PREDICTION IS NOT ACTION: RELIABILITY-GUIDED ANALYTIC REFINEMENT OF MOLECULAR CONFORMERS

Anonymous authors
Paper under double-blind review

ABSTRACT

Molecular conformer generation is an important task in computational chemistry and drug discovery. Despite continued improvements in generative models, generated conformers can still retain localized bond-length and bond-angle deviations. Because local correction depends on the geometric state of the current input conformer, predicting plausible internal geometry alone is insufficient to determine the actual coordinate update. We therefore propose a Prediction-to-Action conformer-refinement framework that explicitly separates graph-conditioned local-geometry prediction from Source-conditioned corrective action. Source-conditioned Reliability modulates local correction signals, analytic Bond/Angle geometry maps these signals to Cartesian atomic motion, and an independent magnitude controller determines the final intervention scale. On the controlled ETFlow test, the fraction of conformers passing strict local-geometric validity criteria increases from 38.62% to $59.30 \pm 0.26\%$, while the refined conformers remain only $0.00724\,\mathring{A}$ RMSD from their Sources. Across a broader evaluation, this local-geometric improvement extends to training-unseen Bemis–Murcko scaffolds in GEOM-DRUGS. Models trained only on ETFlow also improve DiTMC outputs under zero-shot transfer without retraining. The resulting model contains approximately 0.77M trainable parameters.

1 INTRODUCTION

Molecular conformer generation from molecular graphs is an important task in computational chemistry and drug discovery. Early learning-based approaches explored molecular distance geometry, bilevel reconstruction, gradient fields, and internal geometric variables for constructing 3D structures (Simm & Hernández-Lobato, 2019; Xu et al., 2021; Shi et al., 2021; Ganea et al., 2021; Jing et al., 2022). More recent methods increasingly model continuous 3D conformational space directly through equivariant diffusion, flow-based generation, diffusion transformers, and scalable coordinate-space formulations (Xu et al., 2022; Hassan et al., 2024; Wang et al.; Frank et al., 2026), building in part on general generative frameworks such as flow matching (Lipman et al., 2022). Recent work has further improved generation through faster sampling and hierarchical molecular representations (Fan et al., 2024; Cao et al., 2025; Park & Shen, 2024; Reidenbach & Krishnapriyan, 2025; Hu et al., 2026). These efforts primarily address how to generate better conformers. Once generation is complete and the overall conformational state of a molecule has largely been established, however, correcting residual local geometric deviations becomes a distinct post-generation refinement problem.

Computational chemistry already provides well-established routes for improving an existing conformer. Classical molecular force fields can relax a given structure through energy minimization, with MMFF94s as one example, while semiempirical tight-binding quantum-chemical methods such as GFN2-xTB provide another efficient route to geometry optimization (Halgren, 1999; Bannwarth et al., 2019). As learning-based conformer generation has developed, the idea of further improving an existing conformer has also entered the post-processing stage of generative pipelines. For example, Flow-Matching Refiner takes an upstream generated conformer as a new sampling starting point and reorganizes the sampling process to refine that conformer (Xu & Gao, 2025). Optimization or refinement from an existing conformer is therefore a viable technical route. These representative

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.

approaches mainly address energy-driven structural relaxation or errors handled through renewed generative sampling. In contrast, we focus on a local post-generation setting in which the overall conformational state is already largely established but strict local geometric defects remain.

A generated conformer can pass relatively broad structural-validity checks while still failing stricter bond-length or bond-angle criteria. This pattern is directly observed on our fixed ETFlow diagnostic cohort (Appendix C.4). This mismatch motivates a different refinement objective: repair residual local geometric defects while preserving the current conformational state, rather than regenerate the entire 3D conformer or permit unrestricted structural relaxation. This gap between broad structural plausibility and strict local geometric validity provides the direct motivation for the refinement setting considered here.

Identifying a local geometric deviation, however, does not determine how that deviation should be corrected. A model may learn what geometry a particular local structure typically prefers, but for a specific generated conformer, whether that prediction should be used for correction, and to what extent, depends on the state of the current conformer. Predicting what local geometry is plausible and deciding how that prediction should act on the current conformer are related but distinct problems. We summarize this distinction as *Prediction Is Not Action*.

Motivated by this distinction, we formulate conformer refinement as a Prediction-to-Action factorization. The Prediction stage estimates preferred local geometry and its predictive residual scale from the molecular graph without reading the current Source coordinates. The Action stage combines those predictions with the current Source through primitive-level Reliability, analytic Bond/Angle-to-Cartesian mapping, and an independent update-magnitude controller. A predicted local geometric preference is therefore not treated as a Cartesian coordinate update. To our knowledge, we are the first to explicitly formulate post-generation molecular conformer refinement as a Prediction-to-Action factorization that separates graph-conditioned preferred local geometry from Source-conditioned corrective action. The resulting formulation is described in Section 3.

On the controlled ETFlow test, the fraction of conformers passing strict local-geometric validity criteria increases from 38.62% to $59.30 \pm 0.26\%$ with only $0.00724\ \mathring{A}$ Source RMSD. Mechanistic ablations show that smaller structural displacement alone does not yield better refinement: learned Reliability, learned predictive-scale information in action construction, and input-dependent magnitude control materially affect how predictions become corrective action. The same refinement behavior persists across a broader molecular cohort and on training-unseen Bemis–Murcko scaffolds. The ETFlow-trained models also improve DiTMC outputs under zero-shot transfer without retraining.

Our main contributions are as follows:

- **A Prediction-to-Action refinement factorization.** We explicitly separate graph-conditioned preferred-geometry prediction from Source-conditioned corrective action. Primitive-level Reliability controls how local predictions participate in correction, known analytic Bond/Angle geometry converts primitive corrections into Cartesian motion, and an independent magnitude controller determines the final intervention scale.
- **A post-generation local geometric refinement formulation.** We study a regime in which the overall conformational state has already been established while residual strict defects remain in local Bond/Angle geometry. The objective is to repair these deviations while preserving the existing Source conformer as much as possible.
- **Mechanistic and generalization evaluation.** Controlled evaluation and systematic ablations test how corrective action is formed. The ETFlow-trained checkpoints retain local-geometric improvements across a broader set of molecular identities, on Bemis–Murcko scaffolds unseen during training, and under zero-shot transfer to DiTMC outputs.

2 Related Work

Molecular conformer generation. Learning-based conformer generation spans learned geometric constraints, internal-coordinate modeling, and direct 3D coordinate generation. Representative early approaches include GraphDG, ConfVAE, and ConfGF (Simm & Hernández-Lobato, 2019; Xu et al., 2021; Shi et al., 2021), while GeoMol, Torsional Diffusion, and GINGER model local ge-

ometry, torsions, or internal coordinates (Ganea et al., 2021; Jing et al., 2022; Raush et al., 2024). More recent approaches include Cartesian diffusion and consistency models, flow-based generation, hierarchical formulations, and diffusion transformers (Xu et al., 2022; Wang et al.; Fan et al., 2024; Hassan et al., 2024; Cao et al., 2025; Park & Shen, 2024; Reidenbach & Krishnapriyan, 2025; Hu et al., 2026; Frank et al., 2026). These methods address upstream conformer generation; our problem begins after a Source conformer has already been produced.

Post-generation optimization and learned refinement. Existing conformers can be improved through classical geometry optimization or learned post-processing. MMFF94s and GFN2-xTB provide molecular-mechanics and semiempirical optimization (Halgren, 1999; Bannwarth et al., 2019), while Auto3D combines conformer generation with neural-potential-based low-energy structure optimization (Liu et al., 2022). Learned approaches address related objectives in different ways: Energy-Inspired Molecular Conformation Optimization learns gradient fields of an implicit conformational energy landscape, Uni-Mol+ iteratively updates raw coordinates toward DFT equilibrium geometries, and GOLF combines a learned energy model with an external optimizer (Guan et al., 2021; Lu et al., 2024; Tsypin et al., 2024). GeoOpt-Net directly predicts a single-step equivariant coordinate refinement, whereas Flow-Matching Refiner revisits the sampling trajectory from an upstream-generated conformer (Liu et al., 2026; Xu & Gao, 2025).

Across these closest formulations, corrective motion is produced through energy or gradient optimization, direct or iterative coordinate updates, or renewed generative sampling. Our setting instead targets residual Bond/Angle defects in an already generated conformer while remaining close to the current Source. The distinguishing factorization is between graph-conditioned preferred local geometry and Source-conditioned corrective action, rather than between two independent network backbones.

Local geometry and structural validity. Directional and equivariant molecular models such as DimeNet, EGNN, PaiNN, and NequIP encode geometric information while respecting spatial symmetries (Gasteiger et al., 2020; Satorras et al., 2021; Schütt et al., 2021; Batzner et al., 2022). Our analytic Bond/Angle mapping serves a different role: it converts primitive correction signals into Cartesian motion at the current Source geometry. For evaluation, GenBench3D Validity3D combines Bond/Angle statistics with intramolecular-clash and aromatic-ring planarity checks, whereas PoseBusters measures broader chemical and 3D plausibility (Baillif et al., 2024; Buttenschoen et al., 2024). Recent benchmark analyses further motivate careful treatment of local geometric validity (Nikitin et al., 2025; Fan et al., 2026).

