

000 RECON: ROBUST SYMMETRY DISCOVERY VIA 001 EXPLICIT CANONICAL ORIENTATION NORMALIZATION 002

003 **Anonymous authors**

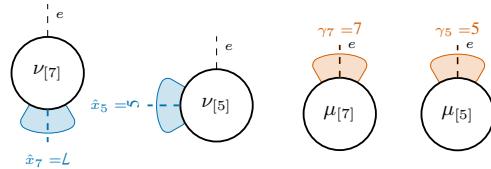
004 Paper under double-blind review

005 ABSTRACT

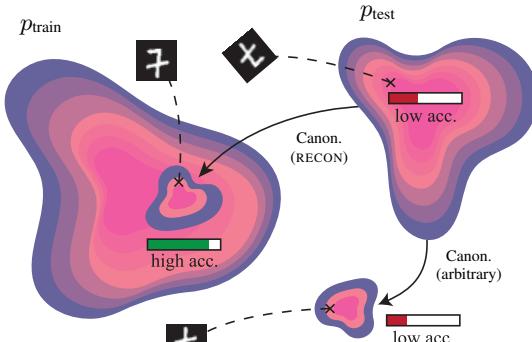
006
007
008
009
010
011 Real world data often exhibits unknown, instance-specific symmetries that rarely
012 exactly match a transformation group \mathcal{G} fixed a priori. Class-pose decompositions
013 aim to create disentangled representations by factoring inputs into invariant features
014 and a pose $g \in \mathcal{G}$ defined relative to a training-dependent, *arbitrary* canonical
015 representation. We introduce RECON, a class-pose agnostic *canonical orientation*
016 *normalization* that corrects arbitrary canonicals via a simple right-multiplication,
017 yielding *natural*, data-aligned canonicalizations. This enables (i) unsupervised
018 discovery of instance-specific *pose* distributions, (ii) detection of out-of-distribution
019 poses and (iii) a *plug-and-play test-time canonicalization layer*. This layer can be
020 attached on top of any pre-trained model to infuse group invariance, improving its
021 performance without retraining. We demonstrate results on 2D image benchmarks
022 and extend unsupervised instance-level pose discovery to 3D groups.

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	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	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
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	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	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
	Class-pose	RECON
Data	<i>Canon. Distrib.</i>	<i>Canon. Distrib.</i>
$\{1, \dots, 7, \dots, 7\}$	L	$\nu_{[7]}$
$\{5, \dots, 5, \dots, 5\}$	ν	$\nu_{[5]}$

(a) Inputs with identical $\pm 30^\circ$ symmetries.



(b) Relative distributions $\nu_{[x]}$ (□) vs normalized distributions $\mu_{[x]}$ (■) via RECON.



(c) Distribution shift (p_{train} vs p_{test}) induced by unseen symmetries at test time, corrected by RECON.

Figure 1: (a) Class-pose methods assign arbitrary (often out-of-distribution) canonicals per class. (b) This leads to distinct relative-pose distributions $\nu_{[x]}$, obscuring the shared $\pm 30^\circ$ symmetries. RECON corrects these offsets, mapping inputs under the same symmetries to the *same* distribution $\mu_{[x]}$ and extracting their natural pose γ . (c) Our data-aligned canonicalization removes symmetry-induced distribution shifts, improving downstream performance of pre-trained backbones without architectural changes or retraining.

1 INTRODUCTION

Symmetry transformations like rotations arise naturally in many domains (Cohen & Welling, 2016; Higgins et al., 2018; Bronstein et al., 2021), with objects appearing in poses related by group transformations $g \in \mathcal{G}$ (e.g., molecules in different orientations) (Weiler et al., 2018; Brandstetter et al., 2022). While \mathcal{G} -equivariant neural networks exploit such structure (Cohen et al., 2019; Romero et al., 2020; Romero & Cordonnier, 2020; Wang et al., 2020; Cohen et al., 2018), they can be overly constrained when there is a mismatch between the pre-fixed group \mathcal{G} and the symmetries in the data (Romero & Lohit, 2022; Weiler & Cesa, 2019). In effect, real-world symmetries are often (i) unknown a priori, (ii) partial (covering only part of the group), or (iii) instance-dependent. This motivates methods that *discover* symmetries from data (van der Linden et al., 2024; Forestano et al., 2023; Benton et al., 2020; Allingham et al., 2024).

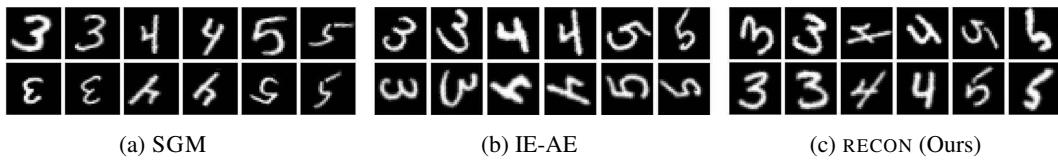


Figure 2: Comparison of canonical representations. Top row: Input MNIST digits under varying rotations. Bottom row: canonicals generated by (a) SGMs (Allingham et al., 2024) and (b) IE-AE (Winter et al., 2022), which yield arbitrary orientations, while RECON consistently produces data-aligned canonical poses (upright digits) (c).

However, existing approaches often require supervision or learn only dataset-level patterns (Sec. 4). Our goal is to discover, in an unsupervised manner, the group transformations inherent to each instance; in particular, we aim to learn probability distributions on \mathcal{G} that describe the *poses in which each instance appears in the data*.¹ We argue that a promising foundations lies on class-pose decomposition methods (Winter et al., 2022; Marchetti et al., 2023; Yokota & Hontani, 2022; Allingham et al., 2024), which disentangle inputs into invariant *class* features and a *pose* $g \in \mathcal{G}$ relative to a learned canonical representation. The canonical’s orientation is generally *arbitrary* (Winter et al., 2022; Allingham et al., 2024), which results in an out-of-distribution (OOD) canonicalization and in arbitrarily shifted relative-pose distributions (Fig. 1).

We address arbitrary canonicals and propose a framework for *Robust unsupervised discovery of intrinsic symmetry distributions via Explicit Canonical Orientation Normalization* (RECON). We prove (Proposition 3.1) that by estimating the centroid (Fréchet mean) of the observed relative poses, which captures the offset induced by the arbitrary canonical, we can approximate symmetry distributions centered at the input’s *natural pose* via a simple right-multiplication. This yields instance-specific symmetry descriptions that are *robust* (independent of the arbitrary canonical), *interpretable* (centered at $e \in \mathcal{G}$, representing the input’s natural pose), and *comparable* across classes (Fig. 1b). RECON is validated on 2D image benchmarks and 3D molecular conformations – beyond typical 2D-only settings of prior work (van der Linden et al., 2024; Romero & Lohit, 2022; Benton et al., 2020; Allingham et al., 2024; Kim et al., 2024). Lastly, we provide practical applications in (i) OOD pose detection and (ii) test-time canonicalization, a drop-in method to grant group invariance to frozen pre-trained models, improving downstream performance (Sec. 5.2).

Contributions

1. We propose RECON, a method for discovery of instance-specific *pose* distributions from unlabeled data leveraging class-pose representation learning methods.
2. We achieve this through *canonical orientation normalization* (Proposition 3.1), an architecture-agnostic correction of arbitrary canonicals, yielding data-aligned natural canonicalizations and well-behaved symmetry distributions.
3. We empirically validate on 2D images and large-scale, real-world 3D data. We offer applications in OOD pose detection and test-time canonicalization, offering performance improvements to pre-trained backbones via a simple plug-in canonicalization step.

Our code is publicly available at [link-hidden-for-double-blind-review](https://github.com/RECON-Group/RECON).

2 PRELIMINARIES

Our approach leverages class-pose decomposition methods, a class of neural networks designed to disentangle input data into an invariant (*class*) component and an equivariant (*pose*) component based on a given transformation group \mathcal{G} . In particular, we build upon Invariant-Equivariant Autoencoders (IE-AEs) (Winter et al., 2022). We briefly review the core concepts here; formal definitions regarding group theory, representations and \mathcal{G} -equivariance are deferred to Appendix A.

IE-AEs aim to learn group invariant and equivariant representations for an input $x \in \mathcal{X}$ of a vector space, e.g., an image or a 3D structure. For a chosen group \mathcal{G} , an IE-AE maps x to a \mathcal{G} -invariant component $z \in \mathbb{R}^n$ and a \mathcal{G} -equivariant group element (pose) $g \in \mathcal{G}$.²

¹Note how this differs from estimating the stabilizer \mathcal{S}_x describing the (self) symmetries of x (cf. Appx D.1).

²Figure 7 in the Appendix, adapted from Winter et al. (2022), visualizes this architecture.

108 **Invariant component z and canonical representation \hat{x}** First, a \mathcal{G} -invariant encoder $\eta : \mathcal{X} \rightarrow$
 109 $\mathcal{Z} \subseteq \mathbb{R}^n$ learns an invariant representation $z = \eta(x)$. This embedding z captures features of the input
 110 that are independent of its pose under \mathcal{G} , meaning z remains unchanged if x is transformed by any
 111 element of the group. A corresponding decoder $\delta : \mathcal{Z} \rightarrow \mathcal{X}$ then reconstructs the input $\hat{x} = \delta(z)$
 112 from this invariant representation. Note that since the decoder only sees the invariant embedding
 113 $z = \eta(x)$, it will produce the *same reconstruction* $\hat{x} = \delta(\eta(x))$ for all group transformations of the
 114 input. This common reconstruction \hat{x} is called the *canonical representation*. Notably, the specific
 115 pose or orientation of \hat{x} is *arbitrary*, influenced by initialization or training dynamics (Winter et al.,
 116 2022). This behaviour is common to other class-pose decomposition methods (Yokota & Hontani,
 117 2022; Allingham et al., 2024).

118 **Pose component g** The canonical \hat{x} has some fixed, arbitrary pose. To recover the original
 119 input x , we have to determine the specific group transformation $g \in \mathcal{G}$ that maps \hat{x} back to x .
 120 This is the role of the *group function* $\psi : \mathcal{X} \rightarrow \mathcal{G}$, which predicts this transformation $g = \psi(x)$.
 121 All the IE-AE components (η, δ, ψ) are then trained jointly by minimizing the reconstruction loss
 122 $d(\rho_{\mathcal{X}}(\psi(x)) \delta(\eta(x)), x)$,³ where $\rho_{\mathcal{X}}$ is the group action of \mathcal{G} on \mathcal{X} .

123 In summary, IE-AE provides an invariant latent vector $z = \eta(x)$, an arbitrary canonical $\hat{x} = \delta(z)$, and
 124 the *relative* group transformation $g = \psi(x)$ that maps the canonical \hat{x} back to x . Our method leverages
 125 these components, specifically the distribution of relative transformations $g = \psi(x)$ and the invariant
 126 latent space \mathcal{Z} , to discover the specific distribution of transformations that appear in the data.

