

STIEFEL FLOW MATCHING FOR MOMENT-CONSTRAINED STRUCTURE ELUCIDATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Molecular structure elucidation is a critical step in understanding chemical phenomena, with applications in identifying molecules in natural products, lab syntheses, forensic samples, and the interstellar medium. We consider the task of elucidating a molecule’s 3D structure given only its molecular formula and moments of inertia, motivated by the ability of rotational spectroscopy to precisely measure these moments. While existing generative models can conditionally sample 3D structures with approximately correct moments, this soft conditioning fails to leverage the many digits of precision afforded by experimental rotational spectroscopy. To address this, we first show that the space of n -atom point clouds with a fixed set of moments of inertia is embedded in the Stiefel manifold $\text{St}(n, 4)$. We then propose *Stiefel Flow Matching* as a generative model for elucidating 3D structure under exact moment constraints. Additionally, we learn simpler and shorter flows by finding approximate solutions for optimal transport on the Stiefel manifold. Empirically, satisfying moment constraints exactly allows Stiefel Flow Matching to achieve higher success rates and faster sampling than Euclidean diffusion models, even on high-dimensional manifolds corresponding to large molecules in the GEOM dataset.

1 INTRODUCTION

Elucidating the structure of unknown molecules is a ubiquitous task in chemistry, important for analyzing environmental samples (Moneta et al., 2023), identifying novel drugs (Sonstrom et al., 2023), and determining potential building blocks of life in the interstellar medium (McGuire et al., 2016). The challenge is to aggregate information from multiple sources of analytical data to unambiguously determine a molecule’s structure. Rotational spectroscopy holds a unique capacity to provide precise measurements about a molecule’s moments of inertia, which in turn has routinely provided the highest quality gas-phase 3D structures attainable from experiment (Domingos et al., 2020). As the moments of inertia are trivial to compute for a given 3D structure, structure elucidation with rotational spectroscopy usually proceeds by confirming whether a *known* structure’s moments matches with experimentally measured moments (Lee & McCarthy, 2019; McCarthy et al., 2020). However, this approach is limited only to molecules whose structures are known *a priori*. This fundamentally limits applications related to identifying molecules that have not been catalogued, such as novel natural products, undiscovered molecules, and key reactive intermediate species that cannot be easily isolated (Womack et al., 2015).

To overcome this, we apply generative modeling to structure elucidation by inferring candidate 3D structures based on the moments and molecular formula *alone*. By themselves, moments of inertia only describe the spread of the molecule’s mass in three directions. Going from moments to 3D structure ($3 \rightarrow 3n$ values) is therefore a heavily underconstrained inverse problem. Nevertheless, deep generative models such as diffusion (Ho et al., 2020; Song et al., 2020) and flow matching (Lipman et al., 2023; Liu et al., 2022) have shown promise in solving various inverse problems (Song et al., 2022; Chung et al., 2023; Song et al., 2023a). Previous work has trained a Euclidean diffusion model to sample 3D structures conditioned on a given set of moments of inertia (Cheng et al., 2024).

However, analytic formulas for the moments of inertia are sufficiently simple that we can define a feasible space where these constraints are always exactly satisfied. Such precise adherence to moments can potentially leverage the many digits of precision provided by rotational spectroscopy

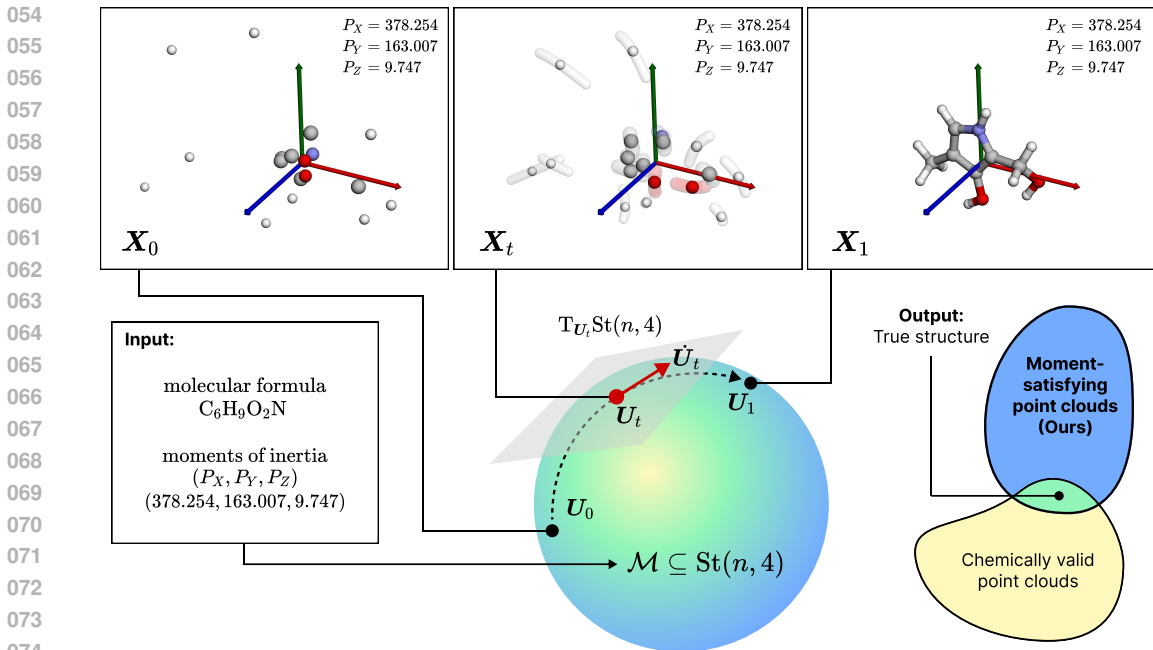


Figure 1: Stiefel Flow Matching learns to elucidate 3D molecular structure from moments and molecular formula alone by transforming uniform Stiefel noise \mathbf{X}_0 into *valid* molecular structures \mathbf{X}_1 . Generative modelling on the Stiefel manifold $\text{St}(n, 4)$ guarantees that samples always have the correct moments of inertia, which allows the network to focus only on generating chemically stable structures. Within the intersection of these spaces lies the true 3D structure.

(Shipman et al., 2011) in order to constrain the space of plausible structures. We first show that \mathcal{M} , the set of n -atom point clouds with fixed moments of inertia, is embedded in the Stiefel manifold $\text{St}(n, 4)$. We then propose *Stiefel Flow Matching* as a generative model on the Stiefel manifold for solving the moment-constrained structure elucidation problem (Figure 1). Stiefel Flow Matching builds on ideas from Riemannian flow matching (Chen & Lipman, 2024) with equivariant optimal transport (Klein et al., 2023; Song et al., 2023c), which simplifies and shortens generation paths.

Concretely, our contributions are:

1. We propose the task of moment-constrained structure elucidation as a challenging generative modelling problem on the Stiefel manifold.
2. To solve this problem, we present *Stiefel Flow Matching*, a Riemannian flow matching approach. Furthermore, we formulate an objective for equivariant optimal transport on the Stiefel manifold, which obtains shorter and simpler flows.
3. Stiefel Flow Matching predicts 3D structure with greater success rate and lower cost than Euclidean diffusion models on both the QM9 (Ramakrishnan et al., 2014) and GEOM (Axelrod & Gomez-Bombarelli, 2022) datasets.

2 BACKGROUND AND APPROACH

We consider a 3D molecule as a point cloud of n atoms with atomic numbers $\mathbf{a} \in \mathbb{N}^n$ and 3D coordinates $\mathbf{X} = (x_i, y_i, z_i)_{i=1}^N \in \mathbb{R}^{n \times 3}$. We also refer to \mathbf{X} as the molecule’s *3D structure*. Molecules have translational, rotational, and permutation symmetry. However, as \mathbf{a} and \mathbf{X} are stored on a computer, they necessarily have a node ordering and orientation. Therefore, when we refer to *structure*, we really mean the equivalence class containing \mathbf{X} under these symmetries.

2.1 PROBLEM STATEMENT

We are given a molecule’s molecular formula and moments of inertia, and wish to infer the 3D structure of the molecule. The molecular formula provides us with the number of atoms n , atomic numbers $\mathbf{a} \in \mathbb{N}^n$, and atomic masses $\mathbf{m} \in \mathbb{R}^n$. The moments of inertia¹ (P_X, P_Y, P_Z) are three nonnegative numbers that summarize the mass distribution of the point cloud. They are defined as the eigenvalues of the *planar dyadic* $\mathbf{P} \in \mathbb{R}^{3 \times 3}$ (Kraitchman, 1953), calculated from masses \mathbf{m} and the 3D atomic coordinates $\mathbf{X} = (x_i, y_i, z_i)_{i=1}^N \in \mathbb{R}^{n \times 3}$ as

$$\mathbf{P} = \mathbf{X}^\top (\text{diag } \mathbf{m}) \mathbf{X} = \sum_{i=1}^n \begin{pmatrix} m_i x_i^2 & m_i x_i y_i & m_i x_i z_i \\ m_i y_i x_i & m_i y_i^2 & m_i y_i z_i \\ m_i z_i x_i & m_i z_i y_i & m_i z_i^2 \end{pmatrix}. \quad (1)$$

This definition assumes a coordinate system whose origin is the weighted center of mass of the molecule, which gives the constraint $\mathbf{m}^\top \mathbf{X} = \mathbf{0}$. We assume that $P_X > P_Y > P_Z > 0$, which eliminates rare edge cases (Appendix A.2). Note that \mathbf{P} is symmetric and positive semidefinite and can be thought of as the mass-weighted covariance matrix of the point cloud, as in principal component analysis (PCA).

Diagonalizing \mathbf{P} yields three eigenvectors, referred to as the *principal axes of rotation*, which orient a molecule in a canonical representation up to sign-flips. The principal axes correspond to directions which “explain” the most molecular mass, in analogy to PCA. The *principal axis system* is then the coordinate system whose origin is the center-of-mass and whose axes are the principal axes of rotation. We fix our coordinate system to be the principal axis system *by construction*, which gives the following constraints on \mathbf{X} :

$$\begin{aligned} P_X &= \sum_{i=1}^n m_i x_i^2, & 0 &= \sum_{i=1}^n m_i y_i z_i, & 0 &= \sum_{i=1}^n m_i x_i, \\ P_Y &= \sum_{i=1}^n m_i y_i^2, & 0 &= \sum_{i=1}^n m_i x_i z_i, & 0 &= \sum_{i=1}^n m_i y_i, \\ P_Z &= \sum_{i=1}^n m_i z_i^2, & 0 &= \sum_{i=1}^n m_i x_i y_i, & 0 &= \sum_{i=1}^n m_i z_i. \end{aligned} \quad (2)$$

These constraints also canonicalize a 3D structure up to sign-flips of the axes (e.g. $x \mapsto -x$). The center-of-mass constraint removes translational degrees of freedom, while the off-diagonal constraints remove rotational degrees of freedom.

The goal of molecular identification from moments and molecular formula is to find all molecular structures \mathbf{X} which are consistent with these constraints *and* are thermodynamically *stable*, i.e., local minima of potential energy surfaces. Structures which satisfy these criteria can then be compared to experimental measurements from rotational spectroscopy.

2.2 THE FEASIBLE SPACE OF MOMENT-CONSTRAINED STRUCTURES

The *Stiefel manifold* is the set of orthonormal $n \times p$ matrices, defined as

$$\text{St}(n, p) = \{\mathbf{U} \in \mathbb{R}^{n \times p} \mid \mathbf{U}^\top \mathbf{U} = \mathbf{I}_p\}, \quad (3)$$

where \mathbf{I}_p is the $p \times p$ identity matrix, and $n \geq p$. Elements of $\text{St}(n, p)$ can be thought of as the first p columns of some n -dimensional (improper) rotation matrix, or can be thought of as orthonormal p -frames residing in n -dimensional space.

Moment-constrained structures \mathbf{X} can then be mapped into $\text{St}(n, 4)$ by scaling rows by masses and columns by moments. Letting $M = \sum_{i=1}^n m_i$ be the total mass, consider the following construction:

$$\mathbf{U} = \begin{pmatrix} \sqrt{\frac{m_1}{P_X}} x_1 & \sqrt{\frac{m_1}{P_Y}} y_1 & \sqrt{\frac{m_1}{P_Z}} z_1 & \sqrt{\frac{m_1}{M}} \\ \vdots & \vdots & \vdots & \vdots \\ \sqrt{\frac{m_n}{P_X}} x_n & \sqrt{\frac{m_n}{P_Y}} y_n & \sqrt{\frac{m_n}{P_Z}} z_n & \sqrt{\frac{m_n}{M}} \end{pmatrix} \in \mathbb{R}^{n \times 4}. \quad (4)$$

¹This overview is a slight simplification of rotational spectroscopy. Full details are outlined in Appendix A.

162 It can be verified by inspection and comparison to the Equation 2 constraints that the columns of \mathbf{U}
 163 are orthonormal. That is, $\mathbf{U} \in \text{St}(n, 4)$. However, we cannot yet freely convert between \mathbf{X} and \mathbf{U} :
 164 the last column of \mathbf{U} is not free as it must equal the unit mass vector $\hat{\mathbf{m}} = \sqrt{m/M}$ to satisfy the zero
 165 center-of-mass constraint. To convert an arbitrary $\mathbf{U} \in \text{St}(n, 4)$ to an \mathbf{X} satisfying all constraints, we
 166 first apply a rigid n -dimensional rotation \mathbf{R} to \mathbf{U} so that its last column is aligned to $\hat{\mathbf{m}}$ (Appendix
 167 B.5), before finally unscaling the rows and columns of \mathbf{U} .

168 The feasible space \mathcal{M} of moment-constrained structures is therefore the subset of $\text{St}(n, 4)$ whose last
 169 column is fixed to $\hat{\mathbf{m}}$, i.e.,

$$170 \mathcal{M} = \{\mathbf{U} \in \text{St}(n, 4) \mid \mathbf{U}_{:,4} = \hat{\mathbf{m}}\}. \quad (5)$$

171 In fact, we show in Appendix B.7 that \mathcal{M} is a totally geodesic submanifold of $\text{St}(n, 4)$, which means
 172 that geodesics between points in \mathcal{M} stay in \mathcal{M} . The first three columns of elements in \mathcal{M} form
 173 the intersection between $\text{St}(n, 3)$ and the orthogonal complement to $\text{span}(\hat{\mathbf{m}})$, which is in turn
 174 equivalent to $\text{St}(n-1, 3)$. Hence, the dimension of \mathcal{M} is $3n-9$, which corresponds to removing
 175 3 translational, 3 rotational, and 3 moment degrees of freedom. Going forward, we assume that
 176 the molecule of interest contains $n \geq 5$ atoms, so that we always deal with Stiefel manifolds of
 177 rectangular matrices.

