

Robust Symmetry Detection via Riemannian Langevin Dynamics

JIHYEON JE*, Stanford University, USA
JIAYI LIU*, Stanford University, USA
GUANDAO YANG*, Stanford University, USA
BOYANG DENG*, Stanford University, USA
SHENGQU CAI, Stanford University, USA
GORDON WETZSTEIN, Stanford University, USA
OR LITANY, Technion, Israel
LEONIDAS GUIBAS, Stanford University, USA




Fig. 1. In this work, we present a novel Langevin dynamics-based symmetry detection method. Our solution is robust to noise, and can detect both partial and global reflective symmetries on diverse shapes. Here we visualize detected reflective symmetries at different noise levels for 2D (top half), and illustrate detected planes along with its supports for selected 3D shapes (bottom half).

*Equal contribution

Authors' Contact Information: Jihyeon Je, jihyeonj@stanford.edu, Stanford University, Stanford, USA; Jiayi Liu, jiayiliu@stanford.edu, Stanford University, Stanford, USA; Guandao Yang, guandao@stanford.edu, Stanford University, Stanford, USA; Boyang Deng, bydeng@stanford.edu, Stanford University, Stanford, USA; Shengqu Cai, shengqu@stanford.edu, Stanford University, Stanford, USA; Gordon Wetzstein, gordonwz@stanford.edu, Stanford University, Stanford, USA; Or Litany, orlitany@gmail.com, Technion, Haifa, Israel; Leonidas Guibas, guibas@cs.stanford.edu, Stanford University, Stanford, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SA Conference Papers '24, December 3–6, 2024, Tokyo, Japan

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1131-2/24/12

<https://doi.org/10.1145/3680528.3687682>

Symmetries are ubiquitous across all kinds of objects, whether in nature or in man-made creations. While these symmetries may seem intuitive to the human eye, detecting them with a machine is nontrivial due to the vast search space. Classical geometry-based methods work by aggregating "votes" for each symmetry but struggle with noise. In contrast, learning-based methods may be more robust to noise, but often overlook partial symmetries due to the scarcity of annotated data. In this work, we address this challenge by proposing a novel symmetry detection method that marries classical symmetry detection techniques with recent advances in generative modeling. Specifically, we apply Langevin dynamics to a redefined symmetry space to enhance robustness against noise. We provide empirical results on a variety of shapes that suggest our method is not only robust to noise, but can also identify both partial and global symmetries. Moreover, we demonstrate the utility of our detected symmetries in various downstream tasks, such as compression and symmetrization of noisy shapes. Please see our project page and code at <https://symmetry-langevin.github.io/>.

CCS Concepts: • **Computing methodologies** → **Shape analysis**.

Additional Key Words and Phrases: Geometry Processing, Generative Modeling, Langevin Dynamics

ACM Reference Format:

Jihyeon Je, Jiayi Liu, Guandao Yang, Boyang Deng, Shengqu Cai, Gordon Wetzstein, Or Litany, and Leonidas Guibas. 2024. Robust Symmetry Detection via Riemannian Langevin Dynamics. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24)*, December 3–6, 2024, Tokyo, Japan. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3680528.3687682>

1 INTRODUCTION

Symmetry, being an aesthetically beautiful concept per se, is fundamental to our understanding of structures in geometry. According to the great mathematician Blaise Pascal, it is essentially “what we see at a glance”. Therefore, the problem of discovering symmetries in shapes has been a central pursuit in geometry processing [Atallah 1985; Mitra et al. 2006; Pauly et al. 2008; Wolter et al. 1985]. The capability of automatically detecting symmetries can consequently enable us to compress [Pauly et al. 2008] and symmetrize [Mitra et al. 2013] shapes for numerous uses in shape modeling, manufacturing, or artistic creation.

Albeit an intuitive concept, symmetry is not trivial to detect, especially when both global and partial symmetries are considered. The challenge of symmetry detection lies in the vast search space of this problem – one has to determine not only the symmetries present but also their support in the shape geometry. The brute force solution has been to search over pairs of points on the shape surface and ascertain whether their neighbouring regions may indicate the presence of a symmetry. To address this complexity, Mitra et al. [2006] propose to map the problem from the shape space to a transformation space, where each point is the transformation between a point pair sampled on the original shape surface. In this transformation space, symmetries can be found by identifying dense clusters, which they do by mean-shift clustering. However, while this largely reduces the search space, it is built on the assumption that the original shape is noise-free. In contrast, in the real world, more often we face scenarios when we need to process noisy or even incomplete shapes, sourced from imperfect sensors. Hence, limited robustness to noisy shapes can significantly restrict the applicability of a symmetry detection algorithm. On the other hand, attempts have been made to design symmetry detection algorithms by learning from data [Gao et al. 2020; Ji and Liu 2019; Zhou et al. 2021]. While this improves robustness to data noise, the lack of annotated data, particularly for partial symmetries, unfortunately limits these methods to detecting only global symmetries.

In this work, we aim to design a symmetry detection algorithm that is robust to noise, while being capable of identifying symmetries of different strengths. We extend the formulation of symmetry detection in the transformation space in Mitra et al. [2006] and provide a fresh perspective of understanding this transformation space. Inspired by the recent success of score-based sampling [Karras et al. 2022; Poole et al. 2023; Song et al. 2020, 2021a] in generative models, we view the transformation space as a score field where the symmetries we seek are the modes in this field. As robust mode seeking algorithms are well-studied in generative modeling, we introduce one, specifically Langevin dynamics, to symmetry detection. We

hypothesize that such formulation of the transformation space reduces sensitivity to noise, due to smoothing over sampled points. Building on these insights, we propose a robust symmetry detection algorithm that marries elegant ideas in geometry processing and generative models, leveraging Langevin dynamics as the key driver. We show empirical results indicating that our proposed symmetry detection algorithm is more robust to noise compared to other baselines. In addition, our method is capable of detecting both global and partial symmetries, and is easily generalizable to a wide range of 2D and 3D shapes.

2 RELATED WORKS

Symmetry. Symmetries play a crucial part in Computer Vision and Computer Graphics. They are widely used in various tasks including 3D reconstruction [Chen et al. 2018; Fawcett et al. 1994; François et al. 2003; Gordon 1990; Huynh 1999; Mukherjee et al. 1995; Phillips et al. 2016; Sinha et al. 2012; Thrun and Wegbreit 2005; Tulsiani et al. 2020; Xu et al. 2020], inverse rendering [Wu et al. 2020], shape refinement [Mitra et al. 2007], and image manipulation [Zhang et al. 2020]. Consequently, symmetry detection has long been an essential challenge in geometry processing. In this regard, we categorize prior methods into two groups, depending on whether data-driven techniques are used. For the non-learning based group, early works focused on finding exact symmetries on planar point sets [Atallah 1985; Wolter et al. 1985]. Later on, the seminal work by Mitra et al. [2006] reformulated this detection problem as clustering in the transformation space, offering a new perspective that allows the detection of partial and approximate symmetries. Follow-up works further refined this framework to detect regular and repeated geometric structures [Pauly et al. 2008] and orbits [Shi et al. 2016]. In contrast, following the revolutionary breakthrough of deep learning, more recent symmetry detection approaches have focused on learning symmetries from data. SymNet [Ji and Liu 2019] devised a supervised training framework to classify whether points are located on the symmetry plane. Alternatively, SymmetryNet [Shi et al. 2020] predicted symmetrical correspondences to avoid overfitting. NeRD [Zhou et al. 2021] further introduced a geometry reconstruction phase to verify predicted symmetries at inference time. However, in these supervised frameworks, annotating symmetries has been an expensive bottleneck. To address this issue, PRSNet [Gao et al. 2020] proposed an unsupervised symmetry prediction framework using a novel symmetry distance loss, while E3Sym [Li et al. 2023] developed an unsupervised correspondence prediction framework. Building upon the aforementioned approaches, our proposed algorithm bridges classical geometry processing techniques with tools from learning-based approaches by interpreting the transformation space as a score field. Using Langevin dynamics, we achieve training-free symmetry detection that is more robust to noisy inputs than heuristic methods, while also enabling partial symmetry prediction—a challenge yet to be addressed by learning-based approaches.