127 3 METHOD

128 3.1 PROBLEM STATEMENT

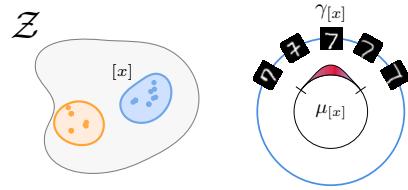
129 Our goal is to discover the characteristic distribution
 130 of symmetry transformations associated with
 131 different classes of objects within unlabeled data \mathcal{X} .
 132 Consider datasets such as GEOM (Axelrod & Gómez-
 133 Bombarelli, 2022), where each molecule class is re-
 134 presented as an ensemble of multiple 3D molecular
 135 conformations; or MNIST, where same digits share
 136 an underlying shape under varying handwriting styles.
 137 Generally, variations *within* these classes are a com-
 138 bination of (a) non-group structural distortions (e.g.,
 139 bond rotations or ring puckering in molecules; style
 140 variations in digits) and (b) transformations from a
 141 symmetry group \mathcal{G} (like $SO(3)$ rotations of molecu-
 142 lar conformations). We specifically aim to model the
 143 distribution of the underlying group transformations.
 144 Therefore, to achieve this goal, we first need to (i) define a way to group objects based on their
 145 intrinsic shape or similarity – independently of their pose or minor deformations and under the
 146 absence of labels, and (ii) model their pose variations.

147 **(i) Modeling pose-invariant similarity** We model pose invariant similarity through equivalence
 148 classes. The \mathcal{G} -invariant latent space typical of class-pose decomposition methods is a promising
 149 candidate for this purpose. In effect, we rely on the principle that structurally similar objects generally
 150 map to nearby points in \mathcal{Z} . This assumption is empirically supported by previous work in class-pose
 151 decompositions methods showing that \mathcal{Z} maps different classes into well-separated connected clusters
 152 (see t-SNE visualizations of \mathcal{Z} in IE-AEs (Winter et al., 2022)). Additionally, it aligns with broader
 153 findings indicating that deep features correctly capture perceptual and semantic similarity, even when
 154 learned without supervision (Zhang et al., 2018).

155 Formally, let $\eta : \mathcal{X} \rightarrow \mathcal{Z}$ be such a \mathcal{G} -invariant encoder, we then define an equivalence relation \sim_{ε} in
 156 \mathcal{X} based on connected proximity in \mathcal{Z} :

157
$$x \sim_{\varepsilon} y \iff \exists \{x_i\}_{i=0}^N \subseteq \mathcal{X} \text{ for some } N \in \mathbb{N} \text{ s.t. } x_0 = x, x_N = y, \text{ and } d_{\mathcal{Z}}(\eta(x_i), \eta(x_{i+1})) < \varepsilon \forall i$$

158
 159 ³Any estimator ψ satisfying the proposed reconstruction loss is \mathcal{G} -equivariant as proved in Winter et al.
 160 (2022).



161 Figure 3: Problem setup. *Left*: A class $[x]$ is defined by inputs clustering together in the
 162 invariant space \mathcal{Z} . *Right*: We model instances $s \in [x]$ as $\rho_{\mathcal{X}}(g)\gamma_{[x]} + \epsilon_s$, where $\gamma_{[x]}$ is a reference
 163 frame and g is drawn from a distribution over rotation angles $\mu_{[x]}$. The objective is to
 164 recover $\mu_{[x]}$ from the unlabeled data.

162 where $d_{\mathcal{Z}}$ is a norm-induced distance in \mathcal{Z} and $\varepsilon > 0$ is a small threshold. Intuitively, this equivalence
 163 relation groups objects whose invariant features form a connected component (Fig. 3, left). We denote
 164 the resulting equivalence class containing x as $[x]$.⁴ We empirically validate this equivalence class
 165 definition in Appendix D.3, confirming effective capture of pose-invariant similarity, and analyze its
 166 limitations in detail.

167 **(ii) Modeling pose variations** Having established the classes $[x]$, we now focus on modeling the
 168 distribution of group transformations responsible for pose variations within each class. We model
 169 this as a probabilistic process.
 170

171 Conceptually, there exists some true underlying probability distribution $\mu_{[x]}$ that governs how in-
 172 stances $s \in [x]$ are generated by transforming some reference pose $\gamma_{[x]} \in \mathcal{X}$ plus residual non-group
 173 variations. Formally, let \mathcal{G} be a Lie group and consider $\mu_{[x]}$ a probability distribution over \mathcal{G} . We
 174 assume that instances $s \in [x]$ can be generated by sampling a transformation $g \sim \mu_{[x]}$ and applying
 175 it to a reference pose $\gamma_{[x]}$, plus a deviation term $\varepsilon_s \sim \mathcal{P}_{\varepsilon'}$:

$$s = \rho_{\mathcal{X}}(g)\gamma_{[x]} + \varepsilon_s, \quad \text{with } g \sim \mu_{[x]}, \varepsilon_s \sim \mathcal{P}_{\varepsilon'}. \quad (1)$$

176 Here, $\mu_{[x]}$ represents the true probability distribution relative to the reference pose $\gamma_{[x]}$. The distribu-
 177 tion $\mathcal{P}_{\varepsilon'}$ models variations that are not explained by the group transformation (e.g., style variations in
 178 digits, internal conformational changes in ensembles of molecules) and is assumed to have zero mean
 179 and small variance (we assume $\|\varepsilon_s\| < \varepsilon'$ almost surely for simplicity). Fig. 3 (right) illustrates this
 180 model.
 181

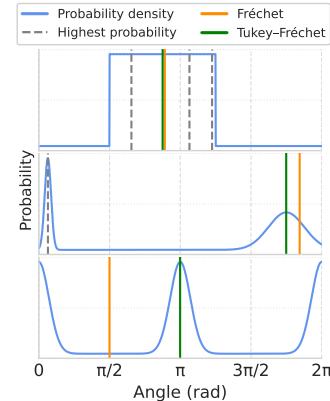
182 Note that under this model, there are several ways of representing the symmetries in the data,
 183 depending on the reference pose $\gamma_{[x]}$. Consider a dataset of handwritten ‘7’s exhibiting rotational
 184 symmetries uniformly between $\pm 30^\circ$: $\{1, \dots, 7, \dots, 7\}$. Using this model, we can describe the
 185 symmetries in this dataset as $\mu_{[7]} = \mathcal{U}([150^\circ, 210^\circ])$ for a reference pose $\gamma_{[7]} = 'L'$, or as $\mu_{[7]} =$
 186 $\mathcal{U}([-30^\circ, 30^\circ])$ for a reference pose $\gamma_{[7]} = '7'$. While both descriptions are *mathematically valid*, the
 187 latter, identity-centered representation is far more desirable. It aligns with our intuition of a *naturally*
 188 *occurring pose* and directly reflects the symmetries as deviations from a neutral reference frame (e,
 189 no transformation), which offers several advantages (as demonstrated in Sec. 5).
 190

191 **Defining a natural pose** Ideally, we aim to obtain such identity-
 192 centered descriptions of the data symmetries. As we have just ex-
 193 exemplified, this boils down to obtaining the symmetry distribution
 194 $\mu_{[x]}$ whose $\gamma_{[x]}$ is the “natural” pose in the data – like the upright
 195 ‘7’ – centered at the group identity. But without labels defining what
 196 a reference or *true* canonical pose is, how can we define a *natural*
 197 *pose*? Specifically, how can we obtain a canonical pose that is ge-
 198 ometrically grounded in the symmetries in the data – rather than just
 199 an arbitrary pose?
 200

201 While the most likely pose (i.e., with maximum probability den-
 202 sity) may seem like a good candidate – for instance in the case of
 203 Gaussian or other unimodal distributions – it can be ill-defined (e.g.,
 204 for uniform distributions, see Fig. 4, top) or not representative (e.g.,
 205 in the case of a distribution with an anomalous sharp spike in its
 206 density, Fig. 4, middle). To establish a general definition, we use
 207 the *Fréchet mean* (Pennec, 2006) or Riemannian center of mass,
 208 a measure of the central tendency of the distribution on the group
 209 manifold. Given any distribution μ over \mathcal{G} , its Fréchet mean $\mathcal{F}(\mu)$ is the unique transformation
 210 minimizing the expected squared Riemannian distance $d_{\mathcal{R}}$ to samples from μ :

$$\mathcal{F}(\mu) = \operatorname{argmin}_{y \in \mathcal{G}} \mathbb{E}_{g \sim \mu} [d_{\mathcal{R}}(y, g)^2]. \quad (2)$$

211 Under mild conditions, the Fréchet mean provides a unique centroid (Afsari, 2011) that aligns with
 212 intuition for common cases: it corresponds to the midpoint for symmetric uniform distributions and
 213 the peak for strongly unimodal ones (e.g. Gaussian). However, in certain multimodal cases, e.g., a
 214 bimodal distribution on $SO(2)$ with opposing identical peaks (Fig. 4, bottom), the Fréchet mean may



215 Figure 4: Estimator comparison for different distributions over $SO(2)$.

⁴While we use this definition throughout our theoretical derivations (see Proofs B), in practice, we approximate the classes $[x]$ by simply computing the k -nearest neighbors of $\eta(x)$ in \mathcal{Z} for efficiency (see Algorithm 1).

fall outside the support of μ . This would lead to an out-of-distribution canonicalization, affecting downstream tasks (Sec. 5.2). Therefore, we propose a robust Fréchet mean extension based on robust *M-estimators* (Shevlyakov et al., 2008) that we coin the *Tukey-Fréchet mean* $\mathcal{F}_r(\mu)$:

$$\mathcal{F}_r(\mu) = \arg \min_{y \in \mathcal{G}} \mathbb{E}_{g \sim \mu} [m(d_{\mathcal{R}}(y, g); c)], \quad (3)$$

where $m(u; c)$ is the Tukey biweight loss (Shin & Oh, 2022; Rousseeuw & Hubert, 2011). This variation extends the Fréchet mean and converges consistently to the same mode in multimodal distributions (cf. Appendix C). For simplicity, we derive our theoretical framework with the classical Fréchet mean, whose properties are well established, but in practice, the robust Tukey-Fréchet can be used as a drop-in replacement.

We thus define our target distribution $\mu_{[x]}$ in equation 1 as the probability distribution over \mathcal{G} that is centered at the group identity e in the Fréchet sense, i.e., satisfying $\mathcal{F}(\mu_{[x]}) = e$. Consequently, the reference pose $\gamma_{[x]}$ (the natural pose) corresponds implicitly to the Fréchet mean (i.e., the upright ‘7’ in the previous example). Our objective is to estimate this well-behaved $\mu_{[x]}$ from unlabeled data. The subsequent section details our method for achieving this by normalizing the arbitrarily offset outputs of IE-AEs.

3.2 RECOVERING SYMMETRIES VIA CANONICAL ORIENTATION NORMALIZATION

The core idea is to leverage the disentangled representation $(z, g) \in \mathcal{Z} \times \mathcal{G}$ of the IE-AE to extract the symmetry distribution of an input, and correct for the arbitrary canonical pose to yield a canonical-pose independent distribution. Our main result provides a way to recover the true symmetries $\mu_{[x]}$ from the relative transformations $\psi(x)$. We provide a proof in Appendix B

Proposition 3.1 (Approximation of $\mu_{[x]}$ via Normalization). *Let \mathcal{X} be a metric space, \mathcal{G} a Lie group and η, δ, ψ an IE-AE where ψ is continuous on a compact domain \mathcal{X} . Suppose that \mathcal{X} exhibits symmetries characterized by $\mu_{[x]}$ where $\mathcal{F}(\mu_{[x]}) = e$ as described above. Consider a random sample $\{s_i\}_{i=1}^N$ of $[x]$ and denote their images by ψ as $\psi([x]) = \{\psi(s_i)\}_{i=1}^N$. Let $\hat{\Gamma}_{[x]} \in \mathcal{G}$ be the empirical Fréchet mean of $\psi([x])$. Then, the empirical distribution $\hat{\mu}_{[x]}$ corresponding to the normalized samples $\psi([x])\hat{\Gamma}_{[x]}^{-1}$ approximates the target distribution $\mu_{[x]}$. Specifically, $\hat{\mu}_{[x]}$ converges in Wasserstein distance to $\mu_{[x]}$ as $\varepsilon' \rightarrow 0$ and $N \rightarrow \infty$.*

Interpretation. Proposition 3.1 provides a practical method for consistently retrieving $\mu_{[x]}$ through *canonical orientation normalization*, outlined in practice in Algorithm 1. The process involves identifying the class $[x]$ and collecting the relative transformations $\psi([x])$ (representing poses relative to the arbitrary canonical \hat{x}). The distribution of transformation in this set is offset w.r.t. $\mu_{[x]}$ by a translation induced by the canonical pose, which can be estimated through the empirical Fréchet mean $\hat{\Gamma}_{[x]}$ of the observed transformations $\psi([x])$. Then, right-multiplying by the inverse $\hat{\Gamma}_{[x]}^{-1}$ corrects this offset and centers the Fréchet mean at the identity consistently for all classes. This removes the influence of the arbitrary choice of the canonical pose \hat{x} , providing a geometrically meaningful canonical (centered at the Fréchet mean) and enabling the retrieval of symmetry distributions with useful properties, as we show in our experiments (Sec. 5).

