning Representations (ICLR)*, 2023.

Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3D. *ICML*, 2022.

Sékou-Oumar Kaba, Arnab Kumar Mondal, Yan Zhang, Yoshua Bengio, and Siamak Ravanbakhsh. Equivariance with learned canonicalization functions. In *International Conference on Machine Learning (ICML)*, 2023.

Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: Exact likelihood generative learning for symmetric densities. In *International Conference on Machine Learning (ICML)*, 2020.

Nick W. Koning. Measuring evidence against exchangeability and group invariance with E-values. *arXiv preprint arXiv:2310.01153*, 2023. URL https://arxiv.org/abs/2310.01153.

Hannah Lawrence, Elyssa Hofgard, Vasco Portilheiro, Yuxuan Chen, Tess Smidt, and Robin Walters. To augment or not to augment? diagnosing distributional symmetry breaking, 2025. ICLR 2026.

Shentong Mo, Enze Xie, Ruihang Chu, Lanqing Hong, Matthias Niessner, and Zhenguo Li. DiT-3D: Exploring plain diffusion transformers for 3d shape generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

David A. Moore. Symmetrized variational inference. Workshop paper presented at the NIPS 2016 Workshop on Advances in Approximate Bayesian Inference, 2016.

Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022.

Dmitri S. Pavlichin. Nearest symmetric distributions. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pp. 1630–1634, 2015. doi: 10.1109/ISIT.2015.7282732.

Omri Puny, Matan Atzmon, Heli Ben-Hamu, Ishan Misra, Aditya Grover, Edward J Smith, and Yaron Lipman. Frame averaging for invariant and equivariant network design. In *International Conference on Learning Representations (ICLR)*, 2022.

Víctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n) equivariant graph neural networks. In *ICML*, 2021.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *ICLR*, 2021.

Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.

Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. PointFlow: 3d point cloud generation with continuous normalizing flows. *International Conference on Computer Vision (ICCV)*, 2019.

Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. LION: Latent point diffusion models for 3d shape generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Leo Zhang, Kianoosh Ashouritaklimi, Yee Whye Teh, and Rob Cornish. SymDiff: Equivariant diffusion via stochastic symmetrisation. In *International Conference on Learning Representations (ICLR)*, 2025.

A NOTATION

Table 2 collects the symbols used in the main text and appendices.

Table 2: Notation used in the main text and appendices.

symbol	meaning
G, $g\cdot X$	symmetry group (finite, or compact with normalized Haar measure), acting on $X \in \mathbb{R}^{3N}$ by measure-preserving linear maps ($\lvert\det g\rvert$=1)
p	G-invariant target density
q	generative-model density with tractable likelihood
$\bar{q}(X) = \mathbb{E}_{g\sim G}[q(g\cdot X)]$	orbit average of the model (its symmetrization; m-projection onto the invariant family)
$D_{\mathrm{sym}}(q) = \mathrm{KL}(q \,\|\, \bar{q})$	symmetry defect; provable lower bound on $\mathrm{KL}(q \,\|\, p)$
$\mathrm{KL}(q \,\|\, p) = D_{\mathrm{sym}}(q) + \mathrm{KL}(\bar{q} \,\|\, p)$	exact decomposition (Theorem 2)
$R = D_{\mathrm{sym}}/\mathrm{KL}(q \,\|\, p) \in [0,1]$	tightness ratio of the lower bound
H	octahedral rotation subgroup ($\lvert H\rvert$=24), enumerated exactly in the inner average
D^H_{sym}, $\widehat{D}^H_{\mathrm{sym}}$	population subgroup defect (a valid, looser lower bound) and its finite outer-MC estimate
m	number of sampled inner rotations for continuous $\mathrm{SO}(3)$
inner / outer	inner: estimate of the orbit average at a fixed X (exact enumeration or m samples); outer: sample mean over $X \sim q$
S_n	permutation group on the n points
β	controlled symmetry-break strength in the known-density experiments
κ	orientation spread of the real point-cloud flow (smaller κ means stronger breaking)
FEE	function-space equivariance error (Gruver et al., 2023), a map-level diagnostic without KL calibration
NLL/pt, SW2	negative log-likelihood per point; sliced Wasserstein-2 distance
MMD-CD, 1-NNA-CD	maximum mean discrepancy and 1-NN accuracy under Chamfer distance

B FULL PROOF OF THEOREM 2

Let G be finite or compact with normalized Haar probability measure, act on $\mathbb{R}^{3N}$ by measure-preserving linear maps ($\lvert\det g\rvert$=1), and let p be G-invariant. Let q have orbit-average $\bar{q}(X) = \mathbb{E}_{g\sim G}[q(g\cdot X)]$. We assume the log ratios below are integrable and the displayed KL terms are finite. Then $\bar{q}$ is a normalized density: $\int \bar{q} = \mathbb{E}_g \int q(g\cdot X)\, dX = \mathbb{E}_g 1 = 1$ by measure preservation. $\bar{q}$ is G-invariant: $\bar{q}(h\cdot X) = \mathbb{E}_g[q(gh\cdot X)] = \mathbb{E}_{g'}[q(g'\cdot X)] = \bar{q}(X)$ by invariance of the group measure.

Step 1 (invariant functions have equal expectation under q and $\bar{q}$). For any G-invariant f,

$$\mathbb{E}_{\bar{q}}[f] = \int f(X)\, \mathbb{E}_g[q(g\cdot X)]\, dX = \mathbb{E}_g \int f(X)\, q(g\cdot X)\, dX$$
$$= \mathbb{E}_g \int f(g^{-1}\cdot Y)\, q(Y)\, dY = \mathbb{E}_g\, \mathbb{E}_q[f] = \mathbb{E}_q[f],$$

using the substitution $Y = g\cdot X$ ($\lvert\det g\rvert$=1) and $f(g^{-1}\cdot Y) = f(Y)$.

Step 2 (decomposition). Both $\log p$ and $\log \bar{q}$ are G-invariant. Hence

$$\mathrm{KL}(q \,\|\, p) - \mathrm{KL}(q \,\|\, \bar{q}) = \mathbb{E}_q[\log \bar{q} - \log p] = \mathbb{E}_{\bar{q}}[\log \bar{q} - \log p] = \mathrm{KL}(\bar{q} \,\|\, p) \geq 0,$$

where the middle equality applies Step 1 to the invariant function $\log \bar{q} - \log p$. This is equation 3, and nonnegativity of $\mathrm{KL}(\bar{q} \,\|\, p)$ gives $D_{\mathrm{sym}} \leq \mathrm{KL}(q \,\|\, p)$, with equality iff $\mathrm{KL}(\bar{q} \,\|\, p)$=0, i.e. $\bar{q}$=p.

Step 3 (entropy form). $D_{\mathrm{sym}} = \mathbb{E}_q[\log q - \log \bar{q}] = -H(q) - \mathbb{E}_q[\log \bar{q}] = -H(q) - \mathbb{E}_{\bar{q}}[\log \bar{q}] = H(\bar{q}) - H(q) \geq 0$ by Step 1 and concavity of entropy under the mixture $\bar{q}$.

Verification. The identity, the entropy form, and the lower bound are checked symbolically (a multi-orbit finite group action) and numerically (finite and continuous Monte-Carlo) by a dedicated checker; the maximum identity residual is $\sim 10^{-16}$ and the bound holds in all trials.