Score-based Generative Models. Score-based generative models have emerged as a powerful framework for estimating and sampling from unknown data distributions, achieving unparalleled success in diverse domains such as image generation [Dhariwal and Nichol

2021; Ho et al. 2020a; Karras et al. 2022; Rombach et al. 2022; Saharia et al. 2022; Song et al. 2020], 3D generation [Poole et al. 2023] and video synthesis [Blattmann et al. 2023a,b; Brooks et al. 2024; Ge et al. 2023; Guo et al. 2023; He et al. 2022; Ho et al. 2022]. Initially inspired by non-equilibrium thermodynamics, Sohl-Dickstein et al. [2015] proposed learning a generative model by reversing a gradual Gaussian diffusion process. One of the seminal contributions to score-based generative models is the denoising diffusion probabilistic model (DDPM) [Ho et al. 2020a], which employs two Markov chains: a forward chain that corrupts original data into noise for training purposes, and a reverse chain that restores this noise to data sample during inference. The forward chain is typically designed to methodically convert any data distribution into a simpler prior distribution, commonly a standard Gaussian distribution. The reverse chain undoes this diffusion process by utilizing deep neural networks to learn the necessary transitions. To generate new data points, one begins with a random noise vector and uses ancestral sampling through the reverse Markov chain. Leveraging the DDPM framework, diffusion models have demonstrated impressive capabilities, particularly in image synthesis [Dhariwal and Nichol 2021]. Score-based generative models [Song and Ermon 2019, 2020] offer an alternative formulation. Similar to DDPMs [Ho et al. 2020a], these models perturb data with progressively increasing Gaussian noise and estimate the score functions across all noisy distributions. A deep neural network, conditioned on varying noise levels, is trained to estimate the score functions. During inference, samples are drawn by sequentially applying score functions at decreasing noise levels. DDPMs and score-based models can also be generalized to infinitely many time steps or noise levels, where the perturbation and denoising are modeled as solutions to SDEs. Score-SDE [Song et al. 2021b] represents a significant advance by employing a diffusion process described by an SDE to perturb data with noise. The reverse process, modeled by a corresponding reverse-time SDE, gradually transforms noise back into data. With the score function parameterized by a deep neural network, samples can be drawn through various numerical techniques, such as Langevin dynamics [Song and Ermon 2019], SDE solvers [Song et al. 2021b], and ODE solvers [Karras et al. 2022; Lu et al. 2022; Song et al. 2020, 2021b]. Subsequent works focus on improving the accuracy and efficiency of score-based generative models. An early work on accelerating diffusion model sampling, Denoising Diffusion Implicit Models (DDIM) [Song et al. 2020] extends the original DDPM [Ho et al. 2020a] formulation, transitioning the sampling process to a fully deterministic reversion of this non-Markov perturbation process. Further, EDM [Karras et al. 2022] demonstrates that Heun’s second-order method [Ascher and Petzold 1998] offers a balance between sample quality and sampling speed. Its proposed higher-order solver minimizes discretization errors at the expense of an additional evaluation of the learned score function per timestep and shows that using Heun’s method [Ascher and Petzold 1998] can produce samples with quality on par with or better than Euler’s method within fewer sampling steps.

3 BACKGROUND

In this section, we introduce the notations and provide the necessary background in symmetry detection to discuss our method.

3.1 Geometric Symmetry

Symmetries, one of the most fundamental shape descriptors, can be categorized as either extrinsic or intrinsic depending on which space the symmetry exists in, and global and partial depending on the extent of the region it impacts. Despite the ubiquity, symmetries are very challenging to define [Mitra et al. 2006; Shi et al. 2016; Tevs et al. 2014]. In this paper, we adapt the geometric view to define symmetries as a transformation under which a significant region of the shape is invariant. Formally, given a shape $S \in \mathcal{S}$ represented as a set of \mathbb{R}^d points, symmetry is a transformation $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ that can be applied to a shape region $R \subseteq S$:

$$\exists R \subset S, \text{ s.t. } \mathcal{D}(T(R), R) \leq \varepsilon, \text{ and } M(R) > \tau \quad (1)$$

where $\mathcal{D} : S \times S \rightarrow \mathbb{R}$ is a distance function between two patches, $T(R)$ denotes the transformed point region $T(R) = \{T(p) : p \in R\}$, $M : S \rightarrow \mathbb{R}$ measures the significance of the shape patch. To make this definition specific, one needs to define the distance function \mathcal{D} between regions, the class of transformations T , the measure of significance M , as well as the thresholds τ and ε . We illustrate the exact formulations below.

An instance of the distance function is Chamfer distance, which leverages the Euclidean distance between points:

$$\mathcal{D}_{CD}(S_1, S_2) = \frac{1}{|S_1|} \sum_{x \in S_1} \min_{y \in S_2} \|x - y\| + \frac{1}{|S_2|} \sum_{y \in S_2} \min_{x \in S_1} \|x - y\| \quad (2)$$

Alternatively, one can use shape properties other than extrinsic shape features to compute the distance. For example, intrinsic shape features such as the Laplace–Beltrami eigenfunctions can give rise to the definition of intrinsic symmetries [Ovsjanikov et al. 2008]. One can also use semantic information or geodesic distance to define \mathcal{D} . For simplicity, we use Chamfer distance as our distance function.

Note that if the transformation class is expressive enough or the tolerance ε is very large, then any arbitrary pair of points can give rise to a symmetry. However, it is often the prominent symmetries, such as the global symmetries that can be applied to the entire shape, that are useful in downstream applications. To ensure the detection of prominent symmetries, we define a measure M for patch significance. One common way to define this measure is the area ratio of the patch region relative to the entire shape, $M(R) = |R|/|S|$. With an adequate threshold τ , we prioritize symmetries that can be applied to a significant enough region of the shape. Finally, we need to decide on the class of transformation T . In this paper, we focus on perhaps the simplest yet the most common type of symmetry: reflective symmetry, while we also show that the pipeline could be extended to other symmetry classes, such as rotation and translation (Section 5.2).

3.2 Detection of Reflective Symmetries

Reflective symmetry can be defined using a hyperplane parameterized by the surface normal $\mathbf{n} \in \mathbb{R}^d$, $\|\mathbf{n}\| = 1$ and a point on the plane $\mathbf{i} \in \mathbb{R}^d$: $H(\mathbf{n}, \mathbf{i}) = \{\mathbf{x} \mid (\mathbf{x} - \mathbf{i})^\top \mathbf{n} = 0\}$. Given




Fig. 2. Reflection across a hyperplane

this definition, reflective symmetry is a transformation:

$$T_{(\mathbf{n}, \mathbf{i})}(\mathbf{x}) = \mathbf{x} + 2(\text{proj}_{(\mathbf{n}, \mathbf{i})}(\mathbf{x}) - \mathbf{x}), \quad (3)$$

where $\text{proj}_{(\mathbf{n}, \mathbf{i})}(\mathbf{x}) = \mathbf{x} + ((\mathbf{i} - \mathbf{x})^\top \mathbf{n})\mathbf{n}$ is the function projecting point \mathbf{x} to the hyperplane parameterized by the normal vector \mathbf{n} and a point on the plane \mathbf{i} . The general idea of symmetry detection is to find a hyperplane (\mathbf{n}, \mathbf{i}) that defines a reflection applicable to a significant part of the shape. One way to achieve this is to formulate symmetry detection as finding all hyperplanes under which the result of following optimization problem returns a significant region:

$$\begin{aligned} &\text{Find all } (\mathbf{n}, \mathbf{i}) \text{ s.t. } M(R^*(\mathbf{n}, \mathbf{i})) > \tau, \\ &\text{where } R^*(\mathbf{n}, \mathbf{i}) = \arg \max_{R \subseteq S, \mathcal{D}(T_{(\mathbf{n}, \mathbf{i})}(R), R) < \varepsilon} M(R) \end{aligned} \quad (4)$$