Algorithm 1 Canonical orientation normalization

Require: Trained IE-AE (η, δ, ψ) , input x , number of neighbors k

- 1: Compute invariant embedding: $z \leftarrow \eta(x)$
- 2: Find k -nearest neighbors s_j of x based on $d_{\mathcal{Z}}(\eta(s_j), z)$ to approximate class $[x]$
- 3: Collect relative transformations:

$$\psi([x]) \leftarrow \{\psi(s_j) \mid s_j \in k\text{-NN}(x)\}$$
- 4: Estimate offset via Fréchet mean:

$$\hat{\Gamma}_{[x]} \leftarrow \operatorname{argmin}_{y \in \mathcal{G}} \sum_{g_i \in \psi([x])} d_{\mathcal{G}}(y, g_i)^2$$
- 5: Compute inverse offset: $\hat{\Gamma}_{[x]}^{-1}$
- 6: Compute normalized transformations:

$$\psi'([x]) \leftarrow \{g_i \hat{\Gamma}_{[x]}^{-1} \mid g_i \in \psi([x])\}$$
- 7: **return** $\psi'([x])$ ▷ Samples approximating $\mu_{[x]}$

3.3 INFERRING SYMMETRIES VIA LEARNED MAPPINGS

Algorithm 1 provides a way to estimate the symmetry distribution and the centering transformation for any class *in the training data*. To enable efficient inference for *unseen* inputs without explicit class computations at test time, we can train learnable mappings. If $\mu_{[x]}$ has a known parametric form

270 (or we approximate it parametrically), its parameters $\theta_{[x]}$ can be estimated as
 271

$$\hat{\theta}_{[x]} = \varphi \left(\psi ([x]) \hat{\Gamma}_{[x]}^{-1} \right), \quad (4)$$

272 where φ is an appropriate estimator for the parameters of the distribution, e.g., maximum likelihood.
 273 Consider the estimates $\hat{\theta}_{[x]}$ and $\hat{\Gamma}_{[x]}$. We can learn two maps using these estimates as pseudo-labels:
 274

- 275 • A map Θ predicting the parameters of the symmetry distribution of an input by minimizing
 $\mathcal{L}_p = d_\theta(\Theta(x), \hat{\theta}_{[x]}),$
- 276 • A map $\Lambda : \mathcal{X} \rightarrow \mathcal{G}$, predicting the centering transformation by minimizing $\mathcal{L}_c = d_G(\Lambda(x), \hat{\Gamma}_{[x]}),$

277 where d_θ, d_G are appropriate distances. The first mapping generalizes the estimation process:
 278 at test time, given x , we can predict its symmetry parameters as $\hat{\theta} = \Theta(x)$ without requiring class
 279 computations. The second mapping allows us to obtain our RECON canonicalizations during inference
 280 as $C(x) = \rho_{\mathcal{X}}(\Lambda(x) \cdot \psi(x)^{-1}) x$. Additionally, we can use these functions in combination to detect
 281 out-of-distribution symmetries Section 5.2. We model Θ and Λ as \mathcal{G} -invariant networks, ensuring
 282 that predictions depend only on the object’s class $[x]$, not its specific input pose.
 283

284 4 RELATED WORK

285 **Class-pose decomposition methods** Unsupervised learning of disentangled invariant and equivariant
 286 representations via autoencoders or other learning paradigms has seen various propositions
 287 (Shu et al., 2018; Guo et al., 2019; Feige, 2019; Kosiorek et al., 2019; Koneripalli et al., 2020;
 288 Winter et al., 2021; 2022; Yokota & Hontani, 2022). In this work, we build on top of the IE-AE
 289 framework (Winter et al., 2022). Our main contribution is the discovery of the transformations in the
 290 data via an invariant latent space search coupled with an explicit canonical orientation normalization
 291 step (Proposition 3.1), which corrects the pose offset introduced by arbitrary canonicals typical of
 292 class-pose decomposition methods like the IE-AE. In principle, however, this approach can be applied
 293 to any method which factors inputs into an invariant component and a symmetry component. For
 294 instance, Quotient Autoencoders (Yokota & Hontani, 2022) also learn canonical representations in a
 295 similar spirit; Marchetti et al. (Marchetti et al., 2023) proposes a class-pose decomposition network
 296 akin to IE-AEs, albeit with a different learning paradigm and advantages; SGMs (Allingham et al.,
 297 2024) also feature a canonical representation or prototype and a relative transformation component.
 298 Our approach remains compatible with such backbones.
 299

300 **Learning symmetries from data** While standard group equivariant networks impose fixed symmetries (Cohen & Welling, 2016; Cohen et al., 2018; Weiler et al., 2018; Weiler & Cesa, 2019; Cohen et al., 2019; Romero et al., 2020; Romero & Cordonnier, 2020; Wang et al., 2020), recent effort have focused on learning symmetries from data. **Approaches like Augerino** (Benton et al., 2020) **learn data augmentations for non-equivariant models**, while others implement relaxed equivariance constraints. **Partial G-CNNs modulate the equivariance per-layer by learning a distribution over the group** (Romero & Lohit, 2022); **Residual Pathway Priors** (Finzi et al., 2021) propose handling partial equivariances through a combination of equivariant and non-equivariant models. These methods require supervision or learn dataset-level transformations rather than instance-specific distributions. Recently, Variational Partial Group Convolutions (VP G-CNNs)(Kim et al., 2024) extended Partial G-CNNs to adapt to instance-level symmetries with a variational inference approach. While they can compute a class-dependent “equivariance error”, this method does not expose a clear symmetry distribution for each input. Equivariance via weight-sharing patterns (Ravanbakhsh et al., 2017; Yeh et al., 2022; Zhou et al., 2021) can also be leveraged to adapt to partial or dataset-level symmetries. Akin to Partial G-CNNs, WSCNNs (van der Linden et al., 2024) introduce layers that can adjust their equivariance based on the data by modulating the weight-sharing pattern. This approach learns a single set of transformations per layer, while our method discovers instance-level symmetry distributions instead.

318 **Another line of work aims to learn a dataset-level symmetry (sub)group $\mathcal{H} \leq \mathcal{G}$ (or generators) of**
 319 **a prescribed ambience group \mathcal{G} from data.** LieGG (Moskalev et al., 2023) extracts infinitesimal
 320 generators from a trained model to reveal learned invariances; LieGAN (Yang et al., 2023) adversarially
 321 learns Lie algebra generators and a dataset-level distribution over coefficients on the group,
 322 discovering subgroups of a prescribed \mathcal{G} without supervision, as well as subsets of the group through
 323 regularization strategies; LaLiGAN (Yang et al., 2024) lifts to a latent space to capture non-linear
 324 actions. Other methods reason through Lie derivatives (Otto et al., 2025) or extend to broader types

324 of symmetries (Forestano et al., 2023; Shaw et al., 2024; Ko et al., 2024). In contrast to these
 325 methods, our approach discovers *instance-specific* distributions of symmetries over a group \mathcal{G} and a
 326 data-aligned canonicalization function without supervision, rather than discovering a dataset-level
 327 (sub)group \mathcal{G} itself. We emphasize that our method is complementary and a continuation to this line
 328 of work: these approaches can be used to infer a suitable \mathcal{G} for all the data, while RECON aims to
 329 infer instance-level distributions and a test-time canonicalization operator.

330 Symmetry-aware Generative Model (SGM) (Allingham et al., 2024) is the most closely related
 331 method in terms of instance-level granularity. SGM uses a flow-based model to learn the relative
 332 distribution of transformations. However, this distribution is relative to an arbitrary canonical, which
 333 presents several disadvantages that RECON addresses (Figures 1, 2, 5b). Note that SGM is a class-
 334 pose decomposition method, and therefore is compatible with RECON. Preliminary work (Urbano &
 335 Romero, 2024) on class-pose methods proposed enforcing a constraint on the group action estimator
 336 via a regularization term, which can be leveraged to learn symmetries, but this often leads to
 337 degenerate solutions where $\psi(x) \approx e$ for all x , limiting symmetry learning. Other methods model
 338 probability distributions over the group. Implicit-PDF (Murphy et al., 2022) models instance-level
 339 distributions over $SO(3)$ with an implicit density and pose supervision; Alignist (Vutukur et al.,
 340 2024) estimates orientation distributions using CAD shape priors and correspondence distributions.
 341 In contrast, RECON requires neither pose labels nor CAD models and operates on general groups
 342 supported by the class-pose backbone.

343 **Test-time canonicalization.** Other works can grant model-agnostic group invariance to pre-trained
 344 models by inserting a small canonicalization in front of them. Spatial Transformer Networks
 345 (STN) (Jaderberg et al., 2015) learn an input-dependent geometric warp via a small network which
 346 can be seen as a group-free form of canonicalization, but it is typically trained end-to-end with
 347 a discriminative objective and requires joint training with the predictor. Equivariant Adaptation
 348 (EquiAdapt) (Mondal et al., 2023) places an equivariant canonicalizer before the frozen predictor
 349 and trains only the canonicalizer using a canonicalization prior; for smaller classifiers they also
 350 propose joint fine-tuning of the predictor. In both cases, training the canonicalization is tied to a
 351 specific pretrained model and its size. In contrast, our canonicalization is trained independently of
 352 the downstream model and can be plugged in front of any model that operates on the same input
 353 domain. In a similar fashion, Affine Steerable EquivarLayer (Li et al., 2025) learn a canonicalizer
 354 trained independently of the pre-trained model, but we note that it is trained using known random
 355 transforms, i.e. pose labels. Our canonicalization, however, is obtained without any supervision.

356 5 EXPERIMENTS

357 We empirically validate RECON on diverse datasets and rotational symmetry distributions. In par-
 358 ticular, we use benchmark image datasets and a large-scale, real-world geometric graphs dataset
 359 exhibiting known ground truth $SO(2)$ and $SO(3)$ rotational symmetries respectively. We focus on
 360 rotation groups since rotations have been recognized as a significant challenge (specially in 3D)
 361 by several studies (Zhao et al., 2020; Shen et al., 2021) even for small ($<15^\circ$) angles (Sun et al.,
 362 2022), but our framework remains general. All our results are on test data, using the learned Θ
 363 and Λ predictors trained with the pseudo-labels obtained through Algorithm 1 (Sec. 3.3). Further
 364 experimental details are provided in Appx. E.