179 2.3 NAVIGATING THE STIEFEL MANIFOLD

180 The Stiefel manifold provides rich structure for navigating the feasible space of molecular structures.
 181 As a *manifold*, it is locally Euclidean but globally curved. This means that every point $\mathbf{U} \in \text{St}(n, p)$
 182 is attached a vector space called its *tangent space* $\text{T}_{\mathbf{U}}\text{St}(n, p)$. For the Stiefel manifold, these tangent
 183 spaces are given as

$$184 \text{T}_{\mathbf{U}}\text{St}(n, p) = \{\Delta \in \mathbb{R}^{n \times p} \mid \mathbf{U}^{\top} \Delta + \Delta^{\top} \mathbf{U} = \mathbf{0}\}. \quad (6)$$

185 Then, equipping every tangent space with an inner product $\langle \cdot, \cdot \rangle_{\mathbf{U}}$ turns $\text{St}(n, p)$ into a *Riemannian*
 186 *manifold*, giving rise to notions of angles and distances. The collection of inner products for each
 187 tangent space is called the *Riemannian metric*. One such metric for the Stiefel manifold is the
 188 canonical metric (Edelman et al., 1998),

$$189 \langle \Delta, \tilde{\Delta} \rangle_{\mathbf{U}} = \text{trace } \Delta^{\top} \left(\mathbf{I}_n - \frac{1}{2} \mathbf{U} \mathbf{U}^{\top} \right) \tilde{\Delta}, \quad (7)$$

190 which we exclusively use for this work. The canonical metric induces a norm $\|\Delta\|_{\mathbf{U}} = \sqrt{\langle \Delta, \Delta \rangle_{\mathbf{U}}}$
 191 on each tangent space, allowing us to define the length of a curve $\gamma: [0, 1] \rightarrow \text{St}(n, p)$ as $L(\gamma) =$
 192 $\int_0^1 \|\dot{\gamma}(t)\|_{\gamma(t)} dt$. Curves that are locally length-minimizing are called *geodesics*, providing a notion
 193 of “straight lines” for efficiently navigating around the manifold. Geodesics are defined by their
 194 starting point and initial velocity. Indeed, the exponential map $\exp_{\mathbf{U}}(\Delta)$ takes in a starting point
 195 $\mathbf{U} \in \text{St}(n, p)$ and an initial velocity $\Delta \in \text{T}_{\mathbf{U}}\text{St}(n, p)$, and outputs the final manifold point after
 196 following this geodesic for unit time. The exponential map is locally invertible, which gives the
 197 existence of the logarithmic map $\log_{\mathbf{U}}(\tilde{\mathbf{U}})$. The logarithmic map takes in a starting point \mathbf{U} and
 198 a target point $\tilde{\mathbf{U}}$, and outputs the tangent vector needed to travel from \mathbf{U} to $\tilde{\mathbf{U}}$. Algorithms for
 199 computing exponential and logarithmic maps under the canonical metric for the Stiefel manifold
 200 are given in Appendix B.3. In addition, points which are not on the manifold can be projected onto
 201 the manifold (Appendix B.6), which can turn a Euclidean generative model into a Stiefel generative
 202 model.
 203
 204

205 3 STIEFEL FLOW MATCHING

206
 207 Having shown that the feasible space of moment-constrained structures is a Stiefel manifold, we
 208 can now formulate the problem of moment-constrained structure elucidation as an *unconstrained*
 209 *generative modeling problem on the Stiefel manifold*. An attractive approach to this is flow matching
 210 (FM), which trains a network as a time-dependent velocity field u_t that transforms samples from a
 211 prior noise distribution $x_0 \sim p_0(x_0)$ into samples which approximately match the data distribution
 212 $x_1 \sim p_1(x_1) \approx p_{\text{data}}(x_1)$ (Lipman et al., 2023). The integration of the velocity field over time is then
 213 a *continuous normalizing flow* ψ_t , which generates marginal probability densities p_t over time by the
 214 pushforward operation $p_t = [\psi_t]_{\#}(p_0)$. In practice, this is realized by sampling initial conditions
 215 $x_0 \sim p_0(x_0)$ and evolving them from time 0 to t according to the ODE $\frac{dx}{dt} = u_t(x)$ (Appendix C.4).
 The goal of training is to approximate this u_t using a neural network $v_{\theta}(t, x)$ parameterized by θ .

Flow matching is readily generalized to distributions on Riemannian manifolds (Chen & Lipman, 2024). When closed-form geodesics are available, Riemannian flow matching provides a simulation-free training objective for learning u_t , called Riemannian conditional flow matching,

$$\mathcal{L}_{\text{RCFM}}(\theta) = \mathbb{E}_{t \sim U(0,1), p(\mathbf{U}_0), p(\mathbf{U}_1)} \left[\|v_\theta(t, \mathbf{U}_t) - \dot{\mathbf{U}}_t\|_{\mathbf{U}_t}^2 \right], \quad (8)$$

To compute this loss, we require (1) sampling $\mathbf{U}_0 \sim p(\mathbf{U}_0)$ from a prior noise distribution, (2) geodesic interpolation between \mathbf{U}_0 and \mathbf{U}_1 to get \mathbf{U}_t , (3) computing the time derivative of the interpolant $\dot{\mathbf{U}}_t$, and (4) evaluating the norm $\|\cdot\|_{\mathbf{U}_t}$.

To sample uniformly from the feasible space \mathcal{M} , we can sample \mathbf{U} uniformly from $\text{St}(n, 4)$ (Appendix B.4) and then rigidly rotate \mathbf{U} so that its last column aligns with the unit mass vector $\hat{\mathbf{m}}$ (Appendix B.5). Then, geodesics can be computed from the exponential (Edelman et al., 1998) and logarithmic (Zimmermann & Hüper, 2022) maps, which are efficient to compute for $\text{St}(n, 4)$,

$$\mathbf{U}_t = \exp_{\mathbf{U}_0}(t \log_{\mathbf{U}_0}(\mathbf{U}_1)), \quad (9)$$

making our training objective simulation-free. Once we have the interpolant \mathbf{U}_t , which will be the input to the neural network, we can then calculate the network’s target $\dot{\mathbf{U}}_t$. Instead of autodifferentiation, which introduces unnecessary overhead, we compute $\dot{\mathbf{U}}_t$ using the logarithmic map,

$$\dot{\mathbf{U}}_t = \frac{1}{1-t} \log_{\mathbf{U}_t}(\mathbf{U}_1), \quad (10)$$

observing that $\dot{\mathbf{U}}_t$ is a unit length tangent vector along the geodesic from \mathbf{U}_t to \mathbf{U}_1 . Finally, we compute the norm of $v_\theta(t, \mathbf{U}_t) - \dot{\mathbf{U}}_t$ following Appendix B.2. Additionally, Appendix Theorem 4 shows how we can compute the logarithm in $\text{St}(n, 3)$, rather than $\text{St}(n, 4)$, which slightly saves time.

3.1 REFLECTION AND PERMUTATION EQUIVARIANCE

As mentioned earlier, we set the coordinate axes as the principal axes of rotation by construction. This canonicalizes the 3D structure, removing translational symmetries and reducing the rotational symmetries to sign-flip symmetries of the eigenvectors of \mathbf{P} , which are the coordinate axes (Puny et al., 2022; Duval et al., 2023). Hence, the flow needs to be equivariant with respect to sign-flips of the coordinate axes, which we call “reflection-equivariance” for brevity (Lim et al., 2023; Cheng et al., 2024). In addition, because 3D structure is invariant under node order permutations, the learned velocity should be equivariant to permutations. Together, the explicit equivariance constraints on the network are given as $v_\theta(t, \mathbf{\Pi U}_t \mathbf{R}) = \mathbf{\Pi} v_\theta(t, \mathbf{U}_t) \mathbf{R}$, for all node permutations $\mathbf{\Pi}$ and reflections $\mathbf{R} = \text{diag}(\pm 1, \pm 1, \pm 1, 1)$. We satisfy these constraints with a reflection-equivariant graph neural network architecture described in Appendix C.1.

3.2 EQUIVARIANT OPTIMAL TRANSPORT

Flow matching learns a velocity field which transports samples from the noise distribution to the data distribution, but there is no guarantee that samples will follow paths γ that are optimal with respect to transport cost $L(\gamma)$. Optimal paths are desired because they afford more efficient training and faster generation (Pooladian et al., 2023; Tong et al., 2024). During training, the model should be encouraged to learn dynamics which follow paths with lower transport cost. Since molecules have permutation and rotational symmetries, optimal paths should connect structures whose equivalence classes are close to each other. This corresponds to finding optimal node permutations and rotations to align noise samples \mathbf{X}_0 to data samples \mathbf{X}_1 . In contrast to the Euclidean case (Klein et al., 2023; Song et al., 2023c), only reflections are needed, as opposed to rotations, because the coordinate axes are already fixed in place by the principal axes of rotation, enforced by the orthogonality condition.

In addition, transport cost on $\text{St}(n, 4)$ must be measured using the Riemannian distance. Thus, the optimal transport map from \mathbf{U}_0 to \mathbf{U}_1 minimizes the following cost over atom-type-preserving node permutations $\mathbf{\Pi}$ (i.e., $\mathbf{\Pi a} = \mathbf{a}$) and reflections $\mathbf{R} = \text{diag}(\pm 1, \pm 1, \pm 1, 1)$:

$$c(\mathbf{U}_0, \mathbf{U}_1) = \|\log_{\mathbf{\Pi U}_0 \mathbf{R}}(\mathbf{U}_1)\|_{\mathbf{\Pi U}_0 \mathbf{R}}. \quad (11)$$

While searching for optimal alignments, we do not compute the full Stiefel logarithm and instead approximately calculate distance using only one iteration of the inner loop described in Appendix

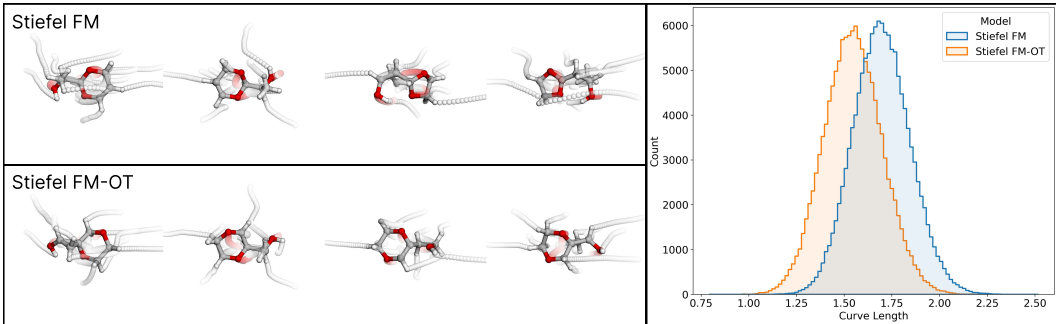


Figure 2: (Left) Learned sampling trajectories for Stiefel FM and Stiefel FM-OT on QM9. Each column begins generation from the same noise. (Right) Histogram of curve lengths of all QM9 sampling trajectories for Stiefel FM and Stiefel FM-OT. Permutation and reflection alignment lead to simpler and shorter paths.

Algorithm 2, and justified in Appendix Figure 11. This approximate distance is then heuristically optimized over atom permutations and reflections with a greedy random local search. Appendix Algorithm 4 outlines this procedure, which samples several U matrices for each reflection to approximately identify the best reflection, and then performs a local search of random atom type-preserving index swaps (Appendix Figure 9).

4 EXPERIMENTS

We evaluate Euclidean diffusion models and Stiefel Flow Matching on the QM9 and GEOM datasets. For each example, the model takes in moments and molecular formula and produces $K = 10$ samples.

Datasets. For QM9 (Ramakrishnan et al., 2014), we use the conformers provided by the GEOM dataset. We abbreviate GEOM-Drugs (Axelrod & Gomez-Bombarelli, 2022) as GEOM. We use the same training, validation, and test splits as Cheng et al. (2024), except that examples which are unstable or have less than 5 atoms are removed, which drops 76 examples from QM9 and drops no examples from GEOM. QM9 has train/val/test splits of 104265/13056/13033 molecules, while GEOM has splits of 233625/29203/29203 molecules, or 5537598/29203/29203 conformers. We only predict 3D structures for the lowest-energy conformers of GEOM. This reflects experimental reality, as the lowest-energy conformer typically has the highest proportion in the population after cooling by supersonic jet expansion, assuming thermodynamic equilibrium (Ruoff et al., 1990).

Structure elucidation. The only measure of a model’s success is its ability to generate the correct 3D structure X with high accuracy at least once, where accuracy is measured by root-mean-squared-deviation (RMSD) of the predicted coordinates to the ground truth. High accuracy is needed because the only confirming evidence available to us is (1) agreement with moments and (2) thermodynamic stability by quantum chemistry. For this reason, we use stringent thresholds of $\text{RMSD} < 0.25 \text{ \AA}$ and $\text{RMSD} < 0.10 \text{ \AA}$ to ensure that the predicted structure is in the same potential energy basin as the true structure. Success is reported if the minimum RMSD over $K = 10$ generated samples satisfies these thresholds. The success rate is the percentage of the test set whose generated samples has a minimum RMSD which satisfies each threshold. Error bars are standard errors of the mean.

Evaluating RMSD is nontrivial because the generated and ground truth structure must first be aligned under same-atom-type permutations and reflections. We use the same RMSD procedure as Cheng et al. (2024): We first align node permutations by solving a linear assignment problem whose cost matrix is squared Euclidean distance, and repeat this for all 8 reflections, taking the minimum. Then we compute RMSD between the ground truth coordinates and the aligned coordinates. This is similar to the alignment procedure used in (Klein et al., 2023; Song et al., 2023c).

Auxiliary metrics. To characterize the samples generated by each model, we report additional metrics. In contrast to success rate, which checks the minimum RMSD of generated samples for each example, these auxiliary metrics are averaged over all generated samples (except diversity). Error measures how much the generated structure violates the moment constraints. If \hat{P} is the computed

Table 1: Experimental results on QM9. Stiefel FM shows no violation of moment constraints as shown in the **Error** metrics, and has the highest success rate for structure elucidation, with the lowest computational cost.

Method	% < RMSD \uparrow		Error \downarrow	Valid \uparrow	Stable \downarrow	Diverse \uparrow	NFE \downarrow
	0.25 Å	0.10 Å					
Stiefel Random	0.00 \pm 0.00	0.00 \pm 0.00	0.00	0.061	nan	2.640	0
KREED	11.22 \pm 0.28	9.55 \pm 0.26	5.18	0.878	-1.335	1.429	1000
KREED-XL	13.65 \pm 0.30	10.94 \pm 0.27	3.64	0.933	-1.048	0.870	1000
KREED-XL-DPS	12.36 \pm 0.29	9.40 \pm 0.26	1.33	0.744	-0.826	1.060	1000
KREED-XL-proj	13.67 \pm 0.30	10.93 \pm 0.27	0.00	0.924	-0.905	0.871	1000
Stiefel FM	15.17\pm0.31	13.82\pm0.30	0.00	0.882	-1.125	1.040	200
Stiefel FM-OT	13.99 \pm 0.30	12.68 \pm 0.29	0.00	0.835	-1.039	1.045	200




Figure 3: Histograms of minimum RMSD for predicted QM9 examples show two distinct clusters for RMSD. The 0.25Å threshold captures molecular structures that are useful for structure elucidation.

planar dyadic of the generated structure, this is computed as $\frac{1}{\sqrt{6}} \|\text{triu}(\hat{\mathbf{P}} - \mathbf{P})\|_2$, where triu takes the upper triangular part of the matrix (contains 6 elements). *Validity* is a heuristic check based on bond detection with `rdDetermineBonds.DetermineConnectivity` (Landrum, 2013; Kim & Kim, 2015). *Stability* is the log norm of the gradient of energy with respect to coordinates, as reported by the `xtb` quantum chemistry program (Bannwarth et al., 2019). Stable structures should have a gradient norm close to zero (log norm very negative), assuming the structure is at a local minimum and not a saddle point. *Diversity* is calculated as the average pairwise RMSD of all generated samples of a single example. *NFE* is the number of function evaluations used during generation, and measures computational cost. We do not set boldface for these metrics because they do not correspond directly to success criteria.