C Properties of the symmetry defect

We collect properties of D_{sym} used in the main text. All reuse the orbit-average identity of Appendix B (for any G-invariant f, $\mathbb{E}_q[f] = \mathbb{E}_{\bar{q}}[f]$) and are checked numerically by the same checker.

Proposition 3 (Non-negativity and exactness)**.** $D_{\mathrm{sym}}(q) \geq 0$, with $D_{\mathrm{sym}}(q) = 0$ if and only if q is G-invariant ($q = \bar{q}$ almost everywhere).

Proof. $D_{\mathrm{sym}} = \mathrm{KL}(q \,\|\, \bar{q}) \geq 0$ with equality iff $q = \bar{q}$. If q is G-invariant then $\bar{q} = \mathbb{E}_g[q(g\cdot)] = q$; conversely $\bar{q}$ is always G-invariant, so $q = \bar{q}$ implies q is invariant. □

Proposition 4 (Information-theoretic identity)**.** Draw g from the group measure and $X \sim q$ independently, and set $Y = g \cdot X$. Then the marginal of Y is $\bar{q}$, and

$$D_{\mathrm{sym}}(q) = I(g; Y),$$

the mutual information between the (uniform) group element and the generated sample. Thus D_{sym} is exactly the amount of group information (orientation under $\mathrm{SO}(3)$, ordering under S_n) that the model's output carries, and is zero iff the output is statistically independent of the group element, i.e. q is invariant.

Proof. The conditional density of Y given g is $q_g(Y) = q(g^{-1} \cdot Y)$ (measure preservation), whose mixture over g is $\mathbb{E}_g[q(g^{-1} \cdot Y)] = \bar{q}(Y)$. Hence $I(g; Y) = \mathbb{E}_g \,\mathrm{KL}(q_g \,\|\, \bar{q})$. For each g, substituting $X = g^{-1} \cdot Y$ and using $\bar{q}(g \cdot X) = \bar{q}(X)$, $\mathrm{KL}(q_g \,\|\, \bar{q}) = \int q(X) \log \frac{q(X)}{\bar{q}(g \cdot X)}\, dX = \mathrm{KL}(q \,\|\, \bar{q}) = D_{\mathrm{sym}}$. Averaging over g gives $I(g; Y) = D_{\mathrm{sym}}$. □

Proposition 5 (Subgroup monotonicity / nesting)**.** For a subgroup $H \leq G$ with orbit-averages $\bar{q}_H$ and $\bar{q}_G$,

$$D^G_{\mathrm{sym}}(q) = D^H_{\mathrm{sym}}(q) + \mathrm{KL}(\bar{q}_H \,\|\, \bar{q}_G) \;\geq\; D^H_{\mathrm{sym}}(q),$$

so any subgroup average is a valid (possibly looser) sound lower bound: $D^H_{\mathrm{sym}} \leq D^G_{\mathrm{sym}} \leq \mathrm{KL}(q \,\|\, p)$.

Proof. Averaging over G factors as averaging over H followed by averaging over coset representatives, so $\bar{q}_G$ is the G-average of $\bar{q}_H$ and in particular is H-invariant. Apply Theorem 2 to the H-action with model q and H-invariant target $\bar{q}_G$: $\mathrm{KL}(q \,\|\, \bar{q}_G) = \mathrm{KL}(q \,\|\, \bar{q}_H) + \mathrm{KL}(\bar{q}_H \,\|\, \bar{q}_G)$. The left side is D^G_{sym} and the first right-hand term is D^H_{sym}; nonnegativity of the second gives the inequality. The chain with Theorem 2 ($D^G_{\mathrm{sym}} \leq \mathrm{KL}(q \,\|\, p)$) completes the bound. □

Proposition 6 (Product groups: separation)**.** If $G = G_1 \times G_2$ (e.g. $\mathrm{SO}(3) \times S_n$ acting on the two factors), then $D^G_{\mathrm{sym}} \geq \max(D^{G_1}_{\mathrm{sym}}, D^{G_2}_{\mathrm{sym}})$, and each single-factor defect isolates one symmetry: a model that breaks only G_2 has $D^{G_1}_{\mathrm{sym}} \approx 0$.

Proof. G_1 and G_2 are subgroups of G; apply Proposition 5 to each. If q is G_1-invariant then $D^{G_1}_{\mathrm{sym}} = 0$ by Proposition 3. □

This is the formal basis for the group-separation column of Table 5: an S_n-breaking bias leaves $D^{\mathrm{SO}(3)}_{\mathrm{sym}} \approx 0$ because the model remains (near-)rotation-invariant.

Proposition 7 (The corrector generalizes to every f-divergence)**.** Let D_f be any f-divergence (f convex, $f(1) = 0$) and let p be G-invariant. Then for any model q, $D_f(\bar{q} \,\|\, p) \leq D_f(q \,\|\, p)$: symmetrizing never increases the divergence to an invariant target, for any f-divergence, so the corrector (Prop. 1) is not special to KL.

Proof. D_f is jointly convex in its argument pair. Writing $q \circ g$ for the pushforward of q by g, each g is measure-preserving and p is G-invariant, so $D_f(q \circ g \,\|\, p) = D_f(q \,\|\, p)$; also $p = \mathbb{E}_g[p \circ g]$ and $\bar{q} = \mathbb{E}_g[q \circ g]$. Joint convexity (Jensen over g) gives $D_f(\bar{q} \,\|\, p) = D_f\big(\mathbb{E}_g[q \circ g] \,\big\|\, \mathbb{E}_g[p \circ g]\big) \le \mathbb{E}_g\, D_f(q \circ g \,\|\, p \circ g) = \mathbb{E}_g\, D_f(q \,\|\, p) = D_f(q \,\|\, p)$. □

Remark 2 (The detection bound is KL-specific). The lower bound $D_{\mathrm{sym}} = \mathrm{KL}(q \,\|\, \bar{q}) \le \mathrm{KL}(q \,\|\, p)$ used for detection rests on the Pythagorean equality equation 3, which is special to KL (the m-projection); for a general f-divergence only the corrector direction (Prop. 7) survives. We therefore state detection in KL, where the orbit average is simultaneously the optimal corrector and the source of the bound.

Proposition 8 (Tightness)**.** The tightness ratio obeys $R := D_{\mathrm{sym}}/\mathrm{KL}(q \,\|\, p) = \big(1 + \mathrm{KL}(\bar{q} \,\|\, p)/D_{\mathrm{sym}}\big)^{-1} \in [0, 1]$. Hence $R \to 1$ (the bound is tight) exactly when the residual non-symmetry error $\mathrm{KL}(\bar{q} \,\|\, p)$ is small relative to the symmetry defect, i.e. when the model's error is dominated by symmetry breaking, and $R \to 0$ when the non-symmetry residual dominates.

Proof. Immediate from equation 3, $\mathrm{KL}(q \,\|\, p) = D_{\mathrm{sym}} + \mathrm{KL}(\bar{q} \,\|\, p)$. □

Proposition 9 (Symmetrization under an invariant target)**.** Let q be a model with orbit-average $\bar{q}$. (i) Invariant target. If the deployment target p is G-invariant, then by equation 2 $\mathrm{KL}(\bar{q} \,\|\, p) = \mathrm{KL}(q \,\|\, p) - D_{\mathrm{sym}}(q) \le \mathrm{KL}(q \,\|\, p)$: symmetrizing never increases the error and removes exactly D_{sym}. (ii) Pose-specific target. If instead the target p' is not G-invariant, such as a fixed canonical pose or ordering, the Pythagorean comparison does not apply and there is no universal ordering between $\mathrm{KL}(q \,\|\, p')$ and $\mathrm{KL}(\bar{q} \,\|\, p')$. Symmetrization may therefore hurt. For example, taking $q = p'$ gives zero error before symmetrization, while $\mathrm{KL}(\bar{q} \,\|\, p') > 0$ whenever that KL is finite and $\bar{q} \neq p'$ almost everywhere.