3 Prediction-to-Action Conformer Refinement

3.1 Overview and Problem Formulation

Post-generation refinement starts from a complete Source conformer $\mathbf{x}_s \in \mathbb{R}^{N\times 3}$ produced by an upstream model. The molecular graph fixes connectivity, and the Source conformer specifies the current 3D conformational state; our objective is to repair residual Bond/Angle defects rather than regenerate the conformer or perform unrestricted relaxation.

We separate refinement into *Prediction* and *Action*. Prediction uses the molecular graph to estimate preferred local geometry and its predictive residual scale. Action determines how those predictions should modify the current Source. Source geometry enters Reliability, analytic Cartesian mapping, and update-magnitude control, but not preferred-geometry prediction. The network therefore does not directly regress the final Cartesian displacement.

Given a molecular graph $G = (V, E)$ with $N = |V|$ atoms, we write the refiner as $\mathbf{x}_{\text{out}} = \mathcal{R}_\theta(G, \mathbf{x}_s)$. Training samples a Reference conformer $\mathbf{x}_{\text{ref}}$ from the same molecule's frozen conformer ensemble for local geometric supervision. Inference requires only G and $\mathbf{x}_s$.

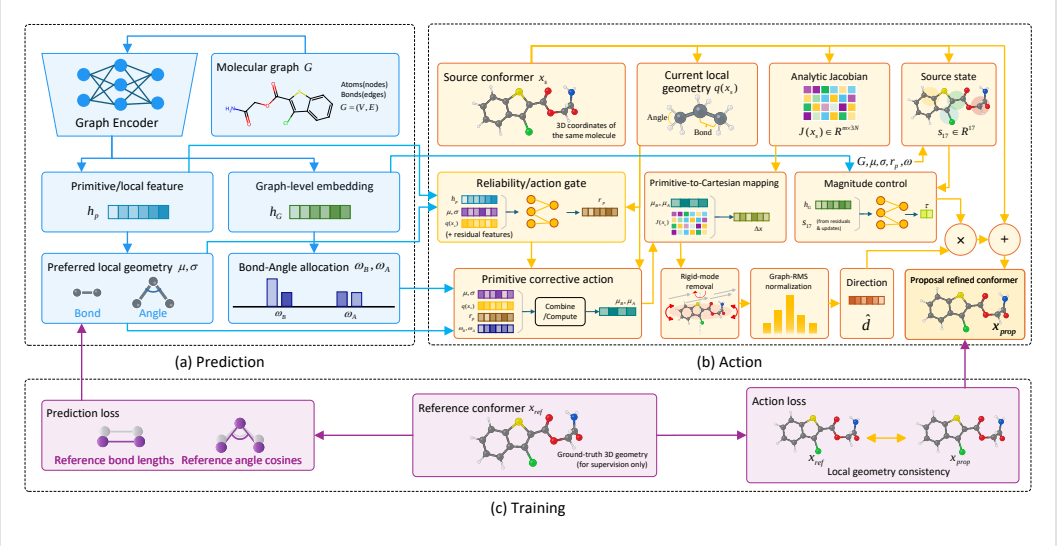

Figure 1: Overview of the Prediction-to-Action conformer refinement framework. Given a molecular graph and a Source conformer, the framework explicitly separates local-geometry prediction from coordinate refinement action. The Prediction stage estimates graph-conditioned preferred local geometry and predictive residual scales, while the Action stage constructs Source-conditioned Cartesian updates through reliability-guided correction, analytic primitive-to-Cartesian mapping, and adaptive magnitude control.

Local geometric primitives. We model Bond primitives $\mathcal{P}_B$ and Angle primitives $\mathcal{P}_A$. For a Bond primitive $p = (i, j)$ and an Angle primitive $p = (i, j, k)$ centered at atom j,

$$q_{B,p}(\mathbf{x}) = \|\mathbf{x}_i - \mathbf{x}_j\|_2,$$
$$q_{A,p}(\mathbf{x}) = \frac{(\mathbf{x}_i - \mathbf{x}_j)^\top (\mathbf{x}_k - \mathbf{x}_j)}{\|\mathbf{x}_i - \mathbf{x}_j\|_2 \|\mathbf{x}_k - \mathbf{x}_j\|_2}. \tag{1}$$

The Angle primitive is the cosine of the corresponding bond angle. We use $f \in \{B, A\}$ for the two primitive families.

3.2 Predicting Preferred Local Geometry

We use a shared message-passing encoder based on the MPNN framework (Gilmer et al., 2017). It produces atom representations $\mathbf{h}_i = [F_\theta(G)]_i$ and primitive representations $\mathbf{h}_{f,p}$. Each primitive predicts a preferred center $\mu_{f,p}$ and residual scale $\sigma_{f,p}$. Bond centers use $\mu_{B,p} = \mathrm{softplus}(a_{B,p})$, whereas Angle centers use $\mu_{A,p} = \tanh(a_{A,p})$. Neither head receives Source coordinates.

Predictive residual scale. To allow local residuals to receive scale-aware corrective weighting, we start from a frozen statistical scale estimated on TRAIN and learn a primitive-dependent graph-conditioned adjustment:

$$\delta_{f,p} = c_\sigma \tanh\left(\frac{b_{f,p}}{c_\sigma}\right), \qquad \sigma_{f,p} = s_{\mathrm{stat},f,p} \exp(\delta_{f,p}), \qquad c_\sigma = 6. \tag{2}$$

Here $s_{\mathrm{stat},f,p}$ is the frozen statistical scale estimated from TRAIN. We interpret $\sigma_{f,p}$ as a *predictive residual scale*, not as calibrated confidence, probability, or epistemic uncertainty. Graph-conditioned mean and residual-scale prediction follows the heteroscedastic-regression setting (Nix & Weigend, 1994). We train the center and scale jointly with β-NLL (Seitzer et al., 2022), which weakens the direct variance-dependent reweighting of mean learning while retaining scale learning. The resulting scale also enters the corrective action.

3.3 From Prediction to Corrective Action

The predicted local geometry does not directly define a Cartesian displacement. Instead, we construct an intermediate corrective action that determines whether and how each primitive-level prediction should modify the current Source conformer.

Reliability-guided correction. For each primitive p, we compute a reliability factor that combines the predicted residual scale and the current Source geometry:

$$r_{f,p} = \sigma(-g_{f,p}), \tag{3}$$

where $g_{f,p}$ is produced by a Source-conditioned reliability head.

The reliability module acts as a primitive-level correction gate. It does not represent a probability of correctness, but instead provides a continuous weighting factor for constructing corrective actions.

The reliability-weighted primitive correction coefficient is defined as:

$$u_{f,p} = \frac{w_f}{n_f} r_{f,p} \frac{q_{f,p}(\mathbf{x}_s) - \mu_{f,p}}{\sigma^2_{f,p}}, \tag{4}$$

where w_f/n_f denotes the normalized contribution of primitive family f, and $\sigma_{f,p}$ provides predictive-scale weighting.

Bond–Angle family weighting. Because different primitive families contribute differently to coordinate updates, the normalized family weights are predicted adaptively:

$$\boldsymbol{\alpha} = \mathrm{softmax}(a_B, a_A), \tag{5}$$

where molecule-conditioned logits from a detached graph representation determine the relative contribution of Bond and Angle corrective signals.

The family weighting is incorporated into w_f/n_f in Eq. 4.

3.4 Analytic Primitive-to-Cartesian Mapping

We transform primitive-level corrections into Cartesian coordinate updates using analytic geometric Jacobians.

For a Bond primitive $p = (i, j)$, the distance derivative with respect to atomic coordinates is:

$$\nabla_{\mathbf{x}_i} q_{B,p} = \frac{\mathbf{x}_i - \mathbf{x}_j}{\|\mathbf{x}_i - \mathbf{x}_j\|_2}, \qquad \nabla_{\mathbf{x}_j} q_{B,p} = -\nabla_{\mathbf{x}_i} q_{B,p}. \tag{6}$$

For an Angle primitive $p = (i, j, k)$, the corresponding analytic angular Jacobian is derived from the normalized vectors:

$$\mathbf{u} = \frac{\mathbf{x}_i - \mathbf{x}_j}{\|\mathbf{x}_i - \mathbf{x}_j\|_2}, \qquad \mathbf{v} = \frac{\mathbf{x}_k - \mathbf{x}_j}{\|\mathbf{x}_k - \mathbf{x}_j\|_2}. \tag{7}$$

The Cartesian correction direction is obtained by accumulating primitive contributions:

$$\mathbf{d} = -\sum_{f,p} \widetilde{J}_{f,p}(\mathbf{x}_s)^\top u_{f,p}, \tag{8}$$

where $\widetilde{J}_{f,p}(\mathbf{x}_s)$ denotes the stabilized analytic primitive Jacobian evaluated at the current Source geometry, and $u_{f,p}$ is the reliability-weighted primitive correction coefficient.

This maps local primitive corrections into Cartesian space without directly predicting atomic displacement vectors.

Removing rigid-body components. The raw Cartesian direction may contain global translation and rotation components. We therefore remove rigid-body modes before applying the update:

$$\mathbf{d}_{\mathrm{eq}} = P_{\mathrm{rigid}}(\mathbf{d}), \tag{9}$$

where P_{rigid} projects out translational and rotational components.