365 5.1 RECOVERING SYMMETRIES

366 **Imaging** We first validate RECON on rotated MNIST and FashionMNIST (Axelrod & Gómez-
 367 Bombarelli, 2022) datasets where different classes exhibit distinct rotational symmetry patterns. For
 368 MNIST, we apply random rotations drawn uniformly from $\pm 60^\circ$ for digits 0-4 and $\pm 90^\circ$ for digits
 369 5-9. For FashionMNIST, we apply rotations drawn from a Gaussian $\mathcal{N}(0, \sigma)$ with $\sigma=0$ for classes
 370 0-2, $\sigma=32$ for classes 3-5 and $\sigma=64$ for classes 6-9. We aim to discover these input-dependent
 371 symmetries from the unlabeled datasets. Canonicalization plots (Fig. 2) show RECON’s consistent,
 372 upright canonicals across classes, compared to other class-pose backbones. In Figure 9 Col.3, we
 373 visually confirm that our normalized distributions match the true sampling regime applied to each
 374 class in the MNIST experiment. Note how this clear distinction and direct interpretability is lost when
 375 analyzing the unnormalized, relative distributions from the class-pose backbone (Figure 9 Col.2).
 376 Figure 5a shows quantitative results for MNIST (top) and FashionMNIST (bottom), confirming that
 377 the predicted distribution parameters (given by $\Theta(x)$) align closely with the ground truth parameters,
 even when these symmetries stem from distributions with different shapes and scales per-class. We

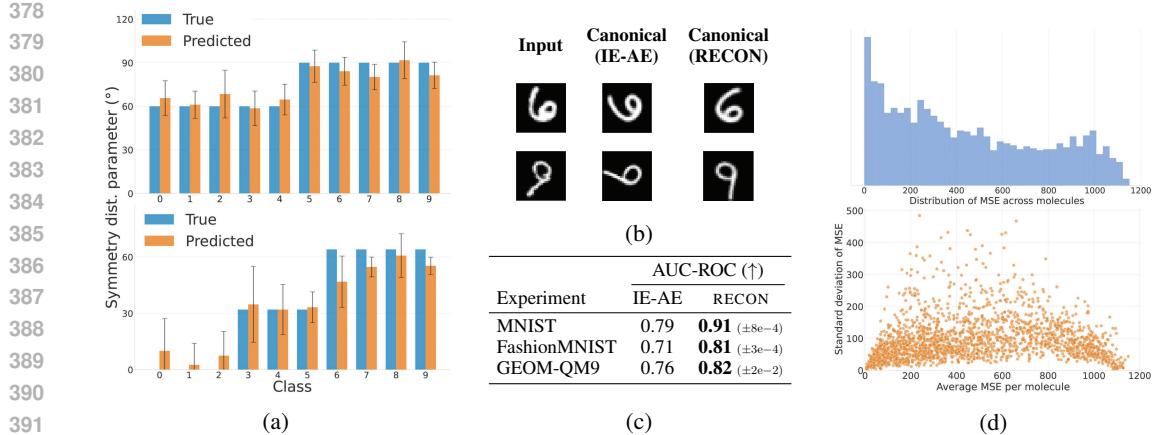


Figure 5: (a) True (■) vs predicted (■) parameter of the label-dependent symmetry distribution in MNIST (top) and FashionMNIST (bottom) experiment. (b) Canonicals for digits of the class 6 (first row) and 9 (second row) obtained with IE-AE and RECON. (c) Out-of-distribution detection performance with IE-AE and RECON. (d) Histogram of average prediction error across molecules and a scatter plot of their standard deviations for GEOM.

obtain these results using $k=10$ neighbors for the class computation $[x]$, which we found to give the best performance (details in App. D.2).

A notable strength of RECON is its ability to pick up on *partial symmetries*. RECON canonicalizations can distinguish between *distinct* classes that are related by a group transformation – such as digits ‘6’ and ‘9’, which are related by a 180° rotation – as opposed to the IE-AE canonicalizations, which collapse both classes into the same canonical⁵ (Fig. 5b). This sensitivity to the data’s contextual orientation allows RECON to handle cases beyond perfect group orbits where full equivariance is inappropriate. We discuss this in detail in Appendix D.4.

Geometric graphs To demonstrate RECON’s ability to discover symmetries in higher-dimensional groups and complex data, we apply it to the Geometric Ensemble of Molecules (GEOM) dataset (Axelrod & Gómez-Bombarelli, 2022), which provides multiple molecular conformations (i.e., unique 3D geometries or conformers) for several classes of molecules. We focus on the the QM9 (Ramakrishnan et al., 2014) subset of GEOM, and to ensure sufficient pose variation in each molecule class, select only molecules with at least 64 distinct low-energy conformers. This results in approximately 175k samples for training, 21k for validation, and 24k for testing across 2, 221 distinct classes of molecules, which can be identified by their SMILES (Anderson et al., 1987) string (more details in Appx. E). We then apply random $SO(3)$ rotations to each molecule class to create ground truth pose variations. Specifically, subsets of molecules are rotated using three distinct matrix-Fisher distributions (Mardia & Jupp, 2009), each centered on a different axis e_1, e_2, e_3 to simulate rotations around a fixed direction. These parameterized distributions $\{\mathcal{M}(F_{true}^1), \mathcal{M}(F_{true}^2), \mathcal{M}(F_{true}^3)\}$ are visualized on the sphere following Mohlin et al. (2020), where brighter regions indicate preferred orientations. For instance, in Fig. 6a, rotations of C[C@H]1(O)C[C@H]1[C@H](O)CO conformers are sampled from $\mathcal{M}(F_{true}^1)$, which perturbs the e_1 axis locally while the orthogonal axes rotate around it freely forming a ring.

Figure 6 evaluates how accurately RECON recovers each molecule’s $SO(3)$ symmetry distribution by measuring per-molecule average MSE between the true and predicted matrix-Fisher parameters. Panels (a-c) visualize randomly sampled molecules from three error regimes – low, average, and high MSE based on percentiles. The panel visualizations show close alignment between the predicted and the true distributions for molecules in the low and average error regime. Figure 5d provides a global view, showing the histogram of MSE and their standard deviations. In general, about 80% of molecules fall within a region of accurate predictions (approx. $MSE < 800$ based on visualizations), demonstrating RECON’s ability to capture diverse $SO(3)$ symmetry patterns from large-scale, unlabeled, realistic 3D data. Interestingly, we observe no clear correlation between MSE and (i) molecular size, (ii) flexibility in conformational variation or (iii) reconstruction error

⁵This limitation can also be observed in SGMs (Allingham et al., 2024).

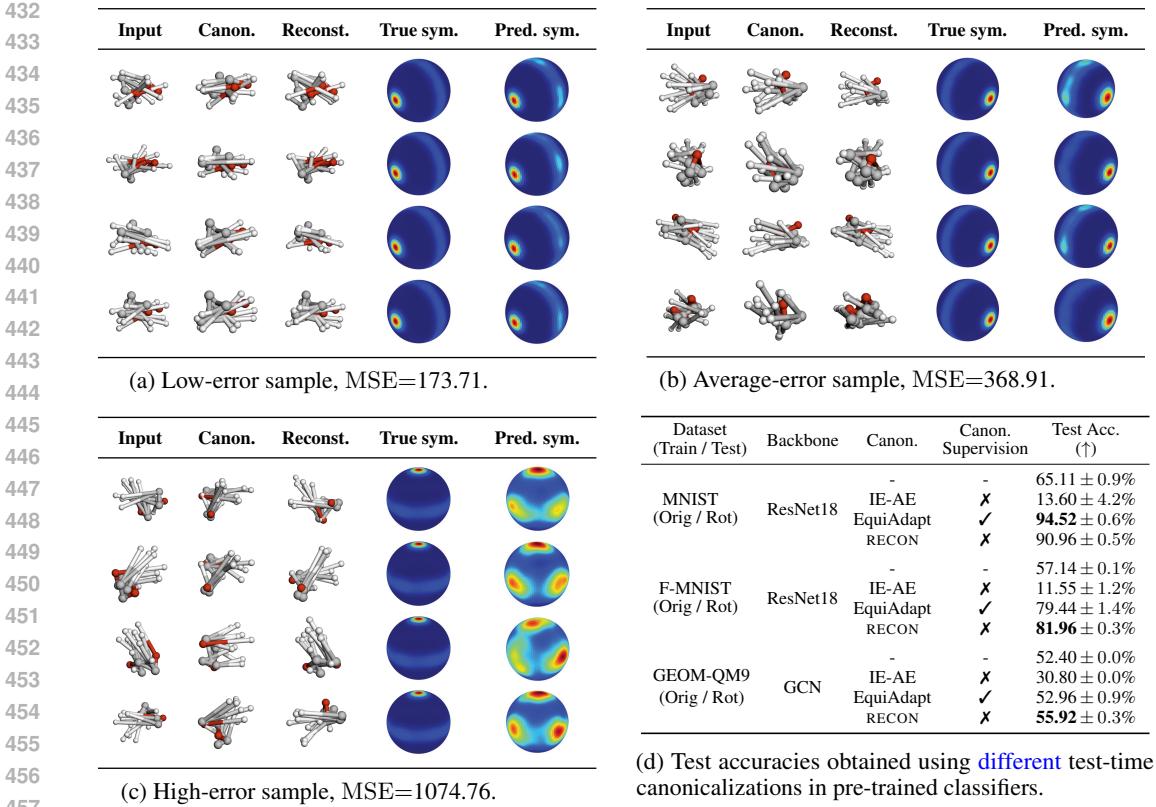


Figure 6: **(a-c)** Test-time comparison for molecules randomly sampled from low error ($\leq p_{10}$ tenth percentile average per-molecule MSE), average error ($p_{45} - p_{55}$), and high error ($\geq p_{90}$) regimes. For each selected molecule, four sample conformers are shown using a ball-and-stick model, displaying: (column 1) input conformer, (2) our canonical reconstruction, (3) final reconstruction, (4) true symmetry distribution (constant for the molecule), and (5) per-conformer predicted symmetry distribution. **(d)** Test-time canonicalization experiment on the per-class rotated dataset variants.

(Figure 14, reconstruction losses defined in Appendix E.2.4). This indicates that our symmetry estimates are robust to conformational complexity, and that the remaining error stems from other sources (see limitations in Sec. 6). RECON also demonstrates remarkable data efficiency, providing accurate symmetry distributions inferences in a dataset where most molecule classes have just a few (<100) number of conformers per class (Fig 13b), highlighting its ability to work with limited per-class data.

5.2 DOWNSTREAM APPLICATIONS

OOD pose detection The centered distributions recovered by RECON yield a natural anomaly score. Given predictors $\Theta(x)$ and $\Lambda(x)$ (Sec. 3.3), consider the absolute pose $g_{\text{abs}} = \psi(x)\Lambda(x)^{-1}$. Then, the likelihood of g_{abs} under the distribution parameterized by $\Theta(x)$ serves as an anomaly score:

$$s(x) = -\log p_{\Theta(x)}(g_{\text{abs}}). \quad (5)$$

Low $s(x)$ indicates in-distribution, while high $s(x)$ flags OOD. We empirically validate this by classifying randomly oriented $SO(2)$ (images) and $SO(3)$ (GEOM) test instances, showing strong OOD detection (AUC-ROC in Table 5, measuring separability between in and OOD predictions). While this score is reference-frame invariant in theory (we provide a proof in Prop. B.2), centering via RECON improves optimization, resulting in consistently higher AUC-ROC scores against IE-AE (full ablation details in Appendix F.1, ROC curves in Fig. 15).

