

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 EQUIVARIANT FLOW MATCHING FOR POINT CLOUD ASSEMBLY

Anonymous authors

Paper under double-blind review

ABSTRACT

The goal of point cloud assembly is to reconstruct a complete 3D shape by aligning multiple point cloud pieces. This work presents a novel equivariant solver for assembly tasks based on flow matching models. We first theoretically show that the key to learning equivariant distributions via flow matching is to learn related vector fields. Based on this result, we propose an assembly model, called equivariant diffusion assembly (Eda), which learns related vector fields conditioned on the input pieces. We further construct an equivariant path for Eda, which guarantees high data efficiency of the training process. Our numerical results show that Eda is highly competitive on practical datasets, and it can even handle the challenging situation where the input pieces are non-overlapped.

1 INTRODUCTION

Point cloud (PC) assembly is a classic machine learning task which seeks to reconstruct 3D shapes by aligning multiple point cloud pieces. This task has been intensively studied for decades and has various applications such as scene reconstruction (Zeng et al., 2017), robotic manipulation (Ryu et al., 2024), cultural relics reassembly (Wang et al., 2021) and protein designing (Watson et al., 2023). A key challenge in this task is to correctly align PC pieces with small or no overlap region, *i.e.*, when the correspondences between pieces are lacking.

To address this challenge, some recent methods (Ryu et al., 2024; Wang and Jörnsten, 2024) utilized equivariance priors for pair-wise assembly tasks, *i.e.*, the assembly of two pieces. In contrast to most of the state-of-the-art methods (Qin et al., 2022; Zhang, 1994) which align PC pieces based on the inferred correspondence, these equivariant methods are correspondence-free, and they are guided by the equivariance law underlying the assembly task. As a result, these methods are able to assemble PCs without correspondence, and they enjoy high data efficiency and promising accuracy. However, the extension of these works to multi-piece assembly tasks remains largely unexplored.

In this work, we develop an equivariant method for multi-piece assembly based on flow matching (Lipman et al., 2023). Our main theoretical finding is that to learn an equivariant distribution via flow matching, one only needs to ensure that the initial noise is invariant and the vector field is related (Thm. 4.2). In other words, instead of directly handling the $SE(3)^N$ -equivariance for N -piece assembly tasks, which can be computationally expensive, we only need to handle the related vector fields on $SE(3)^N$, which is efficient and easy to construct. Based on this result, we present a novel assembly model called equivariant diffusion assembly (Eda), which uses invariant noise and predicts related vector fields by construction. Eda is correspondence-free and is guaranteed to be equivariant by our theory. Furthermore, we construct a short and equivariant path for the training of Eda, which guarantees high data efficiency of the training process. When Eda is trained, an assembly solution can be sampled by numerical integration, *e.g.*, the Runge-Kutta method, starting from a random noise. All proofs can be found in Appx. F. A brief walk-through of our theory using a toy example with minimal terminologies is provided in Appx. C

The contributions of this work are summarized as follows:

- We present an equivariant flow matching framework for multi-piece assembly tasks. Our theory reduces the task of constructing equivariant conditional distributions to the task of constructing related vector fields, thus it provides a feasible way to define equivariant flow matching models.

054 - Based on the theoretical result, we present a simple and efficient multi-piece PC assembly model,
 055 called equivariant diffusion assembly (Eda), which is correspondence-free and is guaranteed to be
 056 equivariant. We further construct an equivariant path for the training of Eda, which guarantees
 057 high data efficiency.
 058 - We numerically show that Eda produces highly accurate results on the challenging 3DMatch and
 059 BB datasets, and it can even handle non-overlapped pieces.
 060

061 **2 RELATED WORK**
 062

064 Our proposed method is based on flow matching (Lipman et al., 2023), which is one of the state-
 065 of-the-art diffusion models for image generation tasks (Esser et al., 2024). Some applications on
 066 manifolds have also been investigated (Chen and Lipman, 2024; Yim et al., 2023). Our model has
 067 two distinguishing features compared to existing methods: it learns conditional distributions instead
 068 of marginal distributions, and it explicitly incorporates equivariance priors.

069 The PC assembly task studied in this work is related to various tasks in the literature, such as PC
 070 registration (Qin et al., 2022; Yu et al., 2023), robotic manipulation (Ryu et al., 2024; 2023) and
 071 fragment reassembly (Wu et al., 2023a). All these tasks aim to align the input PC pieces, but they are
 072 different in settings such as the number of pieces, deterministic or probabilistic, and whether the PCs
 073 are overlapped. More details can be found in Appx. B. In this work, we consider the most general
 074 setting: we aim to align multiple pieces of non-overlapped PCs in a probabilistic way.

075 Recently, diffusion-based methods have been proposed for assembly tasks (Chen et al., 2025; Jiang
 076 et al., 2023; Wu et al., 2023b; Li et al., 2025; Ryu et al., 2024; Scarpellini et al., 2024; Xu et al.,
 077 2024). However, most of these works ignore the manifold structure or the equivariance priors of the
 078 task. One notable exception is Ryu et al. (2024), which developed an equivariant diffusion method
 079 for robotic manipulation, *i.e.*, pair-wise assembly tasks. Compared to Ryu et al. (2024), our method
 080 is conceptually simpler because it does not require Brownian diffusion on $SO(3)$ whose kernel is
 081 computationally intractable, and it solves the more general multi-piece problem. On the other hand,
 082 the invariant flow theory has been studied in Köhler et al. (2020), which can be regarded as a special
 083 case of our theory as discussed in Appx. F.1. Furthermore, the optimal-transport-based method was
 084 explored for invariant flow (Song et al., 2023; Klein et al., 2023).

085 Another branch of related work is equivariant neural networks. Due to their ability to incorporate
 086 geometric priors, this type of networks has been widely used for processing 3D graph data such
 087 as PCs and molecules. In particular, E3NN (Geiger and Smidt, 2022) is a well-known equivariant
 088 network based on the tensor product of the input and the edge feature. An acceleration technique
 089 for E3NN was recently proposed (Passaro and Zitnick, 2023). On the other hand, the equivariant
 090 attention layer was studied in Fuchs et al. (2020); Liao and Smidt (2023); Liao et al. (2024). Our
 091 work is related to this line of approach, because our diffusion network can be seen as an equivariant
 092 network with an additional time parameter.

093 **3 PRELIMINARIES**
 094

095 This section introduces the major tools used in this work. We first define the equivariances in Sec. 3.1,
 096 then we briefly recall the flow matching model in Sec. 3.2.
 097

098 **3.1 EQUIVARIANCES OF PC ASSEMBLY**
 099

100 Consider the action $G = \prod_{i=1}^N SE(3)$ on a set of N ($N \geq 2$) PCs $X = \{X_1, \dots, X_N\}$, where
 101 $SE(3)$ is the 3D rigid transformation group, \prod is the direct product, and X_i is the i -th PC piece
 102 in 3D space. We define the action of $\mathbf{g} = (g_1, \dots, g_N) \in G$ on X as $\mathbf{g}X = \{g_i X_i\}_{i=1}^N$, *i.e.*, each
 103 PC X_i is rigidly transformed by the corresponding g_i . For the rotation subgroup $SO(3)^N$, the
 104 action of $\mathbf{r} = (r_1, \dots, r_N) \in SO(3)^N$ on X is $\mathbf{r}X = \{r_i X_i\}_{i=1}^N$. For $SO(3) \subseteq G$, we denote
 105 $\mathbf{r} = (r, \dots, r) \in SO(3)$ for simplicity, and the action of \mathbf{r} on X is written as $\mathbf{r}X = \{r X_i\}_{i=1}^N$.
 106

107 We also consider the permutations of X . Let S_N be the permutation group of N , the action of $\sigma \in S_N$
 108 on X is $\sigma X = \{X_{\sigma(i)}\}_{i=1}^N$, and the action on \mathbf{g} is $\sigma \mathbf{g} = (g_{\sigma(1)}, \dots, g_{\sigma(N)})$. For group multiplication,

108 we denote $\mathcal{R}_{(\cdot)}$ the right multiplication and $\mathcal{L}_{(\cdot)}$ the left multiplication, *i.e.*, $(\mathcal{R}_r)\mathbf{r}' = \mathbf{r}'\mathbf{r}$, and
 109 $(\mathcal{L}_r)\mathbf{r}' = \mathbf{r}\mathbf{r}'$ for $\mathbf{r}, \mathbf{r}' \in SO(3)^N$.
 110

111 In our setting, for the given input X , the solution to the assembly task is a conditional distribution
 112 $P_X \in \mu(G)$, where $\mu(G)$ is the set of probability distribution on G . We study the following three
 113 equivariances of P_X in this work:

114 **Definition 3.1.** Let $P_X \in \mu(G)$ be a probability distribution on $G = SE(3)^N$ conditioned on X ,
 115 and let $(\cdot)_{\#}$ be the pushforward of measures.

- 116 - P_X is $SO(3)^N$ -equivariant if $(\mathcal{R}_{\mathbf{r}^{-1}})_{\#}P_X = P_{\mathbf{r}X}$ for $\mathbf{r} \in SO(3)^N$.
- 117 - P_X is permutation-equivariant if $\sigma_{\#}P_X = P_{\sigma X}$ for $\sigma \in S_N$.
- 118 - P_X is $SO(3)$ -invariant if $(\mathcal{L}_r)_{\#}P_X = P_X$ for $r \in SO(3)$.

121 As an example, we explicitly show the equivariance in Def. 3.1 for a two-piece deterministic problem.

122 **Example 3.2.** Assume that a solution for point clouds (X_1, X_2) is (r_1, r_2) , meaning r_1X_1 and r_2X_2
 123 are assembled, then

- 124 - $SO(3)^2$ -equivariance: a solution for (r_3X_1, r_4X_2) is $(r_1r_3^{-1}, r_2r_4^{-1})$;
- 125 - Permutation-equivariance: a solution for (X_2, X_1) is (r_2, r_1) ;
- 126 - $SO(3)$ -invariance: another solution for (X_1, X_2) is (rr_1, rr_2) .

129 More discussions on the definition of equivariances can be found in Appx. D

131 We finally recall the definition of $SO(3)$ -equivariant networks, which will be the main computational
 132 tool of this work. We call $F^l \in \mathbb{R}^{2l+1}$ a degree- l $SO(3)$ -equivariant feature if the action of $r \in SO(3)$
 133 on F^l is the matrix-vector production: $rF^l = R^lF^l$, where $R^l \in \mathbb{R}^{(2l+1) \times (2l+1)}$ is the degree- l
 134 Wigner-D matrix of r . We call a network w $SO(3)$ -equivariant if it maintains the equivariance
 135 from the input to the output: $w(rX) = rw(X)$, where $w(X)$ is a $SO(3)$ -equivariant feature. More
 136 detailed introduction of equivariances and the underlying representation theory can be found in Cesa
 137 et al. (2022).

138 3.2 VECTOR FIELDS AND FLOW MATCHING

140 To sample from a data distribution $P_1 \in \mu(M)$, where M is a smooth manifold (we only consider
 141 $M = G$ in this work), the flow matching (Lipman et al., 2023) approach constructs a time-dependent
 142 diffeomorphism $\phi_{\tau} : M \rightarrow M$ satisfying $(\phi_0)_{\#}P_0 = P_0$ and $(\phi_1)_{\#}P_0 = P_1$, where $P_0 \in \mu(M)$
 143 is a fixed noise distribution, and $\tau \in [0, 1]$ is the time parameter. Then the sample of P_1 can be
 144 represented as $\phi_1(g)$ where g is sampled from P_0 .

145 Formally, ϕ_{τ} is defined as a flow, *i.e.*, an integral curve, generated by a time-dependent vector field
 146 $v_{\tau} : M \rightarrow TM$, where TM is the tangent bundle of M :

$$147 \quad \begin{aligned} \frac{\partial}{\partial \tau} \phi_{\tau}(\mathbf{g}) &= v_{\tau}(\phi_{\tau}(\mathbf{g})), \\ 148 \quad \phi_0(\mathbf{g}) &= \mathbf{g}, \quad \forall \mathbf{g} \in M. \end{aligned} \quad (1)$$

150 According to Lipman et al. (2023), an efficient way to construct v_{τ} is to define a path h_{τ} connecting
 151 P_0 to P_1 . Specifically, let \mathbf{g}_0 and \mathbf{g}_1 be samples from P_0 and P_1 respectively, and $h_0 = \mathbf{g}_0$ and
 152 $h_1 = \mathbf{g}_1$. v_{τ} can be constructed as the solution to the following problem:

$$154 \quad \min_v \mathbb{E}_{\tau, \mathbf{g}_0 \sim P_0, \mathbf{g}_1 \sim P_1} \|v_{\tau}(h_{\tau}) - \frac{\partial}{\partial \tau} h_{\tau}\|_F^2. \quad (2)$$

155 When v is learned using (2), we can obtain a sample from P_1 by first sampling a noise \mathbf{g}_0 from P_0
 156 and then taking the integral of (1).

158 In this work, we consider a family of vector fields, flows and paths conditioned on the given PC, and
 159 we use the pushforward operator on vector fields to study their relatedness (Tu, 2011). Formally,
 160 let $F : M \rightarrow M$ be a diffeomorphism, v and w be vector fields on M . w is F -related to v if
 161 $w(F(\mathbf{g})) = F_{*,\mathbf{g}}v(\mathbf{g})$ for all $\mathbf{g} \in M$, where $F_{*,\mathbf{g}}$ is the differential of F at \mathbf{g} . Note that we denote
 162 v_X, ϕ_X and h_X the vector field, flow and path conditioned on PC X respectively.

162 *Remark 3.3.* For readers that are not familiar with this definition, relatedness can be simply regarded
 163 as a transformation, so the above definition simply means w is the transformation of v by F . More
 164 details can be found in Sec.14.6 in the text book Tu (2011).

166 4 METHOD

168 In this section, we provide the details of the proposed Eda model. First, the PC assembly problem
 169 is formulated in Sec. 4.1. Then, we parametrize related vector fields in Sec. 4.2. The training and
 170 sampling procedures are finally described in Sec. 4.3 and Sec. 4.4 respectively.