Baselines. Given the novelty of the problem, the number of available baselines is limited. We compare the performance of Stiefel FM to KREED (Cheng et al., 2024). **KREED is a reflection-equivariant diffusion model trained to generate 3D structure conditioned on molecular formula and moments of inertia, and is a specialization of E(3)-equivariant approaches like EDM (Hoogeboom et al., 2022).** Since our architecture for Stiefel FM is much larger (parameter-wise) than the model architecture used in KREED, **and because KREED is tailored for a slightly different task,** we also train another reflection-equivariant diffusion model with an identical neural network architecture to Stiefel Flow Matching, which we label as KREED-XL. The planar dyadic is computable at every step of the generation process, which means that this task can be treated as a nonlinear inverse problem: On top of KREED-XL, we apply Diffusion Posterior Sampling (DPS) (Chung et al., 2023), **which guides generation with an additional drift term for minimizing the planar dyadic error. As a simple baseline which exactly satisfies moment constraints,** we report performance for uniform random

Table 2: Experimental results on GEOM. Stiefel FM generates few valid structures on its own due to the increased difficulty of manifold-constrained generative modelling. When adjusted to generate the same number of valid molecules, Stiefel FM-OT (filter) obtains the highest success rate, without surpassing baselines in computational cost.

Method	% < RMSD \uparrow		Error \downarrow	Valid \uparrow	Stable \downarrow	Diverse \uparrow	NFE \downarrow
	0.25 Å	0.10 Å					
Stiefel Random	0.00 \pm 0.00	0.00 \pm 0.00	0.00	0.000	nan	4.104	0
KREED	0.05 \pm 0.01	0.02 \pm 0.01	58.36	0.353	-0.583	2.297	1000
KREED-XL	3.54 \pm 0.11	2.02 \pm 0.08	30.71	0.907	-0.900	2.190	1000
KREED-XL-proj	3.54 \pm 0.11	2.04 \pm 0.08	0.00	0.904	-0.752	2.188	1000
Stiefel FM	2.17 \pm 0.08	1.24 \pm 0.07	0.00	0.388	-0.066	2.212	200
Stiefel FM-OT	2.45 \pm 0.09	1.49 \pm 0.07	0.00	0.376	-0.002	2.195	200
Stiefel FM (filter)	3.57 \pm 0.11	2.06 \pm 0.08	0.00	0.889	-0.437	2.208	600
Stiefel FM-OT (filter)	3.79\pm0.11	2.34\pm0.09	0.00	0.869	-0.352	2.193	600




Figure 4: Histograms of minimum RMSD for predicted GEOM examples.

sampling on the Stiefel manifold. We also report the results of KREED-XL after projecting samples to the feasible manifold. Relevant hyperparameters for all methods are provided in Appendix C.3.

Results. Table 1 shows our experimental results on QM9, reporting Stiefel FM with and without optimal transport (OT). We note that the average number of atoms in QM9 is 18, meaning that on average the model must infer $3(18) - 9 = 39$ values from 3 moments. We find that Stiefel FM can generate the correct structure with a greater success rate than all Euclidean diffusion models. We also see that incorporating the analytic formula of the moments via DPS does improve agreement with the moments, but at the cost of accuracy. In contrast, Stiefel FM does not suffer from this tradeoff. Projecting samples onto the manifold does not change success rate, because the projection leaves correct structures untouched, while only distorting incorrect structures. In addition, Stiefel FM uses only 20% of the computation used by diffusion models. Unconstrained diffusion models can produce more valid and stable structures, though they may not necessarily generate the correct structure. Figure 3 reveals that when Stiefel FM’s predictions are correct, it is likely to be extremely accurate, achieving RMSD even below 0.05 Å. Training with equivariant optimal transport helps learn simpler generation paths and reduces the average curve length of generation trajectories from 1.696 to 1.547 (Figure 2). Interestingly, training with optimal transport slightly reduces the success rate of Stiefel FM on QM9, though this trend is reversed for GEOM. Appendix Table 5 reports additional trials experimenting with optimal transport, logit-normal timestep sampling (Esser et al., 2024), and stochasticity (Bose et al., 2024).

The task of structure elucidation poses an even harder challenge on GEOM, with an average of 46 atoms, giving a task of inferring $3(46) - 9 = 129$ values. Naively, one would expect that this problem is hopelessly underconstrained. Nevertheless, Table 2 shows that both diffusion models

432 and flow matching can obtain a nontrivial success rate for structure elucidation. However, the low
433 accuracy, validity, and stability of Stiefel FM(-OT) suggests that the model is underfitting the dataset,
434 even as KREED-XL is able to generate numerous valid samples. We observe that less than half
435 of the structures generated by Stiefel FM are valid, which aligns with the fact that *there are fewer*
436 *valid structures on the feasible manifold \mathcal{M} than in regular Euclidean space*. This suggests that a
437 fairer evaluation utilizing the strength of the manifold constraint should generate a similar number of
438 *valid* structures for each model. This is justified by the fact that validity does not use ground truth
439 labels and can be computed at only nominal cost. Therefore, we generate 30 samples for Stiefel FM
440 and Stiefel FM-OT before filtering to retain up to $K = 10$ valid samples. Note that the combined
441 computational cost of generating 3x as many samples (3 x 200 NFE) is still lower than KREED-XL
442 (1000 NFE). After filtering, Stiefel FM-OT obtains a similar validity rate to KREED-XL, but obtains
443 the highest success rate for structure elucidation. Even after filtering, the mediocre stability of
444 generated samples suggests that success rate can be improved further, though at the slightly higher
445 cost of quantum chemistry calculations. Now, we see that optimal transport does help Stiefel FM to
446 predict accurate structures, while also reducing average curve length from 1.421 to 1.344 (Appendix
447 Figure 8). However, biasing generation by selecting for valid samples seems to also select for longer
448 generation trajectories. *We find that generation paths that land on the correct structure are usually*
449 *longer than generation paths that land on incorrect structures (Appendix Figure 7) This may be*
450 *explained by the fact that initial points are sampled anywhere uniformly on the manifold, but for*
451 *success they must end up on the single true structure. In contrast, there are many incorrect structures*
all over the manifold, which may end up on average closer to random initial points.

452 It should be noted that one should expect success rates to be of this magnitude for solving a heavily
453 underconstrained problem. This is simply due to the fact that there exist many stable structures
454 with the same molecular formula and very similar moments of inertia. Indeed, trained models
455 usually generate realistic-looking molecules – see generated examples in Appendix Figure 5 and
456 Figure 6. Furthermore, only 10 to 30 samples were queried for each molecule, but an actual structure
457 elucidation campaign would have a much larger compute budget for generating thousands of samples.
458 *As a highlight, our results show that it is actually possible at 0.25 Å resolution to elucidate 27.4% of*
459 *the test set of QM9 (3580/13033) when combining KREED-XL, Stiefel FM, and Stiefel FM-OT; and*
460 *7.9% of the test set of GEOM (2297/29203), when combining KREED-XL, Stiefel FM (filter), and*
461 *Stiefel FM-OT (filter).*

462 5 RELATED WORK

463 **Generative models for 3D molecules.** While research in applying generative models for 3D structure,
464 only few works have applied these methods to the task of structure elucidation (Cheng et al., 2024).
465 Adjacent work in applying diffusion and flow matching models to 3D molecules include molecular
466 generation (Hoogeboom et al., 2022; Song et al., 2023c), conformer search (Jing et al., 2022; Xu et al.,
467 2022), docking (Corso et al., 2022), biomolecular assembly (Abramson et al., 2024) and Boltzmann
468 generators (Klein et al., 2023). Generative models on Riemannian manifolds have also been applied
469 to protein design (Bose et al., 2024; Yim et al., 2023) and crystal structure prediction (Jiao et al.,
470 2024). Recent work has applied generative modelling to predict 3D structure from powder X-ray
471 diffraction patterns (Lai et al., 2024; Riesel et al., 2024).

472 **Deep learning for molecular identification using rotational spectroscopy.** A limited number of
473 works have provided parts of the overall puzzle needed for structure elucidation using rotational
474 spectroscopy. Zaleski & Prozument (2018) propose RAINet as a forward modeling approach for
475 assigning rotational spectra, where a set of peaks is fed into a classifier, that categorizes the spectra
476 to an appropriate multilayer perceptron that outputs spectroscopic parameters, including moments.
477 McCarthy & Lee (2020) adopt a “mixture-of-experts” approach that maps spectroscopic parameters
478 and approximate molecular formula into a set of complementary experimental observables and
479 SMILES strings, though with limited success. Cheng et al. (2024) present a Euclidean diffusion
480 model KREED for determining 3D structure from moments, molecular formula, and unsigned
481 substitution coordinates, the latter of which is also measurable from rotational spectroscopy, but
482 can be difficult and expensive to obtain. Most recently, Schwarting et al. (2024) provide a thorough
483 analysis into the inverse problem, examining the frequency with which different molecules have
484 moments of inertia that are very close in value. Stiefel Flow Matching could disambiguate these
485 structures by providing a diversity of structures that satisfy moment constraints exactly.

Statistics on the Stiefel manifold. While the Stiefel manifold is often studied in the context of optimization (Absil et al., 2008; Chen et al., 2021; Kong et al., 2022), a number of works study probability distributions on the Stiefel manifold (Chakraborty & Vemuri, 2019; Chikuse, 1990). One distribution on the Stiefel manifold is called the *matrix von Mises-Fisher distribution* or *matrix Langevin distribution* (Pal et al., 2020; Chikuse, 2003; Jupp & Mardia, 1979). Wang & Solo (2020) propose a particle filtering algorithm on the Stiefel manifold, with the first application of optimal transport on the Stiefel manifold. Yataka et al. (2023) propose a continuous normalizing flow on the Grassmann manifold, which is closely related to the Stiefel manifold.

6 CONCLUSION

We propose Stiefel Flow Matching, a Riemannian generative model for generating samples subject to exact orthogonality constraints, and apply it to the challenging inverse problem of structure elucidation from moments of inertia and molecular formula. Empirically, Stiefel Flow Matching achieves a higher success rate than other Euclidean diffusion approaches. Satisfying the constraints exactly will enable future advances in Riemannian generative modelling to directly transfer to generating more stable molecules, without needing to consider agreement to the moments.

6.1 FUTURE RESEARCH DIRECTIONS

Improving Stiefel generative models. Riemannian flow matching empirically shows degradation compared to Riemannian diffusion (Lou et al., 2024; Zhu et al., 2024). This has been attributed to two pathologies of Riemannian flow matching for compact manifolds: (1) the geodesic-based velocity field is discontinuous at the cut locus (Lou et al., 2024; Zhu et al., 2024), and (2) the probability density has a shrinking support (Stark et al., 2024; Holderrieth et al., 2024). These pathologies may explain the difficulty of Stiefel FM in fitting GEOM, and motivate the development of alternative probability paths for Stiefel flow matching, such as Stiefel diffusion (De Bortoli et al., 2022), or flows which asymptotically land on the Stiefel manifold (Ablin & Peyré, 2022; Gao et al., 2022). Training a Stiefel diffusion model with the denoising score matching objective requires the heat kernel as training targets for the neural network, which can be computed using techniques by Azangulov et al. (2022) and Lou et al. (2024). Alternatively, the matrix Langevin distribution may be amenable to modelling with Star-Shaped Denoising Diffusion Probabilistic Models (Okhotin et al., 2024). Future work can also explore varying the metric used on the Stiefel manifold, since the Stiefel manifold admits a 1-parameter family of metrics generalizing the canonical metric (Hüper et al., 2021). These metrics have efficient numerical algorithms for the exponential and logarithm (Mataigne et al., 2024). Sample-time advances in flow matching, such as corrector sampling (Gat et al., 2024) or enhancing the flow with a jump process (Holderrieth et al., 2024), are also orthogonal avenues for improvement.

A key limitation of Stiefel Flow Matching is the requirement of molecular formula as input. Recent approaches in discrete flow matching (Campbell et al., 2024; Gat et al., 2024) could enable a multimodal flow to simultaneously vary both continuous atom positions and discrete atom types. The jump processes of generator matching (Holderrieth et al., 2024) are particularly natural for this problem, as they could allow the flow to jump between Stiefel manifolds of different sizes.

Modeling on real-world chemical data. Since we require only minimal information in the moments and molecular formula, another direction is to incorporate other conditioning information, such as energy and force information, fragments of 2D graphs, dipole moments, or other sources of analytical chemistry data. We can do so through MCMC sampling (Du et al., 2023), guidance (Song et al., 2023a;b; Mardani et al., 2023), or diversity sampling approaches (Corso et al., 2024). This presents an opportunity in mass spectrometry, as the 3D structural information provided by the moments can distinguish molecules which have the same mass. Controllable diversity can also be leveraged here since, in a molecule identification campaign, we can eliminate candidates once we know they are not correct (by calculating the rovibrational correction and comparing to experiment). We then want to sample structures which are different from these eliminated candidates.

Additional applications of Stiefel Flow Matching. Stiefel generative models can also be applied to other domains of orthogonality constrained data, such as molecular orbitals (Mrovec & Berger, 2021; Aoto & da Silva, 2021), orthogonal neural network weights (Kong et al., 2022), and covariance matrices in neural data (Nejatbakhsh et al., 2024).

REFERENCES

- 540 Pierre Ablin and Gabriel Peyré. Fast and accurate optimization on the orthogonal manifold without
541 retraction. In *International Conference on Artificial Intelligence and Statistics*, pp. 5636–5657.
542 PMLR, 2022.
- 543 Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf
544 Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure
545 prediction of biomolecular interactions with alphafold 3. *Nature*, pp. 1–3, 2024.
- 546 P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*.
547 Princeton University Press, 2008.
- 548 Yuri Alexandre Aoto and Márcio Fabiano da Silva. An optimisation on the grassmannian with
549 applications to quantum chemistry. *Proceeding Series of the Brazilian Society of Computational
550 and Applied Mathematics*, 8(1), 2021.
- 551 Simon Axelrod and Rafael Gomez-Bombarelli. GEOM, energy-annotated molecular conformations
552 for property prediction and molecular generation. *Scientific Data*, 9(1):1–14, 2022.
- 553 Iskander Azangulov, Andrei Smolensky, Alexander Terenin, and Viacheslav Borovitskiy. Stationary
554 kernels and gaussian processes on lie groups and their homogeneous spaces i: the compact case.
555 *arXiv preprint arXiv:2208.14960*, 2022.
- 556 Christoph Bannwarth, Sebastian Ehlert, and Stefan Grimme. Gfn2-xtb—an accurate and broadly
557 parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics
558 and density-dependent dispersion contributions. *Journal of chemical theory and computation*, 15
559 (3):1652–1671, 2019.
- 560 Thomas Bendokat, Ralf Zimmermann, and P-A Absil. A grassmann manifold handbook: Basic
561 geometry and computational aspects. *Advances in Computational Mathematics*, 50(1):1–51, 2024.
- 562 Joey Bose, Tara Akhound-Sadegh, Guillaume Hugué, Kilian FATRAS, Jarrid Rector-Brooks, Cheng-
563 Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael M. Bronstein, and Alexander Tong.
564 SE(3)-stochastic flow matching for protein backbone generation. In *The Twelfth International
565 Conference on Learning Representations*, 2024. URL [https://openreview.net/forum?
566 id=kJFIH23hXb](https://openreview.net/forum?id=kJFIH23hXb).
- 567 Gordon G Brown, Brian C Dian, Kevin O Douglass, Scott M Geyer, and Brooks H Pate. The
568 rotational spectrum of epifluorohydrin measured by chirped-pulse fourier transform microwave
569 spectroscopy. *Journal of Molecular Spectroscopy*, 238(2):200–212, 2006.
- 570 Gordon G Brown, Brian C Dian, Kevin O Douglass, Scott M Geyer, Steven T Shipman, and Brooks H
571 Pate. A broadband fourier transform microwave spectrometer based on chirped pulse excitation.
572 *Review of Scientific Instruments*, 79(5), 2008.
- 573 Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative
574 flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design.
575 *arXiv preprint arXiv:2402.04997*, 2024.
- 576 Rudrasis Chakraborty and Baba C Vemuri. Statistics on the stiefel manifold: theory and applications.
577 2019.
- 578 Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *The Twelfth
579 International Conference on Learning Representations*, 2024. URL [https://openreview.
580 net/forum?id=g7ohD1TITL](https://openreview.net/forum?id=g7ohD1TITL).
- 581 Shixiang Chen, Alfredo Garcia, Mingyi Hong, and Shahin Shahrampour. Decentralized riemannian
582 gradient descent on the stiefel manifold. In *International Conference on Machine Learning*, pp.
583 1594–1605. PMLR, 2021.
- 584 Austin H Cheng, Alston Lo, Santiago Miret, Brooks H Pate, and Alán Aspuru-Guzik. Determining
585 3d structure from molecular formula and isotopologue rotational spectra in natural abundance with
586 reflection-equivariant diffusion. *The Journal of Chemical Physics*, 160(12), 2024.