Proof. (i) is equation 2. For (ii), the invariant-target step in Theorem 2 is unavailable. The example follows from KL exactness: $\mathrm{KL}(p' \,\|\, p') = 0$ while $\mathrm{KL}(\overline{p'} \,\|\, p') > 0$ under the stated conditions. □

D Estimator analysis: bias, debiasing, and concentration

The certificate equation 1 replaces $\log \bar{q}$ by $\mathrm{logmeanexp}$ over an inner orbit sample and the outer expectation by a finite sample mean. We analyze the inner estimator below; exact subgroup enumeration eliminates this inner approximation but never the outer-MC uncertainty. Fix a source point X; let $a_j = q(g_j \cdot X)$ for m i.i.d. group draws g_j, so $\mathbb{E}[a_j] = \bar{q}(X)$ and $\widehat{\bar{q}} = \frac{1}{m}\sum_j a_j$ is unbiased for $\bar{q}(X)$, while $\mathrm{logmeanexp}_j \log a_j = \log \widehat{\bar{q}}$.

Proposition 10 (Upward bias and delta-method debiasing)**.** $\log \widehat{\bar{q}}$ is downward biased,

$$\mathbb{E}[\log \widehat{\bar{q}}] = \log \bar{q}(X) - \frac{\sigma^2(X)}{2m\,\bar{q}(X)^2} + O(m^{-2}), \qquad \sigma^2(X) = \mathrm{Var}_g\big[q(g \cdot X)\big],$$

so the plug-in $\widehat{D}_{\mathrm{sym}}$ over-estimates D_{sym} by $\frac{1}{2m}\mathbb{E}_{X \sim q}[\sigma^2/\bar{q}^2] + O(m^{-2}) \ge 0$. Adding the in-sample estimate $\frac{1}{2m}\,\widehat{\sigma}^2/\widehat{\bar{q}}^{\,2}$ to $\log \widehat{\bar{q}}$ removes this leading $O(1/m)$ bias.

Proof. Second-order Taylor of $\log$ about $\bar{q}$: $\log \widehat{\bar{q}} = \log \bar{q} + \frac{\widehat{\bar{q}} - \bar{q}}{\bar{q}} - \frac{(\widehat{\bar{q}} - \bar{q})^2}{2\bar{q}^2} + O((\widehat{\bar{q}} - \bar{q})^3)$. Take expectations using $\mathbb{E}[\widehat{\bar{q}} - \bar{q}] = 0$ and $\mathbb{E}[(\widehat{\bar{q}} - \bar{q})^2] = \sigma^2/m$. The sign makes $\widehat{D}_{\mathrm{sym}} = \log q - \mathbb{E}[\log \widehat{\bar{q}}]$ an over-estimate. □

Proposition 11 (Concentration and sample complexity)**.** Suppose $q(g \cdot X) \in [0, B]$ and $\bar{q}(X) \ge c > 0$ on the support. For m i.i.d. group draws and any $\delta \in (0, 1)$, with probability at least $1 - \delta$, writing $t = B\sqrt{\log(2/\delta)/(2m)}$, whenever $t \le c/2$,

$$\big|\log \widehat{\bar{q}}(X) - \log \bar{q}(X)\big| \;\le\; \frac{2t}{c} \;=\; \frac{2B}{c}\sqrt{\frac{\log(2/\delta)}{2m}}.$$

Hence $m = O\big(\frac{B^2}{c^2\varepsilon^2}\log\frac{1}{\delta}\big)$ orbit samples give an ε-accurate $\log \bar{q}$ and thus an ε-accurate $\widehat{D}_{\mathrm{sym}}$.

Proof. Hoeffding's inequality for the mean of $[0, B]$-bounded i.i.d. variables gives $|\widehat{\bar{q}} - \bar{q}| \le t$ with probability $\ge 1 - \delta$. Writing $u = (\widehat{\bar{q}} - \bar{q})/\bar{q}$, we have $|u| \le t/c \le \frac{1}{2}$, and $|\log(1 + u)| \le 2|u|$ for $|u| \le \frac{1}{2}$, so $|\log \widehat{\bar{q}} - \log \bar{q}| \le 2t/c$. □

Proposition 12 (Bias-cancelling split-sample estimator)**.** Let $\widehat{\bar{q}}^{(m)}$ and $\widehat{\bar{q}}^{(m/2)}$ be the orbit-sample means over m and $m/2$ draws. Since the bias of $\log \widehat{\bar{q}}^{(k)}$ is $-C(X)/k + O(k^{-2})$ with $C(X) = \sigma^2(X)/(2\bar{q}(X)^2)$ (Prop. 10), the Richardson combination $2\log \widehat{\bar{q}}^{(m)} - \log \widehat{\bar{q}}^{(m/2)}$ has bias $O(m^{-2})$, and therefore

$$\widetilde{D}_{\mathrm{sym}} = \log q(X) - \left(2\log \widehat{\bar{q}}^{(m)} - \log \widehat{\bar{q}}^{(m/2)}\right)$$

estimates D_{sym} with $O(m^{-2})$ bias, requiring no variance estimate.

Proof. The bias of $2\log \widehat{\bar{q}}^{(m)} - \log \widehat{\bar{q}}^{(m/2)}$ is $2(-C/m) - (-2C/m) + O(m^{-2}) = O(m^{-2})$, cancelling the leading term of Prop. 10. □

D.1 Population guarantees and finite-sample reports

The exact population subgroup quantity for finite H (here octahedral, $|H|$=24) is $D^H_{\mathrm{sym}} = \mathrm{KL}(q \,\|\, \bar{q}_H)$ and obeys $D^H_{\mathrm{sym}} \le D^{\mathrm{SO}(3)}_{\mathrm{sym}} \le \mathrm{KL}$. Enumerating H removes inner-orbit approximation, but every displayed sample mean still estimates the outer expectation from finitely many $X_i \sim q$. The sampled-$\mathrm{SO}(3)$ $\mathrm{logmeanexp}$ over m rotations is an estimate: it is Jensen-biased, under-estimating $\log \bar{q}$ and hence over-estimating D_{sym}, and is not guaranteed sound; with more rotations it decreases toward the true full-group defect with Monte-Carlo noise (Figure 3). To give the sampled estimator a principled near-unbiased form we add a delta-method variance correction, adding $\frac{1}{2m}\widehat{\mathrm{Var}}(q)/\bar{q}^2$ to $\log \bar{q}$ computed in log-space from the orbit log-likelihoods, which removes the leading $O(1/m)$ bias. On the known-density β=1 model (KL estimate 2.76, full-$\mathrm{SO}(3)$ defect estimate 1.02), the outer-MC estimate with an exact inner octahedral average is 0.68; the raw sampled estimate is 1.67 ± 0.85 at m=4 and decreases to 1.05 ± 0.04 at m=256, while the debiased estimate is consistently closer to the truth, e.g. 1.07 vs 1.17 at m=16, restoring approximate unbiasedness at small m (mean±std over 8 inner-orbit seeds). Guarantee language therefore refers to the population subgroup quantity; numerical values are identified as finite outer-MC estimates, with intervals only where actually computed.