173 4.1 PROBLEM FORMULATION

174 Given a set X containing N PC pieces, *i.e.*, $X = \{X_i\}_{i=1}^N$ where X_i is the i -th piece, the goal of
 175 assembly is to learn a distribution $P_X \in \mu(G)$, *i.e.*, for any sample \mathbf{g} of P_X , $\mathbf{g}X$ should be the
 176 aligned complete shape. We assume that P_X has the following equivariances:

177 **Assumption 4.1.** P_X is $SO(3)^N$ -equivariant, permutation-equivariant and $SO(3)$ -invariant.

179 We seek to approximate P_X using flow matching. To avoid translation ambiguity, we also assume
 180 that, without loss of generality, the aligned PCs $\mathbf{g}X$ and each input piece X_i are centered, *i.e.*,
 181 $\sum_i \mathbf{m}(g_i X_i) = 0$, and $\mathbf{m}(X_i) = 0$ for all i , where $\mathbf{m}(\cdot)$ is the mean vector.

183 4.2 EQUIVARIANT FLOW

185 The major challenge in our task is to ensure the equivariance of the learned distribution, because a
 186 direct implementation of flow matching (1) generally does not guarantee any equivariance. To address
 187 this challenge, we utilize the following theorem, which claims that when the noise distribution P_0 is
 188 invariant and vector fields v_X are related, the pushforward distribution $(\phi_X) \# P_0$ is guaranteed to be
 189 equivariant.

190 **Theorem 4.2.** *Let G be a smooth manifold, $F : G \rightarrow G$ be a diffeomorphism, and $P \in \mu(G)$. If
 191 vector field $v_X \in TG$ is F -related to vector field $v_Y \in TG$, then*

$$193 \quad F_\# P_X = P_Y, \quad (3)$$

194 where $P_X = (\phi_X) \# P_0$, $P_Y = (\phi_Y) \# (F_\# P_0)$. Here $\phi_X, \phi_Y : G \rightarrow G$ are generated by v_X and v_Y
 195 respectively.

197 Specifically, Thm. 4.2 provides a concrete way to construct the three equivariances required by
 198 Assumption 4.1 as follow.

199 **Assumption 4.3 (Invariant noise).** P_0 is $SO(3)^N$ -invariant, permutation-invariant and $SO(3)$ -
 200 invariant, *i.e.*, $(\mathcal{R}_{\mathbf{r}^{-1}}) \# P_0 = P_0$, $\sigma_\# P_0 = P_0$ and $P_0 = (\mathcal{L}_{\mathbf{r}}) \# P_0$ for $\mathbf{r} \in SO(3)^N$, $\sigma \in S_N$
 201 and $\mathbf{r} \in SO(3)$.

203 **Corollary 4.4.** *Under assumption 4.3,*

- 204 • if v_X is $\mathcal{R}_{\mathbf{r}^{-1}}$ -related to $v_{\mathbf{r}X}$, then $(\mathcal{R}_{\mathbf{r}^{-1}}) \# P_X = P_{\mathbf{r}X}$, where $P_X = (\phi_X) \# P_0$ and $P_{\mathbf{r}X} =$
 205 $(\phi_{\mathbf{r}X}) \# P_0$. Here $\phi_X, \phi_{\mathbf{r}X} : G \rightarrow G$ are generated by v_X and $v_{\mathbf{r}X}$ respectively.
- 207 • if v_X is σ -related to $v_{\sigma X}$, then $\sigma_\# P_X = P_{\sigma X}$, where $P_X = (\phi_X) \# P_0$ and $P_{\sigma X} = (\phi_{\sigma X}) \# P_0$.
 208 Here $\phi_X, \phi_{\sigma X} : G \rightarrow G$ are generated by v_X and $v_{\sigma X}$ respectively.
- 209 • if v_X is $\mathcal{L}_{\mathbf{r}}$ -invariant, *i.e.*, v_X is $\mathcal{L}_{\mathbf{r}}$ -related to v_X , then $(\mathcal{L}_{\mathbf{r}}) \# P_X = P_X$, where $P_X = (\phi_X) \# P_0$.

212 According to Cor. 4.4, if the vector fields v_X are related, then the solution P_X is guaranteed
 213 to be equivariant. Therefore, the problem is reduced to constructing related vector fields. We
 214 start by constructing $(\mathcal{R}_{\mathbf{g}^{-1}})$ -related vector fields, which are $(\mathcal{R}_{\mathbf{g}^{-1}})$ -related by definition, where
 215 $\mathbf{g} \in SE(3)^N$ and $\mathbf{r} \in SO(3)^N$. Specifically, we have the following proposition:

216 **Proposition 4.5.** v_X is $\mathcal{R}_{\mathbf{g}^{-1}}$ -related to $v_{\mathbf{g}X}$ if and only if $v_X(\mathbf{g}) = v_{\mathbf{g}X}(e)\mathbf{g}$ for all $\mathbf{g} \in SE(3)^N$.

Prop. 4.5 suggests that for $(\mathcal{R}_{g^{-1}})$ -related vector fields v_X , $v_X(\mathbf{g})$ is fully determined by the value of v_{gX} at the identity element e . Therefore, to parametrize v_X , we only need to parametrize v_{gX} at one single point e . Specifically, let f be a neural network parametrizing $v_X(e)$ for input X , *i.e.*, $f(X) = v_X(e)$, v_X can then be written as

$$v_X(\mathbf{g}) = f(gX)\mathbf{g}. \quad (4)$$

Here, $f(X) \in \mathfrak{se}(3)^N$ takes the form of

$$f(X) = \bigoplus_{i=1}^N f_i(X) \quad \text{where} \quad f_i(X) = \begin{pmatrix} w_X^i(X) & t^i(X) \\ 0 & 0 \end{pmatrix} \in \mathfrak{se}(3) \subseteq \mathbb{R}^{4 \times 4}. \quad (5)$$

The rotation component $w_X^i(X) \in \mathbb{R}^{3 \times 3}$ is a skew matrix with elements in the vector $w^i(X) \in \mathbb{R}^3$, and $t^i(X) \in \mathbb{R}^3$ is the translation component. For simplicity, we omit the superscript i when the context is clear.

Now we proceed to the other two types of relatedness of v_X . According to the following proposition, when v_X is written as (4), these two relatedness of v_X can be guaranteed if the network f is equivariant.

Proposition 4.6. *For v_X defined in (4),*

- *if f is permutation-equivariant, *i.e.*, $f(\sigma X) = \sigma f(X)$ for $\sigma \in S_N$ and PCs X , then v_X is σ -related to $v_{\sigma X}$.*
- *if f is $SO(3)$ -equivariant, *i.e.*, $w(rX) = rw(X)$ and $t(rX) = rt(X)$ for $r \in SO(3)$ and PCs X , then v_X is \mathcal{L}_r -invariant.*

Finally, we define $P_0 = (U_{SO(3)} \otimes \mathcal{N}(0, \omega I))^N$, where $U_{SO(3)}$ is the uniform distribution on $SO(3)$, \mathcal{N} is the normal distribution on \mathbb{R}^3 with mean zero and isotropic variance $\omega \in \mathbb{R}_+$, and \otimes represents the independent coupling. It is straightforward to verify that P_0 indeed satisfies assumption 4.3.

In summary, with P_0 and v constructed above, the learned distribution is guaranteed to be $SO(3)^N$ -equivariance, permutation-equivariance and $SO(3)$ -invariance.

4.3 TRAINING

To learn the vector field v_X (4) using flow matching (2), we now need to define h_X , and the sampling strategy of τ , \mathbf{g}_0 and \mathbf{g}_1 . A canonical choice (Chen and Lipman, 2024) is $\bar{h}(\tau) = \mathbf{g}_0 \exp(\tau \log(\mathbf{g}_0^{-1} \mathbf{g}_1))$, where \mathbf{g}_0 and \mathbf{g}_1 are sampled independently, and τ is sampled from a predefined distribution, *e.g.*, the uniform distribution $U_{[0,1]}$. However, this definition of h , \mathbf{g}_0 and \mathbf{g}_1 does not utilize any equivariance property of v_X , thus it does not guarantee a high data efficiency.

To address this issue, we construct a “short” and equivariant h_X in the following two steps. First, we independently sample \mathbf{g}_0 from P_0 and $\tilde{\mathbf{g}}_1$ from P_X , and obtain $\mathbf{g}_1 = r^* \tilde{\mathbf{g}}_1$, where $r^* \in SO(3)$ is a rotation correction of $\tilde{\mathbf{g}}_1$:

$$r^* = \arg \min_{r \in SO(3)} \|r \tilde{\mathbf{g}}_1 - \mathbf{g}_0\|_F^2. \quad (6)$$

Then, we define h_X as

$$h_X(\tau) = \exp(\tau \log(\mathbf{g}_1 \mathbf{g}_0^{-1})) \mathbf{g}_0. \quad (7)$$

We call h_X (7) a path generated by \mathbf{g}_0 and $\tilde{\mathbf{g}}_1$. A similar rotation correction in the Euclidean space was studied in Song et al. (2023); Klein et al. (2023). Note that h_X (7) is a well-defined path connecting \mathbf{g}_0 to \mathbf{g}_1 , because $h_X(0) = \mathbf{g}_0$ and $h_X(1) = \mathbf{g}_1$, and \mathbf{g}_1 follows P_X (Prop. F.5).

The advantages of h_X (7) are twofold. First, instead of connecting a noise \mathbf{g}_0 to an independent data sample $\tilde{\mathbf{g}}_1$, h_X connects \mathbf{g}_0 to a modified sample \mathbf{g}_1 where the redundant rotation component is removed, thus it is easier to learn. Second, the velocity fields of h_X enjoy the same relatedness as v_X (4), which leads to high data efficiency. Formally, we have the following observation.

Proposition 4.7 (Data efficiency). *Under assumption 4.3, 4.1, and F.4, we further assume that v_X satisfies the relatedness property required in Cor. 4.4, *i.e.*, v_X is $\mathcal{R}_{r^{-1}}$ -related to v_{rX} , v_X is σ -related to $v_{\sigma X}$, and v_X is \mathcal{L}_r -invariant. Denote $L(X) = \mathbb{E}_{\tau, \mathbf{g}_0 \sim P_0, \tilde{\mathbf{g}}_1 \sim P_X} \|v_X(h_X(\tau)) - \frac{\partial}{\partial \tau} h_X(\tau)\|_F^2$ the training loss (2) of PC X , where h_X is generated by \mathbf{g}_0 and $\tilde{\mathbf{g}}_1$ as defined in (7). Then*

270 - $L(X) = L(rX)$ for $r \in SO(3)^N$.
 271
 272 - $L(X) = L(\sigma X)$ for $\sigma \in S_N$.
 273
 274 - $L(X) = \hat{L}(X)$, where $\hat{L}(X) = \mathbb{E}_{\tau, \mathbf{g}'_0 \sim P_0, \mathbf{g}'_1 \sim (\mathcal{L}_r)_{\#} P_X} \|v_X(h_X(\tau)) - \frac{\partial}{\partial \tau} h_X(\tau)\|_F^2$ is the loss
 275 where the data distribution P_X is pushed forward by $\mathcal{L}_r \in SO(3)$.

276 Prop. 4.7 implies that when h_X (7) is combined with the equivariant components developed in
 277 Sec. 4.2, the following three data augmentations are automatically incorporated into the training
 278 process: 1) random rotation of each input piece X_i , 2) random permutation of the order of the input
 279 pieces, and 3) random rotation of the assembled shape.

281 4.4 SAMPLING VIA THE RUNGE-KUTTA METHOD

283 Finally, when the vector field v_X (4) is learned, we can obtain a sample \mathbf{g}_1 from P_X by numerically
 284 integrating v_X starting from a noise \mathbf{g}_0 from P_0 . In this work, we use the Runge-Kutta (RK) solver
 285 on $SE(3)^N$, which is a generalization of the classical RK solver on Euclidean spaces. For clarity,
 286 we present the formulations below, and refer the readers to Crouch and Grossman (1993) for more
 287 details.

288 To apply the RK method, we first discretize the time interval $[0, 1]$ into I steps, *i.e.*, $\tau_i = \frac{i}{I}$
 289 for $i = 0, \dots, I$, with a step length $\eta = \frac{1}{I}$. For the given input X , denote $f(\mathbf{g}X)$ at time τ
 290 by $f_\tau(\mathbf{g})$ for simplicity. The first-order RK method (RK1), *i.e.*, the Euler method, is to iterate:
 291 $\mathbf{g}_{i+1} = \exp(\eta f_{\tau_i}(\mathbf{g}_i))\mathbf{g}_i$, for $i = 0, \dots, I$. To achieve higher accuracy, we can use the fourth-order
 292 RK method (RK4). More details can be found in E.

294 5 IMPLEMENTATION

296 This section provides the details of the net-
 297 work f (5). Our design principle is to imitate
 298 the standard transformer structure (Vaswani
 299 et al., 2017) to retain its best practices. In ad-
 300 dition, according to Prop. 4.6, we also require
 301 f to be permutation-equivariant and $SO(3)$ -
 302 equivariant.

303 The overall structure of the proposed network is
 304 shown in Fig. 1. In a forward pass, the input PC
 305 pieces $\{X_i\}_{i=1}^N$ are first downsampled using a
 306 few downsampling blocks, and then fed into
 307 the Croco blocks (Weinzaepfel et al., 2022) to
 308 model their relations. Meanwhile, the time step
 309 τ is first embedded using a multi-layer perceptron (MLP)
 310 and then incorporated into the above blocks
 311 via adaptive normalization (Peebles and Xie, 2023). The output is finally obtained by a piece-wise
 312 pooling.

313 Next, we provide details of the equivariant attention layers, which are the major components of both
 314 the downsampling block and the Croco block, in Sec. 5.1. Other layers, including the nonlinear and
 315 normalization layers, are described in Sec. 5.2.

316 5.1 EQUIVARIANT ATTENTION LAYERS

318 The equivariant attention layers are based on e3nn (Geiger and Smidt, 2022). For the input point
 319 cloud, the KNN graph is first built, and the query Q , key K and value V matrices are computed for
 320 each node. Then the dot-product attention is computed where each node attends to its neighbors. We
 321 further use the reduction technique (Passaro and Zitnick, 2023) to accelerate the computation. More
 322 details can be found in Appx. G.