The resulting direction represents internal structural correction rather than global motion.

3.5 Input-Dependent Update Magnitude

The Cartesian correction direction determines where the Source conformer should move, but not how much it should move. We therefore introduce a molecule-level magnitude controller:

$$\tau = \mathrm{softplus}(h_\tau(G, \mathbf{x}_s)), \tag{10}$$

where h_τ is a lightweight controller and softplus parameterization is used to produce a positive update magnitude without an additional hard upper bound.

The final refined conformer is:

$$\mathbf{x}_{\mathrm{out}} = \mathbf{x}_s + \tau \frac{\mathbf{d}_{\mathrm{eq}}}{\|\mathbf{d}_{\mathrm{eq}}\|_{\mathrm{rms}} + \epsilon}. \tag{11}$$

The adaptive magnitude controller allows the intervention scale to vary across molecules, supporting source-preserving refinement through explicit control of the final update magnitude.

3.6 Training Objectives and Gradient Separation

For a minibatch $\mathcal{B}$, let

$$y_{f,p}^{(m)} = q_{f,p}(\mathbf{x}_{\mathrm{ref}}^{(m)}),$$

and let $\mathrm{sg}(\cdot)$ denote stop-gradient. The prediction objective jointly trains the preferred center and predictive residual scale using β-NLL with $\beta = 0.5$:

$$\begin{aligned}\mathcal{L}_{\mathrm{pred}} = \frac{1}{|\mathcal{B}|} \sum_{m \in \mathcal{B}} \frac{1}{2} \sum_{f \in \{B,A\}} \frac{1}{|\mathcal{P}_f^{(m)}|} \sum_{p \in \mathcal{P}_f^{(m)}} \left[\frac{1}{2} \left(\frac{y_{f,p}^{(m)} - \mu_{f,p}^{(m)}}{\sigma_{f,p}^{(m)}} \right)^2 + \log \sigma_{f,p}^{(m)} \right] \\ \times \mathrm{sg}\left(\left[\left(\sigma_{f,p}^{(m)} \right)^2 \right]^\beta \right), \qquad \beta = 0.5.\end{aligned} \tag{12}$$

The action objective supervises local Bond/Angle geometry of the refined conformer rather than a unique Cartesian Reference pose:

$$\mathcal{L}_{\mathrm{act}} = \frac{1}{|\mathcal{B}|} \sum_{m \in \mathcal{B}} \frac{1}{2} \sum_{f \in \{B,A\}} \frac{1}{|\mathcal{P}_f^{(m)}|} \sum_{p \in \mathcal{P}_f^{(m)}} \left(\frac{q_{f,p}(\mathbf{x}_{\mathrm{out}}^{(m)}) - q_{f,p}(\mathbf{x}_{\mathrm{ref}}^{(m)})}{\mathrm{sg}\left(\sigma_{f,p}^{(m)} \right)} \right)^2. \tag{13}$$

The total training objective is

$$\mathcal{L} = \mathcal{L}_{\mathrm{pred}} + \mathcal{L}_{\mathrm{act}}. \tag{14}$$

The μ and σ quantities used in corrective-action construction are detached, so the action objective does not update the preferred-geometry or predictive-scale heads through this path. The prediction objective nevertheless jointly trains both μ and σ.

4 Experiments

We organize the evaluation around three questions: **Q1**, how much strict local geometry can be repaired while keeping the Source nearly unchanged; **Q2**, whether the same behavior persists across broader molecular identities, training-unseen scaffolds, and zero-shot transfer to another upstream generator; and **Q3**, how Reliability, learned predictive-scale information, update magnitude, and Bond–Angle allocation affect the resulting corrective action.

4.1 Experimental Setup

Data. We use GEOM-DRUGS (Axelrod & Gomez-Bombarelli, 2022). TRAIN contains 50,000 molecules and 150,000 ETFlow Source records, with three Source conformers per molecule. Source

and Reference conformers are sampled independently from frozen conformer pools of the same molecule during training. The controlled ETFlow test contains 100 molecular-identity-disjoint molecules and 23,242 Source conformers. Formal ablations use a fixed 2,500-molecule/5,000-record ETFlow primary-final ablation cohort. The large-scale ETFlow and DiTMC evaluations use the same 1,000 molecular identities with 10 fixed Source conformers per molecule; DiTMC is used only for zero-shot evaluation.

Metrics. Validity3D (V3D) measures strict 3D validity including local geometry constraints (Baillif et al., 2024), whereas PoseBusters (PB) measures broader chemical and 3D structural plausibility (Buttenschoen et al., 2024). We also report Bond Reference MAE, Angle-cosine Reference MAE, Source RMSD for structural displacement, and single-point GFN2-xTB energy. The fixed $1{,}000\times10$ evaluation does not reconstruct complete conformer ensembles, so COV/AMR are not reported.

Comparisons and implementation. The controlled test compares Source, MMFF94s, GFN2-xTB optimization, and Ours from the same Source conformers. The final model has 766,288 trainable parameters (approximately 0.77M) and is trained independently with seeds 307, 331, and 353 for 17,500 steps. For Ours, primary V3D results are reported as the three-model arithmetic mean and sample SD. Unless otherwise stated, other single-valued Ours metrics are arithmetic means of the corresponding per-seed metrics. For the controlled-test xTB statistic, the reported median ΔE is the mean of the three seed-specific medians. Further training, baseline, and evaluation details are in Appendix B.

4.2 Controlled Refinement Evaluation

Q1 asks how much local geometry can be repaired relative to the structural adjustment required. Table 1 reports the controlled evaluation on 23,242 Source conformers from 100 molecular identities in the Formal-compatible evaluation set.

Table 1: Main results on the 100-molecule ETFlow test set.

	Validity		Local Geometry		Source Preservation	Energy
Method	V3D ↑	PB ↑	Bond Ref. MAE (Å) ↓	Angle-cos Ref. MAE ↓	Src. RMSD (Å)	med. ΔE (kcal/mol)
Source	38.62%	97.14%	0.00927	0.01981	0	0
MMFF94s	63.47%	98.61%	0.01003	0.02300	0.9961	+5.110
GFN2-xTB	59.08%	97.82%	0.00235	0.01375	0.4122	−4.919
Ours	$59.30 \pm 0.26\%$	97.19%	0.00409	0.01672	0.00724	−1.937

Ours raises V3D from 38.62% to $59.30 \pm 0.26\%$, a gain of about 20.7 percentage points, with only $0.00724\ \mathring{A}$ Source RMSD. Bond and Angle-cosine Reference errors also decrease, while PB remains nearly unchanged. The improvement is therefore concentrated in the strict local-geometric criteria targeted by the refiner.

Compared with GFN2-xTB and MMFF94s, Ours operates at a substantially lower structural-displacement scale while achieving a similar level of strict local validity to GFN2-xTB.

We additionally report xTB energy as a supplementary diagnostic. Ours has a median xTB ΔE of -1.937 kcal/mol, and 99.93% of valid energy pairs are below their corresponding Source energy. Complete energy and convergence statistics are reported in Appendix C.

4.3 Large-Scale Evaluation and Zero-Shot Transfer to DiTMC

On the 1,000-molecule ETFlow evaluation, V3D rises from 24.13% to $51.57 \pm 0.29\%$ with $0.01467\ \mathring{A}$ Source RMSD, extending the aggregate controlled-test pattern to a broader molecular cohort (Table 2).

Among these 1,000 molecules, 438 have non-empty Bemis–Murcko scaffolds (Bemis & Murcko, 1996) absent from TRAIN, giving 4,380 Source records. On this frozen GEOM-DRUGS subset,

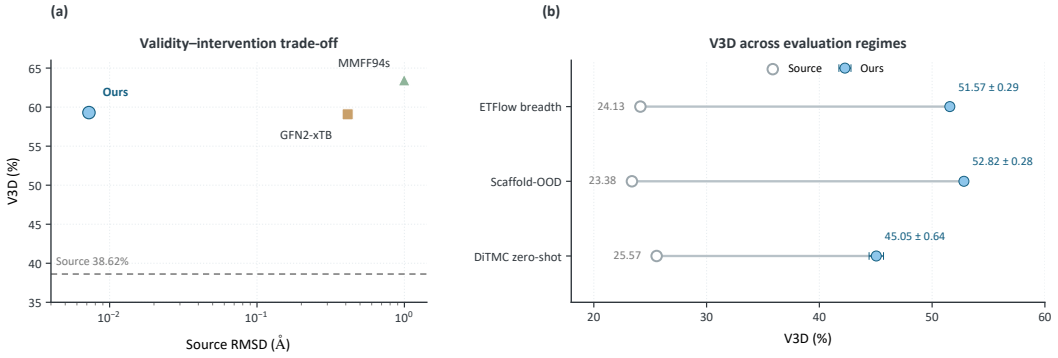

Figure 2: Validity–intervention trade-off and evaluation transfer of the proposed refinement. **(a)** V3D versus Source RMSD for Ours and the optimization baselines on the controlled ETFlow cohort; Source is shown as the zero-displacement reference. **(b)** Source-to-Ours V3D changes across ETFlow breadth, training-unseen Bemis–Murcko scaffolds, and ETFlow→DiTMC zero-shot transfer.

Table 2: Large-scale ETFlow evaluation and zero-shot transfer to DiTMC.