E Additional ablations: group size and orbit estimation

To probe how the certificate behaves as the group and the orbit estimator vary, we run a controlled study over abelian (reflection/sign-flip) and non-abelian (permutation S_n) groups on known-density targets (Figure 8). Group-size / dimension scaling: finite-MC estimates agree with the population inequality in every cell, and the estimated ratio $R = D_{\mathrm{sym}}/\mathrm{KL}$ remains high across group sizes from $|G|$=4 to $|G|$=120 and dimensions up to d=10 ($R \approx 0.94$–0.98 across the sign-flip family and S_3, S_4; $R \approx 0.80$ for the largest case, S_5 at d=10, where the more multimodal target is harder to fit). Orbit estimation: on a group of order 64, both a random m-element orbit sample and a subgroup average approach the full-group defect empirically from below in this run as the number of elements grows (subgroup: $0.20, 0.55, 0.74, 0.99, 1.16, 1.20$ for $|H| = 2, \ldots, 64$; random: $0.35, 0.66, 0.92, 1.03, 1.14, 1.20$ for $m = 2, \ldots, 64$; full-group value 1.20). The population subgroup average is the principled guaranteed object when the full group is too large to enumerate; its displayed value still has outer-sample uncertainty.

F Additional figures and tables

This appendix collects floats deferred from the condensed experiments of Section 6: sample galleries (Figs. 9 and 10), the full known-density diffusion and ShapeNet panels (Figs. 11 and 12), the PointFlow internal diagnosis (Fig. 13), the selection-rule and corrector checks (Figs. 14 and 15), and the accompanying tables (Tables 3, 4, 7, 9, 11, 12, 15, and 16). Nothing here changes any reported number or claim.

G Implementation details

Model. A RealNVP coupling flow (12 affine-coupling layers, hidden width 128) on the flattened point coordinates, trained by maximum likelihood (Adam, 10^{-3}, cosine schedule, 4000 steps, batch

Table 3: Pretrained DiTMC on QM9. The population 24-element-subgroup quantity lower-bounds KL for the likelihood-consistent t=0.97 ODE marginal. Results are finite outer-sample estimates with exact inner enumeration. The displayed apeB interval is clustered over seven molecular conditions; no interval is invented for other rows.

model / check	quantity (nats)	result
apeB (non-equivariant)	$D_{\mathrm{sym}}^{\mathrm{rot},H}$	$\mathbf{0.307}$ $[0.151, 0.462]$
pe3B ($\mathrm{SO}(3)$-equivariant)	$D_{\mathrm{sym}}^{\mathrm{rot},H}$	$-0.0011 \pm 0.0074 \approx 0$
apeB permutation control	$D_{\mathrm{sym}}^{\mathrm{perm}}$	$-0.0021 \pm 0.0090 \approx 0$
apeB reverse ODE steps	$10/20/40$	$0.266/0.265/0.279$

Table 4: Pretrained PointFlow conditional diagnosis (airplane). Top: finite outer estimates of the conditional-decoder defect; octahedral enumeration is exact only for the inner average. Bottom: internal orbit-self-consistency against a rotation-augmented self-reference generated by PointFlow. These are not real-data quality metrics, and the invalid legacy COV result is omitted.

Detection: conditional-decoder symmetry defect		
$\widehat{D}_{\mathrm{sym}}^{\mathrm{rot},H}$ (inner exact; nats/cloud)	$3.18 = \log 24$ (saturated)	
$D_{\mathrm{sym}}^{\mathrm{rot}}$ sampled $\mathrm{SO}(3)$, estimate (nats/cloud)	$\approx 2\times10^{4}$	
$D_{\mathrm{sym}}^{\mathrm{perm}}$ (by construction)	$8\times10^{-4} \approx 0$	
Internal orbit-self-consistency: model-derived reference		
metric	before q	after $\bar{q}$
MMD-CD ↓	1.09	**0.65**
1-NNA-CD ($\rightarrow 0.5$)	0.98	**0.45**
conditional NLL/pt on rotated self-samples ↓	37.5	**16.8**

Table 5: Known-density 3D targets. Every number is a finite outer-MC estimate. "Inner exact" means only that the octahedral or S_n average is enumerated; the corresponding population quantity obeys $D_{\mathrm{sym}} \leq \mathrm{KL}$. In this run each defect estimate lies below its known-density KL estimate.

group	quantity	β=0	β=0.25	β=0.5	β=1.0	β=2.0
$\mathrm{SO}(3)$	$\mathrm{KL}(q\|p)$ (known-density MC)	0.057	0.284	0.846	2.762	9.548
	$\widehat{D}_{\mathrm{sym}}^{H}$ (inner exact)	0.015	0.135	0.348	0.684	0.987
	D_{sym} $\mathrm{SO}(3)$ MC (m=256)	0.017	0.149	0.391	0.929	2.046
	$R = D_{\mathrm{sym}}/\mathrm{KL}$ (MC)	0.29	0.52	0.46	0.34	0.21
	observed $\widehat{D}_{\mathrm{sym}}^{H} \leq \widehat{\mathrm{KL}}$	✓	✓	✓	✓	✓
S_n	$\mathrm{KL}(q\|p)$ (known-density MC)	0.057	0.318	0.982	3.493	12.255
	$\widehat{D}_{\mathrm{sym}}^{S_n}$ (inner exact)	0.012	0.096	0.258	0.628	1.160
	$R = D_{\mathrm{sym}}/\mathrm{KL}$	0.21	0.30	0.26	0.18	0.095
	$\mathrm{SO}(3)$ cross-check (≈ 0)	0.015	0.015	0.013	0.015	0.017
	observed $\widehat{D}_{\mathrm{sym}}^{S_n} \leq \widehat{\mathrm{KL}}$	✓	✓	✓	✓	✓

Table 6: Real-data diffusion diagnostic (ShapeNet airplane; PF-ODE likelihood). No target KL is available. The population octahedral quantity has the lower-bound guarantee, while displayed values are finite outer estimates with exact inner enumeration. The whole-cloud corrector NLL is evaluated on an $\mathrm{SO}(3)$-augmented target. All entries are per point.

quantity	canonical-pose (target not invariant)	rotation-augmented (target genuinely invariant)
held-out NLL / pt	-3.46	-1.88
$\widehat{D}_{\mathrm{sym}}^{\mathrm{rot},H}$ (inner exact)	0.0062	0.0003
$D_{\mathrm{sym}}^{\mathrm{rot}}$ SO(3) MC	$\mathbf{1.764}$ (fires)	$\mathbf{0.0003}$
$D_{\mathrm{sym}}^{\mathrm{perm}}$ (control)	$8\times10^{-8} \approx 0$	≈ 0
corrector NLL/pt on $\mathrm{SO}(3)$-aug. target (before→after)	$5.03 \rightarrow \mathbf{-2.20}$	$-1.88 \rightarrow -1.89$

Table 7: PF-ODE likelihood on known-density 3D targets. All entries are finite outer-MC estimates; "inner exact" denotes subgroup enumeration. The population quantities satisfy the lower bound and exact corrector identity. The run-level estimates agree with those statements within sampling error.