- 594 Yasuko Chikuse. Distributions of orientations on stiefel manifolds. *Journal of multivariate analysis*,
595 33(2):247–264, 1990.
- 596
- 597 Yasuko Chikuse. Concentrated matrix langevin distributions. *Journal of Multivariate Analysis*, 85(2):
598 375–394, 2003.
- 599
- 600 Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye.
601 Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International*
602 *Conference on Learning Representations*, 2023. URL [https://openreview.net/forum?](https://openreview.net/forum?id=OnD9zGAGT0k)
603 [id=OnD9zGAGT0k](https://openreview.net/forum?id=OnD9zGAGT0k).
- 604 Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock:
605 Diffusion steps, twists, and turns for molecular docking. *arXiv preprint arXiv:2210.01776*, 2022.
- 606 Gabriele Corso, Yilun Xu, Valentin De Bortoli, Regina Barzilay, and Tommi S. Jaakkola. Particle
607 guidance: non-i.i.d. diverse sampling with diffusion models. In *The Twelfth International Confer-*
608 *ence on Learning Representations*, 2024. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=KqbCvIFBY7)
609 [KqbCvIFBY7](https://openreview.net/forum?id=KqbCvIFBY7).
- 610
- 611 Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and
612 Arnaud Doucet. Riemannian score-based generative modelling. *Advances in Neural Information*
613 *Processing Systems*, 35:2406–2422, 2022.
- 614 J. Demaison, James E. Boggs, and Attila G. Császár (eds.). *Equilibrium Molecular Structures: From*
615 *Spectroscopy to Quantum Chemistry*. CRC Press, Boca Raton, 2011. ISBN 978-1-4398-1132-0.
- 616
- 617 Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis.
618 In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural*
619 *Information Processing Systems*, 2021.
- 620 Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H
621 Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, et al. Continuous diffusion
622 for categorical data. *arXiv preprint arXiv:2211.15089*, 2022.
- 623
- 624 Sérgio R. Domingos, Cristóbal Pérez, Mark D. Marshall, Helen O. Leung, and Melanie Schnell.
625 Assessing the performance of rotational spectroscopy in chiral analysis. *Chemical Science*, 11(40):
626 10863–10870, October 2020. ISSN 2041-6539. doi: 10.1039/D0SC03752D.
- 627 Yilun Du, Conor Durkan, Robin Strudel, Joshua B Tenenbaum, Sander Dieleman, Rob Fergus, Jascha
628 Sohl-Dickstein, Arnaud Doucet, and Will Sussman Grathwohl. Reduce, reuse, recycle: Composi-
629 tional generation with energy-based diffusion models and mcmc. In *International conference on*
630 *machine learning*, pp. 8489–8510. PMLR, 2023.
- 631
- 632 Alexandre Agm Duval, Victor Schmidt, Alex Hernández-García, Santiago Miret, Fragkiskos D
633 Malliaros, Yoshua Bengio, and David Rolnick. Faenet: Frame averaging equivariant gnn for
634 materials modeling. In *International Conference on Machine Learning*, pp. 9013–9033. PMLR,
635 2023.
- 636
- 637 Alan Edelman, Tomás A Arias, and Steven T Smith. The geometry of algorithms with orthogonality
638 constraints. *SIAM journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.
- 639
- 640 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
641 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for
642 high-resolution image synthesis. *arXiv preprint arXiv:2403.03206*, 2024.
- 643
- 644 Matthias Fey and Jan Eric Lenssen. Fast Graph Representation Learning with PyTorch Geometric,
645 May 2019. URL https://github.com/pyg-team/pytorch_geometric.
- 646
- 647 Bin Gao, Simon Vary, Pierre Ablin, and P-A Absil. Optimization flows landing on the stiefel manifold.
IFAC-PapersOnLine, 55(30):25–30, 2022.
- Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *arXiv preprint arXiv:2407.15595*, 2024.

- 648 Walter Gordy, Robert L Cook, and Arnold Weissberger. *Microwave molecular spectra*, volume 18. Wi-
649 ley New York, 1984. URL [https://app.knovel.com/kn/resources/kpMMSE0001/
650 toc](https://app.knovel.com/kn/resources/kpMMSE0001/toc).
- 651 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
652 neural information processing systems*, 33:6840–6851, 2020.
- 653 Peter Holderrieth, Marton Havasi, Jason Yim, Neta Shaul, Itai Gat, Tommi Jaakkola, Brian Karrer,
654 Ricky TQ Chen, and Yaron Lipman. Generator matching: Generative modeling with arbitrary
655 markov processes. *arXiv preprint arXiv:2410.20587*, 2024.
- 656 Emiel Hoogetboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion
657 for molecule generation in 3d. In *International conference on machine learning*, pp. 8867–8887.
658 PMLR, 2022.
- 659 Knut Hüper, Irina Markina, and Fátima Silva Leite. A lagrangian approach to extremal curves on
660 stiefel manifolds. *AIMS*, 2021.
- 661 Rui Jiao, Wenbing Huang, Peijia Lin, Jiaqi Han, Pin Chen, Yutong Lu, and Yang Liu. Crystal
662 structure prediction by joint equivariant diffusion. *Advances in Neural Information Processing
663 Systems*, 36, 2024.
- 664 Bowen Jing, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. Torsional
665 diffusion for molecular conformer generation. *Advances in Neural Information Processing Systems*,
666 35:24240–24253, 2022.
- 667 Peter E Jupp and Kanti V Mardia. Maximum likelihood estimators for the matrix von mises-fisher
668 and bingham distributions. *The Annals of Statistics*, 7(3):599–606, 1979.
- 669 Yeonjoon Kim and Woo Youn Kim. Universal structure conversion method for organic molecules:
670 from atomic connectivity to three-dimensional geometry. *Bulletin of the Korean Chemical Society*,
671 36(7):1769–1777, 2015.
- 672 Leon Klein, Andreas Krämer, and Frank Noe. Equivariant flow matching. In *Thirty-seventh
673 Conference on Neural Information Processing Systems*, 2023. URL [https://openreview.
674 net/forum?id=eLH2NF001B](https://openreview.net/forum?id=eLH2NF001B).
- 675 Lingkai Kong, Yuqing Wang, and Molei Tao. Momentum stiefel optimizer, with applications to
676 suitably-orthogonal attention, and optimal transport. *arXiv preprint arXiv:2205.14173*, 2022.
- 677 J Kraitchman. Determination of molecular structure from microwave spectroscopic data. *American
678 Journal of Physics*, 21(1):17–24, 1953.
- 679 Qingsi Lai, Lin Yao, Zhifeng Gao, Siyuan Liu, Hongshuai Wang, Shuqi Lu, Di He, Liwei Wang,
680 Cheng Wang, and Guolin Ke. End-to-end crystal structure prediction from powder x-ray diffraction.
681 *arXiv preprint arXiv:2401.03862*, 2024.
- 682 Greg Landrum. Rdkit documentation. *Release*, 1(1-79):4, 2013.
- 683 John M. Lee. *Smooth Manifolds*, pp. 1–29. Springer New York, New York, NY, 2003. ISBN 978-
684 0-387-21752-9. doi: 10.1007/978-0-387-21752-9_1. URL [https://doi.org/10.1007/
685 978-0-387-21752-9_1](https://doi.org/10.1007/978-0-387-21752-9_1).
- 686 Kin Long Kelvin Lee and Michael McCarthy. Study of Benzene Fragmentation, Isomerization,
687 and Growth Using Microwave Spectroscopy. *The Journal of Physical Chemistry Letters*, 10(10):
688 2408–2413, April 2019. ISSN 1948-7185. doi: 10.1021/acs.jpcllett.9b00586.
- 689 Kin Long Kelvin Lee and Michael McCarthy. Bayesian Analysis of Theoretical Rotational Constants
690 from Low-Cost Electronic Structure Methods. *The Journal of Physical Chemistry A*, 124(5):
691 898–910, February 2020. ISSN 1089-5639. doi: 10.1021/acs.jpca.9b09982.
- 692 Derek Lim, Joshua Robinson, Stefanie Jegelka, and Haggai Maron. Expressive sign equivariant
693 networks for spectral geometric learning. In *Thirty-seventh Conference on Neural Information
694 Processing Systems*, 2023. URL <https://openreview.net/forum?id=UWd4ysACo4>.

- 702 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
703 matching for generative modeling. In *The Eleventh International Conference on Learning Repre-*
704 *sentations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.
705
- 706 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
707 transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.
- 708 Aaron Lou, Minkai Xu, Adam Farris, and Stefano Ermon. Scaling riemannian diffusion models.
709 *Advances in Neural Information Processing Systems*, 36, 2024.
710
- 711 Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving
712 inverse problems with diffusion models. *arXiv preprint arXiv:2305.04391*, 2023.
713
- 714 Alan G Marshall and Christopher L Hendrickson. High-resolution mass spectrometers. *Annu. Rev.*
715 *Anal. Chem.*, 1:579–599, 2008.
- 716 Simon Maigne, Ralf Zimmermann, and Nina Miolane. An efficient algorithm for the riemannian log-
717 arithm on the stiefel manifold for a family of riemannian metrics. *arXiv preprint arXiv:2403.11730*,
718 2024.
719
- 720 Michael McCarthy and Kin Long Kelvin Lee. Molecule Identification with Rotational Spectroscopy
721 and Probabilistic Deep Learning. *The Journal of Physical Chemistry A*, 124(15):3002–3017, April
722 2020. ISSN 1089-5639, 1520-5215. doi: 10.1021/acs.jpca.0c01376.
- 723 Michael C. McCarthy, Kin Long Kelvin Lee, P. Brandon Carroll, Jessica P. Porterfield, P. Bryan
724 Changala, James H. Thorpe, and John F. Stanton. Exhaustive Product Analysis of Three Benzene
725 Discharges by Microwave Spectroscopy. *The Journal of Physical Chemistry A*, 124(25):5170–5181,
726 June 2020. ISSN 1089-5639, 1520-5215. doi: 10.1021/acs.jpca.0c02919.
727
- 728 Brett A McGuire. 2018 census of interstellar, circumstellar, extragalactic, protoplanetary disk, and
729 exoplanetary molecules. *The Astrophysical Journal Supplement Series*, 239(2):17, 2018.
730
- 731 Brett A. McGuire, P. Brandon Carroll, Ryan A. Loomis, Ian A. Finneran, Philip R. Jewell, Anthony J.
732 Remijan, and Geoffrey A. Blake. Discovery of the interstellar chiral molecule propylene oxide
733 (CH₃CHCH₂O). *Science*, 352(6292):1449–1452, June 2016. doi: 10.1126/science.aae0328.
- 734 Benedetta Giannelli Moneta, Sara Elsa Aita, Elena Barbaro, Anna Laura Capriotti, Andrea Cerrato,
735 Aldo Laganà, Carmela Maria Montone, Susy Piovesana, Federico Scoto, Carlo Barbante, et al.
736 Untargeted analysis of environmental contaminants in surface snow samples of svalbard islands by
737 liquid chromatography-high resolution mass spectrometry. *Science of the Total Environment*, 858:
738 159709, 2023.
739
- 740 Martin Mrovec and JA Berger. A diagonalization-free optimization algorithm for solving kohn–sham
741 equations of closed-shell molecules. *Journal of Computational Chemistry*, 42(7):492–504, 2021.
- 742 Amin Nejatbakhsh, Isabel Garon, and Alex Williams. Estimating noise correlations across continuous
743 conditions with wishart processes. *Advances in Neural Information Processing Systems*, 36, 2024.
744
- 745 Andrey Okhotin, Dmitry Molchanov, Arkhipkin Vladimir, Grigory Bartosh, Viktor Ohanesian, Aibek
746 Alanov, and Dmitry P Vetrov. Star-shaped denoising diffusion probabilistic models. *Advances in*
747 *Neural Information Processing Systems*, 36, 2024.
- 748 Subhadip Pal, Subhajit Sengupta, Riten Mitra, and Arunava Banerjee. Conjugate priors and posterior
749 inference for the matrix langevin distribution on the stiefel manifold. 2020.
750
- 751 Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman,
752 and Ricky T. Q. Chen. Multisample flow matching: Straightening flows with minibatch couplings.
753 In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and
754 Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*,
755 volume 202 of *Proceedings of Machine Learning Research*, pp. 28100–28127. PMLR, 23–29 Jul
2023. URL <https://proceedings.mlr.press/v202/pooladian23a.html>.

- 756 Omri Puny, Matan Atzmon, Edward J. Smith, Ishan Misra, Aditya Grover, Heli Ben-Hamu, and
757 Yaron Lipman. Frame averaging for invariant and equivariant network design. In *International*
758 *Conference on Learning Representations*, 2022. URL [https://openreview.net/forum?](https://openreview.net/forum?id=zIUyj55nXR)
759 [id=zIUyj55nXR](https://openreview.net/forum?id=zIUyj55nXR).
- 760 Cristina Puzzarini and John F Stanton. Connections between the accuracy of rotational constants and
761 equilibrium molecular structures. *Physical Chemistry Chemical Physics*, 25(3):1421–1429, 2023.
- 762 Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum
763 chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014.
- 764 Eric A Riesel, Tsach Mackey, Hamed Nilforoshan, Minkai Xu, Catherine K Badding, Alison B
765 Altman, Jure Leskovec, and Danna E Freedman. Crystal structure determination from powder
766 diffraction patterns with generative machine learning. *Journal of the American Chemical Society*,
767 2024.
- 768 RS Ruoff, TD Klots, T Emilsson, and HS Gutowsky. Relaxation of conformers and isomers in seeded
769 supersonic jets of inert gases. *The Journal of chemical physics*, 93(5):3142–3150, 1990.
- 770 Marcus Schwarting, Nathan A. Seifert, Michael J. Davis, Ben Blaiszik, Ian Foster, and Kirill
771 Prozument. Twins in rotational spectroscopy: Does a rotational spectrum uniquely identify a
772 molecule?, April 2024.
- 773 Steven T Shipman, Justin L Neill, Richard D Suenram, Matt T Muckle, and Brooks H Pate. Structure
774 determination of strawberry aldehyde by broadband microwave spectroscopy: Conformational
775 stabilization by dispersive interactions. *The Journal of Physical Chemistry Letters*, 2(5):443–448,
776 2011.
- 777 Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion
778 models for inverse problems. In *International Conference on Learning Representations*, 2023a.
779 URL https://openreview.net/forum?id=9_gsMA8MRKQ.
- 780 Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin
781 Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable generation.
782 In *International Conference on Machine Learning*, pp. 32483–32498. PMLR, 2023b.
- 783 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
784 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*
785 *arXiv:2011.13456*, 2020.
- 786 Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging
787 with score-based generative models. In *International Conference on Learning Representations*,
788 2022. URL <https://openreview.net/forum?id=vaRCHVj0uGI>.
- 789 Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou,
790 and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport for 3d molecule
791 generation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023c.
- 792 Reilly E. Sonstrom, Zoua Pa Vang, Haley N. Scolati, Justin L. Neill, Brooks H. Pate, and Joseph R.
793 Clark. Rapid Enantiomeric Excess Measurements of Enantioisotopomers by Molecular Rotational
794 Resonance Spectroscopy. *Organic Process Research & Development*, 27(7):1185–1197, July 2023.
795 ISSN 1083-6160. doi: 10.1021/acs.oprd.3c00028.
- 796 Jean-Thibaut Spaniol, Kin Long Kelvin Lee, Olivier Pirali, Cristina Puzzarini, and Marie-Aline
797 Martin-Drumel. A rotational investigation of the three isomeric forms of cyanoethynylbenzene
798 (HCC-C 6 H 4 -CN): Benchmarking experiments and calculations using the “Lego brick” approach.
799 *Physical Chemistry Chemical Physics*, 25(8):6397–6405, 2023. doi: 10.1039/D2CP04825F.
- 800 Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and
801 Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design. *arXiv preprint*
802 *arXiv:2402.05841*, 2024.