		Validity		Local Geometry		Source Preservation
Generator	Method	V3D ↑	PB ↑	Bond Ref. MAE (Å) ↓	Angle-cos Ref. MAE ↓	Src. RMSD (Å)
ETFlow	Source	24.13%	93.97%	0.01279	0.02116	0
	Ours	$51.57 \pm 0.29\%$	93.96%	0.01059	0.01679	0.01467
DiTMC	Source	25.57%	94.10%	0.01229	0.01866	0
	Ours	$45.05 \pm 0.64\%$	94.15%	0.00675	0.01663	0.00984

the mean V3D gain is 29.44 percentage points (molecule-cluster bootstrap 95% confidence interval: $[26.91, 32.05]$ percentage points), supporting refinement on training-unseen scaffolds. Appendix D.5 gives the complete scaffold construction, identity-overlap audit, and metrics.

DiTMC changes the upstream generator while keeping the ETFlow-trained refiner fixed, with no retraining, fine-tuning, or checkpoint reselection. The molecule-level mean V3D gain is 19.48 percentage points, with a molecule-cluster bootstrap 95% confidence interval of $[17.84, 21.15]$ percentage points. Additional transfer and classical-baseline results are reported in Appendix D.

4.4 Ablation and Mechanism Analysis

Q3 examines how Prediction becomes Action.

All formal ablations use the fixed 2,500-molecule/5,000-record ETFlow primary-final ablation cohort. Parameterization changes use matched retraining, whereas action-construction and inference-rule interventions are applied to the frozen final models.

Action control. Table 3 examines learned Reliability, learned predictive-scale information in action construction, and input-dependent update magnitude.

All three action-control ablations reduce V3D despite smaller Source displacement, showing that moving the Source less does not by itself improve refinement; corrective action must still be selected, weighted, and scaled appropriately.

These results show that smaller Source displacement alone is insufficient; corrective actions must be appropriately selected, weighted, and scaled.

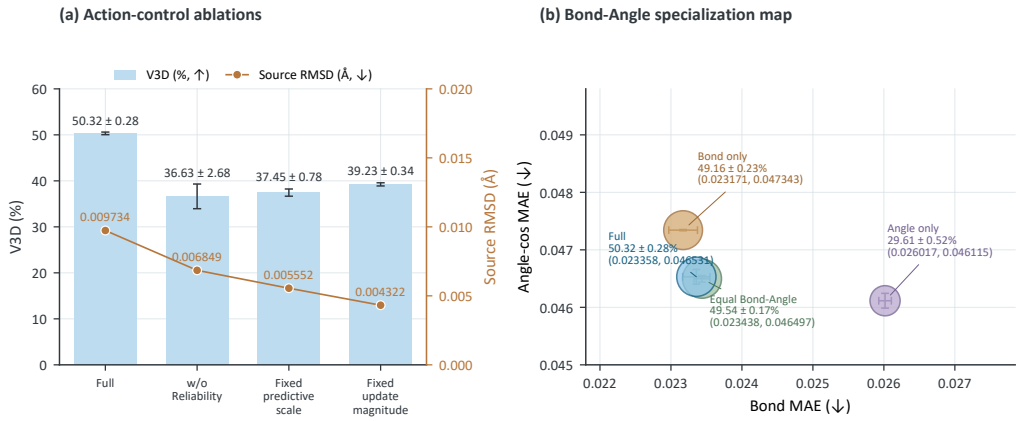

Figure 3: Mechanistic analysis of the Prediction-to-Action refinement process. **(a)** Action-control ablations evaluating the effects of learned Reliability, learned predictive-scale information in action construction, and input-dependent update magnitude. **(b)** Bond–Angle decomposition showing how different primitive allocation strategies affect aggregate V3D and local geometric errors.

Table 3: Ablation of Prediction-to-Action control mechanisms on the fixed ETFlow primary-final ablation cohort.

Variant	Protocol	V3D ↑	Source RMSD (Å)
Full model	Final model	$50.32 \pm 0.28\%$	0.00973
w/o Reliability	Matched retraining	$36.63 \pm 2.68\%$	0.00685
Fixed predictive scale	Inference intervention	$37.45 \pm 0.78\%$	0.00555
Fixed update magnitude	Inference intervention	$39.23 \pm 0.34\%$	0.00432

Bond–Angle decomposition. Equal Bond–Angle allocation remains close to the full model in V3D. Bond-only retains 49.16% V3D and the lowest Bond MAE, whereas Angle-only falls to 29.61% V3D despite the lowest Angle-cosine MAE. Bond-related correction therefore accounts for most of the current aggregate V3D improvement, while the Angle branch provides complementary correction of local angular geometry. The full seven-way ablation and all metrics are reported in Appendix E.

5 Conclusion

We introduced Prediction-to-Action for post-generation conformer refinement. The framework separates graph-conditioned prediction of preferred local geometry from Source-conditioned corrective action, where Reliability, analytic Bond/Angle mapping, and input-dependent magnitude control determine how local predictions are converted into Cartesian updates.

Across controlled ETFlow evaluation, broader ETFlow cohorts, training-unseen Bemis–Murcko scaffolds in GEOM-DRUGS, and zero-shot transfer to DiTMC outputs, the method improves strict local geometric validity while preserving the Source conformational state. Mechanistic ablations further show that refinement quality depends not only on local-geometry prediction accuracy, but also on how predicted geometry is transformed into corrective action.

AI USE STATEMENT

Generative AI tools were used during manuscript preparation for English translation, language editing, LaTeX formatting, and organizational assistance. The authors reviewed and verified the scientific content, mathematical formulations, experimental results, citations, and final manuscript text. Generative AI outputs were not treated as sources of scientific evidence.

REPRODUCIBILITY STATEMENT

Appendices A–G provide additional mathematical details, model and training specifications, fixed evaluation protocols, cross-generator results, final-model ablations, supplementary diagnostics, and development-stage mechanism analyses.

Exact data manifests, checkpoint identifiers, and evaluation artifacts are provided in the anonymous supplementary package.

REFERENCES

Simon Axelrod and Rafael Gomez-Bombarelli. Geom, energy-annotated molecular conformations for property prediction and molecular generation. *Scientific data*, 9(1):185, 2022.

Benoit Baillif, Jason Cole, Patrick McCabe, and Andreas Bender. Benchmarking structure-based three-dimensional molecular generative models using genbench3d: ligand conformation quality matters. *arXiv preprint arXiv:2407.04424*, 2024.

Jon Baker, Alain Kessi, and Bernard Delley. The generation and use of delocalized internal coordinates in geometry optimization. *The Journal of chemical physics*, 105(1):192–212, 1996.

Vebjørn Bakken and Trygve Helgaker. The efficient optimization of molecular geometries using redundant internal coordinates. *The Journal of chemical physics*, 117(20):9160–9174, 2002.

Christoph Bannwarth, Sebastian Ehlert, and Stefan Grimme. Gfn2-xtb—an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *Journal of chemical theory and computation*, 15 (3):1652–1671, 2019.

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1): 2453, 2022.

Guy W Bemis and Mark A Murcko. The properties of known drugs. 1. molecular frameworks. *Journal of medicinal chemistry*, 39(15):2887–2893, 1996.

Martin Buttenschoen, Garrett M Morris, and Charlotte M Deane. Posebusters: Ai-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chemical science*, 15(9):3130–3139, 2024.

Zhonglin Cao, Mario Geiger, Allan Dos Santos Costa, Danny Reidenbach, Karsten Kreis, Tomas Geffner, Franco Pellegrini, Guoqing Zhou, and Emine Kucukbenli. Efficient molecular conformer generation with so (3)-averaged flow matching and reflow. *arXiv preprint arXiv:2507.09785*, 2025.

Fan Fan, Bin Xi, Xianghu Meng, Han Wang, Bowen Zhang, Qingbo Xu, Wei Feng, Wenfeng Gao, Xiaoman Wang, Yuji Wang, et al. Assessing conformation validity and rationality of deep learning-generated 3d molecules. *Nature Communications*, 17(1):2481, 2026.

Zhiguang Fan, Yuedong Yang, Mingyuan Xu, and Hongming Chen. Ec-conf: A ultra-fast diffusion model for molecular conformation generation with equivariant consistency. *Journal of Cheminformatics*, 16(1):107, 2024.

J Thorben Frank, Winfried Ripken, Gregor Lied, Klaus-Robert Müller, Oliver Unke, and Stefan Chmiela. Sampling 3d molecular conformers with diffusion transformers. *Advances in Neural Information Processing Systems*, 38:168881–168931, 2026.

Octavian Ganea, Lagnajit Pattanaik, Connor Coley, Regina Barzilay, Klavs Jensen, William Green, and Tommi Jaakkola. Geomol: Torsional geometric generation of molecular 3d conformer ensembles. *Advances in neural information processing systems*, 34:13757–13769, 2021.

Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. *arXiv preprint arXiv:2003.03123*, 2020.

Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pp. 1263–1272. Pmlr, 2017.

Jiaqi Guan, Wesley Wei Qian, Wei-Ying Ma, Jianzhu Ma, and Jian Peng. Energy-inspired molecular conformation optimization. In *international conference on learning representations*, 2021.