323 Following Croco (Weinzaepfel et al., 2022), we stack two types of attention layers, *i.e.*, the self-
 324 attention layer and the cross-attention layer, into a Croco block to learn the features of each PC

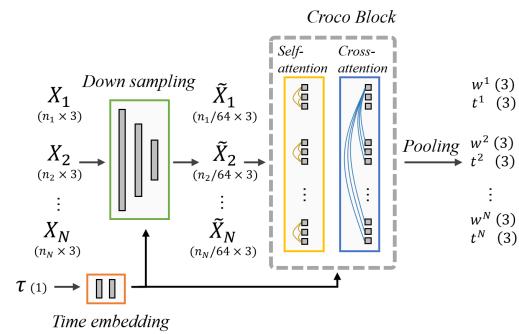


Figure 1: An overview of our model. The shapes of variables are shown in the brackets.

324 piece while incorporating information from other pieces. For self-attention layers, we build KNN
 325 graph where the neighbors are selected from the same pieces, and for cross-attention layers, we build
 326 KNN graph where the neighbors are selected from the different pieces. In addition, to reduce the
 327 computational cost, we use downsampling layers to reduce the number of points before the Croco
 328 layers. Each downsampling layer consists of a farthest point sampling (FPS) layer and a self-attention
 329 layer.

331 5.2 ADAPTIVE NORMALIZATION AND NONLINEAR LAYERS

333 Following the common practice (Devlin et al., 2019), we seek to use the GELU activation function
 334 (Hendrycks and Gimpel, 2016) in our transformer structure. However, GELU in its original form
 335 is not $SO(3)$ -equivariant. To address this issue, we adopt a projection formulation similar to Deng et al.
 336 (2021). Specifically, we define the equivariant GELU (Elu) layer as: $Elu(F^l) = GELU(\langle F^l, \widehat{WF^l} \rangle)$
 337 where $\widehat{x} = x/\|x\|$ is the normalization, $W \in \mathbb{R}^{c \times c}$ is a learnable weight. Note that Elu is a natural
 338 extension of GELU, because when $l = 0$, $Elu(F^0) = GELU(\pm F^0)$.

339 As for the normalization layers, we use RMS-type layer normalization layers (Zhang and Sennrich,
 340 2019) following Liao et al. (2023), and we use the adaptive normalization (Peebles and Xie, 2023)
 341 technique to incorporate the time step τ . Specifically, we use the adaptive normalization layer AN
 342 defined as: $AN(F^l, \tau) = F^l/\sigma \cdot MLP(\tau)$, where $\sigma = \sqrt{\frac{1}{c \cdot l_{max}} \sum_{l=1}^{l_{max}} \frac{1}{2l+1} \langle F^l, F^l \rangle}$, l_{max} is the
 343 maximum degree, and MLP is a multi-layer perceptron that maps τ to a vector of length c .

344 We finally remark that the network f defined in this section is $SO(3)$ -equivariant because each layer
 345 is $SO(3)$ -equivariant by construction. f is also permutation-equivariant because it does not use any
 346 order information of X_i .

349 6 EXPERIMENT

351 This section evaluates Eda on practical assembly tasks. After introducing the experiment settings in
 352 Sec. 6.1, we first evaluate Eda on the pair-wise registration tasks in Sec. 6.2, and then we consider the
 353 multi-piece assembly tasks in Sec. 6.3. An ablation study is finally presented in Sec. 6.4.

355 6.1 EXPERIMENT SETTINGS

357 We evaluate the accuracy of an assembly solution using the averaged pair-wise error. For a predicted
 358 assembly g and the ground truth \hat{g} , the rotation error Δr and the translation error Δt are computed as:
 359 $(\Delta r, \Delta t) = \frac{1}{N(N-1)} \sum_{i \neq j} \tilde{\Delta}(\hat{g}_i, \hat{g}_j g_j^{-1} g_i)$, where the pair-wise error $\tilde{\Delta}$ is computed as $\tilde{\Delta}(g, \hat{g}) =$
 360 $(\frac{180}{\pi} \text{acos}(\frac{1}{2} (\text{tr}(r\hat{r}^T) - 1)), \|\hat{t} - t\|)$. Here $g = (r, t)$, $\hat{g} = (\hat{r}, \hat{t})$, and $\text{tr}(\cdot)$ represents the trace.
 361 This metric is the pair-wise rotation/translation error: it measures the averaged error of g_i w.r.t. g_j
 362 for all (i, j) pairs of pieces.

363 For Eda, we use 2 Croco blocks, and 4 downsampling layers with a downsampling ratio 0.25. We
 364 use $k = 10$ nearest neighbors, $l_{max} = 2$ degree features with $d = 64$ channels and 4 attention
 365 heads. Following Peebles and Xie (2023), we keep an exponential moving average (EMA) with a
 366 decay of 0.99, and we use the AdamW (Loshchilov and Hutter, 2017) optimizer with a learning rate
 367 10^{-4} . Following Esser et al. (2024), we use a logit-normal sampling for time variable τ . For each
 368 experiment, we train Eda on 3 Nvidia A100 GPUs for at most 5 days. We denote Eda with q steps of
 369 RKp as “Eda (RKp, q)”, e.g., Eda (RK1, 10) represents Eda with 10 steps of RK1.

371 6.2 PAIR-WISE REGISTRATION

373 This section evaluates Eda on rotated
 374 3DMatch (Zeng et al., 2017) (3DM) dataset
 375 containing PC pairs from indoor scenes.
 376 Following Huang et al. (2021), we consider the
 377 3DLoMatch split (3DL), which contains PC
 378 pairs with smaller overlap ratios. Furthermore, to highlight the ability of Eda on non-overlapped

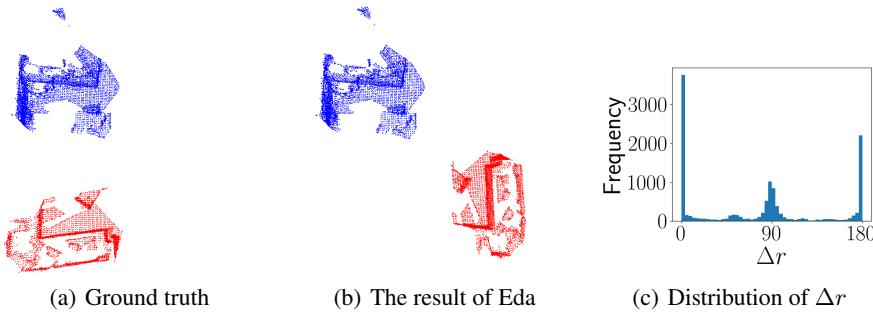
Table 1: The overlap ratio of PC pairs (%).

	3DM	3DL	3DZ
Training set	(10, 100)	0	
Test set	(30, 100)	(10, 30)	0

378 assembly tasks, we consider a new split called 3DZeroMatch (3DZ), which contains non-overlapped
 379 PC pairs. The comparison of these three splits is shown in Tab. 1.
 380

381 We compare Eda against the following
 382 baseline methods: FGR (Zhou et al., 2016),
 383 GEO (Qin et al., 2022), ROI (Yu et al.,
 384 2023), and AMR (Chen et al., 2025),
 385 where FGR is a classic optimization-based
 386 method, GEO and ROI are correspondence-
 387 based methods, and AMR is a recently pro-
 388 posed diffusion-like method based on GEO.
 389 We report the results of the baseline meth-
 390 ods using their official implementations.
 391 Note that the correspondence-free methods
 392 like Ryu et al. (2024); Wang and Jörnsten (2024) do not scale to this dataset.

393 We report the results in Tab 2. On 3DM and 3DL, we observe that Eda outperforms the baseline
 394 methods by a large margin, especially for rotation errors, where Eda achieves more than 50% lower
 395 rotation errors on both 3DL and 3DM. We provide more details of Eda on 3DL in Fig. 5 in the
 396 appendix.



407 Figure 2: More details of Eda on 3DZ. (b): A result of Eda. Cameras are set to look at the room
 408 from above. Two PC pieces are marked by different colors. (c): the distribution of Δr on the test set.
 409

410 As for 3DZ, we only report the results of Eda in Tab 2, because all baseline methods are not applicable
 411 to 3DZ, *i.e.*, their training goal is undefined when the correspondence does not exist. We observe
 412 that Eda’s error on 3DZ is much larger compared to that on 3DL, suggesting that there exists much
 413 larger ambiguity. Nevertheless, as shown in in Fig. 2(b), Eda indeed learned the global geometry of
 414 the indoor scenes instead of just random guessing, because it tends to place large planes, *i.e.*, walls,
 415 floors and ceilings, in a parallel or orthogonal position, and keep a plausible distance between walls
 416 of the assembled room.

417 To show that this behavior is consistent in the whole test set, we present the distribution of Δr of
 418 Eda on 3DZ in Fig. 2(c). A simple intuition is that for rooms consisting of 6 parallel or orthogonal
 419 planes (four walls, a floor and a ceiling), if the orthogonality or parallelism of planes is correctly
 420 maintained in the assembly, then Δr should be 0, 90, or 180. We observe that this is indeed the
 421 case in Fig. 2(c), where Δr is centered at 0, 90, and 180. We remark that the ability to learn global
 422 geometric properties beyond correspondences is a key advantage of Eda, and it partially explains the
 423 superior performance of Eda in Tab. 2

424 6.3 MULTI-PIECE ASSEMBLY

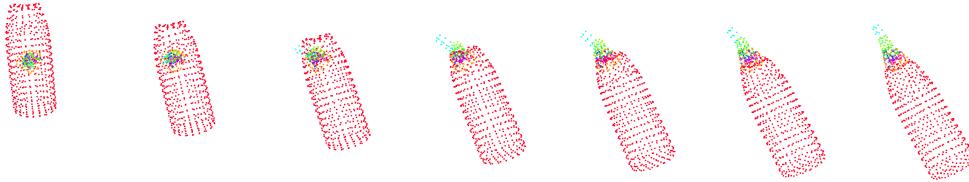
426 This section evaluates Eda on the volume constrained version of BB dataset (Sellán et al., 2022). We
 427 consider the shapes with $2 \leq N \leq 8$ pieces in the “everyday” subset. We compare Eda against the
 428 following baseline methods: DGL (Zhan et al., 2020), LEV (Wu et al., 2023a), GLO (Sellán et al.,
 429 2022), JIG (Lu et al., 2023) and GARF (Li et al., 2025). JIG is correspondence-based, GARF is
 430 diffusion-based, and other baseline methods are regression-based. For Eda, we process all fragments
 431 by grid downsampling with a grid size 0.02. For the baseline methods, we follow their original
 432 preprocessing steps. We do not pretrain GARF for fair comparison,. To reproduce the results of the

Table 2: Quantitative results on rotated 3DMatch. ROI
 (n): ROI with n RANSAC samples.

	3DM		3DL		3DZ	
	Δr	Δt	Δr	Δt	Δr	Δt
FGR	69.5	0.6	117.3	1.3	—	—
GEO	7.43	0.19	28.38	0.69	—	—
ROI (500)	5.64	0.15	21.94	0.53	—	—
ROI (5000)	5.44	0.15	22.17	0.53	—	—
AMR	5.0	0.13	20.5	0.53	—	—
Eda (RK4, 50)	2.38	0.17	8.57	0.4	78.32	2.74

432 baseline methods, we use the implementation of DGL and GLO in the official benchmark suite of
 433 BB, and we use the official implementation of LEV, JIG and GASF.
 434

435 The results are shown in Tab. 3, where we also
 436 report the computation time of all methods on
 437 the test set on a Nvidia T4 GPU except GASF,
 438 which is measured on a A40 GPU because it
 439 does not support the T4 GPU. We observe that
 440 Eda outperforms all baseline methods by a large
 441 margin at a moderate computation cost. We
 442 present some qualitative results in Fig. 7 in the
 443 appendix, where we observe that Eda can gen-
 444 erally reconstruct the shapes more accurately
 445 than the baseline methods. An example of the
 446 assembly process of Eda is presented in Fig. 3.
 447



452 Figure 3: From left to right: the assembly process of a 8-piece bottle by Eda.
 453

454 6.4 ABLATION STUDIES

455 We first investigate the influence of the
 456 number of pieces on the performance
 457 of Eda. We use the kitti odometry
 458 dataset (Geiger et al., 2012) containing PCs
 459 of city road views. For each sequence of
 460 data, we keep pieces that are at least 100
 461 meters apart so that they do not necessarily
 462 overlap, and we downsample them using
 463 grid downsampling with a grid size 0.5.
 464 We train Eda on all consecutive pieces of
 465 length $2 \sim N_{max}$ in sequences $0 \sim 8$. We
 466 call the trained model Eda- N_{max} . We then evaluate Eda- N_{max} on all consecutive pieces of length M
 467 in sequence $9 \sim 10$.

468 The results are shown in Fig. 4. We observe that for Δr , when the length of the test data is seen in
 469 the training set, i.e., $M \leq N_{max}$, Eda performs well, and $M > N_{max}$ leads to worse performance. In
 470 addition, Eda-4 generalizes better than Eda-3 on data of unseen length (5 and 6). The result indicates
 471 the necessity of using training data whose lengths subsume that of the test data. Meanwhile, the
 472 translation errors of Eda-4 and Eda-3 are comparable, and they both increase with the length of data.

473 Then we investigate the influence of the components in our theory.
 474 We compare Eda with Eda- O on the 3DL dataset, where O is a
 475 combination of the following modifications: 1) r : removing r^*
 476 in h_X (7). 2) h : replacing h_X (7) by the canonical path \bar{h} . 3) e :
 477 replacing f by a non-equivariant network. The results are shown in
 478 Tab. 4, where we observe that r leads to a small performance drop,
 479 while h and e lead to large performance drops. In addition, Eda-
 480 (r, h, e) fails to converge. More details can be found in Appx. H.
 481

482 7 CONCLUSION

483 This work studied the theory of equivariant flow matching, and presented a multi-piece assembly
 484 method, called Eda, based on the theory. We show that Eda can accurately assemble PCs on practical
 485 datasets. More discussions can be found in Appx. I.

Table 3: Quantitative results on BB dataset and the total computation time on the test set.