group	quantity	β=0	β=0.5	β=1.0	β=2.0
SO(3)	$\mathrm{KL}(q\|p)$ (known-density MC)	0.057	0.831	2.859	9.749
	$\widehat{D}^{H}_{\mathrm{sym}}$ (inner exact)	0.037	0.324	0.671	0.976
	D_{sym} SO(3) MC (m=64)	0.039	0.409	1.007	1.903
	$R = D_{\mathrm{sym}}/\mathrm{KL}$ (octahedral)	0.65	0.39	0.23	0.10
	$\widehat{\mathrm{KL}}(\bar{q}\|p)$ after	0.019	0.437	1.870	7.716
S_n	$\mathrm{KL}(q\|p)$ (known-density MC)	0.057	1.010	3.553	12.496
	$\widehat{D}^{S_n}_{\mathrm{sym}}$ (inner exact)	0.036	0.264	0.629	1.188
	$R = D_{\mathrm{sym}}/\mathrm{KL}$	0.63	0.26	0.18	0.10
	SO(3) cross-check ($\approx$0)	0.037	0.025	0.033	0.032

Table 8: Matching-group FEE comparison on a known-density target. Rotation-broken rows use rotation FEE and S_n-broken rows use the saved permutation FEE. The population $D^{\mathrm{match}}_{\mathrm{sym}}$ has a KL lower-bound interpretation; displayed values are finite outer estimates with exact inner enumeration. FEE responds to both groups but has no KL lower bound or calibration.

model	matched group	known-density $\widehat{\mathrm{KL}}$	$\widehat{D}^{\mathrm{match}}_{\mathrm{sym}}$	$\mathrm{FEE}^{\mathrm{match}}$
baseline (β=0)	rotation	0.057	0.015	0.150
SO(3)-broken β=0.5	rotation	0.846	0.348	1.175
SO(3)-broken β=1.0	rotation	2.762	0.684	2.867
SO(3)-broken β=2.0	rotation	9.548	0.987	8.873
S_n-broken β=1.0	permutation	3.493	0.628	2.223
S_n-broken β=2.0	permutation	12.255	1.160	5.725

Table 9: Controlled synthetic-shape point-cloud flow. The sampled-$\mathrm{SO}(3)$ defect estimate co-varies with a pooled sliced-Wasserstein proxy. The PCA-canonicalized row is only the orbit spread of the invariant log score $\log q(\mathrm{canon}(X))$; it is not a normalized density, D_{sym}, or KL.

κ (orientation spread)	1.0	0.6	0.3	0.1	0.0
$\widehat{D}_{\mathrm{sym}}$ (SO(3) MC)	1.36	1.73	3.74	16.25	330.60
PCA log-score orbit gap	-9×10^{-8}	-3×10^{-7}	8×10^{-4}	2×10^{-5}	4×10^{-5}
pooled SW2 proxy	0.0007	0.0038	0.0181	0.0455	0.0591

Table 10: Real ShapeNet airplane point-cloud flow (exact-likelihood, permutation-invariant per-point flow). Finite-MC controls are near zero for permutation and large for sampled full-$\mathrm{SO}(3)$ rotation. The bounded-angle sweep is an augmentation-spread proxy: its averaging distribution is not Haar and does not define a group guarantee.

τ (augmentation spread)	0.0	0.1	0.3	0.6	1.0
$A_{\mathrm{spread}}(\tau)$ per cloud (proxy)	0.00	3.99	5.39	5.64	6.59
held-out NLL / point on τ-rotated test	2.47	3.15	4.91	7.35	7.66

Table 11: Finite-MC check of the population corrector identity. On known-density targets, estimated KL decreases in every tested cell and estimated reduction is close to estimated D_{sym}. The exact equality and non-increase are population statements; no unreported confidence interval is implied.

group	quantity	β=0.25	β=0.5	β=1.0	β=2.0
$\mathrm{SO}(3)$	$\widehat{\mathrm{KL}}(q\|p)$ before	0.268	0.833	2.762	9.490
	$\widehat{\mathrm{KL}}(\bar{q}\|p)$ after	0.131	0.418	1.716	7.344
	estimated reduction / $\widehat{D}_{\mathrm{sym}}$	0.138 / 0.147	0.416 / 0.408	1.046 / 1.051	2.146 / 2.171
S_n	$\widehat{\mathrm{KL}}(q\|p)$ before	0.324	1.010	3.521	12.331
	$\widehat{\mathrm{KL}}(\bar{q}\|p)$ after	0.220	0.730	2.831	11.020
	estimated reduction / $\widehat{D}_{\mathrm{sym}}$	0.104 / 0.101	0.280 / 0.272	0.690 / 0.636	1.311 / 1.177

Table 12: Corrector check on a genuinely $\mathrm{SO}(3)$-invariant known-density target. All entries are finite-MC estimates. Their reductions agree with the population identity and with a sample-based SW_2 check; no finite-sample guarantee or unreported interval is asserted.

β	0	0.5	1.0	2.0
$\widehat{\mathrm{KL}}(q\|p)$ before	0.048	0.823	2.764	9.503
$\widehat{\mathrm{KL}}(\bar{q}\|p)$ after	0.038	0.409	1.735	7.480
estimated reduction / $\widehat{D}_{\mathrm{sym}}$	0.010 / 0.010	0.414 / 0.416	1.029 / 1.047	2.023 / 2.105
SW_2 to true samples (before→after)	0.003→0.003	0.061→0.050	0.271→0.236	1.259→1.167

Table 13: Rotation-augmented real-data flow. Defect entries are finite outer estimates; the octahedral inner average is enumerated exactly and its population counterpart has the guarantee. The final column contains a saved per-point augmentation log-score proxy, not a whole-cloud orbit-average corrector and not evidence for Proposition 1.

model	$\widehat{D}^{H}_{\mathrm{sym}}$ (inner exact)	$\widehat{D}^{\mathrm{SO}(3)}_{\mathrm{sym}}$ MC	per-point augmentation proxy
canonical-pose (target not invariant)	3.18= $\log 24$ (saturated)	1.1×10^{3}	8.67 → 4.28
rotation-augmented (target invariant)	**0.77** (graded)	**0.88**	4.191 → 4.190

Table 14: PointFlow conditional diagnosis across categories. Finite outer estimates with exact inner octahedral enumeration saturate at $\log 24$. Other rows are internal comparisons to model-derived rotation-augmented self-references, not real-data generation quality. Legacy COV values are omitted because their nearest-neighbor dimension was wrong.

metric	airplane	chair	car
MMD-CD ↓	1.09 → 0.65	0.80 → 0.70	0.99 → 0.52
1-NNA-CD (→ 0.5)	0.98 → 0.45	0.89 → 0.51	0.99 → 0.49
conditional self-sample NLL/pt ↓	37.5 → 16.8	21.8 → 13.8	32.9 → 13.8
$\widehat{D}^{\mathrm{rot},H}_{\mathrm{sym}}$ (inner exact)	3.18= $\log 24$ (saturated, all three)		

Table 15: Empirical model-selection rule. In this five-candidate pool, finite-MC D_{sym} and known-density KL estimates rank identically and give zero estimated regret. This run supplies no universal selection guarantee.