Mitra et al. [2006] proposes to solve this via a voting algorithm. We provide a brief outline of the 2D algorithm below. For a given 2D shape, any pair of points \mathbf{p} and \mathbf{q} on the boundary of the shape S define a reflective symmetry, represented by the line through the midpoint $\mathbf{m} = \frac{1}{2}(\mathbf{p} + \mathbf{q})$ with a normal in the direction of $\mathbf{n} = (\mathbf{p} - \mathbf{q}) \|\mathbf{p} - \mathbf{q}\|^{-1}$. Now, by sufficiently sampling such point pairs on the shape, we can collect “votes” for potential reflective symmetries, T_i , in the transformation space. Mitra et al. [2006] observe that these votes create a density function in the transformation space, and the modes of the such density function correspond to significant symmetries since they are a transformation that is proximal to most votes. In other words, the optimization problem in Eq. (4) can be formulated as a mode seeking procedure of the density defined in the transformation space.

Building on this concept, Mitra et al. [2006] suggest using mean-shift algorithm [Fukunaga and Hostetler 1975] to identify clusters associated with significant modes. Specifically, given a set of candidate transformations $\{T_i\}_{i=1}^n$, the mean-shift algorithm can be thought as first constructing a density field using kernel density estimation:

$$p(T) = \frac{1}{|N(T)|h^d} \sum_{T_i \in N(T)} K\left(\frac{T - T_i}{h}\right), \quad (5)$$

where K is an appropriate kernel function, h is the bandwidth, and $N(T)$ contains all transformations near the query T . The algorithm first initializes from any of the candidate transformation $T^{(0)} = T_i$ and performs the following iterations:

$$T^{(k+1)} \leftarrow \frac{\sum_{T' \in N_k} K((T' - T^{(k)})h^{-1})T'}{\sum_{T' \in N_k} K((T' - T^{(k)})h^{-1})}, \quad (6)$$

where $N_k = N(T^{(k)})$ is the center of the kernel at the k^{th} iteration. This process repeats until the centroid is converged, *i.e.* $\|T^{(k+1)} - T^{(k)}\| < \gamma$ for certain γ . Once the proposed transformation T^* is found, Mitra et al. [2006] computes R^* in Eq. (4) using a region growing algorithm and filters out R^* if $M(R^*)$ is not significant enough. While the mean-shift voting scheme has shown promising results in detecting varying types and scales of symmetries, they are limited in that it fails under noisy shapes (*i.e.* noisy transformation space) and does not fully leverage modern GPU compute or parallel processing power. To tackle such shortcomings, we establish a connection between Mitra et al. [2006] and the literature

state-of-the-art generative models. This connection allows us to leverage advancements in modern computing and deep learning techniques, applying them to the task of symmetry detection.

4 METHOD

In this section, we first establish the connection between Langevin dynamics, as used in diffusion models, and the mean-shift algorithm used in symmetry detection (Section 4.1). Building on the insight derived from this connection, we propose a symmetry detection method using Langevin dynamics to achieve robustness to noise. We also outline the design of the transformation representation, define the distance between reflective symmetries (Section 4.2), and illustrate how to perform Langevin dynamics in this transformation space (Section 4.3). Finally, we provide a simple algorithm to post-process the resulting modes after Langevin dynamics to extract the final symmetries (Section 4.4).

4.1 Connecting Mean-shift and Langevin Dynamics

Diffusion models [Ho et al. 2020b] and flow matching [Lipman et al. 2022] aim to learn the score function, or the gradient of the log density function. These score functions can be used with Langevin dynamics (or any other SDE solver) to transport samples from the prior distribution to the data distribution. Specifically, given empirical data samples $\mathcal{X} = \{x_i \sim P_{\text{data}}\}_{i=1}^n$, one way to define the objective of the diffusion model is to learn the score for the following density, which convolves the actual data density with a Gaussian to create a blurry version:

$$P_{\sigma}(x) = \int P_{\text{data}}(y)\mathcal{N}(x; y, \sigma^2 I)dy \approx \frac{1}{|\mathcal{X}|} \sum_{i=1}^{|\mathcal{X}|} \mathcal{N}(x; x_i, \sigma^2 I), \quad (7)$$

where the RHS is an approximation using empirical samples \mathcal{X} . This allows us to write the score function [Cai et al. 2020; Karras et al. 2022; Pabbaraju et al. 2023; Song et al. 2021a] as the following:

$$\nabla_x \log P_{\sigma}(x) \approx \left(\frac{\sum_{y \in \mathcal{X}} \mathcal{N}(x; y, \sigma^2 I) \cdot y}{\sum_{y \in \mathcal{X}} \mathcal{N}(x; y, \sigma^2 I)} - x \right) \sigma^{-2}. \quad (8)$$

Once the model learns to approximate this score function, one way to obtain samples is through Langevin dynamics. Specifically, we can initialize Langevin dynamics with samples from a prior distribution $x^{(0)} \sim P_{\text{prior}}$, and then perform the following iterations:

$$x^{(t+1)} \leftarrow x^{(t)} + \alpha_t \nabla_x \log P_{\sigma_t}(x^{(t)}) + \sqrt{2\alpha_t} \beta_t \epsilon_t, \quad (9)$$

where $\epsilon_t \sim \mathcal{N}(0, I)$ is the noise and α_t is the step size. A key connection here is that if we modify the above-mentioned formula slightly, we are able to obtain the updates of mean-shift in Eq. (6). Specifically, if we set $\beta_t = 0$, $\alpha_t = \sigma_t^2$, and $\sigma_t = h$, then the Langevin updating rule becomes:

$$\begin{aligned} x^{(t+1)} &\leftarrow x^{(t)} + h^2 \left(\frac{\sum_{y \in \mathcal{X}} \mathcal{N}(x^{(t)}; y, h^2 I) \cdot y}{\sum_{y \in \mathcal{X}} \mathcal{N}(x^{(t)}; y, h^2 I)} - x \right) / h^2 \\ &= \frac{\sum_{y \in \mathcal{X}} \mathcal{N}(x^{(t)}; y, h^2 I) \cdot y}{\sum_{y \in \mathcal{X}} \mathcal{N}(x^{(t)}; y, h^2 I)}. \end{aligned} \quad (10)$$

The update rule above can be thought as mean-shift algorithm using Gaussian kernel (i.e. $K(x, y) = \exp\left(-\frac{\|x-y\|^2}{h^2}\right)$), but with no stochasticity, infinite neighborhood (i.e. $N(x) = \mathcal{X}$), as well as constant noise value $\sigma_t = h$. Intuitively, this connection roots in the fact that both Langevin dynamics and mean-shift are trying to perform iterative hill climbing in a density function.

Despite the similarities, there are also differences in the choice of hyperparameters. For example, adding stochastic noise in Langevin dynamics is generally considered to bring an advantage in sampling quality, allowing us to obtain symmetry modes of varying strengths. A design choice of mean-shift algorithm is to update with the gradient estimated from a finite neighborhood, which makes the update computationally efficient. The same hyperparameter choice in Langevin dynamics has similar effects to using a small σ_t or truncating the gradients. However, this can be problematic when data is sparse or noisy, potentially creating spurious modes that require post-processing. These observations lead us to ask: *can diffusion models be applied to perform mode-seeking for symmetry detection?* Nevertheless, directly applying Langevin dynamics to the space of reflective transformation is challenging because the transformation space has geometries different from Euclidean space, as it contains an $SO(n)$ rotation. This requires us to construct a space where we can define an appropriate density function, estimate scores, and perform Langevin dynamics.