	Δr	Δt	Time (min)
GLO	126.3	0.3	0.9
DGL	125.8	0.3	0.9
LEV	125.9	0.3	8.1
JIG	106.5	0.24	122.2
GASF	95.6	0.2	(48)
Eda (RK1, 10)	80.64	0.16	19.4
Eda (RK4, 10)	79.2	0.16	76.9

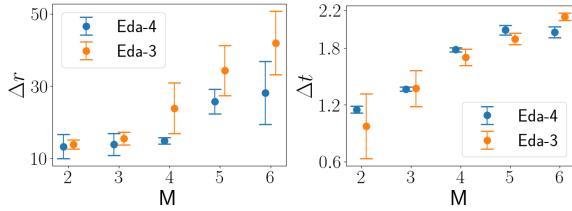


Figure 4: The results of Eda on different number of pieces.

Table 4: Ablation study.

	Δr	Δt
Eda	13.3	0.2
Eda-(r)	15.4	0.23
Eda-(r, h)	79.4	0.51
Eda-(r, e)	86.2	0.37
Eda-(r, h, e)	—	—

486 REFERENCES
487

488 Federica Arrigoni, Beatrice Rossi, and Andrea Fusiello. Spectral synchronization of multiple views
489 in se (3). *SIAM Journal on Imaging Sciences*, 9(4):1963–1990, 2016.

490 K Somani Arun, Thomas S Huang, and Steven D Blostein. Least-squares fitting of two 3-d point sets.
491 *IEEE Transactions on pattern analysis and machine intelligence*, (5):698–700, 1987.

492

493 Gabriele Cesa, Leon Lang, and Maurice Weiler. A program to build e (n)-equivariant steerable cnns.
494 In *International Conference on Learning Representations*, 2022.

495 Ricky TQ Chen and Yaron Lipman. Flow matching on general geometries. In *The Twelfth Interna-
496 tional Conference on Learning Representations*, 2024.

497

498 Yun-Chun Chen, Haoda Li, Dylan Turpin, Alec Jacobson, and Animesh Garg. Neural shape mating:
499 Self-supervised object assembly with adversarial shape priors. In *Proceedings of the IEEE/CVF
500 Conference on Computer Vision and Pattern Recognition*, pages 12724–12733, 2022.

501 Zhi Chen, Yufan Ren, Tong Zhang, Zheng Dang, Wenbing Tao, Sabine Susstrunk, and Mathieu
502 Salzmann. Adaptive multi-step refinement network for robust point cloud registration. *Transactions
503 on Machine Learning Research*, 2025. ISSN 2835-8856. URL [https://openreview.net/
504 forum?id=M3SkSMfWcP](https://openreview.net/forum?id=M3SkSMfWcP).

505 Peter E Crouch and R Grossman. Numerical integration of ordinary differential equations on
506 manifolds. *Journal of Nonlinear Science*, 3:1–33, 1993.

507

508 Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. FlashAttention: Fast and
509 memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing
510 Systems (NeurIPS)*, 2022.

511 Congyue Deng, Or Litany, Yueqi Duan, Adrien Poulenard, Andrea Tagliasacchi, and Leonidas J
512 Guibas. Vector neurons: A general framework for so (3)-equivariant networks. In *Proceedings of
513 the IEEE/CVF International Conference on Computer Vision*, pages 12200–12209, 2021.

514

515 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
516 bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of
517 the North American chapter of the association for computational linguistics: human language
518 technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019.

519

520 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
521 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for
522 high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.

523

524 Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se(3)-transformers: 3d roto-
525 translation equivariant attention networks. *Advances in neural information processing systems*, 33:
526 1970–1981, 2020.

527

528 Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti
529 vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*,
530 pages 3354–3361. IEEE, 2012.

531

532 Mario Geiger and Tess Smidt. e3nn: Euclidean neural networks. *arXiv preprint arXiv:2207.09453*,
533 2022.

534

535 Zan Gojcic, Caifa Zhou, Jan D Wegner, Leonidas J Guibas, and Tolga Birdal. Learning multiview 3d
536 point cloud registration. In *International conference on computer vision and pattern recognition
(CVPR)*, 2020.

537

538 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint
539 arXiv:1606.08415*, 2016.

540

541 Shengyu Huang, Zan Gojcic, Mikhail Usvyatsov, Andreas Wieser, and Konrad Schindler. Predator:
542 Registration of 3d point clouds with low overlap. In *Proceedings of the IEEE/CVF Conference on
543 computer vision and pattern recognition*, pages 4267–4276, 2021.

540 Haobo Jiang, Mathieu Salzmann, Zheng Dang, Jin Xie, and Jian Yang. Se (3) diffusion model-based
 541 point cloud registration for robust 6d object pose estimation. *Advances in Neural Information
 542 Processing Systems*, 36:21285–21297, 2023.

543 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott
 544 Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models.
 545 *arXiv preprint arXiv:2001.08361*, 2020.

546 Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. *Advances in Neural
 547 Information Processing Systems*, 36:59886–59910, 2023.

548 Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative learning for
 549 symmetric densities. In *International conference on machine learning*, pages 5361–5370. PMLR,
 550 2020.

551 Seong Hun Lee and Javier Civera. Hara: A hierarchical approach for robust rotation averaging. In
 552 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages
 553 15777–15786, 2022.

554 Sihang Li, Zeyu Jiang, Grace Chen, Chenyang Xu, Siqi Tan, Xue Wang, Irving Fang, Kristof
 555 Zyskowski, Shannon P McPherron, Radu Iovita, Chen Feng, and Jing Zhang. Garf: Learning
 556 generalizable 3d reassembly for real-world fractures. In *International Conference on Computer
 557 Vision (ICCV)*, 2025.

558 Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3d atomistic
 559 graphs. In *The Eleventh International Conference on Learning Representations*, 2023. URL
 560 <https://openreview.net/forum?id=KwmPfARgOTD>.

561 Yi-Lun Liao, Brandon Wood, Abhishek Das, and Tess Smidt. Equiformerv2: Improved equivariant
 562 transformer for scaling to higher-degree representations. *arXiv preprint arXiv:2306.12059*, 2023.

563 Yi-Lun Liao, Brandon M Wood, Abhishek Das, and Tess Smidt. Equiformerv2: Improved equivariant
 564 transformer for scaling to higher-degree representations. In *The Twelfth International Conference
 565 on Learning Representations*, 2024. URL <https://openreview.net/forum?id=mCOBKZmrzD>.

566 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
 567 matching for generative modeling. In *The Eleventh International Conference on Learning Representations*,
 568 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.

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

571 Jiaxin Lu, Yifan Sun, and Qixing Huang. Jigsaw: Learning to assemble multiple fractured objects.
 572 *Advances in Neural Information Processing Systems*, 36:14969–14986, 2023.

573 Saro Passaro and C Lawrence Zitnick. Reducing so (3) convolutions to so (2) for efficient equivariant
 574 gnns. In *International Conference on Machine Learning*, pages 27420–27438. PMLR, 2023.

575 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of
 576 the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023.

577 Zheng Qin, Hao Yu, Changjian Wang, Yulan Guo, Yuxing Peng, and Kai Xu. Geometric transformer
 578 for fast and robust point cloud registration. In *Proceedings of the IEEE/CVF conference on
 579 computer vision and pattern recognition*, pages 11143–11152, 2022.

580 Hyunwoo Ryu, Hong-in Lee, Jeong-Hoon Lee, and Jongeun Choi. Equivariant descriptor fields: Se
 581 (3)-equivariant energy-based models for end-to-end visual robotic manipulation learning. *arXiv
 582 preprint arXiv:2206.08321*, 2022.

583 Hyunwoo Ryu, Hong in Lee, Jeong-Hoon Lee, and Jongeun Choi. Equivariant descriptor fields:
 584 Se(3)-equivariant energy-based models for end-to-end visual robotic manipulation learning. In
 585 *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=dnjZSPGmY50>.

594 Hyunwoo Ryu, Jiwoo Kim, Hyunseok An, Junwoo Chang, Joohwan Seo, Taehan Kim, Yubin Kim,
 595 Chaewon Hwang, Jongeun Choi, and Roberto Horowitz. Diffusion-edfs: Bi-equivariant denoising
 596 generative modeling on $se(3)$ for visual robotic manipulation. In *Proceedings of the IEEE/CVF*
 597 *Conference on Computer Vision and Pattern Recognition*, pages 18007–18018, 2024.

598 Gianluca Scarpellini, Stefano Fiorini, Francesco Giuliari, Pietro Moreiro, and Alessio Del Bue.
 599 Difassemble: A unified graph-diffusion model for 2d and 3d reassembly. In *Proceedings of the*
 600 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28098–28108, 2024.

601 Silvia Sellán, Yun-Chun Chen, Ziyi Wu, Animesh Garg, and Alec Jacobson. Breaking bad: A dataset
 602 for geometric fracture and reassembly. *Advances in Neural Information Processing Systems*, 35:
 603 38885–38898, 2022.

604 Anthony Simeonov, Yilun Du, Andrea Tagliasacchi, Joshua B Tenenbaum, Alberto Rodriguez, Pulkit
 605 Agrawal, and Vincent Sitzmann. Neural descriptor fields: $Se(3)$ -equivariant object representations
 606 for manipulation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages
 607 6394–6400. IEEE, 2022.

608 Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou,
 609 and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport for 3d molecule
 610 generation. *Advances in Neural Information Processing Systems*, 36:549–568, 2023.

611 Loring W Tu. Manifolds. In *An Introduction to Manifolds*, pages 47–83. Springer, 2011.

612 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 613 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing*
 614 *systems*, 30, 2017.

615 Haiping Wang, Yufu Zang, Fuxun Liang, Zhen Dong, Hongchao Fan, and Bisheng Yang. A
 616 probabilistic method for fractured cultural relics automatic reassembly. *Journal on Computing and*
 617 *Cultural Heritage (JOCCH)*, 14(1):1–25, 2021.

618 Ziming Wang and Rebecka Jörnsten. $Se(3)$ -bi-equivariant transformers for point cloud assembly. In
 619 *The Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

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

623 Philippe Weinzaepfel, Vincent Leroy, Thomas Lucas, Romain Brégier, Yohann Cabon, Vaibhav Arora,
 624 Leonid Antsfeld, Boris Chidlovskii, Gabriela Csurka, and Jérôme Revaud. Croco: Self-supervised
 625 pre-training for 3d vision tasks by cross-view completion. *Advances in Neural Information*
 626 *Processing Systems*, 35:3502–3516, 2022.

627 Ruihai Wu, Chenrui Tie, Yushi Du, Yan Zhao, and Hao Dong. Leveraging $se(3)$ equivariance for
 628 learning 3d geometric shape assembly. In *Proceedings of the IEEE/CVF International Conference*
 629 *on Computer Vision*, pages 14311–14320, 2023a.

630 Yue Wu, Yongzhe Yuan, Xiaolong Fan, Xiaoshui Huang, Maoguo Gong, and Qiguang Miao. Pcrdiff-
 631 fusion: Diffusion probabilistic models for point cloud registration. *CoRR*, 2023b.

632 Qun-Ce Xu, Hao-Xiang Chen, Jiacheng Hua, Xiaohua Zhan, Yong-Liang Yang, and Tai-Jiang
 633 Mu. Fragmentdiff: A diffusion model for fractured object assembly. In *SIGGRAPH Asia 2024*
 634 *Conference Papers*, pages 1–12, 2024.

635 Jason Yim, Andrew Campbell, Andrew YK Foong, Michael Gastegger, José Jiménez-Luna, Sarah
 636 Lewis, Victor Garcia Satorras, Bastiaan S Veeling, Regina Barzilay, Tommi Jaakkola, et al. Fast
 637 protein backbone generation with $se(3)$ flow matching. *arXiv preprint arXiv:2310.05297*, 2023.

638 Hao Yu, Zheng Qin, Ji Hou, Mahdi Saleh, Dongsheng Li, Benjamin Busam, and Slobodan Ilic.
 639 Rotation-invariant transformer for point cloud matching. In *Proceedings of the IEEE/CVF confer-
 640 ence on computer vision and pattern recognition*, pages 5384–5393, 2023.

648 Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas
649 Funkhouser. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In
650 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1802–1811,
651 2017.

652 Guanqi Zhan, Qingnan Fan, Kaichun Mo, Lin Shao, Baoquan Chen, Leonidas J Guibas, Hao Dong,
653 et al. Generative 3d part assembly via dynamic graph learning. *Advances in Neural Information
654 Processing Systems*, 33:6315–6326, 2020.

655 Biao Zhang and Rico Sennrich. Root mean square layer normalization. *Advances in Neural
656 Information Processing Systems*, 32, 2019.

657 Zhengyou Zhang. Iterative point matching for registration of free-form curves and surfaces. *Inter-
658 tional Journal of Computer Vision*, 13(2):119–152, 1994.

659 Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In *Computer Vision–ECCV
660 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings,
661 Part II 14*, pages 766–782. Springer, 2016.

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 A THE USE OF LARGE LANGUAGE MODELS (LLM)
703704 We use an LLM to correct grammar errors.
705706 707 B MORE DETAILS OF THE RELATED TASKS
708709 The registration task aims to reconstruct the scene from multiple overlapped views. A registration
710 method generally consists of two stages: first, each pair of pieces is aligned using a pair-wise
711 method (Qin et al., 2022), then all pieces are merged into a complete shape using a synchronization
712 method (Arrigoni et al., 2016; Lee and Civera, 2022; Gojcic et al., 2020). In contrast to other tasks,
713 the registration task generally assumes that the pieces are overlapped. In other words, it assumes that
714 some points observed in one piece are also observed in the other piece, and the goal is to match the
715 points observed in both pieces, *i.e.*, corresponding points. The state-of-the-art registration methods
716 usually infer the correspondences based on the feature similarity (Yu et al., 2023) learned by neural
717 networks, and then align them using the SVD projection (Arun et al., 1987) or RANSAC.718 The robotic manipulation task aims to move one PC to a certain position relative to another PC. For
719 example, one PC can be a cup, and the other PC can be a table, and the goal is to move the cup onto
720 the table. Since the input PCs are sampled from different objects, they are generally non-overlapped.
721 Unlike the other two tasks, this task is generally formulated in a probabilistic setting, as the solution
722 is generally not unique. Various probabilistic models, such as energy-based models (Simeonov et al.,
723 2022; Ryu et al., 2023), or diffusion models (Ryu et al., 2024), have been used for this task.724 The reassembly task aims to reconstruct the complete object from multiple fragment pieces. This
725 task is similar to the registration task, except that the input PCs are sampled from different fragments,
726 thus they are not necessarily overlapped, *e.g.*, due to missing pieces or the erosion of the surfaces.
727 Most of the existing methods are based on regression, where the solution is directly predicted from
728 the input PCs (Wu et al., 2023a; Chen et al., 2022; Wang and Jörnsten, 2024). Some probabilistic
729 methods, such as diffusion-based methods (Xu et al., 2024; Scarpellini et al., 2024), have also been
730 proposed. Note that there exist some exceptions (Lu et al., 2023) which assume the overlap of the
731 pieces, and they rely on the inferred correspondences as the registration methods.732 A comparison of these three tasks is presented in Tab. 5.
733

734 735 Table 5: Comparison between registration, reassembly and manipulation tasks.