candidate	known-density $\widehat{\mathrm{KL}}$	$\widehat{D}^{H}_{\mathrm{sym}}$ (inner exact)	$\widehat{D}^{\mathrm{SO}(3)}_{\mathrm{sym}}$ MC
aug-full (β=0)	0.057	0.015	0.017
aug-weak (β=0.25)	0.284	0.135	0.147
aug-mid (β=0.5)	0.846	0.348	0.421
aug-low (β=1)	2.762	0.684	1.034
aug-none (β=2)	9.548	0.987	2.060
this pool: pick = aug-full; estimated regret 0.000 nats (worst 9.49); Spearman 1.0; no general guarantee			

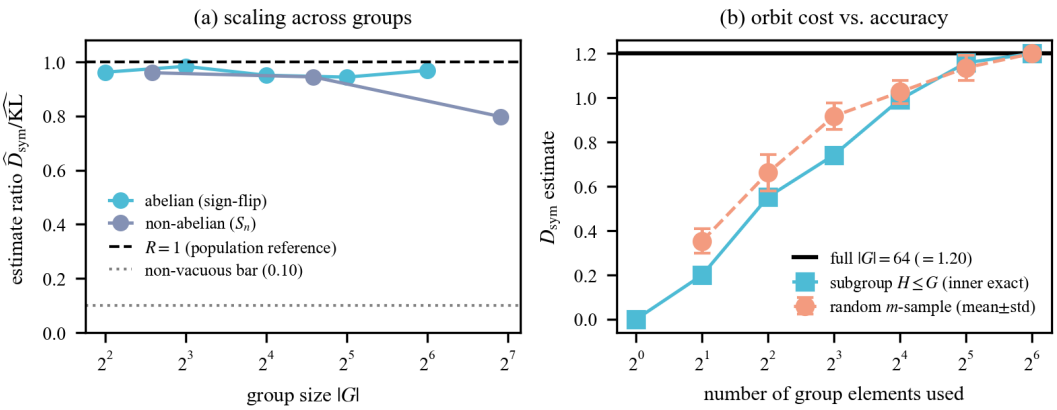

Figure 8: Group-size scaling and orbit estimation (controlled targets). (a) Ratios of finite-MC estimates across groups and dimensions. (b) On a group of order 64, both curves approach the full-inner-group estimate; only the population subgroup quantity has the lower-bound guarantee, and all plotted values retain outer-sample uncertainty.

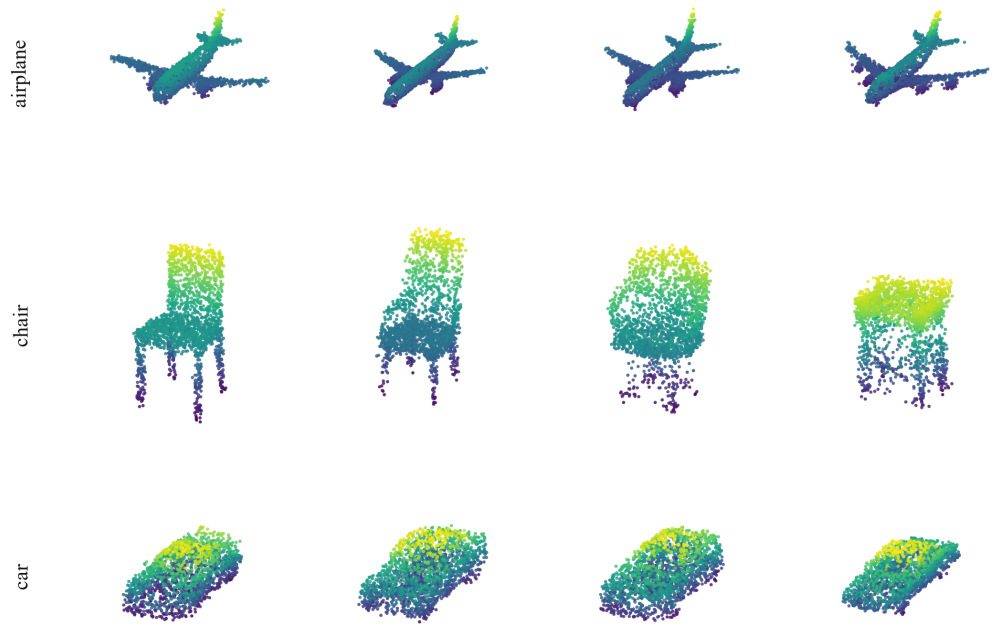

Figure 9: Samples from the pretrained PointFlow generator used only for the conditional-decoder and internal orbit-self-consistency diagnosis.

512) on a single RTX 4080. The known-density targets (N=4, d=12) use a zero-mean Gaussian mixture with exchangeable covariance $M \otimes I_3$ (closed-form density and sampler). The real point-cloud distribution (K=6, d=18) rotates a fixed asymmetric template; orientation spread κ sets the maximum rotation angle.

Groups and orbits. The octahedral group is the 24 signed axis-permutation matrices with determinant $+1$, applied per point. Uniform $\mathrm{SO}(3)$ rotations are sampled from normalized Gaussian quaternions.

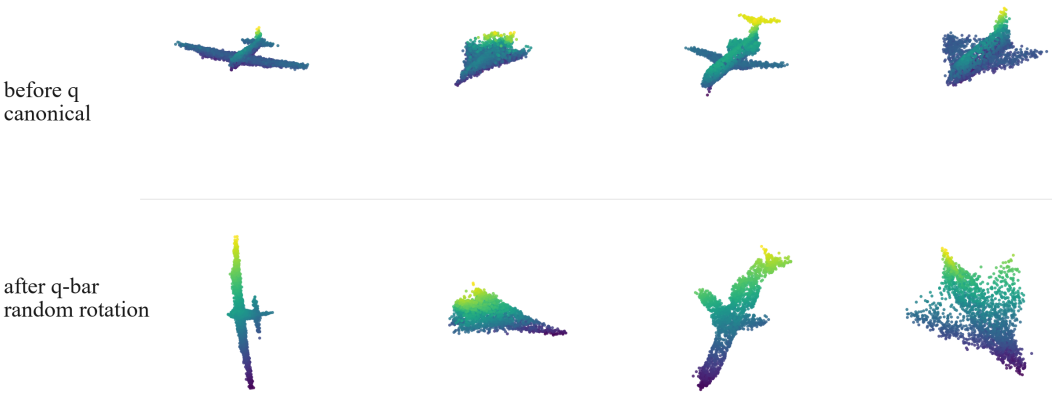

Figure 10: Effect of conditional symmetrization. PointFlow airplane clouds before (top) and after (bottom) one shared random rotation. This visualization documents orbit spreading only; it does not establish improved real-data sample quality.

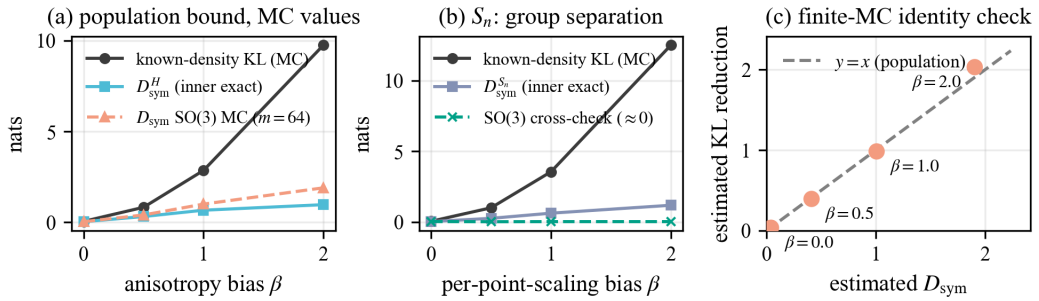

Figure 11: A free-form diffusion model with PF-ODE likelihood. On known-density targets, finite outer-MC estimates using inner subgroup enumeration agree with the population lower bound and corrector identity. Sampled-SO(3) values retain inner Jensen bias; no plotted number is an exact population KL.