Test-time canonicalization RECON’s data-aligned canonicalizations can be used to create a drop-in layer that grants group invariance to frozen backbones (e.g., classifiers, foundation models) at inference, *with no retraining*. Concretely, we canonicalize inputs at test time via a canonicalization

486 layer defined as $C(x) = \rho_{\mathcal{X}}(\Lambda(x) \cdot \psi(x)^{-1}) x$, and feed the canonicalized sample to the pre-trained
 487 model; because RECON aligns inputs to the training data distribution of poses (visualized in Fig. 1d),
 488 it recovers a large fraction of the accuracy lost to test-time transformations.
 489

490 Figure 6d shows considerable performance gains in classifiers. We improve consistently over the base-
 491 line and offer competitive (MNIST) or better performance (FashionMNIST, GEOM-QM9) against
 492 other test-time canonicalization methods such as EquiAdapt (Mondal et al., 2023), while noting that
 493 our canonicalization is trained without supervision, whereas EquiAdapt’s canonicalization is trained
 494 end-to-end using class labels. Notably, in the imaging domain, our method consistently outperforms
 495 equivariance-learning architectures (SE(2) Partial G-CNNs (Romero & Lohit, 2022)) sometimes by a
 496 large margin (67.72% vs 81.96% in FashionMNIST), and stay competitive (MNIST) or even surpass
 497 (FashionMNIST) fully equivariant architectures (SE(2)-equivariant Steerable CNNs (Cesa et al.,
 498 2022)) (cf. Table 3). Note that arbitrary canonicalizations from class-pose methods (e.g., IE-AE,
 499 SGM) map inputs to OOD poses and therefore degrade performance; benefits from these canonical-
 500 izations arise only when they are also applied at training time, thus requiring backbone retraining.
 501 Frame-averaging approaches face the same practical barrier, typically requiring incorporating the
 502 frame-averaging into training to grant invariance (Puny et al., 2022; Kaba et al., 2023), which limits
 503 adoption due to the added training cost. Our canonicalizations work as a plug-in solution during
 504 inference, facilitating its broader adoption. For further details, refer to Appendix F.2.
 505

506 6 LIMITATIONS AND FUTURE WORK

507 We addressed the challenge of learning instance-specific symmetries from unlabeled data by lever-
 508 aging class-pose decompositions and solving the problem of arbitrary canonicals. In our testing,
 509 RECON proved to be successful in recovering pose distributions at the instance-level – including
 510 demonstrations in 3D data. We also showed how RECON enables downstream practical applications,
 511 notably in test-time canonicalization, which yields clear performance gains and has potential for
 512 broader adoption due to its plug-and-play, architecture agnostic, retrain-free nature.
 513

514 **Limitations** A crucial limitation is that RECON relies on class-pose decompositions, where the
 515 relative transformation is defined with respect to the entire input $x \in \mathcal{X}$. This is effective in certain
 516 domains, but when we use class-pose methods as backbones for symmetry discovery e.g. in natural
 517 images, it exposes a key mismatch: the group action $\rho(g)$ moves the whole scene (input), whereas
 518 the object of interest occupies only a subset (e.g., a car within a street scene). Background and
 519 multi-object context make unsupervised association across instances difficult, hindering recovery of
 520 object-level symmetry distributions. A natural next step is moving from instance-level decomposition
 521 to an *object-centric* variant that predicts pose from token or mask-based features from e.g. strong
 522 self-supervised backbones. This preserves our pipeline’s simplicity while targeting in-the-wild
 523 settings where symmetries are local to objects rather than inputs. Complementary improvements
 524 target sampling and modeling; for instance, replacing nearest neighbors with more sophisticated
 525 sampling strategies. On the density side, more flexible estimators for the symmetry distribution – like
 526 flow-based models in the spirit of Allingham et al. (2024), can better capture complex symmetry
 527 patterns than a parametric family. Finally, relaxing the need to predefine \mathcal{G} by inferring group
 528 structure or group-agnostic transformations (van der Linden et al., 2024) is a promising direction.
 529 These extensions preserve RECON’s simplicity while pushing toward more robust performance.
 530

531 REFERENCES

532 Bijan Afsari. Riemannian l^p center of mass: Existence, uniqueness, and convexity. *Proceedings of
 533 the American Mathematical Society*, 139(02):655–655, February 2011. ISSN 0002-9939. doi:
 534 10.1090/S0002-9939-2010-10541-5.

535 James U. Allingham, Bruno K. Mlodzeniec, Shreyas Padhy, Javier Antorán, David Krueger,
 536 Richard E. Turner, Eric Nalisnick, and José M. Hernández-Lobato. A generative model of
 537 symmetry transformations. *Advances in Neural Information Processing Systems*, 37:91091–91130,
 538 December 2024.

539 E. Anderson, Gd Veith, and David Weininger. Smiles: a line notation and computerized interpreter
 540 for chemical structures., 1987. URL <https://api.semanticscholar.org/CorpusID:64884759>.

540 Andrea Aveni and Sayan Mukherjee. Uniform consistency of generalized fréchet means, August
 541 2024. URL <http://arxiv.org/abs/2408.07534>. arXiv:2408.07534 [math].
 542

543 Simon Axelrod and Rafael Gómez-Bombarelli. GEOM, energy-annotated molecular conformations
 544 for property prediction and molecular generation. *Sci. Data*, 9(1):185, April 2022.

545 Gregory W. Benton, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. Learning invariances
 546 in neural networks. *CoRR*, abs/2010.11882, 2020.

547 Lukas Biewald. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.

548 Johannes Brandstetter, Rob Hesselink, Elise van der Pol, Erik J. Bekkers, and Max Welling. Geometric
 549 and physical quantities improve $e(3)$ equivariant message passing, March 2022. URL <http://arxiv.org/abs/2110.02905>. arXiv:2110.02905 [cs].
 550

551 Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning:
 552 Grids, groups, graphs, geodesics, and gauges, May 2021. URL <http://arxiv.org/abs/2104.13478>. arXiv:2104.13478 [cs].
 553

554 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and
 555 Armand Joulin. Emerging properties in self-supervised vision transformers, May 2021. URL
 556 <http://arxiv.org/abs/2104.14294>. arXiv:2104.14294 [cs].
 557

558 Gabriele Cesa, Leon Lang, and Maurice Weiler. A program to build $E(N)$ -equivariant steerable CNNs.
 559 In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=WE4qe9xlnQw>.
 560

561 Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference
 562 on machine learning*, pp. 2990–2999. PMLR, 2016.

563 Taco Cohen, Mario Geiger, and Maurice Weiler. A general theory of equivariant cnns on homogeneous
 564 spaces. *CoRR*, abs/1811.02017, 2018.

565 Taco S. Cohen, Maurice Weiler, Berkay Kicanaoglu, and Max Welling. Gauge equivariant convolutional
 566 networks and the icosahedral CNN. *CoRR*, abs/1902.04615, 2019.

567 Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvassy, Pierre-Emmanuel
 568 Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The faiss library, 2024.
 569

570 Ilya Feige. Invariant-equivariant representation learning for multi-class data, 2019.

571 Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In
 572 *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

573 Marc Finzi, Gregory W. Benton, and Andrew Gordon Wilson. Residual pathway priors for soft
 574 equivariance constraints. *CoRR*, abs/2112.01388, 2021.

575 Roy T Forestano, Konstantin T Matchev, Katia Matcheva, Alexander Roman, Eyup B Unlu, and
 576 Sarunas Verner. Deep learning symmetries and their lie groups, algebras, and subalgebras from first
 577 principles. *Machine Learning: Science and Technology*, 4(2):025027, June 2023. ISSN 2632-2153.
 578 doi: 10.1088/2632-2153/acd989.
 579

580 Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena
 581 Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar,
 582 et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural
 583 information processing systems*, 33:21271–21284, 2020.

584

585 Xifeng Guo, En Zhu, Xinwang Liu, and Jianping Yin. Affine equivariant autoencoder. In *Proceedings
 586 of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pp.
 587 2413–2419. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi:
 588 10.24963/ijcai.2019/335.
 589

594 Irina Higgins, David Amos, David Pfau, Sébastien Racanière, Loïc Matthey, Danilo J. Rezende, and
 595 Alexander Lerchner. Towards a definition of disentangled representations. *CoRR*, abs/1812.02230,
 596 2018.

597 Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer
 598 networks. *CoRR*, abs/1506.02025, 2015.

600 Sékou-Oumar Kaba, Arnab Kumar Mondal, Yan Zhang, Yoshua Bengio, and Siamak Ravanbakhsh.
 601 Equivariance with learned canonicalization functions. In *International Conference on Machine
 602 Learning*, pp. 15546–15566. PMLR, 2023.

603 H. Karcher. Riemannian center of mass and mollifier smoothing. *Communications on Pure and
 604 Applied Mathematics*, 30(5):509–541, 1977. ISSN 1097-0312. doi: 10.1002/cpa.3160300502.

606 Hyunsu Kim, Yegon Kim, Hongseok Yang, and Juho Lee. Variational partial group convolutions for
 607 input-aware partial equivariance of rotations and color-shifts, July 2024. URL <http://arxiv.org/abs/2407.04271>. arXiv:2407.04271 [cs].

609 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.

611 Gyeonghoon Ko, Hyunsu Kim, and Juho Lee. Learning infinitesimal generators of continuous
 612 symmetries from data, 2024.

614 Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch– gordan nets: a fully
 615 fourier space spherical convolutional neural network. In *Advances in Neural In-
 616 formation Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL
 617 https://proceedings.neurips.cc/paper_files/paper/2018/hash/a3fc981af450752046be179185ebc8b5-Abstract.html.

619 Kaushik Koneripalli, Suhas Lohit, Rushil Anirudh, and Pavan Turaga. Rate-invariant autoencoding
 620 of time-series. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and
 621 Signal Processing (ICASSP)*, pp. 3732–3736, 2020. doi: 10.1109/ICASSP40776.2020.9053983.

622 Adam R. Kosiorek, Sara Sabour, Yee Whye Teh, and Geoffrey E. Hinton. Stacked capsule autoen-
 623 coders, 2019.

625 Greg Landrum, Paolo Tosco, Brian Kelley, Ricardo Rodriguez, David Cosgrove, Riccardo Vianello,
 626 sriniker, Peter Gedeck, Gareth Jones, NadineSchneider, Eisuke Kawashima, Dan Nealschneider,
 627 Andrew Dalke, Matt Swain, Brian Cole, Samo Turk, Aleksandr Savelev, tadhurst-cdd, Alain
 628 Vaucher, Maciej Wójcikowski, Ichiru Take, Rachel Walker, Vincent F Scalfani, Kazuya Ujihara,
 629 Daniel Probst, Juuso Lehtivarjo, Hussein Faara, Guillaume Godin, Axel Pahl, and Jeremy Monat.
 630 rdkit/rdkit: 2025_03_1 (q1 2025) beta release, 2025.

631 Yikang Li, Yeqing Qiu, Yuxuan Chen, and Zhouchen Lin. Affine steerable equivariant layer for
 632 canonicalization of neural networks. In *The Thirteenth International Conference on Learning
 633 Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL <https://openreview.net/forum?id=5i6ZZUjCA9>.