- 810 Alexander Tong, Kilian FATRAS, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-
811 Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models
812 with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024. ISSN 2835-
813 8856. URL <https://openreview.net/forum?id=CD9Snc73AW>. Expert Certification.
814
- 815 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
816 Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio,
817 H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information
818 Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- 819 Natalaja Vogt, Jürgen Vogt, and Jean Demaison. Accuracy of the rotational constants. *Journal of
820 molecular structure*, 988(1-3):119–127, 2011.
- 821 Zhichao Wang and Victor Solo. Particle filtering on the stiefel manifold with optimal transport. In
822 *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 4111–4116. IEEE, 2020.
823
- 824 Caroline C Womack, Marie-Aline Martin-Drumel, Gordon G Brown, Robert W Field, and Michael C
825 McCarthy. Observation of the simplest criegee intermediate ch₂oo in the gas-phase ozonolysis of
826 ethylene. *Science advances*, 1(2):e1400105, 2015.
- 827 Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric
828 diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*, 2022.
829
- 830 Ryoma Yataka, Kazuki Hirashima, and Masashi Shiraishi. Grassmann manifold flows for stable
831 shape generation. *Advances in Neural Information Processing Systems*, 36:72377–72411, 2023.
832
- 833 Lia Yeh, Lincoln Satterthwaite, and David Patterson. Automated, context-free assignment of asym-
834 metric rotor microwave spectra. *The Journal of chemical physics*, 150(20), 2019.
- 835 Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay,
836 and Tommi Jaakkola. Se (3) diffusion model with application to protein backbone generation.
837 *arXiv preprint arXiv:2302.02277*, 2023.
- 838 Daniel P. Zaleski and Kirill Prozument. Automated assignment of rotational spectra using artificial
839 neural networks. *The Journal of Chemical Physics*, 149(10):104106, September 2018. ISSN
840 0021-9606. doi: 10.1063/1.5037715.
- 841
- 842 Yuchen Zhu, Tianrong Chen, Lingkai Kong, Evangelos A Theodorou, and Molei Tao. Trivialized mo-
843 mentum facilitates diffusion generative modeling on lie groups. *arXiv preprint arXiv:2405.16381*,
844 2024.
- 845 Ralf Zimmermann and Knut Hüper. Computing the riemannian logarithm on the stiefel manifold:
846 Metrics, methods, and performance. *SIAM Journal on Matrix Analysis and Applications*, 43(2):
847 953–980, 2022.
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

A ROTATIONAL SPECTROSCOPY & MOLECULAR STRUCTURE

A.1 NOTATION

(P_X, P_Y, P_Z) are really called the *planar moments of inertia*. We refer to them as the moments of inertia as shorthand. The actual *moments of inertia* (I_A, I_B, I_C) are uniquely related to the planar moments by a simple linear transformation (Kraitchman, 1953).

A.2 INERTIA EDGE CASES

We assume that the molecule of interest has moments of inertia $P_X > P_Y > P_Z > 0$, or in other words, a nonplanar asymmetric rotor. This assumption holds for the vast majority of molecules. We now discuss edge cases, such as perfectly symmetric, planar, or linear molecules. It is worth noting that, owing to their rarity and symmetries, these edge cases have significant overlap with the set of molecules that have already been studied (McGuire, 2018).

When structures have two equal eigenvalues ($P_Y = P_Z$) in their inertia matrix, there is no longer a unique choice of these two principal axes. These axes now sweep out a plane of possibilities, and numerical diagonalization will arbitrarily pick two orthogonal axes from this plane. But, making an arbitrary choice does not break the mapping between \mathbf{X} and \mathbf{U} in Equation (4). The only issue is that to respect this additional symmetry, the flow should be invariant to in-plane roto-reflections of the molecule. However, these examples are so rare that they can be ignored: 76 examples in QM9 and 8 examples in GEOM. If needed, this symmetry can be handled using data augmentation.

Stiefel Flow Matching cannot handle *exactly* planar ($P_Z = 0$) and *exactly* linear ($P_Y = P_Z = 0$) molecules due to divide by zero in Equation (4). The vast majority of “planar” molecules are actually slightly nonplanar and therefore pose no issue. In rare cases where molecules are truly planar, Stiefel Flow Matching can be reformulated for 1 and 2 dimensions.

A.3 EXPERIMENTAL WORKFLOW

For an unknown molecule, we assume that its molecular formula can be measured by high-resolution mass spectrometry (Marshall & Hendrickson, 2008) and that its moments can be measured from rotational spectroscopy (Gordy et al., 1984).

At a high level, rotational spectroscopy observes how molecules freely rotate in the gas phase. The rotation of molecules is *quantized*, giving rise to a discrete set of rotational states. Molecules can absorb or emit radiation at *characteristic* wavelengths to transition between these energy levels. Rotational spectroscopy measures the energies of these transitions. A broadband microwave spectrometer can simultaneously measure thousands of these transitions, producing a spectrum of many sharp peaks (Brown et al., 2006; 2008). Rotational transition energies are in the microwave to terahertz region, which is why rotational spectroscopy is also known as microwave spectroscopy.

For an asymmetric rigid molecule, and neglecting effects like centrifugal distortion and hyperfine structure, the molecule’s rotational energy levels are essentially determined by three unique *rotational constants*, $A(BC)$, $A > B > C$. Rotational constants are inversely proportional to the principal moments of inertia, i.e., $A = \frac{\hbar^2}{2I_A}$. Each energy level is also indexed by quantum numbers. Transitions are the differences in these energy levels. Transition strength is also affected by the molecule’s dipole moment.

Once a spectrum is measured, the rotational spectroscopist is tasked with assigning each transition to its quantum numbers and ultimately assigning rotational constants $A(BC)$. Spectral assignment is a challenging problem tackled in other works (Zaleski & Prozument, 2018; Yeh et al., 2019). Rotational constants can then simply be inverted to obtain effective moments of inertia.

Using these effective moments, a spectroscopist can now search for the true structure using any of the models developed in this work. In an actual structure elucidation campaign, only a few targets are considered, which permits querying $K > 1000$ samples for each target, and also leaves enough computational resources to evaluate every generated sample using quantum chemistry.

918 A.4 EXPERIMENTAL PRECISION

919
920 While the proposed method generates structures which satisfy the moment constraints *exactly*,
921 and while rotational spectroscopy can measure experimental rotational constants to many digits of
922 precision (Vogt et al., 2011), it is unfortunate that the experimental rotational constants do not directly
923 translate to moments of inertia. This is because molecules are not perfectly rigid: Experiment observes
924 properties that have been *vibrationally* averaged, including the rotational constants. Vibrational
925 averaging is distinguished from conformational fluctuations such as torsions, which can be frozen
926 out by cooling molecules to their ground vibrational state. However, even in the ground vibrational
927 state, a molecule is still vibrating due to zero-point energy. As a result, the experimental rotational
928 constants $A(BC)_0$ are proportional to $\langle \frac{1}{r^2} \rangle$, whereas equilibrium rotational constants $A(BC)_e$
929 are proportional to $\frac{1}{\langle r \rangle^2}$, where $\langle \cdot \rangle$ denotes a vibrational average. Structures in QM9 and GEOM have
930 been geometry optimized to reach equilibrium structures $\langle r \rangle$. Equilibrium rotational constants are on
931 the order of 1% relative error of the experimental rotational constants (Vogt et al., 2011; Puzzarini &
932 Stanton, 2023).

933 Typically, experimental rotational constants $A(BC)_0$ can be corrected into equilibrium rotational
934 constants $A(BC)_e$ by a rovibrational calculation. Experimental and equilibrium rotational constants
935 are related by a perturbative expansion (Demaison et al., 2011)

$$936 A(BC)_0 = A(BC)_e - \frac{1}{2} \sum_i \alpha_i^{A(BC)} + \frac{1}{8} \sum_{ij} \gamma_{ij}^{A(BC)} + \dots, \quad (12)$$

937
938 where i, j index normal modes of vibration ($i, j \in \{1, 2, \dots, 3n - 6\}$ for n atoms), and α and γ are
939 first and second order interaction constants that couple rotational and vibrational motions together.
940 As a Taylor series, α is generally much larger than γ , so γ is usually neglected. The rovibrational
941 correction α can be calculated using an electronic structure method with a good compromise between
942 computational time and accuracy (Puzzarini & Stanton, 2023; Spaniol et al., 2023). However, this
943 calculation requires knowing the structure in the first place.

944 Alternatively, experimental rotational constants can be corrected to equilibrium rotational constants
945 by simple semi-empirical scaling factors under marginalized Bayesian uncertainty (Lee & McCarthy,
946 2020).

947 But, exact moment constraints do allow application in the following sense: if one were to *a priori*
948 guess the equilibrium moments correctly, one could then verify whether they are indeed correct by
949 generating structures, calculating their rovibrational corrections, and then checking their agreement
950 to the experimental rotational constants. Therefore, given experimental rotational constants $A(BC)_0$,
951 one can use this verification procedure in a fine-grid search for the true equilibrium rotational
952 constants $A(BC)_e$. This is feasible since $A(BC)_e$ consist of only 3 numbers. Experimental precision
953 in $A(BC)_0$ is maintained up to the precision in computing α . To verify that a structure is the true
954 structure, we must know that (1) its $A(BC)_e$ and α match $A(BC)_0$ and (2) gradient norm is near 0.
955

956 B THE STIEFEL MANIFOLD

957
958 In this section, we discuss various facts about the Stiefel manifold and computations thereon. We
959 refer readers to (Lee, 2003; Edelman et al., 1998; Bendokat et al., 2024) for further details.

960 The Stiefel manifold $\text{St}(n, p)$ is the set of orthonormal p -frames in \mathbb{R}^n :

$$961 \text{St}(n, p) = \{\mathbf{U} \in \mathbb{R}^{n \times p} \mid \mathbf{U}^\top \mathbf{U} = \mathbf{I}_p\}. \quad (13)$$

962 $\text{St}(n, p)$ is a manifold of dimension $np - \frac{1}{2}p(p + 1)$. Elements of $\text{St}(n, p)$ can be thought of as the
963 first p columns of an $n \times n$ orthogonal matrix. In fact, $\text{St}(n, p)$ generalizes some well-known spaces.
964 For example, $\text{St}(n, n)$ is the orthogonal group $\text{O}(n)$, and $\text{St}(n, n - 1)$ is diffeomorphic to the special
965 orthogonal group $\text{SO}(n)$, and $\text{St}(n, 1)$ recovers the unit n -sphere S^{n-1} .
966

967 B.1 TANGENT SPACE

968
969 The tangent space of $\text{St}(n, p)$ at a point \mathbf{U} is identified with the subspace:

$$970 \text{T}_{\mathbf{U}}\text{St}(n, p) = \{\Delta \in \mathbb{R}^{n \times p} \mid \mathbf{U}^\top \Delta + \Delta^\top \mathbf{U} = \mathbf{0}\}, \quad (14)$$

971 In other words, this is the set of $n \times p$ matrices Δ for which $\mathbf{U}^\top \Delta$ is skew-symmetric.

B.2 CANONICAL METRIC

Smoothly equipping every tangent space with an inner product turns $\text{St}(n, p)$ into a Riemannian manifold. Herein, we will consider the canonical metric:

$$\langle \Delta, \tilde{\Delta} \rangle_{\mathbf{U}} = \text{trace} \Delta^{\top} \left(\mathbf{I}_n - \frac{1}{2} \mathbf{U} \mathbf{U}^{\top} \right) \tilde{\Delta}, \quad (15)$$

Any inner product induces a norm by $\|\Delta\|_{\mathbf{U}} = \sqrt{\langle \Delta, \Delta \rangle_{\mathbf{U}}}$. For the canonical metric, the induced norm is equivalent to:

$$\|\Delta\|_{\mathbf{U}}^2 = \text{trace} \Delta^{\top} \Delta - \frac{1}{2} \text{trace}(\mathbf{U}^{\top} \Delta)^{\top} (\mathbf{U}^{\top} \Delta) = \|\Delta\|_F^2 - \frac{1}{2} \|\mathbf{U}^{\top} \Delta\|_F^2, \quad (16)$$

where $\|\cdot\|_F$ is the Frobenius norm. Computing the canonical norm in this way is much faster for the nodewise-batching approach used by PyTorch Geometric (Fey & Lenssen, 2019).

B.3 EXPONENTIAL AND LOGARITHM

The Stiefel exponential can be computed by the algorithm presented by Edelman et al. (1998), who give the closed-form expression for a geodesic γ with initial conditions $\gamma(0) = \mathbf{U}$ and $\dot{\gamma}(0) = \Delta$. Here, we reproduce the algorithm for computing $\exp_{\mathbf{U}}(\Delta) = \gamma(1)$.

Algorithm 1 Computing a Stiefel geodesic $\gamma(t)$. (Edelman et al., 1998)

Require: Base point $\mathbf{U} \in \text{St}(n, p)$, initial direction $\Delta \in \mathbf{T}_{\mathbf{U}}\text{St}(n, p)$, time $t \in \mathbb{R}$

- 1: $\mathbf{QR} \leftarrow (\mathbf{I}_n - \mathbf{U} \mathbf{U}^{\top}) \Delta$, where $\mathbf{Q} \in \mathbb{R}^{n \times p}$ and $\mathbf{R} \in \mathbb{R}^{p \times p}$ ▷ QR decomposition
- 2: $\mathbf{A} \leftarrow \begin{pmatrix} \mathbf{U}^{\top} \Delta & -\mathbf{R}^{\top} \\ \mathbf{R} & \mathbf{0} \end{pmatrix} \in \mathbb{R}^{2p \times 2p}$
- 3: $\begin{pmatrix} \mathbf{M}(t) & | & \cdots \\ \mathbf{N}(t) & | & \cdots \end{pmatrix} \leftarrow \exp_m(t\mathbf{A})$, where $\mathbf{M}(t), \mathbf{N}(t) \in \mathbb{R}^{p \times p}$ ▷ take submatrices
- 4: **return** $\mathbf{U} \mathbf{M}(t) + \mathbf{Q} \mathbf{N}(t) \in \text{St}(n, p)$

An efficient algorithm for computing the Stiefel logarithm is reproduced and simplified here from Zimmermann & Hüper (2022). We implement the logarithm in C++ for speed on the CPU, as it must be called on every fetch of a data example. This does not present a bottleneck: computing logarithmic maps takes an average of 0.1 ms for molecules in both QM9 and GEOM. The logarithm has a cost of $O(np^2)$. The thin QR decomposition and initial matrix multiplications have a cost of $O(np^2)$, while the inner loop’s Schur decomposition, Sylvester solve, matrix multiplications, and matrix exponential have a cost of $O(p^3)$.