Thomas A Halgren. Mmff vi. mmff94s option for energy minimization studies. *Journal of computational chemistry*, 20(7):720–729, 1999.

Majdi Hassan, Nikhil Shenoy, Jungyoon Lee, Hannes Stark, Stephan Thaler, and Dominique Beaini. Equivariant flow matching for molecular conformer generation. In *ICML 2024 Workshop on Structured Probabilistic Inference $\{\backslash\&\}$ Generative Modeling*, 2024.

Jiapeng Hu, Weizhi Gao, Zhichao Hou, and Xiaorui Liu. Hierarchical multi-scale molecular conformer generation. In *The Fourteenth International Conference on Learning Representations*, 2026.

Bowen Jing, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. Torsional diffusion for molecular conformer generation. *Advances in neural information processing systems*, 35:24240–24253, 2022.

Roland Lindh, Anders Bernhardsson, and Martin Schütz. Force-constant weighted redundant coordinates in molecular geometry optimizations. *Chemical physics letters*, 303(5-6):567–575, 1999.

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

Chengchun Liu, Wendi Cai, Boxuan Zhao, and Fanyang Mo. A cross-domain graph learning protocol for single-step molecular geometry refinement. *Journal of Chemical Theory and Computation*, 2026.

Zhen Liu, Tetiana Zubatiuk, Adrian Roitberg, and Olexandr Isayev. Auto3d: Automatic generation of the low-energy 3d structures with ani neural network potentials. *Journal of chemical information and modeling*, 62(22):5373–5382, 2022.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

Shuqi Lu, Zhifeng Gao, Di He, Linfeng Zhang, and Guolin Ke. Data-driven quantum chemical property prediction leveraging 3d conformations with uni-mol+. *Nature communications*, 15(1): 7104, 2024.

Filipp Nikitin, Ian Dunn, David Ryan Koes, and Olexandr Isayev. Geom-drugs revisited: toward more chemically accurate benchmarks for 3d molecule generation. *Digital Discovery*, 4(11): 3282–3291, 2025.

David A Nix and Andreas S Weigend. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 ieee international conference on neural networks (ICNN'94)*, volume 1, pp. 55–60. IEEE, 1994.

Jiwoong Park and Yang Shen. Equivariant blurring diffusion for hierarchical molecular conformer generation. *Advances in Neural Information Processing Systems*, 37:131645–131675, 2024.

Chunyang Peng, Philippe Y Ayala, H Bernhard Schlegel, and Michael J Frisch. Using redundant internal coordinates to optimize equilibrium geometries and transition states. *Journal of Computational Chemistry*, 17(1):49–56, 1996.

Peter Pulay and Geza Fogarasi. Geometry optimization in redundant internal coordinates. *The Journal of chemical physics*, 96(4):2856–2860, 1992.

Eugene Raush, Ruben Abagyan, and Maxim Totrov. Efficient generation of conformer ensembles using internal coordinates and a generative directional graph convolution neural network. *Journal of Chemical Theory and Computation*, 20(9):4054–4063, 2024.

Danny Reidenbach and Aditi S Krishnapriyan. Coarsenconf: Equivariant coarsening with aggregated attention for molecular conformer generation. *Journal of Chemical Information and Modeling*, 65(1):22–30, 2025.

Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021.

Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *International conference on machine learning*, pp. 9377–9388. PMLR, 2021.

Maximilian Seitzer, Arash Tavakoli, Dimitrije Antic, and Georg Martius. On the pitfalls of heteroscedastic uncertainty estimation with probabilistic neural networks. *arXiv preprint arXiv:2203.09168*, 2022.

Chence Shi, Shitong Luo, Minkai Xu, and Jian Tang. Learning gradient fields for molecular conformation generation. In *International conference on machine learning*, pp. 9558–9568. PMLR, 2021.

Gregor NC Simm and José Miguel Hernández-Lobato. A generative model for molecular distance geometry. *arXiv preprint arXiv:1909.11459*, 2019.

Andrew Stirn, Harm Wessels, Megan Schertzer, Laura Pereira, Neville Sanjana, and David Knowles. Faithful heteroscedastic regression with neural networks. In *International conference on artificial intelligence and statistics*, pp. 5593–5613. PMLR, 2023.

Artem Tsypin, Leonid A Ugadiarov, Kuzma Khrabrov, Alexander Telepov, Egor Rumiantsev, Alexey Skrynnik, Aleksandr Panov, Dmitry P Vetrov, Elena Tutubalina, and Artur Kadurin. Gradual optimization learning for conformational energy minimization. In *International Conference on Learning Representations*, volume 2024, pp. 36639–36656, 2024.

Y Wang, AA Elhag, N Jaitly, JM Susskind, and MA Bautista. Swallowing the bitter pill: Simplified scalable conformer generation, arxiv, 2023. *arXiv preprint arXiv:2311.17932*, 10.

Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2):513–530, 2018.

Minkai Xu, Wujie Wang, Shitong Luo, Chence Shi, Yoshua Bengio, Rafael Gomez-Bombarelli, and Jian Tang. An end-to-end framework for molecular conformation generation via bilevel programming. In *International conference on machine learning*, pp. 11537–11547. PMLR, 2021.

Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*, 2022.

Xiangyang Xu and Hongyang Gao. Flow-matching based refiner for molecular conformer generation. *arXiv preprint arXiv:2510.04878*, 2025.

A Additional Method and Mathematical Details

A.1 Analytic Bond and Angle Derivatives

For a Bond primitive $p = (i, j)$, let

$$q_p(\mathbf{x}) = \|\mathbf{x}_i - \mathbf{x}_j\|_2. \tag{15}$$

On a nondegenerate branch,

$$\nabla_{\mathbf{x}_i} q_p = \frac{\mathbf{x}_i - \mathbf{x}_j}{\|\mathbf{x}_i - \mathbf{x}_j\|_2}, \qquad \nabla_{\mathbf{x}_j} q_p = -\nabla_{\mathbf{x}_i} q_p. \tag{16}$$

Thus, a Bond primitive contributes Cartesian directions of equal magnitude and opposite sign to its two endpoint atoms.

For an Angle primitive $p = (i, j, k)$ centered at atom j, let

$$\mathbf{u} = \mathbf{x}_i - \mathbf{x}_j, \qquad \mathbf{v} = \mathbf{x}_k - \mathbf{x}_j, \tag{17}$$

and define

$$c = \frac{\mathbf{u}^\top \mathbf{v}}{\|\mathbf{u}\|_2 \|\mathbf{v}\|_2}. \tag{18}$$

On a nondegenerate branch,

$$\frac{\partial c}{\partial \mathbf{u}} = \frac{\mathbf{v}}{\|\mathbf{u}\|_2 \|\mathbf{v}\|_2} - c \frac{\mathbf{u}}{\|\mathbf{u}\|_2^2}, \tag{19}$$

and

$$\frac{\partial c}{\partial \mathbf{v}} = \frac{\mathbf{u}}{\|\mathbf{u}\|_2 \|\mathbf{v}\|_2} - c \frac{\mathbf{v}}{\|\mathbf{v}\|_2^2}. \tag{20}$$

The corresponding atom-wise gradients are

$$\nabla_{\mathbf{x}_i} c = \frac{\partial c}{\partial \mathbf{u}}, \qquad \nabla_{\mathbf{x}_k} c = \frac{\partial c}{\partial \mathbf{v}}, \qquad \nabla_{\mathbf{x}_j} c = -\nabla_{\mathbf{x}_i} c - \nabla_{\mathbf{x}_k} c. \tag{21}$$

The executed implementation uses the stabilized analytic form described in Appendix A.4.

For family $f \in \{B, A\}$ and primitive p, the correction coefficient used in the main text is

$$u_{f,p} = \frac{w_f}{n_f} r_{f,p} \frac{q_{f,p}(\mathbf{x}_s) - \mu_{f,p}}{\sigma_{f,p}^2}. \tag{22}$$

Accumulating primitive contributions gives

$$\mathbf{d}_{\mathrm{raw}} = -\widetilde{\mathbf{J}}_B(\mathbf{x}_s)^\top \mathbf{u}_B - \widetilde{\mathbf{J}}_A(\mathbf{x}_s)^\top \mathbf{u}_A. \tag{23}$$

The primitive correction coefficients are therefore converted into a Cartesian displacement direction through a stabilized analytic Jacobian-transpose mapping. Bond, angle, and related primitive derivatives form the basis of internal-coordinate Jacobian constructions widely used in molecular geometry optimization (Pulay & Fogarasi, 1992; Peng et al., 1996; Baker et al., 1996; Lindh et al., 1999; Bakken & Helgaker, 2002). Our method uses learned correction coefficients together with this analytic geometric relation rather than performing a conventional internal-coordinate geometry optimization.

A.2 Local Surrogate Interpretation

Fix μ, σ, Reliability, and the Bond–Angle family weights at the current Source $\mathbf{x}_s$, and assume that no norm floor, cosine clamp, or active mask is triggered.

Define the local surrogate

$$\mathcal{L}_{\mathrm{geo}}(\mathbf{x}) = \frac{1}{2} \sum_{f,p} \alpha_{f,p} \left(q_{f,p}(\mathbf{x}) - \mu_{f,p} \right)^2, \tag{24}$$

where

$$\alpha_{f,p} = \frac{w_f}{n_f} \frac{r_{f,p}}{\sigma_{f,p}^2}. \tag{25}$$

By the chain rule, the Jacobian-transpose action equals the first-order negative gradient of this surrogate at $\mathbf{x}_s$ on the nondegenerate active branch.