635 Giovanni Luca Marchetti, Gustaf Tegnér, Anastasiia Varava, and Danica Kragic. Equivariant rep-
 636 resentation learning via class-pose decomposition. In *Proceedings of The 26th International
 637 Conference on Artificial Intelligence and Statistics*, pp. 4745–4756. PMLR, April 2023. URL
 638 <https://proceedings.mlr.press/v206/marchetti23b.html>.

640 K.V. Mardia and P.E. Jupp. *Directional Statistics*. Wiley Series in Probability and Statistics.
 641 Wiley, 2009. ISBN 978-0-470-31781-5. URL <https://books.google.de/books?id=PTNiCm4Q-M0C>.

643 David Mohlin, Josephine Sullivan, and Gérald Bianchi. Probabilistic orientation estimation with
 644 matrix fisher distributions. In *Advances in Neural Information Processing Systems*, volume 33, pp.
 645 4884–4893. Curran Associates, Inc., 2020.

647 Arnab Kumar Mondal, Siba Smarak Panigrahi, Sekouba Kaba, Sai Rajeswar, and Siamak Ravan-
 648 bakhsh. Equivariant adaptation of large pretrained models. *ArXiv*, abs/2310.01647, 2023.

648 Artem Moskalev, Anna Sepliarskaia, Ivan Sosnovik, and Arnold Smeulders. Liegg: Studying
 649 learned lie group generators, January 2023. URL <http://arxiv.org/abs/2210.04345>.
 650 arXiv:2210.04345 [cs].

651 Kieran Murphy, Carlos Esteves, Varun Jampani, Srikumar Ramalingam, and Ameesh Makadia.
 652 Implicit-pdf: Non-parametric representation of probability distributions on the rotation manifold,
 653 2022. URL <http://arxiv.org/abs/2106.05965>. arXiv:2106.05965 [cs].

654 Samuel E. Otto, Nicholas Zolman, J. Nathan Kutz, and Steven L. Brunton. A unified framework
 655 to enforce, discover, and promote symmetry in machine learning, 2025. URL <http://arxiv.org/abs/2311.00212>. arXiv:2311.00212 [cs].

656 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
 657 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward
 658 Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang,
 659 Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning
 660 library, 2019.

661 Xavier Pennec. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements.
 662 *Journal of Mathematical Imaging and Vision*, 25(1):127–154, July 2006. ISSN 0924-9907, 1573-
 663 7683. doi: 10.1007/s10851-006-6228-4.

664 Omri Puny, Matan Atzmon, Heli Ben-Hamu, Ishan Misra, Aditya Grover, Edward J. Smith, and
 665 Yaron Lipman. Frame averaging for invariant and equivariant network design, March 2022. URL
 666 <http://arxiv.org/abs/2110.03336>. arXiv:2110.03336 [cs].

667 Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole von Lilienfeld. Quantum
 668 chemistry structures and properties of 134 kilo molecules. *Sci. Data*, 1(1):140022, August 2014.

669 Siamak Ravanbakhsh, Jeff Schneider, and Barnabas Poczos. Equivariance through parameter-sharing,
 670 June 2017. URL <http://arxiv.org/abs/1702.08389>. arXiv:1702.08389 [stat].

671 David W. Romero and Jean-Baptiste Cordonnier. Group equivariant stand-alone self-attention for
 672 vision. *CoRR*, abs/2010.00977, 2020.

673 David W Romero and Suhas Lohit. Learning partial equivariances from data. *Advances in Neural
 674 Information Processing Systems*, 35:36466–36478, 2022.

675 David W. Romero, Erik J. Bekkers, Jakub M. Tomczak, and Mark Hoogendoorn. Attentive group
 676 equivariant convolutional networks. *CoRR*, abs/2002.03830, 2020.

677 Peter J. Rousseeuw and Mia Hubert. Robust statistics for outlier detection. *WIREs Data Mining and
 678 Knowledge Discovery*, 1(1):73–79, January 2011. ISSN 1942-4787, 1942-4795. doi: 10.1002/
 679 widm.2.

680 Ben Shaw, Abram Magner, and Kevin R Moon. Symmetry discovery beyond affine transformations,
 681 2024.

682 Wen Shen, Qihan Ren, Dongrui Liu, and Quanshi Zhang. Interpreting representa-
 683 tion quality of dnns for 3d point cloud processing. In *Advances in Neural Infor-
 684 mation Processing Systems*, volume 34, pp. 8857–8870. Curran Associates, Inc.,
 685 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/4a3e00961a08879c34f91ca0070ea2f5-Abstract.html>.

686 Georgy L. Shevlyakov, Stephan Morgenthaler, and Alexander M. Shurygin. Redescending m-
 687 estimators. *Journal of Statistical Planning and Inference*, 138:2906–2917, 2008. URL <https://api.semanticscholar.org/CorpusID:121151493>.

688 Ha-Young Shin and Hee-Seok Oh. Robust geodesic regression, January 2022. URL <http://arxiv.org/abs/2007.04518>. arXiv:2007.04518 [stat].

689 Zhixin Shu, Mihir Sahasrabudhe, Riza Alp Güler, Dimitris Samaras, Nikos Paragios, and Iasonas
 690 Kokkinos. Deforming autoencoders: Unsupervised disentangling of shape and appearance. *CoRR*,
 691 abs/1806.06503, 2018.

702 Jiachen Sun, Qingzhao Zhang, Bhavya Kailkhura, Zhiding Yu, Chaowei Xiao, and Z. Morley Mao.
 703 Benchmarking robustness of 3d point cloud recognition against common corruptions, January
 704 2022. URL <http://arxiv.org/abs/2201.12296>. arXiv:2201.12296 [cs].
 705

706 Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley.
 707 Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds,
 708 May 2018. URL <http://arxiv.org/abs/1802.08219>. arXiv:1802.08219 [cs].
 709

710 Alonso Urbano and David W. Romero. Self-supervised detection of perfect and partial input-
 711 dependent symmetries. In *Proceedings of the Geometry-grounded Representation Learning*
 712 and *Generative Modeling Workshop (GRaM)*, volume 251 of *Proceedings of Machine Learning*
 713 *Research*, pp. 113–131. PMLR, 29 Jul 2024. URL <https://proceedings.mlr.press/v251/urbano24a.html>.
 714

715 Putri A. van der Linden, Alejandro García-Castellanos, Sharvaree Vadgama, Thijs P. Kuipers, and
 716 Erik J. Bekkers. Learning symmetries via weight-sharing with doubly stochastic tensors. *Advances*
 717 in *Neural Information Processing Systems*, 37:42738–42764, December 2024.
 718

719 Shishir Reddy Vutukur, Rasmus Laurvig Haugaard, Junwen Huang, Benjamin Busam, and Tolga
 720 Birdal. Alignist: Cad-informed orientation distribution estimation by fusing shape and correspon-
 721 dences, 2024. URL <http://arxiv.org/abs/2409.06683>. arXiv:2409.06683 [cs].
 722

723 Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant
 724 attention mechanism for weakly supervised semantic segmentation. *CoRR*, abs/2004.04581, 2020.
 725

726 Maurice Weiler and Gabriele Cesa. General e(2)-equivariant steerable cnns. *CoRR*, abs/1911.08251,
 727 2019.
 728

729 Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco Cohen. 3d steerable cnns:
 730 Learning rotationally equivariant features in volumetric data. *CoRR*, abs/1807.02547, 2018.
 731

732 Robin Winter, Frank Noé, and Djork-Arné Clevert. Permutation-invariant variational autoencoder for
 733 graph-level representation learning. *CoRR*, abs/2104.09856, 2021.
 734

735 Robin Winter, Marco Bertolini, Tuan Le, Frank Noé, and Djork-Arné Clevert. Unsupervised learning
 736 of group invariant and equivariant representations. *Advances in Neural Information Processing*
 737 *Systems*, 35:31942–31956, 2022.
 738

739 Andrew Wood. Estimation of the concentration parameters of the fisher matrix distribution on 50(3)
 740 and the bingham distribution on sq, q >= 2. *Australian Journal of Statistics*, 35:69 – 79, 06 2008.
 741 doi: 10.1111/j.1467-842X.1993.tb01313.x.
 742

743 Jianke Yang, Robin Walters, Nima Dehmamy, and Rose Yu. Generative adversarial symmetry
 744 discovery, 2023. URL <http://arxiv.org/abs/2302.00236>. arXiv:2302.00236 [cs].
 745

746 Jianke Yang, Nima Dehmamy, Robin Walters, and Rose Yu. Latent space symmetry discovery,
 747 August 2024. URL <http://arxiv.org/abs/2310.00105>. arXiv:2310.00105 [cs].
 748

749 Raymond A. Yeh, Yuan-Ting Hu, Mark Hasegawa-Johnson, and Alexander Schwing. Equivariance
 750 discovery by learned parameter-sharing. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel
 751 Valera (eds.), *Proceedings of The 25th International Conference on Artificial Intelligence and*
 752 *Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 1527–1545. PMLR,
 753 28–30 Mar 2022. URL <https://proceedings.mlr.press/v151/yeh22b.html>.
 754

755 Tatsuya Yokota and Hidekata Hontani. Manifold modeling in quotient space: Learning an invariant
 756 mapping with decodability of image patches, 2022.
 757

758 Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 759 effectiveness of deep features as a perceptual metric, April 2018. URL <http://arxiv.org/abs/1801.03924>. arXiv:1801.03924 [cs].

756 Yue Zhao, Yuwei Wu, Caihua Chen, and Andrew Lim. On isometry robustness of deep 3d point
757 cloud models under adversarial attacks. In *2020 IEEE/CVF Conference on Computer Vision*
758 and *Pattern Recognition (CVPR)*, pp. 1198–1207, Seattle, WA, USA, June 2020. IEEE. ISBN
759 978-1-7281-7168-5. doi: 10.1109/CVPR42600.2020.00128. URL <https://ieeexplore.ieee.org/document/9157546/>.

760

761 Allan Zhou, Tom Knowles, and Chelsea Finn. Meta-learning symmetries by reparameterization,
762 2021. URL <https://arxiv.org/abs/2007.02933>.

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810
811

A BACKGROUND

812
813
814
815

Groups and group actions. A group \mathcal{G} is a set equipped with a closed, associative binary operation \cdot such that \mathcal{G} contains an identity element $e \in \mathcal{G}$ and every element $g \in \mathcal{G}$ has an inverse $g^{-1} \in \mathcal{G}$. For a given set \mathcal{X} and group \mathcal{G} , the (left) group action of \mathcal{G} on \mathcal{X} is a map $\rho : \mathcal{G} \times \mathcal{X} \rightarrow \mathcal{X}$ that preserves the group structure. Intuitively, it describes how set elements transform by group elements.

816
817
818
819
820
821
822

Group representations. In this work, we focus on cases where \mathcal{X} is a vector space. In such scenarios, the group acts on it by means of *group representations*. A representation of a group \mathcal{G} on a vector space \mathcal{X} is a homomorphism $\rho_{\mathcal{X}} : \mathcal{G} \rightarrow \text{GL}(\mathcal{X})$ mapping each $g \in \mathcal{G}$ to an invertible linear operator on \mathcal{X} . We consider our dataset to be of the form $\mathcal{X} = \{f \mid f : \mathcal{V} \rightarrow \mathcal{W}\}$, where \mathcal{V} is a set on which \mathcal{G} acts, and \mathcal{W} is a vector space on which \mathcal{G} may also act. As an example, molecular conformations can be interpreted as functions $f : \mathcal{V} \rightarrow \mathbb{R}^3$ that map atoms to their spatial coordinates. Following this definition, a group element acts in a data sample as

823
824
$$[\rho_{\mathcal{X}}(g)f](x) \equiv \rho_{\mathcal{W}}(g)f(\rho_{\mathcal{V}}(g^{-1})x). \quad (6)$$

825
826
827
828

For instance, a rigid transformation g might shift or rotate node coordinates via $\rho_{\mathcal{W}}(g)$ while leaving the nodes unchanged or transforming them via $\rho_{\mathcal{V}}(g)$. Whenever we speak of a representation $\rho_{\mathcal{X}}$ on \mathcal{X} , it is understood that we are implicitly referring to the previous equation to understand the transformation of each component.