Table 16: Empirical selection on a heterogeneous, multi-axis pool. Twelve separately-trained flows varying in augmentation β (symmetry axis) and capacity+seed (non-symmetry fit axis); rows average two seeds. This pool gives Spearman 0.97 and estimated regret 0.003 nats (worst 2.89), but no universal guarantee. Same-β rows expose fit variation invisible to D_{sym}.

β	capacity	known-density $\widehat{\mathrm{KL}}$	$\widehat{D}^{H}_{\mathrm{sym}}$ (inner exact)	$\widehat{D}^{\mathrm{SO}(3)}_{\mathrm{sym}}$ MC
0	small	0.100	0.034	0.036
0	big	0.047	0.009	0.010
0.5	small	0.920	0.365	0.432
0.5	big	0.820	0.343	0.406
1.0	small	2.906	0.712	1.054
1.0	big	2.763	0.689	1.024
12 candidates: Spearman(D_{sym}, KL)=0.97; regret($\arg\min_q D_{\mathrm{sym}}$)=0.003 nats (worst 2.89)				

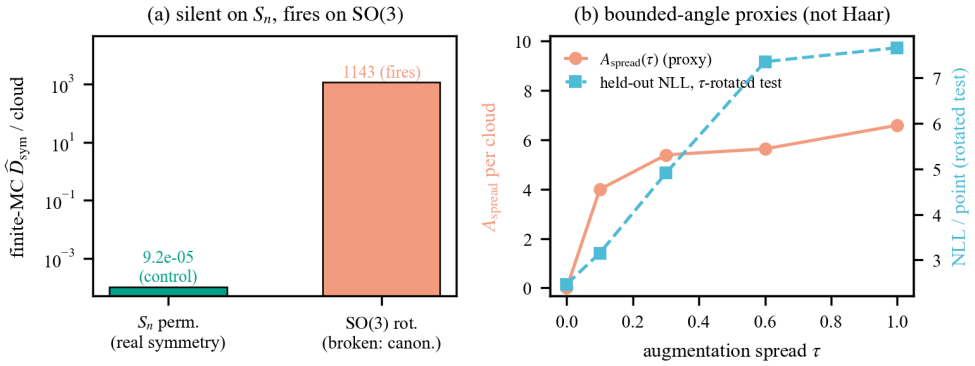

Figure 12: Real ShapeNet airplane point-cloud flow. (a) Finite-MC permutation and sampled-SO(3) diagnostics. (b) A bounded-angle augmentation-spread proxy co-varies with rotated-test NLL. It is not a Haar average, subgroup quantity, or population guarantee.

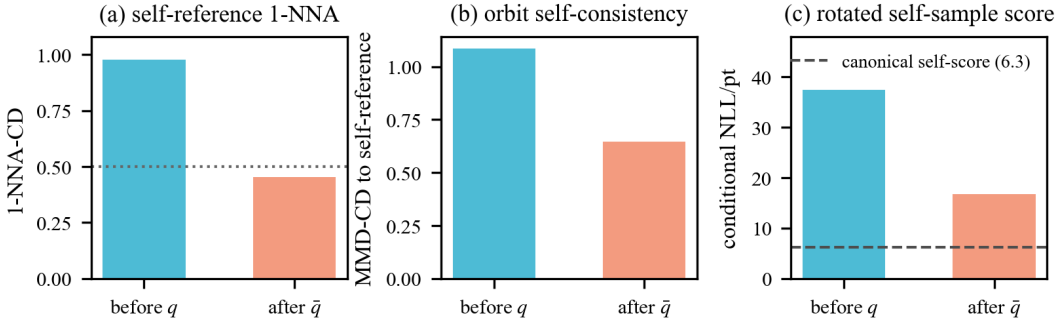

Figure 13: PointFlow internal diagnosis. Conditional-decoder 1-NNA, MMD, and log-score quantities use a PointFlow-derived rotation-augmented self-reference. They diagnose orbit self-consistency only, not real-data quality; no standard COV number is reported. The subgroup population guarantee is conditional on z, while its displayed value is a finite outer estimate with exact inner enumeration.

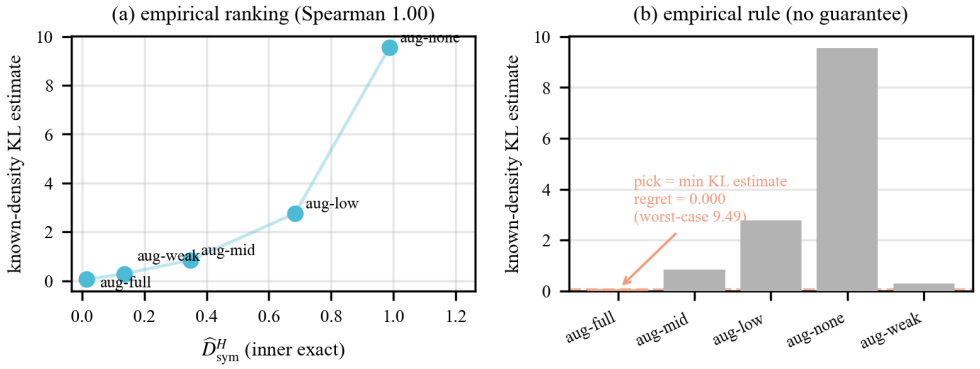

Figure 14: An empirical reference-free rule. In this five-model pool, finite-MC score and known-density KL estimates have Spearman 1.0 and the selected model has zero estimated regret. This is not a universal model-selection guarantee.

Permutations of S_4 are enumerated. The certificate estimator is equation 1 with logmeanexp over the orbit set.

Canonicalized log-score sanity check. We align each cloud to its PCA frame with an invariant sign convention and evaluate $s_c(X) = \log q(\text{canon}(X))$. Because canonicalization is many-to-one, $\exp s_c$ is not established as a normalized density on the original space. We therefore use only the

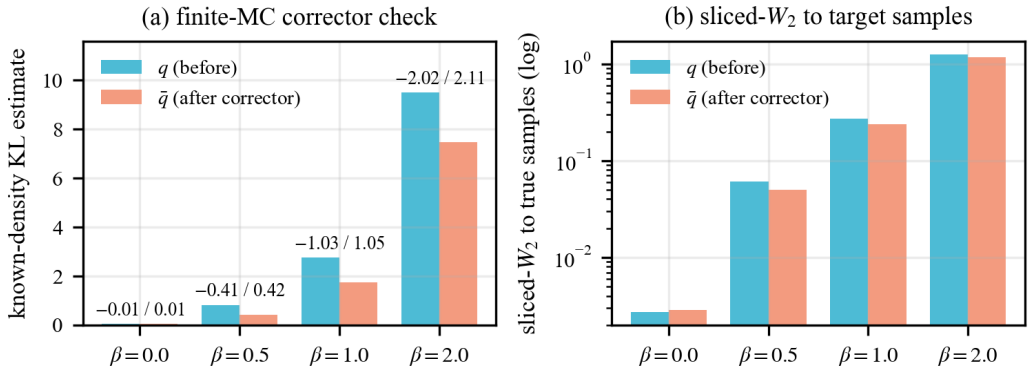

Figure 15: Corrector estimates on an invariant known-density target. The target is $\mathrm{SO}(3)$-invariant by construction. Finite-MC KL and sliced-Wasserstein estimates decrease in the tested biased cells; the exact KL reduction identity holds at population level.