This relation is only a local first-order interpretation. The method does not solve

$$(\mathbf{J}^\top \mathbf{W} \mathbf{J})^{-1} \tag{26}$$

or any equivalent normal equation. Accordingly, the update is not a Gauss–Newton step and does not imply sample-wise monotonicity of geometric error or energy for a finite step. When a stabilization branch is activated, the executed direction is the stabilized analytic direction and is not unconditionally identical to the exact surrogate gradient above.

A.3 Rigid-Mode Projection and Direction Normalization

Bond and Angle primitives are insensitive to global translation and rotation, whereas the refinement target is internal geometry rather than global rigid pose. We construct translation and infinitesimal-rotation bases at the current Source geometry, apply QR decomposition with rank filtering, and obtain an effective orthonormal basis $\mathbf{Q}$.

On the normal full-rank branch, rigid components are removed by

$$\mathbf{d}_{\text{int}} = (\mathbf{I} - \mathbf{Q}\mathbf{Q}^\top)\mathbf{d}_{\text{raw}}. \tag{27}$$

We then compute the graph-level RMS magnitude

$$\rho = \left[N^{-1} \sum_i \|\mathbf{d}_{\text{int},i}\|_2^2 \right]^{1/2}, \tag{28}$$

and, on the normal branch,

$$\widehat{\mathbf{d}} = \frac{\mathbf{d}_{\text{int}}}{\rho}. \tag{29}$$

This preserves the relative atomic-motion pattern while removing the global scale of the raw direction. The final movement is controlled independently by τ:

$$\mathbf{x}_{\text{out}} = \mathbf{x}_s + \tau \widehat{\mathbf{d}}. \tag{30}$$

A.4 Numerical Stabilization

Bond and Angle derivative calculations use norm-floor protection. Angle cosines are clipped inside their valid domain and combined with active masking for stabilized branches; angle-value evaluation also uses norm protection. The rigid-mode basis is rank-filtered after QR decomposition, and RMS direction normalization protects its denominator. For rank-deficient geometries, this rank-filtered QR path is treated as a numerical safeguard; exact recovery of the rigid-motion subspace is asserted only on the normal full-rank branch. Nonfinite or near-zero directions enter a zero-direction stabilization branch.

These numerical guards prevent division by zero and unstable derivatives in degenerate geometries without changing the functional definition of Prediction-to-Action. The corresponding safeguards use fixed constants; the exact configuration is provided with the anonymous supplementary implementation.

A.5 Source-State Representation for Update Magnitude

The final magnitude predictor reads both a detached graph-level representation and a state dynamically computed for the current Source conformer. For each refinement operation, the 17-dimensional Source state is computed on the fly from $\mathbf{x}_s$, local predictions, primitive residuals, and the analytic direction rather than being stored as a fixed per-molecule vector.

Let $n_f = \max(1, |\mathcal{P}_f|)$ and define the standardized Source residual

$$z_{f,p} = \frac{q_{f,p}(\mathbf{x}_s) - \mu_{f,p}}{\sigma_{f,p}}.$$

The family-normalized standardized residual is then

$$\rho_{f,p} = \frac{z_{f,p}}{\sqrt{2n_f}}.$$

The 17-dimensional state is ordered as follows: four Bond summaries of ρ_B (mean absolute value, RMS, maximum absolute value, and 90th percentile absolute value); the corresponding four Angle summaries of ρ_A; the Bond and Angle Gaussian-NLL summaries, denoted E_B and E_A, together with their balanced average $\frac{1}{2}(E_B + E_A)$; projected-direction RMS, mean per-atom direction norm, and maximum per-atom direction norm; and finally the atom, Bond, and Angle counts. Each component is standardized using a frozen mean and standard deviation estimated on TRAIN, with numerical safeguards applied during standardization.

The magnitude head receives a 128-dimensional detached mean-pooled graph embedding together with the 17-dimensional state, giving 145 input dimensions in total. The MLP architecture is

$$145 \to 64 \to 32 \to 1,$$

and softplus is applied to produce positive τ. No additional hard upper bound or inference-time intervention is applied beyond the update rule defined in the main text.

B Architecture, Training, and Reproducibility Details

B.1 Model Architecture

The final model uses a shared graph encoder to produce graph-conditioned atomic representations and constructs Bond and Angle primitive representations from them. Prediction and Action share graph representations rather than using two independently trained backbones. The final model contains 766,288 trainable parameters, approximately 0.77M.

The preferred-geometry branch reads only molecular-graph information. The Bond mean head predicts positive bond lengths, the Angle mean head predicts angle cosines within their valid range, and the predictive-scale head produces the primitive-dependent residual scale σ. The frozen statistical factor $s_{\mathrm{stat},f,p}$ in Eq. 2 is obtained from a typed-context calibration table constructed from TRAIN Reference conformers with hierarchical backoff, rather than being fitted independently for each primitive. We interpret σ as predictive residual scale rather than calibrated uncertainty probability.

B.2 Prediction and Action Heads

For each primitive, the graph-derived representation produces μ_p and σ_p. The current source provides the measured primitive geometry $q_p(\mathbf{x}_s)$ and residual

$$e_p = q_p(\mathbf{x}_s) - \mu_p.$$

The Reliability head combines the primitive graph feature, $q_p(\mathbf{x}_s)$, μ_p, $\log \sigma_p$, raw and absolute residuals, standardized residual, and Bond/Angle type information. A sigmoid produces

$$r_p \in (0, 1).$$

The value r_p represents the relative strength with which the primitive participates in corrective action for the current source, rather than calibrated confidence.

Bond–Angle family allocation predicts w_B and w_A from a detached mean-pooled graph representation and applies softmax so that

$$w_B + w_A = 1.$$

The magnitude predictor reads the detached pooled graph representation together with the 17-dimensional source state to obtain molecule-level τ. Reliability controls primitive-level action suitability, whereas τ controls the overall magnitude applied to the normalized Cartesian direction.

B.3 Gradient Routing and Functional Separation

The separation in Prediction-to-Action refers to functional roles and direct gradient paths rather than parameter-independent subnetworks. The shared graph encoder serves both prediction and Reliability. The pooled graph inputs used by family weighting and magnitude are detached before entering the corresponding heads.

Table 4: Direct gradient routing and functional separation.

Loss / path	Backbone	μ heads	σ heads	Reliability	Bond–Angle weighting	τ
Prediction objective	Yes	Yes	Yes	No	No	No
Refinement objective	Indirect via Reliability	Stop-grad	Stop-grad	Yes	Yes	Yes
Bond–Angle pooled input	Detached	–	–	–	Yes	–
Magnitude pooled/state input	Detached	–	–	–	–	Yes

The prediction objective directly trains the shared backbone and the μ and σ heads. The μ/σ paths in corrective-action construction are stop-gradient; consequently, the refinement loss does not update the μ/σ heads through the analytic action construction. Primitive graph features entering Reliability remain connected to the shared encoder, so the refinement loss can still update the shared encoder indirectly through this path.

B.4 Training Configuration

TRAIN contains 50,000 molecules and 150,000 ETFlow source records, with three source records per molecule. During training, source and reference conformers are sampled independently from frozen conformer pools of the same molecule rather than being paired one-to-one.

Models are independently trained with seeds 307, 331, and 353 for 17,500 steps with batch size 64. We use AdamW (Loshchilov & Hutter, 2017) with a backbone learning rate of 1.5×10^{-4}, head learning rate of 3×10^{-4}, and weight decay of 1×10^{-6}. The cosine schedule uses $T_{\max} = 22{,}500$, and gradient clipping is set to 1.0.

The prediction objective uses $\beta = 0.5$ in β-NLL. The multiplicative $(\sigma^2)^\beta$ reweighting is stop-gradient, while σ itself remains learned by the prediction objective. All three models use the final-step checkpoint; checkpoints are not reselected based on DEV results.

B.5 Evaluation and Baseline Protocol

The main ETFlow test set contains 100 molecules with no molecular-identity overlap with TRAIN and 23,242 source conformers in total. Development-stage mechanism studies use the fixed 2,500-molecule, 5,000-record ETFlow DEV cohort. Final formal ablations use a separate fixed 2,500-molecule, 5,000-record ETFlow primary-final ablation cohort.

Large-scale evaluation uses the same 1,000 molecular identities, with 10 fixed source conformers per molecule, to construct matched ETFlow and DiTMC $1{,}000 \times 10$ sets of 10,000 records each.

All three final checkpoints are trained only on ETFlow sources. The DiTMC evaluation uses no retraining, fine-tuning, or checkpoint reselection. MMFF94s and GFN2-xTB optimization start from exactly the same source conformers as Ours. Formal evaluation preserves a fixed denominator, and optimization failures use a predefined source-fallback policy.

Table 5: Per-seed V3D results on the 100-molecule ETFlow test set.

Model	V3D
seed307	59.5990%
seed331	59.1128%
seed353	59.1817%
Three-seed mean $\pm$ sample SD	$59.30 \pm 0.26\%$

C Extended Main Evaluation

C.1 Per-Seed Results on the 100-Molecule ETFlow Test Set

The three independently trained checkpoints produce results within a relatively narrow range. Accordingly, the main text reports the arithmetic mean across the three models and additionally reports the sample standard deviation for the principal V3D result.