829
830

Orbits. The orbit of x , $\mathcal{O}_x = \{\rho_{\mathcal{X}}(g)x\}_{g \in \mathcal{G}}$ captures all possible transformations of x resulting from the action of all elements of \mathcal{G} .

831
832
833
834
835

Equivalence classes and quotient sets. Our analysis relies on the definition of equivalence classes and their quotient sets. Let \sim be an equivalence relation on \mathcal{X} and consider the *equivalence classes* $[x] = \{y \in \mathcal{X}, \text{ s.t. } x \sim y\}$ of \mathcal{X} . The quotient set \mathcal{X}/\sim is defined as the collection of all equivalent classes in \mathcal{X} under the relation \sim .

836
837
838
839
840
841
842

Group equivariance and group invariance. A map $h : \mathcal{V} \rightarrow \mathcal{W}$ is \mathcal{G} -equivariant with respect to the representations $\rho_{\mathcal{V}}, \rho_{\mathcal{W}}$ if $h(\rho_{\mathcal{V}}(g)x) = \rho_{\mathcal{W}}(g)h(x) \forall g \in \mathcal{G}, \forall x \in \mathcal{X}$. In the context of neural networks, G-CNNs (Cohen & Welling, 2016) are designed to be \mathcal{G} -equivariant by using only \mathcal{G} -equivariant layers in their constructions. This ensures that applying a transformation $g \in \mathcal{G}$ before a layer yields an equivalently transformed output. Analogously, a map h is \mathcal{G} -invariant with respect to $\rho_{\mathcal{V}}$ if $h(\rho_{\mathcal{V}}(g)x) = h(x) \forall g \in \mathcal{G}, \forall x \in \mathcal{X}$. That is, if \mathcal{G} -transformations of the input yield the same result.

843
844

B PROOFS

845
846
847
848
849
850
851
852
853

Proposition B.1 (Approximation of $\mu_{[x]}$ via Normalization). *Let \mathcal{X} be a metric space, \mathcal{G} a Lie group and η, δ, ψ an IE-AE where ψ is continuous on a compact domain \mathcal{X} . Suppose that \mathcal{X} exhibits symmetries characterized by $\mu_{[x]}$ where $\mathcal{F}(\mu_{[x]}) = e$ as described above. Consider a random sample $\{s_i\}_{i=1}^N$ of $[x]$ and denote their images by ψ as $\psi([x]) = \{\psi(s_i)\}_{i=1}^N$. Let $\hat{\Gamma}_{[x]} \in \mathcal{G}$ be the empirical Fréchet mean of $\psi([x])$. Then, the empirical distribution $\hat{\mu}_{[x]}$ corresponding to the normalized samples $\psi([x])\hat{\Gamma}_{[x]}^{-1}$ approximates the target distribution $\mu_{[x]}$. Specifically, $\hat{\mu}_{[x]}$ converges in Wasserstein distance to $\mu_{[x]}$ as $\varepsilon' \rightarrow 0$ and $N \rightarrow \infty$.*

854
855
856
857
858
859
860
861

Proof. Consider N independent samples $\{s_i\}_{i=1}^N$ drawn from class $[x]$ according to the model $s_i = \rho_{\mathcal{X}}(g_i)\gamma_{[x]} + \varepsilon_i$, where $g_i \sim \mu_{[x]}$ i.i.d. and $\varepsilon_i \sim P_{\varepsilon}$ i.i.d, with $\|\varepsilon_i\| < \varepsilon'$ almost surely. Let $\psi([x]) = \{\psi(s_i)\}_{i=1}^N$ be the set of observed relative transformations and $\nu_{[x]}$ be the distribution over \mathcal{G} from which $\psi(s_i)$ are sampled. Our first goal is to show that $\nu_{[x]}$ and $\mu_{[x]}$ differ by some translation $\Gamma_{[x]} \in \mathcal{G}$.

862
863

Let $s'_i = \rho_{\mathcal{X}}(g_i)\gamma_{[x]}$ be the noise-free (i.e., $\varepsilon_i = 0$) counterparts of s_i , and consider their images $\psi(s'_i)$. Denote $\nu'_{[x]}$ as the corresponding distribution over \mathcal{G} from which $\psi(s'_i)$ are sampled. Then,

864
865
$$\psi(s'_i) = g_i\psi(\gamma_{[x]}) = g_i\Gamma_{[x]} \quad (7)$$

\downarrow
 $\psi \text{ } \mathcal{G}-\text{eq}$

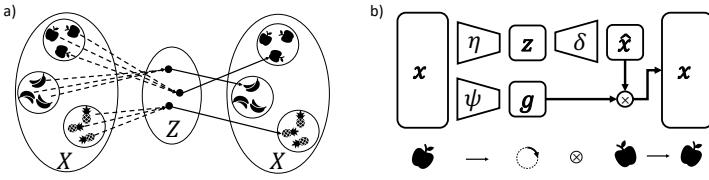


Figure 7: IE-AE architecture visualization. Image taken from (Winter et al., 2022).

for $\Gamma_{[x]} := \psi(\gamma_{[x]}) \in \mathcal{G}$. Since $g_i \sim \mu_{[x]}$, the distribution $\nu'_{[x]}$ is exactly the right-translate $\mu_{[x]}\Gamma_{[x]}$. That is, $\nu'_{[x]} = \mu_{[x]}\Gamma_{[x]}$. This relates the noise-free $\nu'_{[x]}$ and $\mu_{[x]}$. Now, we will show that this relation also holds approximately for $\nu_{[x]}$.

By the uniform continuity of ψ (which holds since \mathcal{X} is compact and ψ is continuous), there exists a modulus of continuity ω such that $\lim_{\delta \rightarrow 0} \omega(\delta) = 0$ and for $s_i, s'_i \in \mathcal{X}$, it holds:

$$d_{\mathcal{G}}(\psi(s_i), \psi(s'_i)) \leq \omega(d_{\mathcal{X}}(s_i, s'_i)) = \omega(\|s_i - s'_i\|) < \omega(\varepsilon') = K, \quad (8)$$

for some small $K > 0$, since ε' is presumably small and $\lim_{\delta \rightarrow 0} \omega(\delta) = 0$. This establishes pointwise closeness almost surely.

Consider the p -Wasserstein distance between both probability distributions, defined as

$$W_p(\nu_{[x]}, \nu'_{[x]}) = \inf_{\pi \in \Pi(\nu_{[x]}, \nu'_{[x]})} \left(\int_{\mathcal{G} \times \mathcal{G}} d_{\mathcal{G}}(g_1, g_2)^p d\pi(g_1, g_2) \right)^{1/p}, \quad (9)$$

where $\Pi(\nu_{[x]}, \nu'_{[x]})$ is the set of all possible couplings (joint probability measures) on $\mathcal{G} \times \mathcal{G}$ with marginals $\nu_{[x]}$ and $\nu'_{[x]}$. Consider the natural coupling π_{nat} determined by the joint random variable pair $(\psi(s_i), \psi(s'_i))$. π_{nat} is a valid coupling in $\Pi(\nu_{[x]}, \nu'_{[x]})$, since its marginals are $\nu_{[x]}$ and $\nu'_{[x]}$ respectively. Since the Wasserstein metric is defined as the minimum over all possible couplings, it must hold that

$$W_p(\nu_{[x]}, \nu'_{[x]}) \leq \left(\int_{\mathcal{G} \times \mathcal{G}} d_{\mathcal{G}}(g_1, g_2)^p d\pi_{nat}(g_1, g_2) \right)^{1/p} = \quad (10)$$

$$= \left(\int_{\mathcal{G} \times \mathcal{G}} d_{\mathcal{G}}(\psi(s_i), \psi(s'_i))^p d\pi_{nat}(g_1, g_2) \right)^{1/p} < \left(\int_{\mathcal{G} \times \mathcal{G}} \omega(\varepsilon')^p d\pi_{nat}(g_1, g_2) \right)^{1/p} = \omega(\varepsilon'). \quad (11)$$

This inequality shows that the distribution $\nu_{[x]}$ generating our samples $\psi([x])$ is close to the distribution $\nu'_{[x]} = \mu_{[x]}\Gamma_{[x]}$ in Wasserstein distance, with proximity controlled by the noise bound ε' via the modulus of continuity ω . Let's show now how to estimate the unknown $\Gamma_{[x]}$.

First, note that we can calculate the Fréchet mean of $\nu'_{[x]}$ as follows:

$$\mathcal{F}(\nu'_{[x]}) = \mathcal{F}(\mu_{[x]})\Gamma_{[x]} = e\Gamma_{[x]} = \Gamma_{[x]}, \quad (12)$$

where we used the property that the Fréchet mean is equivariant under isometries (such as right-translation by $\Gamma_{[x]}$ when using a right-invariant metric), $\mathcal{F}(\mu\Gamma) = \mathcal{F}(\mu)\Gamma$ (Karcher, 1977). This means that the unknown $\Gamma_{[x]}$ is the Fréchet mean of $\nu'_{[x]}$. However, we just proved that $\nu'_{[x]}$ and $\nu_{[x]}$ are close, bounded by $\omega(\varepsilon')$ (equation 10). We want to prove now that their Fréchet means are also close.

Consider the Fréchet mean $\mathcal{F}(\nu_{[x]})$ of the actual data distribution. The Fréchet functional is known to be strictly convex within geodesic balls of a certain radius r_0 determined by the manifold's geometry (Karcher, 1977). Within such regions, the Fréchet mean map $\mu \mapsto \mathcal{F}(\mu)$ is Lipschitz continuous with respect to the Wasserstein distance (Afsari, 2011). Therefore, exists a constant $C_L > 0$ such that

$$d_{\mathcal{G}}(\mathcal{F}(\nu_{[x]}), \mathcal{F}(\nu'_{[x]})) \leq C_L W_p(\nu_{[x]}, \nu'_{[x]}) < C_L \omega(\varepsilon'). \quad (13)$$

Substituting $\mathcal{F}(\nu'_{[x]}) = \Gamma_{[x]}$, we have $d_{\mathcal{G}}(\mathcal{F}(\nu_{[x]}), \Gamma_{[x]}) < C_L \omega(\varepsilon')$. This shows the true Fréchet mean of the distribution we sample from is close to $\Gamma_{[x]}$. Therefore, we can estimate $\Gamma_{[x]}$ by obtaining the population Fréchet mean of $\nu_{[x]}$.

Let's show that using the sample Fréchet mean of $\psi([x])$ estimates the true Fréchet mean. In effect, the sample Fréchet mean is a statistically consistent estimator of the population Fréchet mean $\mathcal{F}(\nu_{[x]})$ (Aveni & Mukherjee, 2024):

$$\hat{\Gamma}_{[x]} \xrightarrow{P} \mathcal{F}(\nu_{[x]}) \quad \text{as } N \rightarrow \infty. \quad (14)$$

Combining all the above, we see that for large N and small ε' , $\hat{\Gamma}_{[x]}$ provides a good approximation of $\Gamma_{[x]}$.