Task	Number of pieces	Probabilistic/Deterministic	Overlap
Registration	≥ 2	Deterministic	Overlapped
Reassembly	≥ 2	Deterministic	Non-overlapped
Manipulation	2	Probabilistic	Non-overlapped
Assembly (this work)	≥ 2	Probabilistic	Non-overlapped

741 742 743 C A WALK-THROUGH OF THE MAIN THEORY
744745 This section provides a walk-through of the theory using the two-piece deterministic example. We
746 follow the notation in example 3.2: let (r_1, r_2) be the solution for the input point clouds (X_1, X_2) ,
747 meaning $r_1 X_1$ and $r_2 X_2$ are assembled.748 Our theory addresses the following equivariance question. Assume that a diffusion model works
749 for the input (X_1, X_2) , *i.e.*, the predicted vector field $v_{(X_1, X_2)}$ flows to the correct solution (r_1, r_2) .750 **How to ensure it also works for the perturbed input?** For example, for $SO(3)^2$ -equivariance, the
751 question is how to ensure the model also works for $(r_3 X_1, r_4 X_2)$. *i.e.*, to ensure the predicted vector
752 field $v_{(r_3 X_1, r_4 X_2)}$ flows to $(r_1 r_3^{-1}, r_2 r_4^{-1})$.753 754 **Corollary 4.4** shows that the goal can be achieved if $v_{(r_3 X, r_4 X_2)}$ is a proper "transformation" of
755 $v_{(X_1, X_2)}$ (relatedness), and the noise is invariant.

756 Then, the next question is how to satisfy the relatedness requirement. **Proposition 4.5** suggests that
 757 this can be simply done by parametrizing the vector fields as
 758

759 $v_{(X_1, X_2)}(r_7, r_8) = f(r_7 X_1, r_8 X_2)(r_7 \oplus r_8)$, where $f(X_1, X_2) = (w_1, t_1) \oplus (w_2, t_2)$ (8)
 760 is a neural network mapping (X_1, X_2) to their respective rotation/translation velocity components
 761 w and t , and \oplus is the concatenation. In summary, we can now answer the question from the last
 762 paragraph: if the diffusion model predicts the vector field as in (8) and it works for (X_1, X_2) , then it
 763 also works for $(r_3 X_1, r_4 X_2)$.

764 Further more, **Proposition 4.6** suggests that, to ensure the other two requirements (permutation
 765 equivariance and $SO(3)$ -invariance) of the model, f needs to satisfy

$$766 \quad f(X_2, X_1) = (w_2, t_2) \oplus (w_1, t_1) \quad \text{and} \quad f(rX_1, rX_2) = (rw_1, rt_1) \oplus (rw_2, rt_2) \quad (9)$$

767 Finally, **Proposition 4.7** suggests that some data augmentations are not needed when all the above
 768 requirements are satisfied. For example, for data (X_1, X_2) we learn a vector field $v_{(X_1, X_2)}$. We can
 769 use randomly augmented data $(r_3 X_1, r_4 X_2)$ and learn $v_{(r_3 X_1, r_4 X_2)}$. However, this is not necessary
 770 because $v_{(r_3 X_1, r_4 X_2)}$ is already guaranteed to be a transformation of $v_{(X_1, X_2)}$ as described above,
 771 and the loss for them is the same, *i.e.*, learning $v_{(X_1, X_2)}$ alone is enough. Similar results hold for the
 772 other two types of augmentations.

774 D CONNECTIONS WITH BI-EQUIVARIANCE

775 This section briefly discusses the connections between Def. 3.1 and the equivariances defined in Ryu
 776 et al. (2024) and Wang and Jörnsten (2024) in pair-wise assembly tasks.

777 We first recall the definition of the probabilistic bi-equivariance.

778 **Definition D.1** (Eqn. (10) in Ryu et al. (2024) and Def. (1) in Ryu et al. (2022)). $\hat{P} \in \mu(SE(3))$ is
 779 bi-equivariant if for all $g_1, g_2 \in SO(3)$, PCs X_1, X_2 , and a measurable set $A \subseteq SE(3)$,

$$780 \quad \hat{P}(A|X_1, X_2) = \hat{P}(g_2 A g_1^{-1} | g_1 X_1, g_2 X_2). \quad (10)$$

781 Note that we only consider $g_1, g_2 \in SO(3)$ instead of $g_1, g_2 \in SE(3)$ because we require all input
 782 PCs, *i.e.*, $X_i, g_i X_i, i = 1, 2$, to be centered.

783 Then we recall Def. 3.1 for pair-wise assembly tasks:

784 **Definition D.2** (Restate $SO(3)^2$ -equivariance and $SO(3)$ -invariance in Def. 3.1 for pair-wise prob-
 785 lems). Let X_1, X_2 be the input PCs and $P \in \mu(SE(3) \times SE(3))$.

- 786 • P is $SO(3)^2$ -equivariant if $P(A|X_1, X_2) = P(A(g_1^{-1}, g_2^{-1})|g_1 X_1, g_2 X_2)$ for all $g_1, g_2 \in SO(3)$
 787 and $A \subseteq SO(3) \times SO(3)$, where $A(g_1^{-1}, g_2^{-1}) = \{(a_1 g_1^{-1}, a_2 g_2^{-1}) : (a_1, a_2) \in A\}$.
- 788 • P is $SO(3)$ -invariant if $P(A|X_1, X_2) = P(rA|X_1, X_2)$ for all $r \in SO(3)$ and $A \subseteq SO(3) \times$
 789 $SO(3)$.

790 Intuitively, both Def. D.1 and Def. D.2 describe the equivariance property of an assembly solution, and
 791 the only difference is that Def. D.1 describes the special case where X_1 can be rigidly transformed and
 792 X_2 is fixed, while Def. D.2 describes the solution where both X_1 and X_2 can be rigidly transformed.
 793 In other words, a solution satisfying Def. D.2 can be converted to a solution satisfying Def. D.1 by
 794 fixing X_2 . Formally, we have the following proposition.

795 **Proposition D.3.** *Let P be $SO(3)^2$ -equivariant and $SO(3)$ -invariant. If $\tilde{P}(A|X_1, X_2) \triangleq P(A \times$
 796 $\{e\}|X_1, X_2)$ for $A \subseteq SO(3)$, then \tilde{P} is bi-equivariant.*

797 *Proof.* We prove this proposition by directly verifying the definition.

$$800 \quad \tilde{P}(g_2 A g_1^{-1} | g_1 X_1, g_2 X_2) = P(g_2 A g_1^{-1} \times \{e\} | g_1 X_1, g_2 X_2) \quad (11)$$

$$801 \quad = P(g_2 A \times \{e\} | X_1, g_2 X_2) \quad (12)$$

$$802 \quad = P(A \times \{g_2^{-1}\} | X_1, g_2 X_2) \quad (13)$$

$$803 \quad = P(A \times \{e\} | X_1, X_2) \quad (14)$$

$$804 \quad = \tilde{P}(A|X_1, X_2). \quad (15)$$

810 Here, the second and the fourth equation hold because P is $SO(3)^2$ -equivariant, the third equation
 811 holds because P is $SO(3)$ -invariant, and the first and last equation are due to the definition. \square
 812

813 We note that the deterministic definition of bi-equivariance in Wang and Jörnsten (2024) is a special
 814 case of Def. D.1, where \hat{P} is a Dirac delta function. In addition, as discussed in Appx. E in Wang and
 815 Jörnsten (2024), a major limitation of the deterministic definition of bi-equivariance is that it cannot
 816 handle symmetric shapes. In contrast, it is straightforward to see that the probabilistic definition, *i.e.*,
 817 both Def. D.1 and Def. D.2 are free from this issue. Here, we consider the example in Wang and
 818 Jörnsten (2024). Assume that X_1 is symmetric, *i.e.*, there exists $g_1 \in SO(3)$ such that $g_1 X_1 = X_1$.
 819 Under Def. D.1, we have $P(A|X_1, X_2) = P(A|g_1 X_1, X_2) = P(Ag_1|X_1, X_2)$, which simply means
 820 that $P(A|X_1, X_2)$ is \mathcal{R}_{g_1} -invariant. Note that this will not cause any contradiction, *i.e.*, the feasible
 821 set is not empty. For example, a uniform distribution on $SO(3)$ is \mathcal{R}_{g_1} -invariant.

822 As for the permutation-equivariance, the swap-equivariance in Wang and Jörnsten (2024) is a
 823 deterministic pair-wise version of the permutation-equivariance in Def. D.2, and they both mean that
 824 the assembled shape is independent of the order of the input pieces.

826 E THE RK4 FORMULATION

$$829 \quad k_1 = f_{\tau_i}(\mathbf{g}_i), \quad k_2 = f_{\tau_i + \frac{1}{2}\eta}(\exp(\frac{1}{2}\eta k_1)\mathbf{g}_i), \quad k_3 = f_{\tau_i + \frac{1}{2}\eta}(\exp(\frac{1}{2}\eta k_2)\mathbf{g}_i), \quad k_4 = f_{\tau_i + \eta}(\exp(\eta k_3)\mathbf{g}_i), \\ 830 \\ 831 \quad \mathbf{g}_{i+1} = \exp(\frac{1}{6}\eta k_4) \exp(\frac{1}{3}\eta k_3) \exp(\frac{1}{3}\eta k_2) \exp(\frac{1}{6}\eta k_1)\mathbf{g}_i. \quad (16)$$

833 Note that RK4 (16) is more computationally expensive than RK1, because it requires four evaluations
 834 of v_X at different points at each step, *i.e.*, four forward passes of network f , while the Euler method
 835 only requires one evaluation per step.

837 F PROOFS

839 F.1 PROOF IN SEC. 4.2

841 To prove Thm. 4.2, which established the relations between related vector fields and equivariant
 842 distributions, we proceed in two steps: first, we prove lemma F.1, which connects related vector
 843 fields to equivariant mappings; then we prove lemma F.2, which connects equivariant mappings to
 844 equivariant distributions.

845 **Lemma F.1.** *Let G be a smooth manifold, $F : G \rightarrow G$ be a diffeomorphism. If vector field v_τ is
 846 F -related to vector field w_τ for $\tau \in [0, 1]$, then $F \circ \phi_\tau = \psi_\tau \circ F$, where ϕ_τ and ψ_τ are generated by
 847 v_τ and w_τ respectively.*

849 *Proof.* Let $\tilde{\psi}_\tau \triangleq F \circ \phi_\tau \circ F^{-1}$. We only need to show that $\tilde{\psi}_\tau$ coincides with ψ_τ .

850 We consider a curve $\tilde{\psi}_\tau(F(\mathbf{g}_0))$, $\tau \in [0, 1]$, for a arbitrary $\mathbf{g}_0 \in G$. We first verify that $\tilde{\psi}_0(F(\mathbf{g}_0)) =$
 851 $F \circ \phi_0 \circ F^{-1} \circ F(\mathbf{g}_0) = F(\mathbf{g}_0)$. Note that the second equation holds because $\phi_0(\mathbf{g}_0) = \mathbf{g}_0$, *i.e.*, ϕ_τ
 852 is an integral path. Then we verify

$$854 \quad \frac{\partial}{\partial \tau}(\tilde{\psi}_\tau(F(\mathbf{g}_0))) = \frac{\partial}{\partial \tau}(F \circ \phi_\tau(\mathbf{g}_0)) \quad (17)$$

$$855 \quad = F_{*, \phi_\tau(\mathbf{g}_0)} \circ \frac{\partial}{\partial \tau}(\phi_\tau(\mathbf{g}_0)) \quad (18)$$

$$856 \quad = F_{*, \phi_\tau(\mathbf{g}_0)} \circ v_\tau(\phi_\tau(\mathbf{g}_0)) \quad (19)$$

$$857 \quad = w_\tau(F \circ \phi_\tau(\mathbf{g}_0)) \quad (20)$$

$$858 \quad = w_\tau(\tilde{\psi}_\tau(F(\mathbf{g}_0))) \quad (21)$$

859 where the 2-nd equation holds due to the chain rule, and the 4-th equation holds becomes v_τ is
 860 F -related to w_τ . Therefore, we can conclude that $\tilde{\psi}_\tau(F(\mathbf{g}_0))$ is an integral curve generated by w_τ

starting from $F(\mathbf{g}_0)$. However, by definition of ψ_τ , $\psi_\tau(F(\mathbf{g}_0))$ is also the integral curve generated by w_τ and starts from $F(\mathbf{g}_0)$. Due to the uniqueness of integral curves, we have $\tilde{\psi}_\tau = \psi_\tau$. \square

Lemma F.2. *Let $\phi, \psi, F : G \rightarrow G$ be three diffeomorphisms satisfying $F \circ \phi = \psi \circ F$. We have $F_\#(\phi_\# \rho) = \psi_\#(F_\# \rho)$ for all distribution ρ on G .*

Proof. Let $A \subseteq G$ be a measurable set. We first verify that $\phi^{-1}(F^{-1}(A)) = F^{-1}(\psi^{-1}(A))$: If $x \in \phi^{-1}(F^{-1}(A))$, then $(F \circ \phi)(x) \in A$. Since $F \circ \phi = \psi \circ F$, we have $(\psi \circ F)(x) \in A$, which implies $x \in F^{-1}(\psi^{-1}(A))$, i.e., $\phi^{-1}(F^{-1}(A)) \subseteq F^{-1}(\psi^{-1}(A))$. The other side can be verified similarly. Then we have

$$(F_\#(\phi_\# \rho))(A) = \rho(\phi^{-1}(F^{-1}(A))) = \rho(F^{-1}(\psi^{-1}(A))) = (\psi_\#(F_\# \rho))(A), \quad (22)$$

which proves the lemma. \square

Now, we can prove Thm. 4.2 using the above two lemmas.