4.2 Transformation Space of Reflective Symmetry

As we directly walk on the transformation manifold, choosing the optimal transformation representation was of utmost importance. Previous works have explored representing transformations as Lie groups [Shi et al. 2016] and Euclidean transformation groups [Mitra et al. 2006], yet these representations are often much higher dimensional – typically in \mathbb{R}^7 . Such complexity makes identification of high-density regions in the space challenging.

A good space for diffusion model should be bounded and easy to compute. Therefore, we resort to the Hesse normal form utilized in Hough transforms [Duda and Hart 1972] to represent the hyperplane that defines a reflective symmetry. Hough transform

represents each hyperplane $\{x | (x - i)^T \mathbf{n} = 0\}$ using a rotation parameterized by the surface normal vector $\mathbf{n} \in \mathbb{R}^d$, oriented away from the origin, and a scalar $l \in \mathbb{R}$




Fig. 3. Transformation space construction

tance to the plane from the origin. Given this representation (\mathbf{n}, l) , the reflective plane can be written as: $\{x | x^T \mathbf{n} = l\}$. A convenient property of Hough transform is that both \mathbf{n} and l are bounded as long as the input shape S is bounded. Moreover, it's easy to compute the potential reflective plane for a pair of points in the shape $(\mathbf{p}, \mathbf{q}) \in S$:

$$\mathbf{n}(\mathbf{p}, \mathbf{q}) = (\mathbf{p} - \mathbf{q}) \|\mathbf{p} - \mathbf{q}\|^{-1}, \quad l(\mathbf{p}, \mathbf{q}) = \frac{1}{2} (\mathbf{p} + \mathbf{q})^T \mathbf{n}(\mathbf{p}, \mathbf{q}) \quad (11)$$

Such computation is very efficient and parallelizable, allowing us to scale to a large number of points. Instead of doing walks in




Fig. 4. Geodesic distance to the black point in transformation space. Blue indicates proximity, and red indicates larger distance. Shortest distance can be computed traveling through the origin as well.

the concatenate space (\mathbf{n}, l) , which has one additional degree of freedom than the actual space of reflective symmetry, we embed the symmetry into \mathbb{R}^d using the following:

$$T(\mathbf{p}, \mathbf{q}) = \mathbf{n}(\mathbf{p}, \mathbf{q}) \cdot (\text{sign}(l(\mathbf{p}, \mathbf{q})) \cdot k + l(\mathbf{p}, \mathbf{q})) \quad (12)$$

where k is the hyperparameter denoting the radius of the invalid region. This parameter is introduced to deal with ambiguities when $l(\mathbf{p}, \mathbf{q}) = 0$, where multiple symmetry planes are mapped to the origin. Specifically, we shift all planes in the direction of the normal by k . This representation is numerically stable in most cases except for when $\mathbf{p} = \mathbf{q}$, which we discard as they do not create a valid vote for reflective symmetry. Each reflective hyperplane thus gets mapped to a single point in the transformation space (Figure 3).

4.3 Langevin Dynamics in Transformation Space

One challenge of performing the walk in the transformation space defined as $T(\mathbf{p}, \mathbf{q})$ is that Euclidean distance in the space does not reflect the actual distance between symmetry planes. This is due to the fact that reflective symmetry contains $SO(n)$ rotation symmetry. Fortunately, the score function can be computed in a Riemannian manifold where only distances are defined [Chen and Lipman 2024; Huang et al. 2022]. Below, we first cover the background of score estimation in a Riemannian manifold, which enables us to compute the geodesic function for the reflective transformation space. Next, using the score function derived from the geodesic, we define the process of walking in the transformation space.

Langevin Dynamics in Riemannian Geometry. Let \mathcal{M} be a Riemannian manifold where we define a premetrics measuring the closeness of points $\mathbf{d} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ such that \mathbf{d} is (1) non-negative (i.e. $\mathbf{d}(x, y) \geq 0, \forall x, y \in \mathcal{M}$), (2) positive (i.e. $\mathbf{d}(x, y) = 0$ iff $x = y$), and (3) non-degenerate (i.e. $\nabla_x \mathbf{d}(x, y) \neq 0$ iff $x \neq y$). Given the premetrics function, Chen and Lipman [2024] suggest that one way to compute the gradient of the probability density flow is using the following formula:

$$g_\sigma(x) = \int_{\mathcal{M}} \mathbf{d}(x, y) \frac{\nabla_x \mathbf{d}(x, y)}{\|\nabla_x \mathbf{d}(x, y)\|^2} \frac{p_\sigma(x|y)}{p_\sigma(x)} P_{\text{data}}(y) d\text{vol}_y, \quad (13)$$

where $p_\sigma(x|y)$ is a Gaussian distribution on the geodesic distance: $p_\sigma(x|y) \propto \exp(-\mathbf{d}(x, y)^2 \sigma^{-2})$. The gradient field of the probability density flow can be used in place of the score function in Langevin dynamics. This means as long as we are able to create a premetrics (or geodesic distance) for the transformation space \mathcal{M} to reflect the closeness of reflective symmetries, we are able to compute g_σ to perform Langevin dynamics.

Distance of Reflective Symmetry. Intuitively, the distance between two reflective symmetries T_1, T_2 should capture how hard it is to transform one hyperplane to the other. A seemingly straightforward approach is to directly compute the Euclidean distance between T_1 and T_2 . While this gives some measure of geodesic, it is problematic near the origin because when $l = 0$, both \mathbf{n} and $-\mathbf{n}$ represent the same hyperplane, yet $(0, \mathbf{n})$ and $(0, -\mathbf{n})$ is far away in our representation (Eq. (12)). To address this issue, we propose the following geodesic distance:

$$d(x, y) = \min \begin{cases} \min_z \|x - z\| + \|y + z\| \\ \min_r \int_0^1 \text{valid}(r(t)) |r'(t)| dt \end{cases}, \quad (14)$$

where $\|z\| = k$, $r(0) = x$, $r(1) = y$, and $\text{valid}(x) = 1 + \infty \cdot \mathbf{1}(\|x\| > k)$ marks the valid region of the path allowed. Intuitively, the first term captures the shortest path if the hyperplane moves past the origin. The second term captures the shortest path if we do not go through the origin. We visualize this geodesic distance at different points in Figure 4. We implement the geodesic with JAX and compute the gradient $\nabla_x d(x, y)$ using automatic differentiation.

Stepping in the Transformation space. Given the definition of the vector field g_σ which maps a point in the geodesic to the tangent space, we can replace the score function in Eq. (10). However, naively doing so will lead to an issue where the points might end up in invalid areas that don't represent any valid reflective symmetry plane (e.g. $\|x^{(t)}\| < k$). This occurs if the gradient field g_σ is trying to move points across the origin to find symmetries with more votes. One way to counter this is to properly define the behavior of moving a hyperplane across the origin with certain velocity. Specifically, suppose T is a point with $l = 0$ and normal \mathbf{n} (so it's on the sphere with radius k). If we would like to move T with velocity v_d and v_n , this is equivalent to moving $-T$ (i.e. $l = 0$ and $-\mathbf{n}$) with v_d and v_n . If we encounter such a case during the walk, we consider moving both T and $-T$, then only send the point to the valid region.

This operation is performed both during the gradient step (i.e. walking along g_σ) as well as the noise perturbation step (i.e. walking along ϵ). We illustrate the full Langevin dynamics trajectory as well as the corresponding hyperplane in the shape space upon convergence in Figure 6. We show the detailed implementation and hyperparameter setup in the appendix.




Fig. 5. Stepping in the transformation space. We move both T and $-T$, and only send the point to the valid destination.