Algorithm 2 Computing the Stiefel logarithm. (Zimmermann & Hüper, 2022)

Require: Base point $\mathbf{U} \in \text{St}(n, p)$, target point $\tilde{\mathbf{U}} \in \text{St}(n, p)$

- 1: $\mathbf{M} \leftarrow \mathbf{U}^{\top} \tilde{\mathbf{U}} \in \mathbb{R}^{p \times p}$
- 2: $\mathbf{QN} \leftarrow \tilde{\mathbf{U}} - \mathbf{U} \mathbf{M} \in \mathbb{R}^{n \times p}$ ▷ thin QR
- 3: $\mathbf{OR} \leftarrow \begin{pmatrix} \mathbf{M} \\ \mathbf{N} \end{pmatrix} \in \mathbb{R}^{2p \times p}$, $\mathbf{O} \in \mathbb{R}^{2p \times 2p}$ ▷ compute orthogonal completion via full QR
- 4: $\mathbf{V} \leftarrow \begin{pmatrix} \mathbf{M} & \mathbf{O}_{:,p:2p} \\ \mathbf{N} & \mathbf{O}_{:,p:2p} \end{pmatrix} \in \text{SO}(2p)$ ▷ flip sign if $\det < 0$
- 5: **for** $k = 1, \dots, 20$ **do**
- 6: $\begin{pmatrix} \mathbf{A} & -\mathbf{B}^{\top} \\ \mathbf{B} & \mathbf{C} \end{pmatrix} \leftarrow \log_m(\mathbf{V})$ ▷ use Schur matrix logarithm
- 7: $\mathbf{S} \leftarrow \frac{1}{12} \mathbf{B} \mathbf{B}^{\top} - \frac{1}{2} \mathbf{I}_p$
- 8: solve $\mathbf{C} = \mathbf{S} \mathbf{\Gamma} + \mathbf{\Gamma} \mathbf{S}$ for $\mathbf{\Gamma} \in \mathbb{R}^{p \times p}$ ▷ symmetric Sylvester equation
- 9: $\mathbf{\Phi} \leftarrow \exp_m(\mathbf{\Gamma}) \in \mathbb{R}^{p \times p}$ ▷ matrix exponential, or use Cayley transform
- 10: $\mathbf{V}_{:,p:2p} \leftarrow \mathbf{V}_{:,p:2p} \mathbf{\Phi}$ ▷ rotate last p columns of \mathbf{V}
- 11: **return** $\mathbf{U} \mathbf{A} + \mathbf{Q} \mathbf{B} \in \mathbf{T}_{\mathbf{U}}\text{St}(n, p)$

1026 B.4 UNIFORM SAMPLING
1027

1028 To sample uniformly on $\text{St}(n, p)$ with respect to the Haar measure, we can compute $\mathbf{U} =$
1029 $\mathbf{Z}(\mathbf{Z}^\top \mathbf{Z})^{-1/2}$ for a random matrix $\mathbf{Z} \in \mathbb{R}^{n \times p}$ whose elements are drawn i.i.d. from a standard
1030 Gaussian (i.e. $\mathbf{Z} = \text{randn}(n, p)$).
1031

1032 B.5 HOUSEHOLDER REFLECTIONS
1033

1034 To align the final column of a matrix $\mathbf{U} \in \text{St}(n, p)$ to a unit vector \mathbf{y} , we can rotate its columns under
1035 a transformation $\mathbf{R} \in \text{SO}(n)$ such that $\mathbf{R}\mathbf{U}_{:,4} = \mathbf{y}$. Let $\mathbf{U}_{:,4} = \mathbf{x}$. For any unit vector $\mathbf{v} \in \mathbb{R}^3$, the
1036 Householder matrix $\mathbf{H}(\mathbf{v}) = \mathbf{I}_n - 2\mathbf{v}\mathbf{v}^\top$ has determinant -1 . Then,
1037

$$1038 \mathbf{R} = \mathbf{H}(\mathbf{y})\mathbf{H}\left(\frac{\mathbf{x} + \mathbf{y}}{\|\mathbf{x} + \mathbf{y}\|}\right) \in \text{SO}(n) \quad (17)$$

1039 is our desired rotation. In particular, we can verify that
1040

$$1041 \mathbf{H}\left(\frac{\mathbf{x} + \mathbf{y}}{\|\mathbf{x} + \mathbf{y}\|}\right)\mathbf{x} = \mathbf{x} - \frac{(2 + 2\mathbf{x}^\top \mathbf{y})}{\|\mathbf{x} + \mathbf{y}\|^2}(\mathbf{x} + \mathbf{y}) = \mathbf{x} - (\mathbf{x} + \mathbf{y}) = -\mathbf{y} \quad (18)$$

1042 so that $\mathbf{R}\mathbf{x} = \mathbf{H}(\mathbf{y})(-\mathbf{y}) = \mathbf{y}$, as desired.
1043
1044
1045

1046 B.6 ORTHOGONAL PROJECTIONS
1047

1048 The orthogonal projection of a point $\mathbf{U}_0 \in \mathbb{R}^{n \times p}$ onto the $\text{St}(n, p)$ is a special case of the well-known
1049 orthogonal Procrustes problem:
1050

$$1051 \arg \min_{\mathbf{U}} \|\mathbf{U}\mathbf{I}_p - \mathbf{U}_0\|_F, \quad \text{subject to } \mathbf{U}^\top \mathbf{U} = \mathbf{I}_p. \quad (19)$$

1052 Letting $\mathbf{U}_0 = \mathbf{A}\mathbf{\Sigma}\mathbf{B}^\top$ under a singular value decomposition, the solution to this problem is $\mathbf{A}\mathbf{B}^\top$.
1053 The orthogonal projection of a point $\mathbf{Z} \in \mathbb{R}^{n \times p}$ onto the tangent space $\text{T}_{\mathbf{U}}\text{St}(n, p)$ is computed by
1054 $\mathbf{Z} - \mathbf{U} \text{sym}(\mathbf{U}^\top \mathbf{Z})$, where $\text{sym}(\mathbf{A}) = \frac{1}{2}(\mathbf{A} + \mathbf{A}^\top)$ is a symmetrization operation.
1055

1056 B.7 THE ZERO CENTER-OF-MASS SUBMANIFOLD
1057

1058 For $n \geq 5$, we are interested in the subset of $\text{St}(n, 4)$ obtained by fixing the final column to a fixed
1059 unit vector $\mathbf{a} \in S^{n-1}$:

$$1060 \mathcal{M} = \{\mathbf{U} \in \text{St}(n, p) \mid \mathbf{U}_{:,4} = \mathbf{a}\}. \quad (20)$$

1061 Note that $S = \pi^{-1}(\{\mathbf{a}\})$ is a level set of the projection map:
1062

$$1063 \pi: \text{St}(n, 4) \rightarrow S^{n-1}, \quad \mathbf{U} \mapsto \mathbf{U}_{:,4}. \quad (21)$$

1064 In fact, π is a smooth surjective map of constant rank, so it is a submersion by the global rank theorem.
1065 Hence, \mathcal{M} is an embedded submanifold by the submersion level set theorem. The tangent vectors to
1066 \mathcal{M} are exactly those whose final column is zero, since
1067

$$1068 \text{T}_{\mathbf{U}}\mathcal{M} = \ker d\pi_{\mathbf{U}} = \{\Delta \in \text{T}_{\mathbf{U}}\text{St}(n, 4) \mid \Delta_{:,4} = \mathbf{0}\}. \quad (22)$$

1069 If \mathcal{M} further inherits the canonical metric from $\text{St}(n, 4)$, then it becomes a Riemannian manifold.
1070 Note that \mathcal{M} is homeomorphic to $\text{St}(n-1, 3)$, so it is connected and compact. Hence, it is geodesi-
1071 cally complete and any two points in \mathcal{M} can be connected with a length-minimizing geodesic on \mathcal{M} .
1072 Theorem 1 shows that \mathcal{M} is totally geodesic, i.e., any geodesic on \mathcal{M} is a geodesic on $\text{St}(n, 4)$. This
1073 allows us to perform simpler computations in the ambient space. Theorem 5 computes projections of
1074 arbitrary matrices onto $\text{T}_{\mathbf{U}}\mathcal{M}$.

1075 **Theorem 1.** *As defined above, \mathcal{M} is totally geodesic.*
1076

1077 *Proof.* A sufficient condition is that for any $\mathbf{U} \in \mathcal{M}$ and $\Delta \in \text{T}_{\mathbf{U}}\mathcal{M}$, the geodesic γ with initial
1078 conditions $\gamma(0) = \mathbf{U}$ and $\dot{\gamma}(0) = \Delta$ stays within \mathcal{M} . Fortunately, γ is given in closed-form by
1079 Edelman et al. (1998) in Algorithm 1. Since $\Delta_{:,4} = \mathbf{0}$, the fourth columns of \mathbf{R} and $\mathbf{U}^\top \Delta$ are zero.
Since $\mathbf{U}^\top \Delta$ is skew-symmetric by Equation 14, its fourth row is also zero. Then, the fourth row and

column of \mathbf{A} and $\exp_m(t\mathbf{A})$ are zero, but $\exp_m(t\mathbf{A})_{4,4} = 1$. It follows that $\mathbf{M}(t)_{\cdot,4} = (0, 0, 0, 1)$ and $\mathbf{N}(t)_{\cdot,4} = \mathbf{0}$, so that

$$\gamma(t)_{\cdot,4} = \mathbf{U}(\mathbf{M}(t)_{\cdot,4}) + \mathbf{Q}(\mathbf{N}(t)_{\cdot,4}) = \mathbf{U}_{\cdot,4} = \mathbf{a}. \quad (23)$$

Hence, $\gamma(t) \in \mathcal{M}$ as desired. \square

We now provide some lemmas which are useful for proving that the Stiefel logarithm on \mathcal{M} can be computed using $\text{St}(n, 3)$ rather than $\text{St}(n, 4)$ (Theorem 4).

Lemma 2. *Let \mathbf{X} be a square matrix whose i^{th} row and column are 0. Positive powers of \mathbf{X} retain zeros in the i^{th} row and column. Suppose*

$$\mathbf{X} = \begin{pmatrix} \mathbf{E} & \mathbf{0} & \mathbf{F} \\ \mathbf{0} & 0 & \mathbf{0} \\ \mathbf{G} & \mathbf{0} & \mathbf{H} \end{pmatrix}, \quad \begin{pmatrix} \mathbf{E} & \mathbf{F} \\ \mathbf{G} & \mathbf{H} \end{pmatrix}^k = \begin{pmatrix} \hat{\mathbf{E}} & \hat{\mathbf{F}} \\ \hat{\mathbf{G}} & \hat{\mathbf{H}} \end{pmatrix}. \quad (24)$$

Then we have

$$\mathbf{X}^k = \begin{pmatrix} \hat{\mathbf{E}} & \mathbf{0} & \hat{\mathbf{F}} \\ \mathbf{0} & 0 & \mathbf{0} \\ \hat{\mathbf{G}} & \mathbf{0} & \hat{\mathbf{H}} \end{pmatrix}. \quad (25)$$

Proof. We show this by induction over k . It is clear this is true for $k = 1$. Then, we assume

$\mathbf{X}^k = \begin{pmatrix} \hat{\mathbf{E}} & \mathbf{0} & \hat{\mathbf{F}} \\ \mathbf{0} & 0 & \mathbf{0} \\ \hat{\mathbf{G}} & \mathbf{0} & \hat{\mathbf{H}} \end{pmatrix}$ and aim to show for $k + 1$. We know that

$$\begin{pmatrix} \mathbf{E} & \mathbf{F} \\ \mathbf{G} & \mathbf{H} \end{pmatrix}^{k+1} = \begin{pmatrix} \hat{\mathbf{E}} & \hat{\mathbf{F}} \\ \hat{\mathbf{G}} & \hat{\mathbf{H}} \end{pmatrix} \begin{pmatrix} \mathbf{E} & \mathbf{F} \\ \mathbf{G} & \mathbf{H} \end{pmatrix} = \begin{pmatrix} \hat{\mathbf{E}}\mathbf{E} + \hat{\mathbf{F}}\mathbf{G} & \hat{\mathbf{E}}\mathbf{F} + \hat{\mathbf{F}}\mathbf{H} \\ \hat{\mathbf{G}}\mathbf{E} + \hat{\mathbf{H}}\mathbf{G} & \hat{\mathbf{G}}\mathbf{F} + \hat{\mathbf{H}}\mathbf{H} \end{pmatrix}. \quad (26)$$

At the same time, we have that

$$\mathbf{X}^{k+1} = \begin{pmatrix} \hat{\mathbf{E}} & \mathbf{0} & \hat{\mathbf{F}} \\ \mathbf{0} & 0 & \mathbf{0} \\ \hat{\mathbf{G}} & \mathbf{0} & \hat{\mathbf{H}} \end{pmatrix} \begin{pmatrix} \mathbf{E} & \mathbf{0} & \mathbf{F} \\ \mathbf{0} & 0 & \mathbf{0} \\ \mathbf{G} & \mathbf{0} & \mathbf{H} \end{pmatrix} = \begin{pmatrix} \hat{\mathbf{E}}\mathbf{E} + \hat{\mathbf{F}}\mathbf{G} & \mathbf{0} & \hat{\mathbf{E}}\mathbf{F} + \hat{\mathbf{F}}\mathbf{H} \\ \mathbf{0} & 0 & \mathbf{0} \\ \hat{\mathbf{G}}\mathbf{E} + \hat{\mathbf{H}}\mathbf{G} & \mathbf{0} & \hat{\mathbf{G}}\mathbf{F} + \hat{\mathbf{H}}\mathbf{H} \end{pmatrix}, \quad (27)$$

which completes the proof. \square

Lemma 3. *Let \mathbf{X} be a square matrix whose i^{th} row and column are 0, and consider the matrix exponential of its nonzero portions:*

$$\mathbf{X} = \begin{pmatrix} \mathbf{E} & \mathbf{0} & \mathbf{F} \\ \mathbf{0} & 0 & \mathbf{0} \\ \mathbf{G} & \mathbf{0} & \mathbf{H} \end{pmatrix}, \quad \exp_m \begin{pmatrix} \mathbf{E} & \mathbf{F} \\ \mathbf{G} & \mathbf{H} \end{pmatrix} = \begin{pmatrix} \bar{\mathbf{E}} & \bar{\mathbf{F}} \\ \bar{\mathbf{G}} & \bar{\mathbf{H}} \end{pmatrix}. \quad (28)$$

Then we have

$$\exp_m(\mathbf{X}) = \begin{pmatrix} \bar{\mathbf{E}} & \mathbf{0} & \bar{\mathbf{F}} \\ \mathbf{0} & 1 & \mathbf{0} \\ \bar{\mathbf{G}} & \mathbf{0} & \bar{\mathbf{H}} \end{pmatrix}. \quad (29)$$

Proof. The matrix exponential is given by

$$\exp_m(\mathbf{X}) = \sum_{k=0}^{\infty} \frac{1}{k!} \mathbf{X}^k. \quad (30)$$

By Lemma 2, all non-identity terms of this series are identical to terms in the series for $\exp_m \begin{pmatrix} \mathbf{E} & \mathbf{F} \\ \mathbf{G} & \mathbf{H} \end{pmatrix}$ but with zeros inserted in the i^{th} row and column. The identity term then contributes the extra 1 in the main diagonal. \square

1134 **Theorem 4.** Let $\mathbf{U} \in \mathcal{M}$, and let $\tilde{\mathbf{U}}$ be its first three columns so that $\mathbf{U} = (\tilde{\mathbf{U}} \quad \mathbf{a})$.