Proof of Thm. 4.2. Since v_X is F -related to v_Y , according to lemma F.1, we have $F \circ \phi_X = \phi_Y \circ F$. Then according to lemma F.2, we have $F_\#(\phi_X \# P_0) = \phi_Y \#(F_\# P_0)$. The proof is complete by letting $P_X = \phi_X \# P_0$ and $P_Y = \phi_Y \#(F_\# P_0)$. \square

We remark that our theory extends the results in Köhler et al. (2020), where only invariance is considered. Specifically, we have the following corollary.

Corollary F.3 (Thm 2 in Köhler et al. (2020)). *Let G be the Euclidean space, F be a diffeomorphism on G , and v_τ be a F -invariant vector field, i.e., v_τ is F -related to v_τ , then we have $F \circ \phi_\tau = \phi_\tau \circ F$, where ϕ_τ is generated by v_τ .*

Proof. This is a direct consequence of lemma F.1 where G is the Euclidean space and $w_\tau = v_\tau$. \square

Note that the terminology used in Köhler et al. (2020) is different from ours: The F -invariant vector fields in our work is called F -equivariant vector field in Köhler et al. (2020), and Köhler et al. (2020) does not consider general related vector fields.

Finally, we present the proof of Prop. 4.5 and Prop. 4.6.

Proof of Prop. 4.5. If v_X is $\mathcal{R}_{g^{-1}}$ -related to v_{gX} , we have $v_{gX}(\hat{\mathbf{g}}\mathbf{g}^{-1}) = (\mathcal{R}_{g^{-1}})_{*,\hat{\mathbf{g}}}v_X(\hat{\mathbf{g}})$ for all $\hat{\mathbf{g}}, \mathbf{g} \in SE(3)^N$. By letting $\mathbf{g} = \hat{\mathbf{g}}$, we have

$$v_X(\mathbf{g}) = (\mathcal{R}_{\mathbf{g}})_{*,e}v_{gX}(e) \quad (23)$$

where $(\mathcal{R}_{\mathbf{g}})_{*,e} = ((\mathcal{R}_{g^{-1}})_{*,\mathbf{g}})^{-1}$ due to the chain rule of $\mathcal{R}_{\mathbf{g}}\mathcal{R}_{g^{-1}} = e$.

On the other hand, if Eqn. (23) holds, we have

$$(\mathcal{R}_{g^{-1}})_{*,\hat{\mathbf{g}}}v_X(\hat{\mathbf{g}}) = (\mathcal{R}_{g^{-1}})_{*,\hat{\mathbf{g}}}(\mathcal{R}_{\hat{\mathbf{g}}})_{*,e}v_{\hat{\mathbf{g}}X}(e) = (\mathcal{R}_{\hat{\mathbf{g}}\mathbf{g}^{-1}})_{*,e}v_{\hat{\mathbf{g}}X}(e) = v_{gX}(\hat{\mathbf{g}}\mathbf{g}^{-1}), \quad (24)$$

which suggests that v_X is $\mathcal{R}_{g^{-1}}$ -related to v_{gX} . Note that the second equation holds due to the chain rule of $\mathcal{R}_{g^{-1}}\mathcal{R}_{\hat{\mathbf{g}}} = \mathcal{R}_{\hat{\mathbf{g}}\mathbf{g}^{-1}}$, and the first and the third equation are the result of Eqn. (23). \square

Proof of Prop. 4.6. 1) Assume v_X is σ -related to $v_{\sigma X}$: $(\sigma)_{*,g}v_X(g) = V_{\sigma X}(\sigma(g))$. By inserting Eqn. (5) to this equation, we have

$$(\sigma)_{*,g}(\mathcal{R}_{\mathbf{g}})_{*,e}f(\mathbf{g}X) = (\mathcal{R}_{\sigma\mathbf{g}})_{*,e}f(\sigma(\mathbf{g})\sigma(X)). \quad (25)$$

Since $\sigma \circ \mathcal{R}_{\mathbf{g}} = \mathcal{R}_{\sigma\mathbf{g}} \circ \sigma$, by the chain rule, we have $\sigma_*(\mathcal{R}_{\mathbf{g}})_* = (\mathcal{R}_{\sigma\mathbf{g}})_*\sigma_*$. In addition, $\sigma(\mathbf{g})\sigma(X) = \sigma(\mathbf{g}X)$. Thus, this equation can be simplified as

$$(\mathcal{R}_{\sigma\mathbf{g}})_*\sigma_*f(\mathbf{g}X) = (\mathcal{R}_{\sigma\mathbf{g}})_{*,e}f(\sigma(\mathbf{g}X)) \quad (26)$$

which suggests

$$\sigma_*f = f \circ \sigma. \quad (27)$$

The first statement in Prop. 4.6 can be proved by reversing the discussion.

918 2) Assume v_X is \mathcal{L}_r -related to v_X : $(\mathcal{L}_r)_{*,g}v_X(\mathbf{g}) = V_X(r\mathbf{g})$. By inserting Eqn. (5) to this equation,
919 we have

$$(\mathcal{L}_r)_{*,g}(\mathcal{R}_g)_{*,e}f(\mathbf{g}X) = (\mathcal{R}_{r\mathbf{g}})_{*,e}f(r\mathbf{g}X). \quad (28)$$

920 Since $\mathcal{R}_{r\mathbf{g}} = \mathcal{R}_g \circ \mathcal{R}_r$, by the chain rule, we have $(\mathcal{R}_{r\mathbf{g}})_{*,e} = (\mathcal{R}_g)_{*,r}(\mathcal{R}_r)_{*,e}$. In addition,
921 $(\mathcal{L}_r)(\mathcal{R}_g) = (\mathcal{R}_g)(\mathcal{L}_r)$, by the chain rule, we have $(\mathcal{L}_r)_{*,g}(\mathcal{R}_g)_{*,e} = (\mathcal{R}_g)_{*,r}(\mathcal{L}_r)_{*,e}$. Thus the
923 above equation can be simplified as

$$(\mathcal{L}_r)_{*,e}f(\mathbf{g}X) = (\mathcal{R}_r)_{*,e}f(r\mathbf{g}X) \quad (29)$$

924 which implies

$$f \circ r = (\mathcal{R}_{r^{-1}})_{*,r} \circ (\mathcal{L}_r)_{*,e} \circ f. \quad (30)$$

925 By representing f in the matrix form, we have

$$w_X^i(rX) = rw_X^i(X)r^T \quad (31)$$

$$t^i(rX) = rt^i(X) \quad (32)$$

926 for all i , where r on the right hand side represents the matrix form of the rotation r . Here the first
927 equation can be equivalently written as $w^i(rX) = rw^i(X)$. The second statement in Prop. 4.6 can
928 be proved by reversing the discussion. \square

936 F.2 PROOFS IN SEC. 4.3

937 To establish the results in this section, we need to assume the uniqueness of r^* (6):

938 **Assumption F.4.** The solution to (6) is unique.

939 Note that this assumption is mild. A sufficient condition (Wang and Jörnsten, 2024) of assumption F.4
940 is that the singular values of $\tilde{\mathbf{g}}_1^T \mathbf{g}_0 \in \mathbb{R}^{3 \times 3}$ satisfy $\sigma_1 \geq \sigma_2 > \sigma_3 \geq 0$, i.e., σ_2 and σ_3 are not equal.
941 We leave the more general treatment without requiring the uniqueness of r^* to future work.

942 We first justify the definition of $\mathbf{g}_1 = r^* \tilde{\mathbf{g}}_1$ by showing that \mathbf{g}_1 follows P_1 in the following proposition.

943 **Proposition F.5.** Let P_0 and P_1 be two $SO(3)$ -invariant distributions, and $\mathbf{g}_0, \tilde{\mathbf{g}}_1$ be independent
944 samples from P_0 and P_1 respectively. If r^* is given by (6) and assumption F.4 holds, then $\mathbf{g}_1 = r^* \tilde{\mathbf{g}}_1$
945 follows P_1 .

946 *Proof.* Define $A_{\tilde{\mathbf{g}}_1} = \{\mathbf{g}_0 | r^*(\mathbf{g}_0, \tilde{\mathbf{g}}_1) = e\}$, where we write r^* as a function of $\tilde{\mathbf{g}}_1$ and \mathbf{g}_0 . Then we
947 have $P(r^* = e | \tilde{\mathbf{g}}_1) = P_0(A_{\tilde{\mathbf{g}}_1})$ by definition. In addition, due to the uniqueness of the solution to (6),
948 for an arbitrary $\hat{r} \in SO(3)$, we have $P(r^* = \hat{r} | \tilde{\mathbf{g}}_1) = P_0(\hat{r} A_{\tilde{\mathbf{g}}_1})$. Since P_0 is $SO(3)$ -invariant, we
949 have $P_0(\hat{r} A_{\tilde{\mathbf{g}}_1}) = P_0(A_{\tilde{\mathbf{g}}_1})$, thus, $P(r^* = \hat{r} | \tilde{\mathbf{g}}_1) = P(r^* = e | \tilde{\mathbf{g}}_1)$. In other words, for a given $\tilde{\mathbf{g}}_1$, r^*
950 follows the uniform distribution $U_{SO(3)}$.

951 Finally we compute the probability density of \mathbf{g}_1 :

$$P(\mathbf{g}_1) = \int P(r^* = \hat{r}^{-1} | \hat{r} \mathbf{g}_1) P_1(\hat{r} \mathbf{g}_1) d\hat{r} \quad (33)$$

$$= \int U_{SO(3)}(\hat{r}) P_1(\mathbf{g}_1) d\hat{r} \quad (34)$$

$$= P_1(\mathbf{g}_1), \quad (35)$$

952 which suggests that \mathbf{g}_1 follows P_1 . Here the second equation holds because P_1 is $SO(3)$ -invariant.
953 \square

954 Then we discuss the equivariance of the constructed h_X (7).

955 **Proposition F.6.** Given $\mathbf{r} \in SO(3)^N$, $\mathbf{g}_0, \tilde{\mathbf{g}}_1 \in SE(3)^N$, $\sigma \in S_N$, $r \in SO(3)$ and $\tau \in [0, 1]$. Let
956 h_X be a path generated by \mathbf{g}_0 and $\tilde{\mathbf{g}}_1$. Under assumption F.4,

- 957 • if $h_{\mathbf{r}X}$ is generated by $\mathbf{g}_0 \mathbf{r}^{-1}$ and $\tilde{\mathbf{g}}_1 \mathbf{r}^{-1}$, then $h_{\mathbf{r}X}(\tau) = \mathcal{R}_{\mathbf{r}^{-1}} h_X(\tau)$.
- 958 • if $h_{\sigma X}$ is generated by $\sigma(\mathbf{g}_0)$ and $\sigma(\tilde{\mathbf{g}}_1)$, then $h_{\sigma X}(\tau) = \sigma(h_X(\tau))$.
- 959 • if \hat{h}_X is generated by $r\mathbf{g}_0$ and $r\tilde{\mathbf{g}}_1$, then $\hat{h}_X(\tau) = \mathcal{L}_r(h_X(\tau))$.

972 *Proof.* 1) Due to the uniqueness of the solution to (6), we have $r^*(\mathbf{g}_0 \mathbf{r}^{-1}, \tilde{\mathbf{g}}_1 \mathbf{r}^{-1}) = r^*(\mathbf{g}_0, \tilde{\mathbf{g}}_1)$.
 973 Thus, we have

$$974 h_{\mathbf{r}X}(\tau) = \exp(\tau \log(\mathbf{g}_1 \mathbf{g}_0^{-1})) \mathbf{g}_0 \mathbf{r}^{-1} = \mathcal{R}_{\mathbf{r}^{-1}}(h_{\mathbf{r}X}(\tau)). \quad (36)$$

976 2) Due to the uniqueness of the solution to (6), we have $r^*(\sigma(\mathbf{g}_0), \sigma(\tilde{\mathbf{g}}_1)) = \sigma(r^*(\mathbf{g}_0, \tilde{\mathbf{g}}_1))$. Thus,
 977 we have $\sigma(h_X) = h_{\sigma X}$.