4.4 Symmetry Extraction

After sufficient steps of Langevin dynamics, randomly initialized points at the beginning of the walk will eventually converge at the local maxima of the density field if guided properly. However, this still gives us a set of extremely narrow distributions centered at each mode. In order to extract a single symmetry for every peak, we perform Density-Based Spatial Clustering of Applications with Noise (DBSCAN), a density-based clustering algorithm that groups points based on neighborhood density [Ester et al. 1996; Khan et al. 2014; Schubert et al. 2017]. After Langevin we already expect the modes to be well-separated and dense, thus, DBSCAN is only utilized




Fig. 6. Visualizing Langevin dynamics trajectory. For each converged mode, we show the full trajectory (top row), as well as the corresponding symmetry in the shape space (bottom row). Note that due to our Langevin construction, we directly jump across the invalid region as opposed to stepping within it. When converged modes are near the boundary of the invalid region ($l \approx k$), we see oscillations through the origin as the two modes are equivalent ($H(0, \mathbf{n}) = H(0, -\mathbf{n})$).

to identify the centroids and to remove potentially remaining low-density points. Therefore, we enforce DBSCAN with a very high density threshold and a small distance threshold. Once the centroid at each local maxima is identified, we map it back to shape space to find the corresponding symmetry. We provide experiment results quantifying the impact of DBSCAN in the appendix.

4.5 Datasets

We evaluate our symmetry detection pipeline on both 2D and 3D shapes. In order to evaluate the performance of the proposed symmetry detection pipeline, there has to be a good amount of symmetries present in the dataset; yet common baselines such as MNIST datasets [LeCun et al. 1998] do not contain much symmetries. Therefore, for 2D shapes, we construct a dataset of English alphabet letters of varying fronts from Google Fonts¹. We extract 10 fonts for each letter, and extract the contour. For 3D shapes, we utilize 20 selected object meshes from ShapeNetCore [Chang et al. 2015] and Adobe Assets. We normalize all shapes by shifting the center of mass to the origin and scaling them such that all points are within $[-1, 1]^d$. To assess the noise-robustness of the method, we introduce varying degrees of noise to the constructed dataset. For 2D, we directly add scaled Gaussian noise $\alpha \cdot \mathcal{N}(0, 1)$ for $\alpha = [0\%, 1\%, 3\%, 5\%]$ to all sampled points on the normalized shape. For 3D, we add up to 3% noise to all vertices in the direction of the normal. This results in multiple variations of each shape at different noise levels.

However, a key challenge remains in data annotation, as there exists no prior symmetry ground truth. Therefore, we develop a pipeline for generating ground truth symmetry annotations on collected shapes. We first take a brute force approach and sample point pairs to construct a symmetry transformation space. For each proposal from a point pair we calculate the number of points that would vote for this proposal, and only keep the proposals that have more than a threshold number of votes. We then apply clustering to group similar transformations together to collect a set of symmetry proposals. Thresholds in this process are set fairly low such that we always propose more symmetries than there actually are. Given

¹<https://fonts.google.com/>

these loose proposals, a human annotator can either decide whether to include the proposed symmetry in the ground truth set or not.

4.6 Evaluation metrics

In our analysis, we employ precision and recall as our primary evaluation metrics. By comparing the set of predicted symmetries and the set of ground truth symmetries, the two metrics could be directly obtained. During our evaluation, we consider two symmetries to be the same if they are within some δ -radius of each other in the transformation space. However, due to the labor-intensiveness of the human-labeling of ground truth data, we propose two additional metrics to evaluate the quality of a predicted set of symmetries: association and compression rate.

To obtain the association measure, we first define the association of a point with a particular symmetry as the condition where the projection of the point through the symmetry falls within an epsilon radius of another point. Then for each symmetry, we calculate the proportion of points satisfying this criterion. We then determine the distribution of these proportions across the set of symmetries. The association measure of the symmetry set is quantified as the integral of this distribution, representing the area under the curve across all ratios of total number of points. This metric serves as a precision measure, capturing the degree to which the set of symmetries accurately associate points of the shape within a defined tolerance.

On the other hand, for the compression measure, we use a heuristic-based compression algorithm to apply symmetries for reducing the number of points required to represent a shape, given the suspected NP-hard nature of an optimal compression algorithm. This algorithm iteratively applies the best symmetries to minimize the representation space, and the resultant compression ratio is calculated by dividing the original memory needed to store all points by themselves with the memory required for the compressed object. Note that the compressed object includes spatial locations of retained points, the symmetries utilized, and the sequence of their application. A lower compression ratio indicates a higher efficacy and relevance of the symmetries used, which serves as a proxy for recall.

5 RESULTS AND APPLICATIONS

In this section, we provide results indicating the robustness of our method, which can identify reflective symmetries of different strengths even on noisy shapes. We first show demonstrative examples of detected symmetry planes, evaluate our method quantitatively with other baselines, and provide potential downstream applications.

5.1 Evaluation on Symmetry Detection Performance

We show quantitative comparisons with neural baselines and Mitra et al. [2006] for 3D (Table 1) and 2D shapes (Table 2). Results indicate that our method outperforms the baseline methods for F1




Fig. 7. Symmetry groups found within a castle wall. The shape patches that support the detected symmetries are colored the same as the corresponding plane. We identify both local and global reflective symmetries.

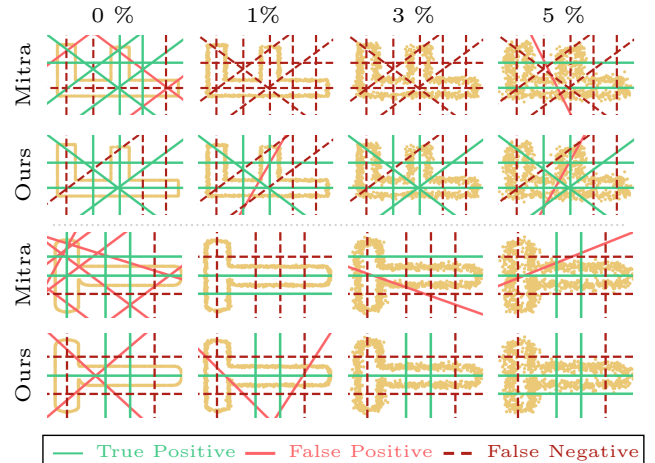


Fig. 8. Qualitative comparison with Mitra et al. [2006] for 2D. We visualize the input 2D font shapes and the predicted reflective symmetries for varying levels of noise. We show that our solution is more robust to noise, capable of detecting key symmetries even at high levels of noise.




Fig. 9. Qualitative comparison with baselines for 3D. All detected symmetries are shown for 0, 1, 3% noise. Our method is capable of detecting both local and global symmetries accurately even at high noise levels.

score, compression, and association, especially by a large margin at high noise levels. This is potentially because our method works under a Gaussian-smoothed transformation space, where Langevin dynamics is utilized to help travel through the potentially noisy parameter space and extract the modes, bypassing certain noisy information. On the other hand, Mitra et al. [2006] directly operates under the noisy transformation space and further uses stochastic

Table 2. Quantitative results for 2D symmetry detection.

Methods	F1 \uparrow				Compression \downarrow				Association \uparrow			
	0%	1%	3%	5%	0%	1%	3%	5%	0%	1%	3%	5%
Mitra	0.602	0.208	0.287	0.347	0.483	0.793	0.783	0.757	0.366	0.303	0.234	0.246
Ours	0.576	0.540	0.481	0.443	0.389	0.413	0.478	0.530	0.476	0.484	0.477	0.459




Fig. 10. Qualitative results for 3D. We visualize the input 3D shape and selected reflective planes from our detection results. We only highlight the largest connected region that agree to the symmetry, and color each side of the reflective plane differently. Results indicate that our method is capable of detecting symmetries of varying strengths on a wide array of shapes.




Fig. 11. We show that our pipeline can be extended to translation and rotational symmetries. For translation, we identify global negative (a) and local positive shifts (b, c). For rotation, we identify global (a) and local rotational axis (b, c, d). The transformation space (left) and the identified modes (right) for both symmetry types are shown.

clustering to extract the modes. As both procedures are not robust to noise, this likely contributes to its weaker performances on noisy

shapes. It is also worth noting that our method has significantly improved compression and association rates across all noise levels, which further demonstrates its advantages in potential downstream applications.