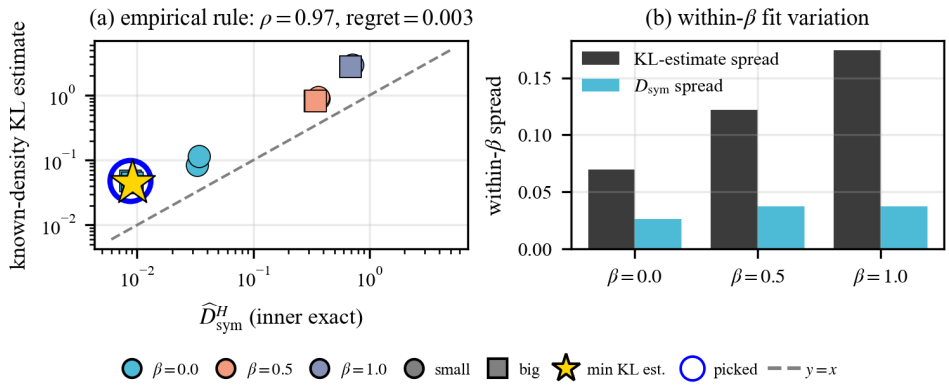

Figure 16: Empirical reference-free selection. The rule performs well on this heterogeneous candidate pool; no general recovery guarantee is implied. The right panel shows within-symmetry fit variation invisible to D_{sym}.

orbit identity $s_c(gX) = s_c(X)$ as a numerical sanity check and compute neither KL nor D_{sym} from it.

Quality proxy. The pooled sliced-Wasserstein proxy compares the pooled 3D points of the generated clouds against a rotation-invariant reference set; it is rotation-sensitive (anisotropic, orientation-biased generations score high) and serves as a stand-in for the unknown true error on the real model.

Metrics. The theorem concerns population $D_{\mathrm{sym}} \leq \mathrm{KL}$. Tables compare finite outer-MC estimates; exact group enumeration refers only to the inner average. Ratios and rank correlations are descriptive statistics of saved runs, not finite-sample guarantees.

H Additional Related Work

This appendix expands Sec. 2.

Exact vs. approximate equivariance in 3D generation. Equivariant architectures bake symmetry into the model: tensor-field networks, Vector Neurons, $E(n)$-GNNs, and equivariant flows and diffusion (Thomas et al., 2018; Deng et al., 2021; Satorras et al., 2021; Köhler et al., 2020; Klein et al., 2023; Hoogeboom et al., 2022) give exact guarantees at the cost of specialized, harder-to-scale machinery. Scalable 3D generators instead pair free-form backbones with augmentation (Nichol

et al., 2022; Mo et al., 2023; Zeng et al., 2022; Abramson et al., 2024; Ding & Hofmann, 2025) and reach only approximate equivariance at sampling time.

Equivariance-error metrics measure a different object. The Local Equivariance Error and related finite-transformation mismatches evaluate the deviation $\|f(Rx) - \rho(R)f(x)\|$ of the network map itself (Gruver et al., 2023). Map-level mismatch neither upper- nor lower-bounds the divergence of the generated distribution from the target: small mismatch can coexist with large distributional error and vice versa. D_{sym} is instead defined directly on distributions and provably lower-bounds $\mathrm{KL}(q \,\|\, p)$; Sec. 6.2 separates the two empirically.

Direct statistical and information-theoretic precedents. Pavlichin (2015) ask for the nearest G-invariant distribution and show that the arithmetic orbit average solves the relevant KL projection on finite spaces. Moore (2016) form the same symmetrized density, derive the same $\mathrm{KL}(q \,\|\, \bar{q})$ gain term in the ELBO, and discuss its Monte-Carlo evaluation. Koning (2023) studies compact-group invariance and derives the log-optimal e-value $q/\bar{q}$; its expected log under q is our D_{sym}. In quantum resource theory, Gour et al. (2009) define relative entropy to the nearest invariant state and connect it to group asymmetry. Together with the information-geometric Pythagorean theorem (Amari, 2016), these are direct precedents for the mathematical core.

Neighboring diagnostics and symmetrization methods. Chiu & Bloem-Reddy (2023) test distributional invariance via distance to a group average, and Lawrence et al. (2025) use a two-sample classifier to diagnose dataset symmetry. Stochastic symmetrization, frame averaging, learned canonicalization, and invariant discriminators use related operators to enforce or exploit symmetry (Zhang et al., 2025; Puny et al., 2022; Kaba et al., 2023; Birrell et al., 2022). We do not claim novelty for orbit averaging, the KL projection, the $q/\bar{q}$ statistic, the entropy identity, or randomized symmetrized sampling. Our novelty claim is an audited 3D-generator deployment study: population certificate/corrector interpretation, explicit inner-versus-outer Monte-Carlo accounting, and controlled and pretrained likelihood-model diagnostics. Model selection remains an empirical rule.

I Limitations and Scope

Certification scope of the symmetry defect. Population D_{sym} is a lower bound that quantifies one symmetry component, but $D_{\mathrm{sym}} \approx 0$ proves only near-invariance, not closeness to the target, since the residual leg $\mathrm{KL}(\bar{q} \,\|\, p)$ is invisible to it. The corrector inherits this boundary, removing the certified component D_{sym} while leaving $\mathrm{KL}(\bar{q} \,\|\, p)$ uncertified: invariance becomes exact while correctness stays open. Selection by D_{sym} has no universal guarantee because it cannot rank invariant fit error.

Requirements, estimation, and scope. The theory assumes a finite or compact group with normalized Haar probability measure, measure-preserving actions, integrable log ratios, and finite displayed KL terms. Evaluating q on the orbit needs a tractable likelihood: flows directly and diffusion via the probability-flow ODE likelihood (both validated here, Sections 6.2 and 6.6), with likelihood-free models such as GANs out of scope. For continuous $\mathrm{SO}(3)$, the plug-in defect estimate has inner Jensen upward bias. Exact octahedral enumeration removes inner-orbit error, but its reported sample mean still has outer-MC uncertainty; no CI is implied unless displayed. We validate on controlled known-density targets, a free-form diffusion model through its PF-ODE likelihood (on both controlled targets and real ShapeNet points), a synthetic-shape flow, a real-ShapeNet per-point flow, pretrained DiTMC, and pretrained PointFlow at a single-GPU budget. DiTMC supplies a coherent, non-saturated subgroup certificate but, because its released clean-endpoint sampler is not invertible, our likelihood statement is deliberately for its $t{=}0.97$ ODE marginal rather than its final discrete prediction. PointFlow supplies only a conditional-decoder population guarantee and internal orbit-self-consistency diagnostics; its outer-MC subgroup estimate saturates. Its legacy COV values are removed because of a nearest-axis implementation error. The Table 13 post column is a saved per-point augmentation proxy, not a whole-cloud orbit average. Larger latent-variable generators without a tractable marginal likelihood remain out of scope.