978 3) Due to the uniqueness of the solution to (6), we have $r^*(r\mathbf{g}_0, r\tilde{\mathbf{g}}_1) = rr^*(\mathbf{g}_0, \tilde{\mathbf{g}}_1)r^{-1}$. Thus,

$$979 \hat{h}_{\mathbf{r}X}(\tau) = \exp(\tau \log(rr^*\tilde{\mathbf{g}}_1\mathbf{g}_0^{-1}r^{-1}))r\mathbf{g}_0 = r \exp(\tau \log(r^*\tilde{\mathbf{g}}_1\mathbf{g}_0^{-1}))\mathbf{g}_0 = \mathcal{L}_r(h_X(\tau)). \quad (37)$$

□

982 With the above preparation, we can finally prove Prop. 4.7.

984 *Proof of Prop. 4.7.* 1) By definition

$$986 L(\mathbf{r}X) = \mathbb{E}_{\tau, \mathbf{g}'_0 \sim P_0, \tilde{\mathbf{g}}'_1 \sim P_{\mathbf{r}X}} \|v_{\mathbf{r}X}(h_{\mathbf{r}X}(\tau)) - \frac{\partial}{\partial \tau} h_{\mathbf{r}X}(\tau)\|_F^2, \quad (38)$$

988 where $h_{\mathbf{r}X}$ is the path generated by \mathbf{g}'_0 and $\tilde{\mathbf{g}}'_1$. Since $P_0 = (\mathcal{R}_{\mathbf{r}^{-1}})_\# P_0$ and $P_{\mathbf{r}X} = (\mathcal{R}_{\mathbf{r}^{-1}})_\# P_X$ by
 989 assumption, we can write $\mathbf{g}'_0 = \mathbf{g}_0 \mathbf{r}^{-1}$ and $\tilde{\mathbf{g}}'_1 = \tilde{\mathbf{g}}_1 \mathbf{r}^{-1}$, where $\mathbf{g}_0 \sim P_0$ and $\tilde{\mathbf{g}}_1 \sim P_X$. According to
 990 the first part of Prop. F.6, we have $h_{\mathbf{r}X}(\tau) = \mathcal{R}_{\mathbf{r}^{-1}} h_X(\tau)$, where h_X is a path generated by \mathbf{g}_0 and $\tilde{\mathbf{g}}_1$.
 991 By taking derivative on both sides of the equation, we have $\frac{\partial}{\partial \tau} h_{\mathbf{r}X}(\tau) = (\mathcal{R}_{\mathbf{r}^{-1}})_{*, h_X(\tau)} \frac{\partial}{\partial \tau} h_X(\tau)$.
 992 Then we have

$$994 L(\mathbf{r}X) = \mathbb{E}_{\tau, \mathbf{g}'_0 \sim P_0, \tilde{\mathbf{g}}'_1 \sim P_{\mathbf{r}X}} \|v_{\mathbf{r}X}(\mathcal{R}_{\mathbf{r}^{-1}} h_X(\tau)) - (\mathcal{R}_{\mathbf{r}^{-1}})_{*, h_X(\tau)} \frac{\partial}{\partial \tau} h_X(\tau)\|_F^2 \quad (39)$$

995 by inserting these two equations into Eqn. (38). Since v_X is $\mathcal{R}_{\mathbf{r}^{-1}}$ -related to $v_{\mathbf{r}X}$ by assumption, we
 996 have $v_{\mathbf{r}X}(\mathcal{R}_{\mathbf{r}^{-1}} h_X(\tau)) = (\mathcal{R}_{\mathbf{r}^{-1}})_{*, h_X(\tau)} v_X(h_X(\tau))$. Thus, we have

$$998 \|v_{\mathbf{r}X}(\mathcal{R}_{\mathbf{r}^{-1}} h_X(\tau)) - (\mathcal{R}_{\mathbf{r}^{-1}})_{*, h_X(\tau)} \frac{\partial}{\partial \tau} h_X(\tau)\|_F^2 = \|(\mathcal{R}_{\mathbf{r}^{-1}})_{*, h_X(\tau)}(v_{\mathbf{r}X}(h_X(\tau)) - \frac{\partial}{\partial \tau} h_X(\tau))\|_F^2 \\ 1000 = \|(v_{\mathbf{r}X}(h_X(\tau)) - \frac{\partial}{\partial \tau} h_X(\tau))\|_F^2 \quad (40)$$

1001 where the second equation holds because $(\mathcal{R}_{\mathbf{r}^{-1}})_{*, h_X(\tau)}$ is an orthogonal matrix. The desired result
 1002 follows.

1004 2) The second statement can be proved similarly as the first one, where σ -equivariance is considered
 1005 instead of $\mathcal{R}_{\mathbf{r}^{-1}}$ -equivariance.

1006 3) Denote $\mathbf{g}'_0 = r\mathbf{g}_0$ and $\tilde{\mathbf{g}}'_1 = r\tilde{\mathbf{g}}_1$, where $\mathbf{g}_0 \sim P_0$ and $\tilde{\mathbf{g}}_1 \sim P_X$. According to the third part of
 1007 Prop. F.6, we have $\hat{h}_X(\tau) = \mathcal{L}_r(h_X(\tau))$. By taking derivative on both sides of the equation, we have
 1008 $\frac{\partial}{\partial \tau} \hat{h}_X(\tau) = (\mathcal{L}_r)_{*, h_X(\tau)} \frac{\partial}{\partial \tau} h_X(\tau)$. Then the rest of the proof can be conducted similarly to the first
 1009 part of the proof. □

1012 G MODEL DETAILS

1014 Let $F_u^l \in \mathbb{R}^{c \times (2l+1)}$ be a channel- c degree- l feature at point u . The equivariant dot-product attention
 1015 is defined as:

$$1016 A_u^l = \sum_{v \in KNN(u) \setminus \{u\}} \frac{\exp(\langle Q_u, K_{vu} \rangle)}{\sum_{v' \in KNN(u) \setminus \{u\}} \exp(\langle Q_u, K_{v'u} \rangle)} V_{vu}^l, \quad (41)$$

1018 where $\langle \cdot, \cdot \rangle$ is the dot product, $KNN(u) \subseteq \bigcup_i X_i$ is a subset of points u attends to, $K, V \in \mathbb{R}^{c \times (2l+1)}$
 1019 take the form of the e3nn (Geiger and Smidt, 2022) message passing, and $Q \in \mathbb{R}^{c \times (2l+1)}$ is obtained
 1020 by a linear transform:

$$1022 Q_u = \bigoplus_l W_Q^l F_u^l, \quad K_v = \bigoplus_l \sum_{l_e, l_f} c_K^{(l, l_e, l_f)} (|uv|) Y^{l_e}(\widehat{vu}) \otimes_{l_e, l_f}^l F_v^{l_f}, \quad (42)$$

$$1024 V_v^l = \sum_{l_e, l_f} c_V^{(l, l_e, l_f)} (|uv|) Y^{l_e}(\widehat{vu}) \otimes_{l_e, l_f}^l F_v^{l_f}. \quad (43)$$

1026 Here, $W_Q^l \in \mathbb{R}^{c \times c}$ is a learnable weight, $|vu|$ is the distance between point v and u , $\widehat{vu} = \vec{vu}/|vu| \in$
 1027 \mathbb{R}^3 is the normalized direction, $Y^l : \mathbb{R}^3 \rightarrow \mathbb{R}^{2l+1}$ is the degree- l spherical harmonic function,
 1028 $c : \mathbb{R}_+ \rightarrow \mathbb{R}$ is a learnable function that maps $|vu|$ to a coefficient, and \otimes is the tensor product with
 1029 the Clebsch-Gordan coefficients.
 1030

1031 To accelerate the computation of K and V , we use the $SO(2)$ -reduction technique (Passaro and
 1032 Zitnick, 2023), which rotates the edge uv to the y -axis, so that the computation of spherical harmonic
 1033 function, the Clebsch-Gordan coefficients, and the iterations of l_e are no longer needed.
 1034

1035 The main idea of $SO(2)$ -reduction (Passaro and Zitnick, 2023) is to rotate the edge uv to the y -axis,
 1036 and then update node feature in the rotated space. Since all 3D rotations are reduced to 2D rotations
 1037 about the y -axis in the rotated space, the feature update rule is greatly simplified.
 1038

1039 Here, we describe this technique in the matrix form to facilitates better parallelization. Let $F_v^l \in$
 1040 $\mathbb{R}^{c \times (2l+1)}$ be a c -channel l -degree feature of point v , and $L > 0$ be the maximum degree of features.
 1041 We construct $\hat{F}_v^l \in \mathbb{R}^{c \times (2L+1)}$ by padding F_v^l with $L - l$ zeros at the beginning and the end of
 1042 the feature, then we define the full feature $F_v \in \mathbb{R}^{c \times L \times (2L+1)}$ as the concatenate of all \hat{F}_v^l with
 1043 $0 < l \leq L$. For an edge vu , there exists a rotation r_{vu} that aligns uv to the y -axis. We define
 1044 $R_{vu} \in \mathbb{R}^{L \times (2L+1) \times (2L+1)}$ to be the full rotation matrix, where the l -th slice $R_{vu}[l, :, :]$ is the l -th
 1045 Wigner-D matrix of r_{vu} with zeros padded at the boundary. K_v defined in (42) can be efficiently
 1046 computed as
 1047

$$1046 \quad K_v = R_{vu}^T \times_{1,2} (W_K \times_3 (D_K \times_{1,2} R_{vu} \times_{1,2} F_v)), \quad (44)$$

1048 where $M_1 \times_i M_2$ represents the batch-wise multiplication of M_1 and M_2 with the i -th dimen-
 1049 sion of M_2 treated as the batch dimension. $W_K \in \mathbb{R}^{(cL) \times (cL)}$ is a learnable weight, $D_K \in$
 1050 $\mathbb{R}^{c \times (2L+1) \times (2L+1)}$ is a learnable matrix taking the form of 2D rotations about the y -axis, *i.e.*, for
 1051 each i , $D_K[i, :, :]$ is
 1052

$$1053 \quad \begin{bmatrix} a_1 & & & & -b_1 \\ a_2 & & & & -b_2 \\ \ddots & & & & \ddots \\ & a_{L-1} & & -b_{L-1} & \\ & b_{L-1} & a_L & a_{L-1} & \\ & \ddots & & & \ddots \\ b_2 & & & a_2 & \\ b_1 & & & & a_1 \end{bmatrix}, \quad (45)$$

1063 where $a_1, \dots, a_L, b_1, \dots, b_{L-1} : \mathbb{R}_+ \rightarrow \mathbb{R}$ are learnable functions that map $|vu|$ to the coefficients.
 1064 V_v defined in (42) can be computed similarly. Note that (44) does not require the computation of
 1065 Clebsch-Gordan coefficients, the spherical harmonic functions, and all computations are in the matrix
 1066 form where no for-loop is needed, so it is much faster than the computations in (42).
 1067

1069 H MORE DETAILS OF SEC. 6

1071 H.1 A REMARK ON METRICS

1073 When $N = 2$, *i.e.*, the input point clouds are X_1 and X_2 , the metric we used is simply the averaged
 1074 RRE/RTE of (X_1, X_2) and (X_2, X_1) . Other popular metrics used in registration papers (Huang
 1075 et al., 2021), like FMR and IR, are not suitable, because they measure the quality of the estimated
 1076 correspondences, which our method does not compute. In addition, we do not use RR, because it
 1077 is similar to the RRE and RTE metric (mean value v.s. threshold percentage), and it depends on a
 1078 manually selected threshold for indoor scene, *i.e.*, 5 degrees and 2cm. This metric does not apply to
 1079 other dataset than 3DMatch, e.g., the translation threshold is way too strict for outdoor dataset like
 KITTI, and is not applicable for dataset with unknown scale like BB.

1080

H.2 AN VERIFICATION OF EQUIVARIANCES

1081

1082 This subsection directly checks the relatedness of the learned vector field v . To verify the $SO(3)$ -
 1083 relatedness, *i.e.*, $v_X(r\mathbf{g}) = rv_X(\mathbf{g})$, we compute loss $L_R = \|v_X(r\mathbf{g}) - rv_X(\mathbf{g})\|_F$ where r is a
 1084 random rotation, and \mathbf{g} is a random rigid transformation. This error will be close to zero if the
 1085 equivariance holds.

1086 On the other hand, we can also verify the σ -relatedness directly by computing $L_P = \|v_{\sigma X}(\sigma\mathbf{g}) -$
 1087 $\sigma v_X(\mathbf{g})\|_F$, where σ is a random permutation and \mathbf{g} is a random rigid transformation. Similarly, this
 1088 error will be close to zero if the equivariance holds.

1089 We compute L_P and L_R 3 times and present the mean results in Table 6. We observe that both errors
 1090 are close to 0 with only floating point errors, suggesting these two equivariances hold.
 1091

1092

1093

1094

1095

1096

1097

H.3 MORE RESULTS ON 3DMATCH, BB AND KITTI

1098

1099 We present more details of Eda on 3DL in Fig. 5. We observe that the vector field is gradually
 1100 learned during training, *i.e.*, the training error converges. On the test set, RK4 outperforms the RK1,
 1101 and they both benefit from more time steps, especially for rotation errors.

1102

1103

1104

1105

1106

1107

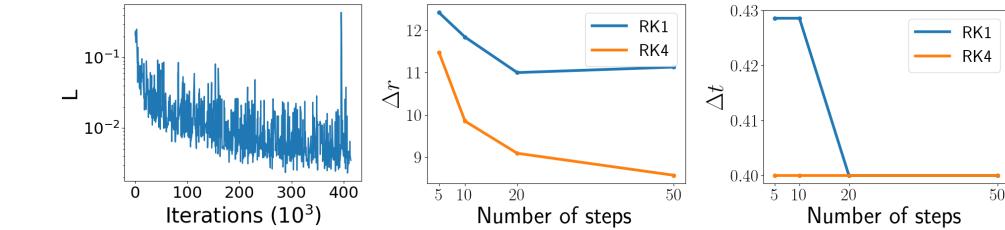
1108

1109

1110

1111

1112



1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1111 Figure 5: More details of Eda on 3DL. Left: the training curve. Middle and right: the influence of
 1112 RK4/RK1 and the number of time steps on Δr and Δt .

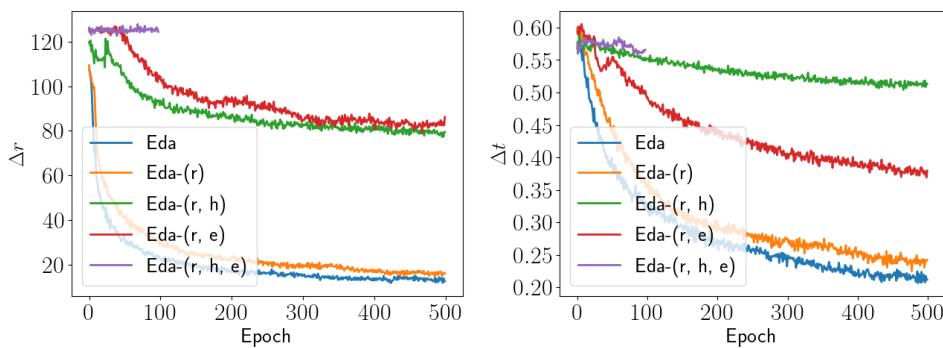
1129

1130

1131

1132

1133



1126

1127

1128

1129

1130

1131

1132

1133

1126 Figure 6: Validation error curves of all methods in Tab. 4. The training of Eda-(r, h, e) is unstable
 1127 and produces NaN value at the early stage.

We now provide more details for the ablation study reported in Tab. 4. The curve of validation errors of all methods are presented in Fig. 6. All methods use (RK1, 10) for sampling. Eda-(r) satisfies all equivariances. Eda-(r, h) breaks the first and third part of Prop. 4.7. Eda-(r, e) and Eda-(r, h, e) further break the second part of Prop. 4.6. The non-equivariant network is obtained by replacing the matrix (45) by a linear transformation with exactly the same number of parameters. All methods considered in this study contain exactly the same number of trainable parameters.