We show qualitative results of our symmetry detection algorithm in comparison with baselines for 2D and 3D (Figure 8, Figure 9), and results on additional 3D shapes (Figure 10). While we miss smaller symmetries as noise is introduced, our method retains higher recall across all levels when compared to baseline methods. Particularly, we see that we are able to constantly identify the biggest modes (i.e. global reflective symmetry) throughout. Whereas previous neural methods were limited to global symmetries, we identify symmetries of varying strengths on a wide range of unseen shapes.




Fig. 12. Applications in symmetry fixing for noisy shapes. We take an imperfect mesh generated from a single view image, then symmetrize the shape based on the detected symmetry planes. Shapes fixed with our detection is of higher quality than with Mitra et al. [2006] as our method is more robust to noise.




Fig. 13. Applications in shape compression. We take a shape with symmetries at many different scales, and perform sequential compression based on the detected modes. Region to be compressed is colored blue at each stage.

5.2 Applications

In this section, we demonstrate potential downstream applications utilizing the detected symmetries. We first show an extension to different types of symmetries - translation and rotation, then demonstrate applications in shape fixing and compression.

Extension to additional symmetry types. As our approach is training-free, extending the pipeline to other symmetry types is possible as long as a representative and bounded transformation space can be constructed. We show proof-of-concept examples of extensions to rotation and translation symmetries, where we only modify the transformation space we walk on (Figure 11). For constructing the transformation space for translation, we directly extract vector displacement for every point pair (instead of the Hough transform representation). For rotational symmetries, we utilize a special case of the Cartan-Dieudonné theorem, which states that every rotation in $SO(n)$ is a composition of at most n -reflections about hyperplanes. Given two reflective planes, the rotational axis corresponds to the line of intersection between two planes, where the angle of rotation is twice the angle of intersection. With this idea, we directly extend the reflective transformation space to rotational space.

Extension to shape fixing. Utilizing the method’s robustness to noise, we show potential applications to downstream shape fixing tasks (Figure 12). We implement the shape symmetrization pipeline

proposed in Mitra et al. [2007]. We demonstrate that given an imperfect shape generated from SOTA image-to-mesh model [Xu et al. 2024], we can accurately detect symmetry planes and fix the shapes accordingly. As our method can be directly applied to unseen noisy shapes, it could be further utilized in correcting the output of 3D generative models.

Extension to shape compression. Another downstream application for symmetry detection is in shape compression. We perform sequential model compression based on the detected local and global reflective planes (Figure 13) with the pipeline. Using the first four significant modes, a reduction to only 22% of the original model size is achieved by removing redundant shape patches. This demonstrates the potential of symmetries as a complexity-reducing concept to simplify shapes.

6 FUTURE WORK AND LIMITATIONS

In this work, we bring a novel perspective to the symmetry detection problem and propose a noise-robust algorithm. Building on classical approaches, we introduce modern concepts from generative models, particularly Langevin dynamics, to address prior shortcomings. However, there are a few limitations that are worth exploring in the future.

Neural gradient field. We show that with sufficiently sampling the shape space, we are able to construct a transformation space that is representative of the symmetries present within the shape. Yet, the quality of the transformation space as well as the gradient field constructed from this space is heavily dependent on the sample size. Therefore, one may learn this gradient field to better capture the transformation space without sampling.

Collecting ground truth evaluation data. As there exists no prior symmetry dataset, the evaluation of the method is greatly limited by the lack of fine-grained symmetry annotations. In this work we propose a semi-supervised method to generate a base set of ground truth data, although expanding this to a larger dataset that captures different types of symmetries will help perform a more comprehensive evaluation.

Covering a larger symmetry space. While we show that our pipeline is capable of identifying reflective symmetries both in partial and global scales, there are other types of symmetries that we did not cover extensively in the scope of this work. Further extending the work to continuous symmetries (e.g. rotational and helical

surfaces), translational symmetries, or a combination of multiple symmetry types would be a promising next step.

7 CONCLUSIONS

We have introduced a novel symmetry detection algorithm based on Riemannian Langevin dynamics. We formulate the transformation space as a score field, on which we perform mode-seeking with Langevin dynamics.

Our experiments validate the robustness and generalizability of the approach, showing improved mode detection even in the presence of noise. The proposed method is capable of detecting both partial and global reflective symmetries in 2D and 3D shapes, which addresses limitations with previous neural approaches. Finally, we show that our detection method generalizes easily to a wide range of shapes and can be integrated into various downstream tasks, highlighting the practical value of our algorithm.

ACKNOWLEDGMENTS

This work was supported by ARL grant W911NF-21-2-0104, Vannevar Bush Faculty Fellowship, and Azrieli Foundation Early Career Faculty Fellowship. We would also like to thank Eric Chan, Marcel Rød, Lucia Zheng, Chenglei Si, Zoe Wefers, George Nakayama, and Maolin Gao for all the helpful discussions and feedback.

REFERENCES

Uri M. Ascher and Linda R. Petzold. 1998. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. Society for Industrial and Applied Mathematics, Philadelphia, PA. <https://doi.org/10.1137/1.9781611971392> arXiv:<https://epubs.siam.org/doi/pdf/10.1137/1.9781611971392>

Atallah. 1985. On symmetry detection. *IEEE Trans. Comput.* 100, 7 (1985), 663–666.

Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. 2023a. Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets. In *arXiv*.

Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. 2023b. Align your Latents: High-Resolution Video Synthesis with Latent Diffusion Models. In *CVPR*.

Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. 2024. Video generation models as world simulators. (2024). <https://openai.com/research/video-generation-models-as-world-simulators>

Ruojin Cai, Guandao Yang, Hadar Averbuch-Elor, Zekun Hao, Serge Belongie, Noah Snaveley, and Bharath Hariharan. 2020. Learning gradient fields for shape generation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III* 16. Springer, 364–381.

Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. 2015. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012* (2015).

Ricky T. Q. Chen and Yaron Lipman. 2024. Flow Matching on General Geometries. arXiv:2302.03660 [cs.LG]

Xin Chen, Yuwei Li, Xi Luo, Tianjia Shao, Jingyi Yu, Kun Zhou, and Youyi Zheng. 2018. Autosweep: Recovering 3d editable objects from a single photograph. *IEEE transactions on visualization and computer graphics* 26, 3 (2018), 1466–1475.

Prafulla Dhariwal and Alexander Quinn Nichol. 2021. Diffusion Models Beat GANs on Image Synthesis. In *NeurIPS*.

Richard O. Duda and Peter E. Hart. 1972. Use of the Hough transformation to detect lines and curves in pictures. *Commun. ACM* 15 (1972), 11–15. <https://api.semanticscholar.org/CorpusID:1105637>

M Ester, H P Kriegel, J Sander, and Xu Xiaowei. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. (12 1996). <https://www.osti.gov/biblio/421283>

Roger Fawcett, Andrew Zisserman, and J Michael Brady. 1994. Extracting structure from an affine view of a 3D point set with one or two bilateral symmetries. *Image and Vision Computing* 12, 9 (1994), 615–622. [https://doi.org/10.1016/0262-8856\(94\)90015-9](https://doi.org/10.1016/0262-8856(94)90015-9)

Alexandre R. J. François, Gérard G. Medioni, and Roman Waupotitsch. 2003. Mirror symmetry to 2-view stereo geometry. *Image and Vision Computing* 21 (2003), 137–143.

K. Fukunaga and L. Hostetler. 1975. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Transactions on Information Theory* 21, 1 (1975), 32–40. <https://doi.org/10.1109/TIT.1975.1055330>

Lin Gao, Ling-Xiao Zhang, Hsien-Yu Meng, Yi-Hui Ren, Yu-Kun Lai, and Leif Kobbelt. 2020. PRS-Net: Planar reflective symmetry detection net for 3D models. *IEEE transactions on visualization and computer graphics* 27, 6 (2020), 3007–3018.