Table 6: Complete results on the 100-molecule ETFlow test set.

	Validity		Local Geometry		Source Preservation	Energy	
Method	V3D ↑	PB ↑	Bond Ref. MAE (Å) ↓	Angle-cos Ref. MAE ↓	Src. RMSD (Å)	med. ΔE (kcal/mol)	Lower-energy ↑
Source	38.62%	97.14%	0.00927	0.01981	0	0	–
MMFF94s	63.47%	98.61%	0.01003	0.02300	0.9961	+5.110	13.82%
GFN2-xTB	59.08%	97.82%	0.00235	0.01375	0.4122	-4.919	99.96%
Ours	$59.30 \pm 0.26\%$	97.19%	0.00409	0.01672	0.00724	-1.937	99.93%

C.2 Baseline Convergence and Fallback Statistics

Table 7: Baseline convergence and source-fallback statistics.

Method	Records	Successful / converged	Source fallback	Timeout
MMFF94s	23,242	23,186	56	0
GFN2-xTB optimization	23,242	23,233	9	0

All fallback records remain in the fixed 23,242-record evaluation denominator. The 56 MMFF94s fallbacks correspond to records that did not satisfy the optimization convergence criterion, while the nine GFN2-xTB fallbacks correspond to unsuccessful optimization convergence.

Subsequent GFN2-xTB single-point energy pairing succeeds for all 23,242 MMFF94s outputs and all 23,242 GFN2-xTB outputs. For Ours, energy evaluation is valid for 23,242/23,242 records for seed307, while seeds 331 and 353 each contain one single-point failure.

C.3 Energy Diagnostics

xTB energy is not part of the training objective and is therefore used only as a supplementary diagnostic after refinement. On the 100-molecule ETFlow test set, Ours obtains a median xTB ΔE of -1.937 kcal/mol and a lower-energy fraction of 99.93%. GFN2-xTB optimization obtains -4.919 kcal/mol, whereas MMFF94s gives $+5.110$ kcal/mol.

These energy statistics should be interpreted together with displacement from the source conformer. GFN2-xTB performs substantially more extensive geometry relaxation and correspondingly produces much larger Source RMSD, whereas Ours remains in a source-near local-refinement regime.

C.4 Motivating Local-Geometry Diagnosis

Table 8: Motivating local-geometry diagnosis on the fixed ETFlow cohort.

Diagnostic	Count / rate
PoseBusters pass	4,657 / 5,000 (93.14%)
V3D pass	1,113 / 5,000 (22.26%)
PB-pass & V3D-fail	3,596
PB-pass & V3D-fail with Bond/Angle failure	3,584 / 3,596

Failure categories can overlap. This diagnostic characterizes the fixed ETFlow cohort in which structures pass broad 3D plausibility checks while failing strict local geometry checks.

D Large-Scale and Transfer Evaluation

D.1 Large-Scale ETFlow Evaluation

Table 9: Large-scale ETFlow evaluation on 1,000 molecules $\times$ 10 Source conformers.

	Validity		Local Geometry		Source Preservation	Energy	
Method	V3D ↑	PB ↑	Bond Ref. MAE (Å) ↓	Angle-cos Ref. MAE ↓	Src. RMSD (Å)	med. ΔE (kcal/mol)	Lower-energy ↑
Source	24.13%	93.97%	0.01279	0.02116	0	0	–
Ours	$51.57 \pm 0.29\%$	93.96%	0.01059	0.01679	0.01467	-3.537	99.87%

At the molecule level, 75.9% of molecules improve in V3D, 17.2% remain unchanged, and 6.9% worsen. This result shows that the improvement observed in the 100-molecule controlled evaluation is not concentrated in only a small number of molecular identities.

D.2 Zero-Shot Transfer to DiTMC

The DiTMC evaluation uses exactly the same 1,000 molecular identities as the ETFlow $1{,}000 \times 10$ evaluation, again with 10 Source conformers per molecule. All three checkpoints are trained only on ETFlow Sources. DiTMC is not used for training, fine-tuning, or checkpoint reselection.

Table 10: Zero-shot transfer to DiTMC with classical optimization baselines.

	Validity		Local Geometry		Source Preservation	Energy	
Method	V3D ↑	PB ↑	Bond Ref. MAE ↓	Angle-cos Ref. MAE ↓	Src. RMSD (Å)	med. ΔE (kcal/mol)	Lower-energy ↑
Source	25.57%	94.10%	0.012289	0.018659	0	0	–
Ours	$45.05 \pm 0.64\%$	94.15%	0.006750	0.016631	0.009837	-2.947673	99.8532%
MMFF94s	59.76%	96.04%	0.011214	0.022273	0.766020	+3.194447	19.1515%
GFN2-xTB	56.77%	94.08%	0.002039	0.011260	0.144382	-5.521589	99.9800%

The classical baselines are included to characterize distinct optimization and refinement regimes.

D.3 Molecule-Level Transfer Statistics

Among the 1,000 DiTMC molecules, 642 improve, 245 remain unchanged, and 113 worsen, corresponding to 64.2%, 24.5%, and 11.3%, respectively.

The molecule-level mean V3D improvement is 19.48 percentage points, and the median improvement is 10 percentage points. Using molecular identity as the resampling unit, the cluster-bootstrap 95% confidence interval is $[17.84, 21.15]$ percentage points.

This experiment uses ETFlow-trained checkpoints for DiTMC zero-shot evaluation.

D.4 Scope of the $1{,}000 \times 10$ Evaluation

The $1{,}000 \times 10$ design fixes 10 Source conformers per molecule. Its purpose is to broaden the range of molecular identities and provide matched cross-generator evaluation rather than reconstruct complete conformer ensembles.

Accordingly, this design does not report COV/AMR and is not mixed with coverage benchmarks that require complete conformer ensembles.

D.5 Scaffold-OOD Evaluation

Scaffold membership is frozen before reading Source/Ours performance artifacts. We use the standard Bemis–Murcko scaffold (Bemis & Murcko, 1996) implemented in RDKit 2026.03.4. Scaffold-based partitioning also has precedent in molecular machine-learning benchmarks such as MoleculeNet (Wu et al., 2018). Our criterion here is more specific: membership is determined by whether the canonical non-empty Bemis–Murcko scaffold occurs in TRAIN, rather than by applying a standard benchmark train/validation/test scaffold split.

TRAIN contains 50,000 molecules, of which 49,838 contain rings, yielding 28,915 unique non-empty scaffolds. In the ETFlow 1,000-molecule breadth cohort, 997 molecules contain rings and three are acyclic.

For an evaluation molecule containing rings, it is assigned to the unseen-scaffold subset if its canonical non-empty Bemis–Murcko scaffold does not occur in TRAIN. Acyclic molecules are not automatically assigned to this subset because of their empty scaffold. This procedure yields 438 unseen-scaffold molecules and 4,380 records, corresponding to 43.8% of the full cohort.

The 438-molecule subset contains 437 unique unseen scaffolds. The median and 90th percentile of the number of molecules per unseen scaffold are both one, with a maximum of two.

Table 11: Scaffold-OOD evaluation on frozen ETFlow artifacts.

	Validity		Local Geometry		Source Preservation
Method	V3D ↑	PB ↑	Bond Ref. MAE ↓	Angle-cos Ref. MAE ↓	Src. RMSD ($\mathring{A}$)
Source	23.38%	93.01%	0.013124	0.021210	0
Ours	$52.82 \pm 0.28\%$	$93.14 \pm 0.01\%$	0.005601 ± 0.000013	0.016690 ± 0.000267	0.01009 ± 0.00026

Relative to Source, Ours improves mean V3D by 29.44 percentage points. At the molecule level, 332 molecules improve, 78 remain unchanged, and 28 worsen, corresponding to a 75.8% improved fraction.

Using molecular identity as the resampling unit, a 10,000-replicate cluster bootstrap gives a 95% confidence interval of $[26.91, 32.05]$ percentage points. PB changes only slightly, from 93.01% to 93.14%.

Within the full 1,000-molecule breadth cohort, two canonical molecular identities overlap with TRAIN. Both belong to the seen-scaffold subset and neither enters the 438-molecule unseen-scaffold subset.

The unseen-scaffold subset therefore has 0/438 molecular-identity overlap with TRAIN, whereas the complete 1,000-molecule breadth cohort is not treated as fully train-disjoint.

The result therefore characterizes generalization to training-unseen Bemis–Murcko scaffolds in GEOM-DRUGS.

E Additional Ablations and Mechanism Analysis

E.1 Complete Ablation Summary

All final matched ablations use the same fixed 2,500-molecule, 5,000-record ETFlow primary-final ablation cohort. Reliability-off and Equal Bond–Angle allocation use matched three-seed retraining. Fixed predictive scale, Fixed update magnitude, Bond-only, and Angle-only are inference-time interventions applied to the frozen final checkpoints. In the Reliability-off variant, primitive Reliability is fixed to $r_{f,p} = 1$, the Reliability head is removed, and each seed is retrained under the matched protocol.