1134 We provide the complete version of Table 2 in Table 7, where we additionally report the standard
 1135 deviations of Eda.
 1136

1137 Table 7: The complete version of Table 2 with stds of Eda reported in brackets.

	3DM	3DL	3DZ			
	Δr	Δt	Δr	Δt	Δr	Δt
FGR	69.5	0.6	117.3	1.3	—	—
GEO	7.43	0.19	28.38	0.69	—	—
ROI (500)	5.64	0.15	21.94	0.53	—	—
ROI (5000)	5.44	0.15	22.17	0.53	—	—
AMR	5.0	0.13	20.5	0.53	—	—
Eda (RK4, 50)	2.38 (0.16)	0.16 (0.01)	8.57 (0.08)	0.4 (0.0)	78.74 (0.6)	0.96 (0.01)

1147 We provide some qualitative results on BB datasets in Fig. 7. Eda can generally recover the shape of
 1148 the objects.
 1149

1150 A complete version of Tab. 3 is provided in Tab. 8, where we additionally report the standard
 1151 deviations of Eda.
 1152

1153 Table 8: The complete version of Table 3 with stds of Eda reported in brackets.

	Δr	Δt	Time (min)
GLO	126.3	0.3	0.9
DGL	125.8	0.3	0.9
LEV	125.9	0.3	8.1
Eda (RK1, 10)	80.64	0.16	19.4
Eda (RK4, 10)	79.2 (0.58)	0.16 (0.0)	76.9

1160
 1161 We provide a few examples of the reconstructed road views in Fig. 8.
 1162

1163 I LIMIATION AND FUTURE WORKS

1164 Eda in its current form has several limitations. First, Eda is slow when using a high order RK solver
 1165 with a large number of steps. Besides its iterative nature, another cause is the lack of CUDA kernel
 1166 level optimization like FlashAttention (Dao et al., 2022) for equivariant attention layers. We expect to
 1167 see acceleration in the future when such optimization is available. Second, Eda always uses all input
 1168 pieces, which is not suitable for applications like archeology reconstruction, where the input data
 1169 may contain pieces from unrelated objects. Finally, in the future research, we plan to make Eda a
 1170 foundation model by scaling up the training, so that it can handle different types of data and achieve
 1171 higher precision. In particular, the scaling law (Kaplan et al., 2020) of Eda worths investigation,
 1172 where we expect to see that an increase in model/data size leads to an increase in performance similar
 1173 to image generation applications (Peebles and Xie, 2023).
 1174

1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187

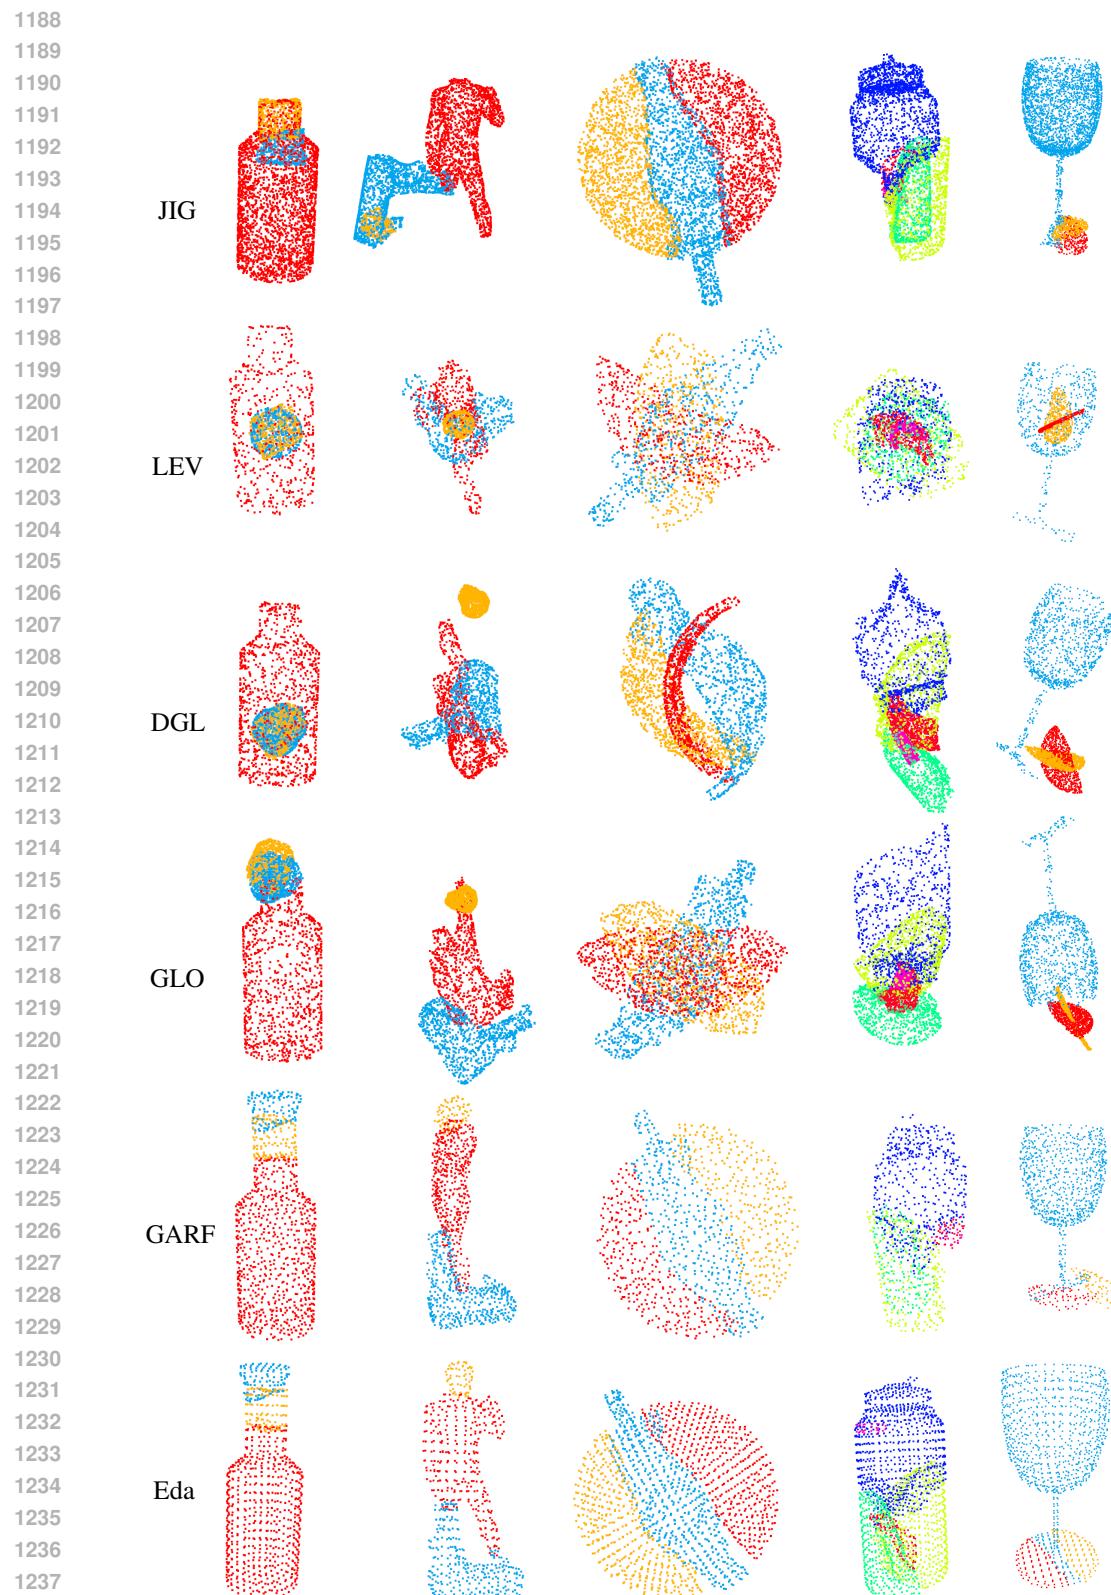


Figure 7: Qualitative results on BB. Zoom in to see details.

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

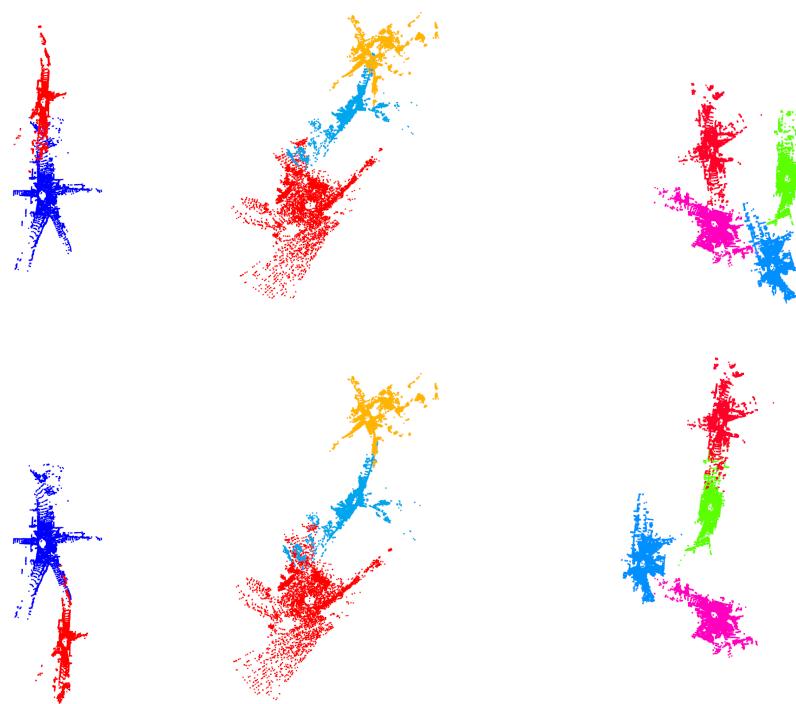
1277

1278

1279

1280

1281



1282 Figure 8: Qualitative results of Eda on kitti. We present the results of Eda (1-st row)
1283 and ground truth (2-nd row). For each assembly, Eda correctly places the input road views on the same plane.
1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

1296

1297

1298

1299

1300

1301

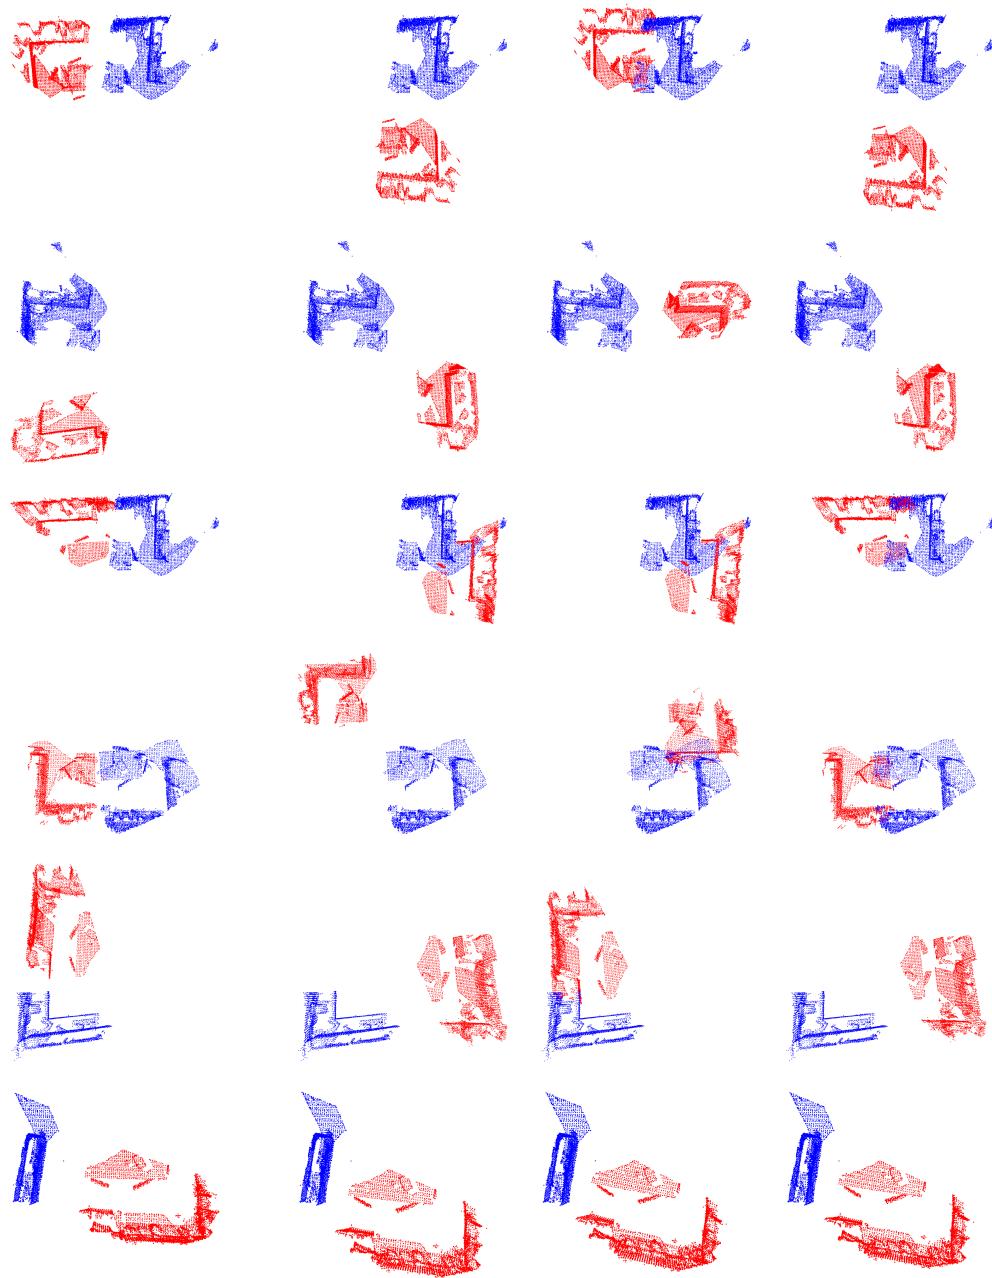


Figure 9: Qualitative results of Eda on 3DZ. Cameras are set to look at the room from above.

1344

1345

1346

1347

1348

1349