Songwei Ge, Seungjun Nah, Guilin Liu, Tyler Poon, Andrew Tao, Bryan Catanzaro, David Jacobs, Jia-Bin Huang, Ming-Yu Liu, and Yogesh Balaji. 2023. Preserve Your Own Correlation: A Noise Prior for Video Diffusion Models. In *ICCV*.

Gaile G. Gordon. 1990. Shape From Symmetry. In *Intelligent Robots and Computer Vision VIII: Algorithms and Techniques*, David P. Casasent (Ed.), Vol. 1192. International Society for Optics and Photonics, SPIE, 297–308. <https://doi.org/10.1117/12.969742>

Yuwei Guo, Ceyuan Yang, Anyi Rao, Yaohui Wang, Yu Qiao, Dahua Lin, and Bo Dai. 2023. AnimateDiff: Animate Your Personalized Text-to-Image Diffusion Models without Specific Tuning. In *arXiv*.

Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. 2022. Latent Video Diffusion Models for High-Fidelity Long Video Generation. In *arXiv*.

Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P. Kingma, Ben Poole, Mohammad Norouzi, David J. Fleet, and Tim Salimans. 2022. Imagen Video: High Definition Video Generation with Diffusion Models. In *arXiv*.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020a. Denoising Diffusion Probabilistic Models. In *NeurIPS*.

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

Chin-Wei Huang, Milad Aghajohari, Joey Bose, Prakash Panangaden, and Aaron C Courville. 2022. Riemannian diffusion models. *Advances in Neural Information Processing Systems* 35 (2022), 2750–2761.

D.Q. Huynh. 1999. Affine reconstruction from monocular vision in the presence of a symmetry plane. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, Vol. 1. 476–482 vol.1. <https://doi.org/10.1109/ICCV.1999.791259>

Penglei Ji and Xinguo Liu. 2019. A fast and efficient 3D reflection symmetry detector based on neural networks. *Multimedia Tools and Applications* 78 (2019), 35471–35492. <https://api.semanticscholar.org/CorpusID:201103305>

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. 2022. Elucidating the Design Space of Diffusion-Based Generative Models. In *NeurIPS*.

Kamran Khan, Saif Ur Rehman, Kamran Aziz, Simon Fong, and Sababady Sarasvady. 2014. DBSCAN: Past, present and future. In *The fifth international conference on the applications of digital information and web technologies (ICADIWT 2014)*. IEEE, 232–238.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.

Ren-Wu Li, Ling-Xiao Zhang, Chungpeng Li, Yu-Kun Lai, and Lin Gao. 2023. E3Sym: Leveraging E (3) invariance for unsupervised 3D planar reflective symmetry detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14543–14553.

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

Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. 2022. DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps. *arXiv preprint arXiv:2206.00927* (2022).

Niloy J Mitra, Leonidas J Guibas, and Mark Pauly. 2006. Partial and approximate symmetry detection for 3d geometry. *ACM Transactions on Graphics (TOG)* 25, 3 (2006), 560–568.

Niloy J Mitra, Leonidas J Guibas, and Mark Pauly. 2007. Symmetrization. *ACM Transactions on Graphics (TOG)* 26, 3 (2007), 63–es.

Niloy J Mitra, Mark Pauly, Michael Wand, and Duygu Ceylan. 2013. Symmetry in 3d geometry: Extraction and applications. In *Computer graphics forum*, Vol. 32. Wiley Online Library, 1–23.

Dipti Prasad Mukherjee, Andrew Zisserman, and Michael Brady. 1995. Shape from Symmetry: Detecting and Exploiting Symmetry in Affine Images. *Philosophical Transactions: Physical Sciences and Engineering* 351, 1695 (1995), 77–106. <http://www.jstor.org/stable/54462>

Maks Ovsjanikov, Jian Sun, and Leonidas Guibas. 2008. Global intrinsic symmetries of shapes. In *Computer graphics forum*, Vol. 27. Wiley Online Library, 1341–1348.

Chirag Pabbaraju, Dhruv Rohatgi, Anish Sevekari, Holden Lee, Ankur Moitra, and Andrej Risteski. 2023. Provable benefits of score matching. arXiv:2306.01993 [cs.LG]

Mark Pauly, Niloy J Mitra, Johannes Wallner, Helmut Pottmann, and Leonidas J Guibas. 2008. Discovering structural regularity in 3D geometry. In *ACM SIGGRAPH 2008 papers*. 1–11.

Cody J. Phillips, Matthieu Lecce, and Kostas Daniilidis. 2016. Seeing Glassware: from Edge Detection to Pose Estimation and Shape Recovery. In *Robotics: Science and Systems*. <https://api.semanticscholar.org/CorpusID:18915780>

Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. 2023. DreamFusion: Text-to-3D using 2D Diffusion. In *ICLR*.

- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR*.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. In *arXiv*.
- Erich Schubert, Jörg Sander, Martin Ester, Hans Peter Kriegel, and Xiaowei Xu. 2017. DBSCAN revisited, revisited: why and how you should (still) use DBSCAN. *ACM Transactions on Database Systems (TODS)* 42, 3 (2017), 1–21.
- Yifei Shi, Junwen Huang, Hongjia Zhang, Xin Xu, Szymon Rusinkiewicz, and Kai Xu. 2020. SymmetryNet: Learning to Predict Reflectional and Rotational Symmetries of 3D Shapes from Single-View RGB-D Images. *ACM Transactions on Graphics (SIGGRAPH Asia 2020)* 39, 6 (2020).
- Zeyun Shi, Pierre Alliez, Mathieu Desbrun, Hujun Bao, and Jin Huang. 2016. Symmetry and orbit detection via lie-algebra voting. In *Computer graphics forum*, Vol. 35. Wiley Online Library, 217–227.
- Sudipta N. Sinha, Krishnan Ramnath, and Richard Szeliski. 2012. Detecting and Reconstructing 3D Mirror Symmetric Objects. In *Computer Vision – ECCV 2012*, Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 586–600.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 37)*, Francis Bach and David Blei (Eds.). PMLR, Lille, France, 2256–2265. <https://proceedings.mlr.press/v37/sohl-dickstein15.html>
- Jiaming Song, Chenlin Meng, and Stefano Ermon. 2020. Denoising Diffusion Implicit Models. In *ICLR*.
- Yang Song and Stefano Ermon. 2019. Generative modeling by estimating gradients of the data distribution. In *NeurIPS*.
- Yang Song and Stefano Ermon. 2020. Improved techniques for training score-based generative models. In *NeurIPS*.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021a. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=PXTIG12RRHS>
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021b. Score-Based Generative Modeling through Stochastic Differential Equations. In *ICLR*.
- Art Tevs, Qixing Huang, Michael Wand, Hans-Peter Seidel, and Leonidas Guibas. 2014. Relating shapes via geometric symmetries and regularities. *ACM Trans. Graph.* 33, 4, Article 119 (jul 2014), 12 pages.
- Sebastian Thrun and Ben Wegbreit. 2005. Shape from symmetry. *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1 2* (2005), 1824–1831 Vol. 2.
- Shubham Tulsiani, Or Litany, Charles R Qi, He Wang, and Leonidas J Guibas. 2020. Object-centric multi-view aggregation. *arXiv preprint arXiv:2007.10300* (2020).
- Jan D Wolter, Tony C Woo, and Richard A Volz. 1985. Optimal algorithms for symmetry detection in two and three dimensions. *The Visual Computer* 1 (1985), 37–48.
- Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. 2020. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1–10.
- Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. 2024. InstantMesh: Efficient 3D Mesh Generation from a Single Image with Sparse-view Large Reconstruction Models. arXiv:2404.07191 [cs.CV] <https://arxiv.org/abs/2404.07191>
- Yifan Xu, Tianqi Fan, Yi Yuan, and Gurprit Singh. 2020. Ladybird: Quasi-Monte Carlo Sampling for Deep Implicit Field Based 3D Reconstruction with Symmetry. arXiv:2007.13393 [cs.CV]
- Xuaner Zhang, Jonathan T. Barron, Yun-Ta Tsai, Rohit Pandey, Xiuming Zhang, Ren Ng, and David E. Jacobs. 2020. Portrait Shadow Manipulation. *ACM Transactions on Graphics (TOG)* 39, 4.
- Yichao Zhou, Shichen Liu, and Yi Ma. 2021. NeRD: Neural 3d reflection symmetry detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15940–15949.