1135
1136 The exponential on \mathcal{M} can be computed as the exponential on $\text{St}(n, 3)$ after discarding the last
1137 column of $\Delta = (\tilde{\Delta} \quad \mathbf{0})$.

$$1138 \quad \exp_{\mathbf{U}}(\Delta) = (\exp_{\tilde{\mathbf{U}}}(\tilde{\Delta}) \quad \mathbf{a}) \quad (31)$$

1139
1140 Similarly, the logarithm on \mathcal{M} can be computed as the logarithm on $\text{St}(n, 3)$ followed by concatenating
1141 a zero column:

$$1142 \quad \log_{\mathbf{U}_0}(\mathbf{U}_1) = (\log_{\tilde{\mathbf{U}}_0}(\tilde{\mathbf{U}}_1) \quad \mathbf{0}) \quad (32)$$

1143
1144 *Proof.* We go line-by-line through Algorithm 1 to show that the Stiefel exponential on \mathcal{M} is equivalent
1145 to the exponential for $\text{St}(n, 3)$.

1146 The last column of $(\mathbf{I}_n - \mathbf{U}\mathbf{U}^\top)\Delta$ is 0 because the last column of Δ is 0. Then,

$$1148 \quad \mathbf{Q}\mathbf{R} = (\mathbf{I}_n - \mathbf{U}\mathbf{U}^\top)\Delta = (\tilde{\mathbf{Q}} \quad \mathbf{q}) \begin{pmatrix} \tilde{\mathbf{R}} & \mathbf{0} \\ \mathbf{0} & 0 \end{pmatrix}, \quad (33)$$

1150 where \mathbf{q} is some orthogonal vector to the 3 columns of $\tilde{\mathbf{Q}}$, and $\tilde{\mathbf{R}}$ is the first 3 columns and rows of
1151 \mathbf{R} . This corresponds to the QR decomposition $\tilde{\mathbf{Q}}\tilde{\mathbf{R}} = (\mathbf{I}_n - \tilde{\mathbf{U}}\tilde{\mathbf{U}}^\top)\tilde{\Delta}$.

1153 In addition, we have that $\mathbf{U}^\top\Delta = \begin{pmatrix} \tilde{\mathbf{U}}^\top\tilde{\Delta} & \mathbf{0} \\ \mathbf{0} & 0 \end{pmatrix}$.

1155 Thus, the block matrix \mathbf{A} in Algorithm 1 is given as

$$1157 \quad \begin{pmatrix} \tilde{\mathbf{U}}^\top\tilde{\Delta} & \mathbf{0} & -\tilde{\mathbf{R}}^\top & \mathbf{0} \\ \mathbf{0} & 0 & \mathbf{0} & 0 \\ \tilde{\mathbf{R}} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & 0 & \mathbf{0} & 0 \end{pmatrix}. \quad (34)$$

1162 By Lemma 3, the matrix exponential of \mathbf{A} containing extra rows and columns of zeros is equivalent
1163 to the matrix exponential with these rows and columns of zeros removed.

1164 Therefore,

$$1165 \quad \exp_m(t\mathbf{A}) = \left(\begin{array}{c|c} \mathbf{M}(t) & \cdots \\ \mathbf{N}(t) & \cdots \end{array} \right) = \left(\begin{array}{c|c} \tilde{\mathbf{M}}(t) & \mathbf{0} \mid \cdots \\ \mathbf{0} & 1 \mid \cdots \\ \tilde{\mathbf{N}}(t) & \mathbf{0} \mid \cdots \\ \mathbf{0} & 0 \mid \cdots \end{array} \right). \quad (35)$$

1170 Finally, the output of Algorithm 1 is

$$1172 \quad \mathbf{U}\mathbf{M}(t) + \mathbf{Q}\mathbf{N}(t) = (\tilde{\mathbf{U}} \quad \mathbf{a}) \begin{pmatrix} \tilde{\mathbf{M}}(t) & \mathbf{0} \\ \mathbf{0} & 1 \end{pmatrix} + (\tilde{\mathbf{Q}} \quad \mathbf{q}) \begin{pmatrix} \tilde{\mathbf{N}}(t) & \mathbf{0} \\ \mathbf{0} & 0 \end{pmatrix} \quad (36)$$

$$1173 \quad = (\tilde{\mathbf{U}}\tilde{\mathbf{M}}(t) \quad \mathbf{a}) + (\tilde{\mathbf{Q}}\tilde{\mathbf{N}}(t) \quad \mathbf{0}) \quad (37)$$

$$1174 \quad = (\tilde{\mathbf{U}}\tilde{\mathbf{M}}(t) + \tilde{\mathbf{Q}}\tilde{\mathbf{N}}(t) \quad \mathbf{a}) \quad (38)$$

$$1175 \quad = (\exp_{\tilde{\mathbf{U}}}(\tilde{\Delta}) \quad \mathbf{a}). \quad (39)$$

1178 Since the Stiefel exponential is locally invertible, this also shows that the Stiefel logarithm can be
1179 computed using $\tilde{\mathbf{U}}_0$ and $\tilde{\mathbf{U}}_1$.

1181 □

1182 **Theorem 5.** Let $\mathbf{U} \in \mathcal{M}$, and let $\tilde{\mathbf{U}}$ be its first three columns so that $\mathbf{U} = [\tilde{\mathbf{U}} \quad \mathbf{a}]$. Given $\mathbf{Z} \in \mathbb{R}^{n \times 3}$,
1183 the minimum-norm projection of $[\mathbf{Z} \quad \mathbf{0}]$ onto $\mathbb{T}_{\mathbf{U}}\mathcal{M}$ is given by $[\pi(\mathbf{Z}) \quad \mathbf{0}]$, where

$$1185 \quad \pi(\mathbf{Z}) = \tilde{\mathbf{U}} \text{skew}(\tilde{\mathbf{U}}^\top \mathbf{Z}) + (\mathbf{I}_n - \mathbf{U}\mathbf{U}^\top)\mathbf{Z} \in \mathbb{R}^{n \times 3} \quad (40)$$

$$1186 \quad = \tilde{\mathbf{U}} \text{skew}(\tilde{\mathbf{U}}^\top \mathbf{Z}) + (\mathbf{I}_n - \mathbf{U}\mathbf{U}^\top)\mathbf{Z} \in \mathbb{R}^{n \times 3} \quad (41)$$

1187 Note there is no tilde in the second term.

1188 *Proof.* We can rewrite

$$1189 \quad \mathbb{T}_{\tilde{U}}\mathcal{M} = \{[\tilde{\Delta} \mathbf{0}] \mid \tilde{U}^\top \tilde{\Delta} + \tilde{\Delta}^\top \tilde{U} = \mathbf{0}, \tilde{\Delta}^\top \mathbf{a} = \mathbf{0}\} \quad (42)$$

$$1191 \quad = \{[\tilde{\Delta} \mathbf{0}] \mid \tilde{\Delta} \in \mathbb{T}_{\tilde{U}}\text{St}(n, 3) \cap \text{span}(\mathbf{a})^\perp\}, \quad (43)$$

1192 Edelman et al. (1998) give an orthogonal projection of \mathbf{Z} onto $\mathbb{T}_{\tilde{U}}\text{St}(n, 3)$ as:

$$1193 \quad \pi_1(\mathbf{Z}) = \tilde{U} \text{skew}(\tilde{U}^\top \mathbf{Z}) + (\mathbf{I}_n - \tilde{U} \tilde{U}^\top) \mathbf{Z}, \quad (44)$$

1194 where $\text{skew}(\mathbf{A}) = \frac{1}{2}(\mathbf{A} - \mathbf{A}^\top)$. The orthogonal projection of \mathbf{Z} onto $\text{span}(\mathbf{a})^\perp$ is

$$1196 \quad \pi_2(\mathbf{Z}) = (\mathbf{I}_n - \mathbf{a} \mathbf{a}^\top) \mathbf{Z}. \quad (45)$$

1197 We can check that π_1 and π_2 commute with $\pi = \pi_1 \circ \pi_2 = \pi_2 \circ \pi_1$. Hence, π is an orthogonal
1198 projection onto $\mathbb{T}_{\tilde{U}}\text{St}(n, 3) \cap \text{span}(\mathbf{a})^\perp$, as desired. \square

1200 C EXPERIMENTAL DETAILS

1202 C.1 ARCHITECTURE

1204 Explicitly, the neural network takes in moments (P_X, P_Y, P_Z) , time t , atom types \mathbf{a} , and coordinates
1205 \mathbf{X} , and outputs a Stiefel tangent vector. Moments are embedded using sinusoidal features with
1206 wavelength geometrically spaced from 0.0001 to 10,000. Time is similarly embedded but with
1207 a wavelength range from 0.001 to 1. We use a reflection-equivariant network. Note that given a
1208 reflection-*invariant* function f , the mapping $\mathbf{X} \mapsto \text{sign}(\mathbf{X}) \odot f(\mathbf{X})$ is reflection-equivariant, where
1209 $\text{sign}(\mathbf{X})$ gives the element-wise signs of \mathbf{X} , with $\text{sign}(0) = 0$. Thus, the problem is reduced to
1210 constructing a network that is reflection-invariant with respect to the input coordinates \mathbf{X} .

1211 As input to our network, we begin by featurizing the molecule in a reflection-invariant manner. We
1212 obtain invariant node features $\mathbf{h}_i \in \mathbb{R}^{d_{\text{node}}}$ by using the molecule’s unsigned coordinates and atom
1213 types, and edge features $\mathbf{e}_{ij} \in \mathbb{R}^{d_{\text{edge}}}$ are computed from the unsigned differences between pairs of
1214 atomic coordinates. These features are passed through a Transformer backbone (Vaswani et al., 2017).
1215 We use the PreLN layout with an adaptive version of LayerNorm (Dieleman et al., 2022; Dhariwal &
1216 Nichol, 2021) that conditions on the timestep and molecule’s moments. In addition, the attention
1217 module is replaced with a message-passing block that jointly updates the node and edge features:

$$1218 \quad (\mathbf{v}_{ij}, a_{ij}, \mathbf{e}'_{ij}) \leftarrow \text{MLP}(\mathbf{h}_i, \mathbf{h}_j, \mathbf{e}_{ij}), \quad \text{for all } i, j, \quad (46)$$

$$1219 \quad \mathbf{e}_{ij} \leftarrow \mathbf{e}_{ij} + \mathbf{e}'_{ij}, \quad \text{for all } i, j, \quad (47)$$

$$1221 \quad \mathbf{y}_i \leftarrow \sum_{j=1}^n \left(\frac{\exp(a_{ij})}{\sum_{k=1}^n \exp(a_{ik})} \right) \mathbf{v}_{ij}, \quad \text{for all } i, \quad (48)$$

$$1223 \quad \mathbf{h}_i \leftarrow \mathbf{h}_i + \text{Linear}(\mathbf{y}_i) \quad \text{for all } i. \quad (49)$$

1224 The last two equations are reminiscent of self-attention in Transformers, and along the same lines, we
1225 use a multi-headed extension of them. Across all experiments, we use $d_{\text{node}} = 768$ and $d_{\text{edge}} = 192$.
1226 We use 16 Transformer-like blocks, 12 update heads, SiLU/Swish activations, and a $4 \times$ expansion in
1227 each block’s feed-forward module. In total, our model has 154M trainable parameters.

1228 This architecture is distinct from the pretrained network of KREED (Cheng et al., 2024), which was
1229 tailored for predicting 3D structure given molecular formula, moments of inertia, and *substitution*
1230 *coordinates*. KREED was trained using random dropout of input substitution coordinates. For
1231 QM9, the model sometimes observed examples with no substitution coordinates during training.
1232 However, for GEOM, the model always received at least some substitution coordinates during training.
1233 Therefore, KREED is operating out-of-distribution when provided with no substitution coordinates at
1234 all.

1236 C.2 STOCHASTICITY

1237 Diversity is important for identifying unknown molecules outside the training set. Therefore, we
1238 experiment with adding stochasticity to the dynamics of the flow, leading individual paths to be
1239 stochastic. To do so during training, we apply an exponential map to a Gaussian variable with noise
1240 scale $\gamma \cdot \frac{\cos(\pi t) + 1}{2}$ in the tangent space of the interpolant before calculating \tilde{U} again. During sampling,
1241 we add Gaussian noise of the same noise schedule to the tangent vector at every time step.

Table 3: General training and sampling hyperparameters.

	Hyperparameter	QM9	GEOM
Training	Epochs	1000	60
	Effective batch size	1024	96
	Optimizer	AdamW	AdamW
	Learning rate	10^{-4}	10^{-4}
	Learning rate warmup steps	2000	2000
	Weight decay	0.01	0.01
	Gradient clipping	yes	yes
	EMA decay	0.9995	0.9995
KREED	Timesteps	1000	1000
	Schedule	polynomial	polynomial
Stiefel FM	Timesteps	200	200

Table 4: Training and sampling hyperparameters for Stiefel Flow Matching.

Dataset	Model	Timestep sampling	OT	stochasticity γ
QM9	Stiefel FM	uniform	no	0.00
	Stiefel FM-OT	uniform	yes	0.00
	Stiefel FM-OT-stoch	uniform	yes	0.10
	Stiefel FM-ln	logit-normal	no	0.00
	Stiefel FM-ln-OT	logit-normal	yes	0.00
GEOM	Stiefel FM	uniform	no	0.00
	Stiefel FM-OT	uniform	yes	0.00

Table 5: Extended ablation study for QM9.

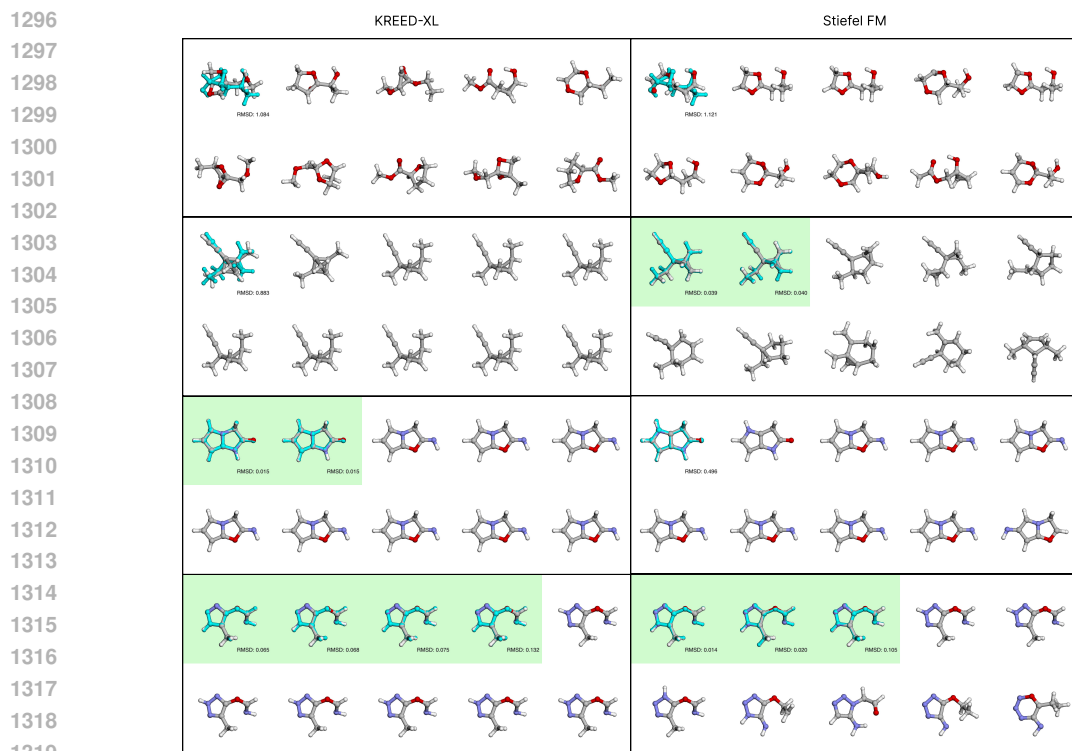
Method	% < RMSD \uparrow		Error \downarrow	Valid \uparrow	Stable \downarrow	Diverse \uparrow	NFE \downarrow
	0.25 Å	0.10 Å					
Stiefel FM	15.17 \pm 0.31	13.82 \pm 0.30	0.00	0.882	-1.125	1.040	200
Stiefel FM-OT	13.99 \pm 0.30	12.68 \pm 0.29	0.00	0.835	-1.039	1.045	200
Stiefel FM-stoch	15.13 \pm 0.31	13.83 \pm 0.30	0.00	0.877	-1.116	1.045	500
Stiefel FM-ln	15.74 \pm 0.32	11.45 \pm 0.28	0.00	0.880	-0.600	0.982	200
Stiefel FM-ln-OT	14.90 \pm 0.31	12.45 \pm 0.29	0.00	0.875	-0.687	1.026	200