For Fixed predictive scale, the learned predictive scale used by action construction is replaced at inference time by the frozen TRAIN typed-context statistical scale. This replacement propagates through the corresponding Reliability features and Source-state quantities used by the magnitude controller.

For Fixed update magnitude, the learned molecule-dependent magnitude is replaced by the frozen constant $\tau = 0.004321542759325059\,\text{Å}$.

Equal Bond–Angle allocation fixes $w_B = w_A = 0.5$ and uses matched three-seed retraining.

Bond-only and Angle-only are inference-time interventions that mask the opposite learned family contribution without renormalizing the remaining family weight.

Table 12: Complete ablation summary on the fixed ETFlow primary-final ablation cohort.

		Performance		Local Geometry		Source Preservation
Variant	Protocol	V3D mean ± SD ↑	PB mean ↑	Bond MAE ↓	Angle-cos MAE ↓	Src. RMSD (Å)
Full model	Final model	$50.32 \pm 0.28\%$	93.213%	0.023358	0.046531	0.009734
w/o Reliability	Matched retraining	$36.63 \pm 2.68\%$	93.180%	0.024096	0.047016	0.006849
Fixed predictive scale	Inference only	$37.45 \pm 0.78\%$	93.193%	0.024322	0.046861	0.005552
Fixed update magnitude	Inference only	$39.23 \pm 0.34\%$	93.180%	0.024725	0.047048	0.004322
Equal Bond–Angle	Matched retraining	$49.54 \pm 0.17\%$	93.220%	0.023438	0.046497	0.009601
Bond only	Inference only	$49.16 \pm 0.23\%$	93.213%	0.023171	0.047343	0.009705
Angle only	Inference only	$29.61 \pm 0.52\%$	93.167%	0.026017	0.046115	0.009647

The primary-final ablation cohort differs from the 100-molecule ETFlow test set. Absolute Bond/Angle MAE values should therefore not be directly compared across these two cohorts. This appendix is intended to characterize the relative behavior of action mechanisms within the same fixed primary-final ablation cohort.

E.2 Action-Control Ablations

Removing learned Reliability with matched retraining reduces V3D from 50.32% to 36.63%. Fixing predictive scale gives 37.45%, while fixing update magnitude gives 39.23%.

The corresponding Source RMSD values are $0.006849\,\text{Å}$, $0.005552\,\text{Å}$, and $0.004322\,\text{Å}$, respectively, all below the Full model value of $0.009734\,\text{Å}$.

Thus, smaller movement can coincide with substantially worse V3D. A limited coordinate update must still be selected and applied with appropriate primitive weighting and overall magnitude. Moving less is not, by itself, equivalent to refining better.

E.3 Bond–Angle Decomposition

Equal Bond–Angle allocation obtains 49.54% V3D, close to the 50.32% of the Full model. Bond-only obtains 49.16% and the lowest Bond MAE, 0.023171. Angle-only has lower aggregate V3D, 29.61%, but achieves the lowest Angle-cosine MAE, 0.046115.

These results indicate that Bond-related correction accounts for most of the current aggregate V3D improvement, while the Angle branch provides complementary correction of local angular geometry. Despite its lower aggregate V3D, Angle-only achieves the lowest Angle-cosine MAE, demonstrating direct improvement of the corresponding angular objective.

E.4 Interpretation of Ablation Protocols

Reliability-off and Equal Bond–Angle allocation are matched-retraining experiments because they alter the model parameterization and training process.

Fixed predictive scale, Fixed update magnitude, Bond-only, and Angle-only are inference-time interventions. They address the question: given the same learned predictions, what changes when the rule that converts prediction into action is modified?

Accordingly, Fixed predictive scale probes predictive-scale information in action construction rather than isolating only the explicit $1/\sigma^2$ weighting term. Likewise, Fixed update magnitude directly tests the role of input-dependent magnitude control in the inference-time action.

F Additional Diagnostics and Limitations

F.1 Reference Conformers and Local Supervision

Reference conformers are used during training to provide local geometric supervision rather than to define a unique global target pose.

Accordingly, global Reference RMSD is not treated as the primary metric for source-preserving local refinement. Evaluation instead focuses on local bond/angle geometry and structural displacement relative to the input source conformer.

F.2 Energy Diagnostics and Limitations

GFN2-xTB single-point energy is used to examine energy changes after refinement, but energy is not part of the main optimization objective.

The current results show that local geometric repair can coincide with favorable cohort-level energy changes. More extensive energy relaxation remains within the scope of force-field or quantum-chemical geometry optimization.

Accordingly, these energy statistics should be interpreted as supplementary diagnostics rather than evidence of sample-wise energy monotonicity or complete energy minimization.

G Development-Stage Mechanism Studies

This section reports development-stage experiments used to understand information routing and training-design choices in Prediction-to-Action. Both studies use the fixed 2,500-molecule, 5,000-record ETFlow DEV cohort and seed307. They are included to interpret design choices and do not replace the final-model three-seed matched ablations in Appendix E.

G.1 J0/J1/J2 × R0/R1 Six-Arm Factorial

This matched factorial compares three predictive-scale/training-semantic designs and their interaction with Reliability.

J0 does not learn a primitive-dependent σ and instead uses a fixed/statistical source-free predictive scale.

J1 learns graph-conditioned σ using β-NLL (Seitzer et al., 2022). Detached variance-dependent reweighting reduces the strong effect of scale estimation on the mean-prediction gradient. J1 is the design used in the final model.

J2 also learns graph-conditioned σ, but uses a Faithful-style mean/variance gradient separation inspired by Stirn et al. (2023) as an alternative heteroscedastic training semantic.

R0 does not use learned Reliability, whereas R1 uses Source-conditioned Reliability. All six variants use the same movement protocol, with Source RMSD approximately $0.004322\,\mathring{A}$. The principal differences therefore arise from predictive-scale/training semantics and Reliability.

Table 13: Development-stage J0/J1/J2 × R0/R1 factorial on ETFlow DEV (seed307).

		Performance		Local Geometry		Source Preservation
Variant	Reliability	V3D ↑	PB ↑	Bond MAE ↓	Angle-cos MAE ↓	Src. RMSD ($\mathring{A}$)
J0–R0	R0 / no learned Reliability	44.76%	93.16%	0.006571	0.015974	0.004322
J0–R1	R1 / Source-conditioned	48.02%	93.20%	0.005074	0.015230	0.004322
J1–R0	R0 / no learned Reliability	47.94%	93.22%	0.005422	0.015352	0.004322
J1–R1	R1 / Source-conditioned	53.54%	93.22%	0.004300	0.014606	0.004322
J2–R0	R0 / no learned Reliability	45.06%	93.16%	0.006566	0.015968	0.004322
J2–R1	R1 / Source-conditioned	48.94%	93.22%	0.004972	0.015183	0.004322

Within this fixed DEV screen, R1 improves V3D relative to the corresponding R0 by 3.26, 5.60, and 3.88 percentage points under J0, J1, and J2, respectively. J1–R1 reaches 53.54% V3D.

This factorial is a seed307 development study used to analyze the interaction between Reliability and predictive-scale/training design. The final-model three-seed matched ablation is reported in Appendix E.

G.2 B/R/W/T Source-Routing Diagnostic

To examine where current-source information should enter learned components, we perform a complete 16-mask B/R/W/T routing screen.

The four bits indicate whether Source-derived conditioning is provided to the belief/prediction branch (B), Reliability (R), Bond–Angle family weighting (W), and update magnitude τ (T). A value of 1 enables the corresponding Source-conditioning route, whereas 0 masks the audited Source-derived input slots for that component.

The 0000 configuration does not mean that Source information is absent from refinement execution. The measured geometry $q(\mathbf{x}_s)$, residual $q(\mathbf{x}_s) - \mu$, and analytic Jacobian $\mathbf{J}(\mathbf{x}_s)$ remain necessary inputs to the refinement operation.

All 16 routing variants are independently trained with seed307 on the same DEV cohort. Their differences therefore include both routing effects and variation arising from independent training.

Table 14: Complete 16-mask Source-routing diagnostic on ETFlow DEV (seed307).

Mask	V3D ↑	Source RMSD (Å)	Mask	V3D ↑	Source RMSD (Å)
0000	56.92%	0.006168	1000	56.64%	0.006088
0001	56.44%	0.006094	1001	56.82%	0.006002
0010	56.30%	0.006230	1010	56.28%	0.006111
0011	57.28%	0.006065	1011	57.08%	0.006055
0100	56.52%	0.006218	1100	56.28%	0.006172
0101	56.78%	0.006118	1101	57.60%	0.006056
0110	56.86%	0.006180	1110	57.20%	0.006087
0111	56.82%	0.006076	1111	56.54%	0.006166

Differences in V3D across the 16 routing masks are limited, and no single Source-routing rule summarizes all observed behavior. The four first-order average effects on V3D are small: approximately $+0.065$, $+0.105$, $+0.045$, and $+0.295$ percentage points for B, R, W, and T, respectively. The corresponding molecule-cluster 95% intervals all cross zero. Complete secondary diagnostics involving Bond/Angle proposals and τ are provided in the anonymous supplementary artifact/configuration report.

The final model uses routing B0R1W0T1: preferred-geometry prediction and Bond–Angle family allocation do not directly read the Source, whereas Reliability and τ use Source-conditioned state.