APPENDIX

A METHOD DETAILS

We provide algorithm blocks to better illustrate the Langevin dynamics process on the transformation manifold. See Algorithm 1 for the overall Langevin dynamics setup, Algorithm 2 for geodesic distance compute, and Algorithm 3 for walking on the manifold. Please refer to Appendix B for detailed hyperparameter setup.

ALGORITHM 1: Langevin Dynamics.

```

Require: Noise levels  $\{\sigma_i\}_{i=1}^k$ ; Step size  $\alpha$ ; Number of steps  $T$ 
Initialize  $x_0$ 
for  $i \leftarrow 1$  to  $k$  do
  for  $t \leftarrow 0$  to  $T - 1$  do
     $\epsilon_t \sim \mathcal{N}(0, 1)$ 
     $x'_{t+1} \leftarrow x_t + \frac{\sqrt{\alpha}\sigma_i\epsilon_t}{\sigma_k}$ 
     $g_\sigma(x) \leftarrow \text{jax.grad}(\text{geodesic}(x))$ 
     $x_{t+1} \leftarrow \text{walk}(x'_{t+1}, \frac{\alpha\sigma_i^2}{2\sigma_k^2} g_\sigma(x'_{t+1}, \sigma_i))$ 
  end
   $x_0 \leftarrow x_T$ 
end

```

ALGORITHM 2: Computing the Geodesic Distance.

```

Require:  $x$ ;  $y$ ; Radius of the invalid region  $R$ 
Function  $\text{geodesic}(x, y, R)$ :
   $x_{\text{valid}}, y_{\text{valid}} \leftarrow \neg\text{inside\_invalid}(x, R), \neg\text{inside\_invalid}(y, R)$ 
  if  $x_{\text{valid}} \wedge y_{\text{valid}}$  then
     $d_{\text{through}} \leftarrow \text{distance\_jump\_invalid}(x, y, R)$ 
     $d_{\text{outside}} \leftarrow \text{distance\_outside\_invalid}(x, y, R)$ 
    return  $\min(d_{\text{through}}, d_{\text{outside}})$ 
  else
    return  $-1$ 
  end

```

ALGORITHM 3: Walking in the Transformation Space.

```

Require:  $x$ ; Gradient  $g$ ; Radius of the invalid region  $R$ 
Function  $\text{walk}(x, g, R)$ :
   $\text{target} \leftarrow x + g$ 
   $\text{inside} \leftarrow \text{True}$  if  $\|\text{target}\| < R$ ; else False
   $t_M \leftarrow -x \cdot g$ 
   $M \leftarrow x + t_M * g$ 
   $\theta_{\text{rot}} \leftarrow \arccos(\frac{\|M\|}{R})$ 
   $Y_{\text{pos}}, Y_{\text{neg}} \leftarrow \text{rotation\_matrix}(\pm\theta_{\text{rot}}) \cdot \frac{M}{\|M\|} \cdot R$ 
   $T_{Y_{\text{pos}}}, T_{Y_{\text{neg}}} \leftarrow (Y_{\text{pos}}, Y_{\text{neg}} - x) \cdot g$ 
   $Y \leftarrow x + \min(T_{Y_{\text{pos}}}, T_{Y_{\text{neg}}}) * g$ 
   $v_y \leftarrow \text{target} - Y$ 
   $v_N \leftarrow (v_y \cdot Y) \cdot Y/R$ 
   $\theta_{\text{sign}} \leftarrow \text{sign}(\frac{v_y - v_N}{\|v_y - v_N\|} \times \frac{Y}{R})$ 
   $Y_{\text{refl}} \leftarrow -\text{rotation\_matrix}(\frac{\|v_y - v_N\|}{R} * \theta_{\text{sign}}) \cdot Y$ 
   $\text{target}_{\text{refl}} \leftarrow Y_{\text{refl}} + \frac{Y_{\text{refl}}}{\|Y_{\text{refl}}\|} \cdot \|v_N\|$ 
  if  $\text{inside}$  then
    return  $\text{target}_{\text{refl}}$ 
  else
    return  $\text{target}$ 
  end

```

B EXPERIMENT DETAILS

The parameters we tune for transformation space construction are: k (the radius of the invalid region) and the number of sampled points. The parameter k is introduced to deal with ambiguities when $l(p, q) = 0$, where multiple symmetry planes are mapped to the origin. Specifically, we shift all planes in the direction of the normal by k (Eq. 12). We find that 50K point pairs sampled from both 2D and 3D shapes are sufficient to construct a dense transformation space to capture local modes.

For mode seeking, we utilize the annealed version of Langevin dynamics [Song and Ermon 2019]. We begin with random samples and perform recursive updates toward the gradient while gradually decreasing noise. The gradient is computed from the kernel-approximated geodesic distances [Chen and Lipman 2024; Huang et al. 2022]. This requires tuning kernel size, step size, and number of steps. A smaller kernel size allows finer modes to be discovered but may require more steps for convergence, while a larger kernel size may overlook these finer modes. With hyperparameter tuning, we achieve the best performance with: kernel size = 0.025, $k = 0.3$, step size = 0.06 for 2D, and kernel size = 0.08, $k = 0.5$, step size = 0.02 for 3D. We use 200 points for Langevin initialization and take 50K steps for both. We show the effects of altering such parameters below, where we profile end-to-end runtime for varying number of steps, for which we measure the wall-clock time on a 12G RTX 2080 GPU (Table A1). We also quantify the importance of the hyperparameter k as well as kernel size (Table A2). We report the F1 score for all.

Table A1. Runtime analysis

	Number of steps			
	1K	5K	10K	50K
F1	0.759	0.813	0.774	0.857
Runtime	29.3s	100s	4min	15min

Table A2. Hyperparameter setup

	k				Kernel size			
	0.3	0.5	0.8	1	0.2	0.1	0.08	0.05
F1	0.733	0.857	0.710	0.733	0.700	0.846	0.857	0.560

C ABLATION STUDIES

As we assume the converged modes after Langevin steps to be well-separated and dense, we only utilize DBSCAN as a post-processing step to extract final mode centroids. With this assumption, we enforce it with a very high density and a small distance threshold. In order to confirm that we are not relying on DBSCAN for clustering, we perform an ablation study quantifying the impact of DBSCAN. We replace mean-shift clustering in [Mitra et al. 2006] with DBSCAN and report the F1 score for all noise levels. We determine the hyperparameters for DBSCAN following the approaches in Sander et al. [1998]. We observe that DBSCAN struggles with noisy shapes even more so than mean-shift, likely due to its reliance on fixed-density thresholds.

Table A3. Ablation study on the impact of DBSCAN

	0%	1%	3%	5%
Mitra	0.602	0.208	0.287	0.347
Mitra + DBSCAN	0.517	0.182	0.120	0.097
Ours	0.576	0.540	0.481	0.443

D ROBUSTNESS TO SPARSE RECONSTRUCTION

We acknowledge that directly adding noise to the shape might not be realistically capturing noises encountered in real-life applications. Therefore, we evaluate the method further on sparsely-reconstructed point clouds, which are generated from back-projecting 4, 6, and 8-views randomly taken from the input 3D shape. We show that the pipeline is capable of consistently detecting modes even in sparse conditions.




Fig. A1. Robustness to sparse input. We show detection results on sparse-view reconstructed input point cloud, as well as the top 3 planes for the 8-view case.