Method	Dataset	Training (min / epoch)	Training (it / s)	Sampling (seconds / K=10 samples)
KREED-XL	QM9	1.02	6.7	13.9
Stiefel FM	QM9	1.32	5.1	2.9
Stiefel FM-OT	QM9	3.48	1.9	2.9
KREED-XL	GEOM	225.6	17.0	71.3
Stiefel FM	GEOM	224.4	17.1	15.0
Stiefel FM-OT	GEOM	229.8	16.7	15.0

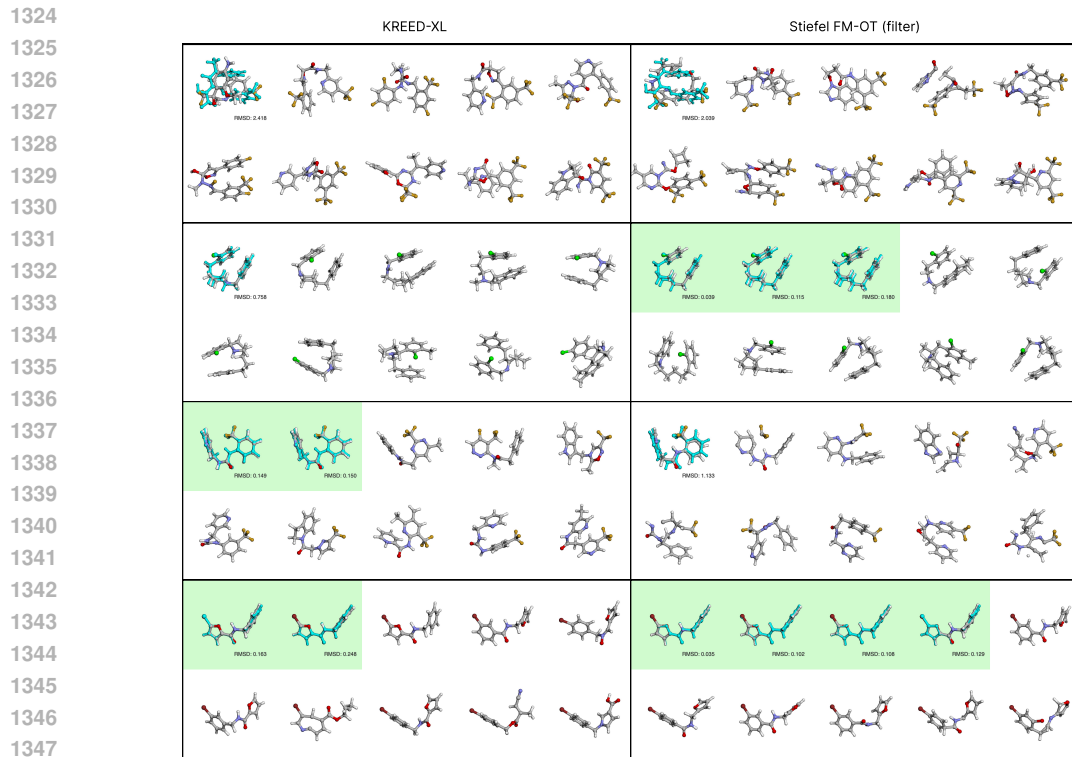
Table 6: Training and sampling timing on QM9 and GEOM. A CPU bottleneck of computing logarithms and optimal transport exists for QM9 due to a large batch size of 256, but this CPU bottleneck disappears for GEOM due to a small batch size of 24.

C.3 TRAINING

Models were trained on 4 NVIDIA A100 40GB GPUs. Tables 3 and 4 gives our hyperparameters. We use the adaptive gradient clipping strategy from Hoogetboom et al. (2022).



1320 Figure 5: Selected QM9 examples. Best viewed zoomed in. Examples are sorted by RMSD to ground
 1321 truth, which is shown in cyan. Green panels indicate meeting the threshold of 0.25 Å.
 1322



1348 Figure 6: Selected GEOM examples. Best viewed zoomed in. Examples are sorted by RMSD to
 1349 ground truth, which is shown in cyan. Green panels indicate meeting the threshold of 0.25 Å.




Figure 7: Normalized histogram of curve lengths of success (0.25 Å) and failure examples of Stiefel FM on the QM9 (left) and GEOM (right) datasets. Success cases are more often the result of longer generation paths.

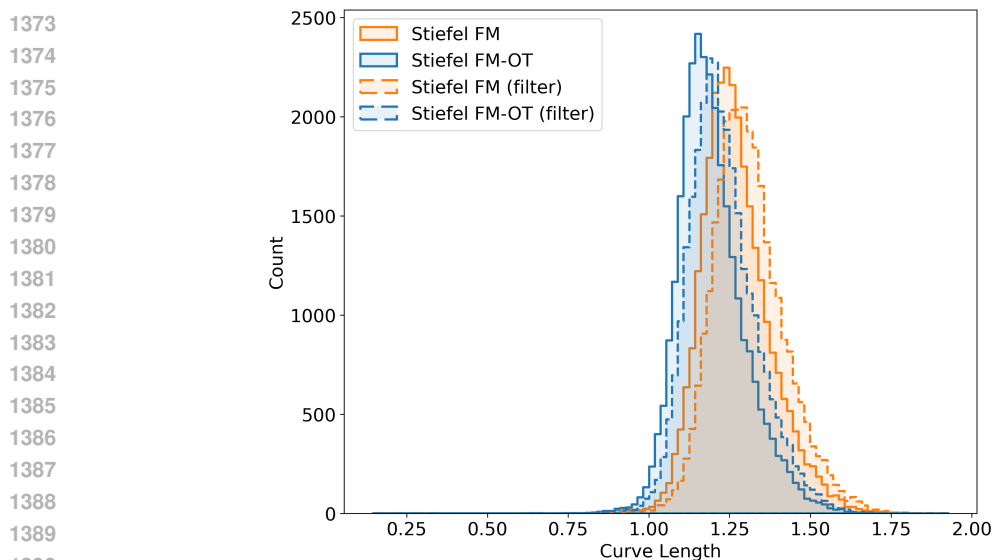


Figure 8: Equivariant optimal transport shortens generation trajectories for GEOM. However, generating more samples and filtering by validity has a slight bias towards *longer* generation trajectories.

C.4 SAMPLING

During sampling, we sample uniformly from \mathcal{M} and iteratively query the trained model for a tangent vector Δ at every step. If stochasticity is turned on, Gaussian noise with scale $\gamma \cdot \frac{\cos(\pi t) + 1}{2}$ is added to this tangent vector. At every step, \mathbf{U}_t is projected onto the manifold (Appendix B.6) before projecting the tangent vector to the tangent space on the manifold $T_{\mathbf{U}_t}\text{St}(n, 4)$ (Theorem 5). Integration proceeds by applying $\exp_{\mathbf{U}_t}$ to this tangent vector, scaled by dt .

Algorithm 3 Sampling under Stiefel Flow Matching.**Require:** Size n , atom types \mathbf{a} , moments $P_{XYZ} = (P_X, P_Y, P_Z)$, stochasticity γ , timesteps T

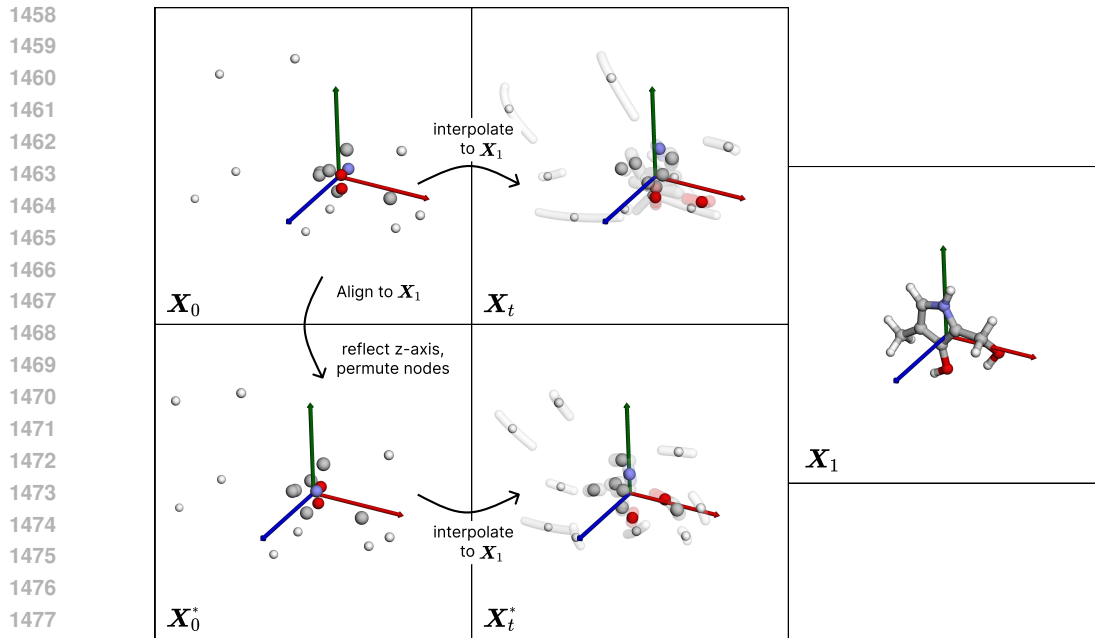
- 1: $\mathbf{U} \sim \text{Uniform}(\text{St}(n, 4))$ using Appendix B.4
- 2: Project \mathbf{U} to \mathcal{M} (defined in Equation 5) using Appendix B.5
- 3: $t \leftarrow 0$
- 4: $t_\Delta \leftarrow 1/T$
- 5: **for** step in $1, \dots, T$ **do**
- 6: $\mathbf{X} \leftarrow \mathbf{X}(\mathbf{U}, \mathbf{m}, P_{XYZ})$ by inverting Equation 4
- 7: $\Delta \leftarrow v_\theta(t, \mathbf{X}, P_{XYZ}, \mathbf{a}) \in \mathbb{R}^{n \times 3}$
- 8: **if** step $< T$ **then**
- 9: $\Delta \leftarrow t_\Delta \Delta + \frac{1}{2} \gamma \sqrt{t_\Delta} (\cos(\pi t) + 1) \boldsymbol{\varepsilon}$, where $\varepsilon_{ij} \sim \mathcal{N}(0, 1)$
- 10: **else**
- 11: $\Delta \leftarrow t_\Delta \Delta$
- 12: Project Δ to $\text{T}_U \mathcal{M}$ using Theorem 5
- 13: $\mathbf{U} \leftarrow \exp_{\mathbf{U}}(\Delta)$
- 14: $t \leftarrow t + t_\Delta$
- 14: **return** $\mathbf{X} \leftarrow \mathbf{X}(\mathbf{U}, \mathbf{m}, P_{XYZ})$ by inverting Equation 4

D GREEDY RANDOM OPTIMAL ASSIGNMENT ALGORITHM

The computational cost of computing the optimal transport map is $O(np^2)$, as it relies on a fixed number of computations of the logarithm.

Algorithm 4 Heuristic alignment algorithm.**Require:** $\mathbf{U}_0, \mathbf{U}_1 \in \text{St}(n, p)$, atom types \mathbf{a} , number of restarts R , local search budget L

- 1: $c^* \leftarrow \infty$ ▷ lowest cost seen so far
- 2: **for** reflections $\mathbf{R} \in \{\text{diag}(\mathbf{s}) \mid \mathbf{s} \in \{-1, +1\}^3\}$ **do** ▷ find a good reflection
- 3: **for** $k = 0, 1, \dots, R$ **do**
- 4: sample a random atom-type-preserving node permutation $\mathbf{\Pi}$
- 5: $\mathbf{U}_{\text{cand}} \leftarrow \mathbf{\Pi} \mathbf{U}_0 \mathbf{R}^\top$
- 6: $c \leftarrow \tilde{d}(\mathbf{U}_{\text{cand}}, \mathbf{U}_1)$, ▷ approximate distance
- 7: **if** $c < c^*$ **then**
- 8: $c^* \leftarrow c$
- 9: $\mathbf{U}_{\text{best}} \leftarrow \mathbf{U}_{\text{cand}}$
- 10: **for** $k = 0, 1, \dots, L$ **do** ▷ local search over node swaps
- 11: sample (i, j) such that $\mathbf{a}_i = \mathbf{a}_j$, without repeating pairs
- 12: $\mathbf{U}_{\text{cand}} \leftarrow \text{swap}(\mathbf{U}_{\text{best}}, i, j)$
- 13: $c \leftarrow \tilde{d}(\mathbf{U}_{\text{cand}}, \mathbf{U}_1)$, ▷ approximate distance
- 14: **if** $c < c^*$ **then**
- 15: $c^* \leftarrow c$
- 16: $\mathbf{U}_{\text{best}} \leftarrow \mathbf{U}_{\text{cand}}$
- 17: **return** \mathbf{U}_{best}



1479 Figure 9: Equivariant optimal transport aligns noise samples X_0 to training examples X_1 over
1480 reflections and permutations, leading to smoother and shorter paths shown to the model during
1481 training.

1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492

1493 E STIEFEL LOGARITHM EMPIRICAL ANALYSIS

1494
1495
1496
1497
1498

1499 We sample 100k training examples from QM9 or GEOM to be used as U_1 and for each example
1500 sample one random point U_0 . We compute the true logarithm $\Delta = \log_{U_0}(U_1)$ using a large number
1501 of iterations and compare it to the 20-iteration truncated logarithm. We record error as the infinity
1502 norm of the difference between the true and approximate logarithms. We set the convergence
1503 threshold to be $1e-6$.

1504
1505
1506
1507
1508
1509

1504 For QM9, the 20-iteration logarithm converges 97.9% of the time (median 9 iterations to converge),
1505 and the median error in case of nonconvergence is $2.4e-4$. The Spearman correlation between the
1506 1-iteration approximate distance and the true distance is $\rho = 0.88$. For GEOM, the 20-iteration
1507 logarithm converges 99.7% of the time (median 9 iterations to converge), and the median error in
1508 case of nonconvergence is $8.8e-5$. The Spearman correlation between the 1-iteration approximate
1509 distance and the true distance is $\rho = 0.82$.

1510
1511

1510 These results empirically validate that the 1-iteration approximate distance used in Algorithm 4 is an
1511 upper bound on the true distance, and that it is a valid heuristic which generally maintains the same
relative ordering as the true Stiefel distance.




Figure 10: Histograms of the number of inner iterations of the Stiefel logarithm required to converge to an error of $1e-6$ for 100k random training examples of QM9 (left) and GEOM (right).




Figure 11: Parity plots comparing the 1-iteration approximate Stiefel distance to the true Stiefel distance of QM9 (left) and GEOM (right).