Finally, consider the Fréchet-normalized empirical measure $\hat{\mu} = \frac{1}{N} \sum_{i=1}^N \delta_{\psi(s_i)\hat{\Gamma}_{[x]}^{-1}}$. By the continuous mapping theorem, and because group inversion and multiplication are continuous operations in the Lie group, it follows that $\hat{\mu}$ converges to $\mu_{[x]}$ in W_p distance. \square

Proposition B.2 (Right-translation invariance of the log density). *Let \mathcal{G} be a Lie group with a right Haar measure λ . Let μ be a probability measure on \mathcal{G} absolutely continuous w.r.t. λ , with density $p_\mu = \frac{d\mu}{d\lambda} \in L^1(\lambda)$. Fix $\gamma \in \mathcal{G}$ and let $r_\gamma(g) = g\gamma$. If $\nu = (r_\gamma)_*\mu$, then*

$$p_\nu(h) = p_\mu(h\gamma^{-1}) \quad \text{for almost every } h \in \mathcal{G}. \quad (15)$$

Consequently,

$$-\log p_\nu(h) = -\log p_\mu(h\gamma^{-1}) \quad \text{for almost every } h \in \mathcal{G}, \quad (16)$$

interpreted in the extended reals with the convention $-\log 0 = +\infty$.

Proof. For any bounded measurable $f : \mathcal{G} \rightarrow \mathbb{R}$,

$$\int f(h) d\nu(h) = \int f(r_\gamma(g)) d\mu(g) = \int f(r_\gamma(g)) p_\mu(g) d\lambda(g). \quad (17)$$

Set $h = g\gamma$. Right invariance of λ gives $d\lambda(h) = d\lambda(g)$, hence

$$\int f(h) d\nu(h) = \int f(h) p_\mu(h\gamma^{-1}) d\lambda(h). \quad (18)$$

By uniqueness of Radon-Nikodym derivatives, $p_\nu(h) = p_\mu(h\gamma^{-1})$ for almost every h . Taking $-\log$ yields equation 16 (with $-\log 0 = +\infty$). \square

C ROBUST FRÉCHET MEAN EXTENSION FOR MULTIMODAL DISTRIBUTIONS

While the Fréchet mean provides a robust and well-defined centroid for many distributions, as discussed above, it may fall outside the support in certain multimodal cases, leading to non-natural canonical poses that do not align with the data's intrinsic symmetries. For instance, consider a bimodal distribution on $SO(2)$ with identical modes at 0 and π radians. Here, the Fréchet mean minimizes the sum of squared geodesic distances and, due to the symmetry, converges to a point midway between the modes ($\pm\pi/2$), which lies in a region of near-zero probability density, outside the effective support of the distribution. This violates the desired natural (in-distribution) reference pose, as it may correspond to an orientation not observed in the data, potentially degrading downstream applications like invariance granting or interpretability.

To address this limitation while preserving the generality of the Fréchet mean, we propose an extension using *redescending M-estimators* (Shevlyakov et al., 2008; Rousseeuw & Hubert, 2011). These estimators replace the squared loss with a bounded, redescending function that caps the influence of distant points, providing robustness to outliers or separated modes. We specifically propose using the Tukey biweight loss (Rousseeuw & Hubert, 2011; Shin & Oh, 2022):

$$m(u; c) = \begin{cases} \frac{c^2}{6} \left(1 - \left(1 - \left(\frac{u}{c} \right)^2 \right)^3 \right) & \text{if } u \leq c \\ \frac{c^2}{6} & \text{if } u > c \end{cases} \quad (19)$$

In the standard Fréchet mean, we minimize $\sum_i d(y, g_i)^2$, which is unbounded quadratic in distances; large distances (outliers or far modes) have disproportionate pull because their squared penalty grows without limit. However, we can provide a robust version by minimizing $\sum_i m(d(y, g_i); c)$ instead.

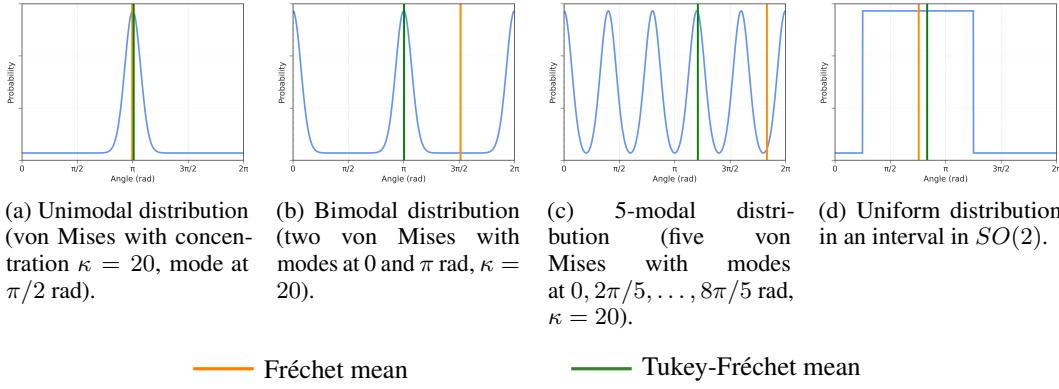


Figure 8: Comparison of Fréchet mean and robust Tukey-Fréchet mean variation on $SO(2)$ distributions. Plots generated by sampling $n = 2000$ points from each distribution, computing the means, and overlaying on the theoretical density. The robust estimator localizes to high-density regions in multimodal cases, while converging to the standard Fréchet mean in unimodal/uniform scenarios.

In effect, the robustness arises from the *bounded influence* of the Tukey biweight loss; m is such that for small u , it holds that $m(u) \approx u^2/2$, behaving like the Fréchet mean locally (quadratic). However, for large distances $u > c$, m is constant at $c^2/6$, assigning zero gradient to distant points:

$$\frac{dm}{du}(u; c) = \begin{cases} u \left(1 - \left(\frac{u}{c}\right)^2\right)^2 & \text{if } |u| \leq c \\ 0 & \text{if } |u| > c \end{cases} \quad (20)$$

The function dm/du , determines how much each sample affects the final estimate during optimization (we use gradient descent to solve the minimization problem as in (Shin & Oh, 2022)). Therefore, in a multimodal distribution with well-separated clusters, large distances ($|u| > c$) have zero influence on the gradient/update, and the objective effectively ignores points from secondary modes when evaluated near a primary mode, minimizing only over the local cluster. To mitigate ambiguity, initialization at a common starting point can guide convergence to a “principal” consistent mode. In contrast, the Fréchet mean’s quadratic penalty (u^2) grows unbounded, pulling the minimizer toward a global compromise, often outside any cluster. We provide some simulations in $SO(2)$ in Fig. 8.

The parameter $c > 0$ controls the robustness threshold (e.g., $c = \pi/4$ for $SO(2)$, half the maximum geodesic distance π ; more generally, c can be set adaptively, but intuitively, smaller values of c induce a “shorter-sight” on the estimator).

Formally, given a distribution μ over the Lie group \mathcal{G} with geodesic distance $d_{\mathcal{R}}$, the robust centroid or *Tukey-Fréchet mean* $\mathcal{F}_r(\mu)$ is defined as:

$$\mathcal{F}_r(\mu) = \arg \min_{y \in \mathcal{G}} \mathbb{E}_{g \sim \mu} [m(d_{\mathcal{R}}(y, g); c)], \quad (21)$$

where $m(u; c)$ is the Tukey biweight loss. For empirical samples $\{g_i\}_{i=1}^n$, the estimator becomes:

$$\hat{\mathcal{F}}_r = \arg \min_{y \in \mathcal{G}} \sum_{i=1}^n m(d_{\mathcal{R}}(y, g_i); c). \quad (22)$$

C.1 EXAMPLES

We show comparison via simulations between the Fréchet mean and the robust variation for common distributions (Fig. 8). In $SO(2)$, the Fréchet mean has a closed-form expression (circular mean). The robust variation lacks a closed form and requires optimization, e.g., scalar minimization over $[0, 2\pi]$, but as a one-time computation for pseudo-labels (Algorithm 1), this is feasible (e.g., <1s for $n = 2000$ on standard hardware). In multimodal cases, it converges to a mode (e.g., π for bimodal, as shown in 8b) while Fréchet converges to $\pm\pi/2$. In $SO(3)$, there is no closed form for either Fréchet mean or the robust variation. In this case, the Fréchet mean can be found via singular value decomposition or Riemannian gradient descent.

1026
1027

D ADDITIONAL INSIGHTS AND EXPERIMENTS DETAILS

1028
1029

D.1 ORBITS, STABILIZERS, AND RELATION TO INSTANCE-LEVEL POSE DISTRIBUTIONS

1030
1031
1032
1033
1034
1035
1036

We clarify how our notion of instance-level pose/orbit distributions relates to the group-theoretic notions of orbits and stabilizers. Given $x \in \mathcal{X}$ and a group \mathcal{G} acting on \mathcal{X} , the stabilizer $\mathcal{S}_x = \{g \in \mathcal{G} \text{ s.t. } gx = x\}$ collects the (self) symmetries of x . For example, on planar rotations, a perfect circle has $\mathcal{S}_{\text{circle}} = SO(2)$, whereas a perfect square has $\mathcal{S}_{\text{square}}$ equal to the group of 90° rotations.

1037
1038
1039
1040
1041
1042
1043
1044
1045
1046

In contrast, our instance-level distributions $\mu_{[x]}$ do not model stabilizers explicitly; instead, they approximate the part of the orbit \mathcal{O}_x that is actually observed in the data. When x has a non-trivial stabilizer \mathcal{S}_x , many group elements act identically on x , so the pose distribution is only identifiable modulo \mathcal{S}_x and becomes broad or multimodal along the corresponding symmetry directions. This makes canonical poses inherently ambiguous for highly symmetric inputs, but our Tukey-Fréchet estimator is explicitly designed to remain robust in this regime.

1047
1048
1049
1050
1051
1052
1053
1054
1055
1056

This is visible for the digit “8” in the per-class rotated MNIST experiment: for a small fraction of very symmetric handwritten 8s, upright and upside down configurations are effectively indistinguishable, so the learned pose distribution shows a non-negligible mass near 180° (pose ambiguity), unlike digits such as 3, 4, or 7 whose histograms are sharply concentrated. In such cases, our Tukey-Fréchet objective does not average the modes into an unstable in-between pose; it admits minimizers aligned with one of the symmetric modes, yielding stable in-distribution canonicalizations (cf. Appendix C) and enabling test-time canonicalization.

1057

D.2 IMPACT OF NEIGHBORHOOD SIZE

1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068

Our symmetry estimation relies on approximating the class $[x]$ via k -nearest neighbors in the invariant space \mathcal{Z} . Figure 10 shows the impact of k on the MAE of the predicted per-class symmetry parameter θ for the MNIST experiment. While Proposition 3.1 suggests convergence as $N \rightarrow \infty$ (larger k), practical datasets have finite class separation and noisy samples. Large k can include samples from different underlying classes in the k -NN approximation, increasing noise and impacting performance. Conversely, very small k may not provide a representative sample. Overall, $k = 25$ offers the best balance for GEOM-QM9 and $k = 10$ for MNIST and FashionMNIST, and the performance on the symmetry discovery task is not overly sensitive to the neighborhood size within a reasonable range of neighbors, indicating that our method is stable w.r.t. the choice of k .

1069
1070

D.3 ANALYSIS OF EQUIVALENCE CLASS DEFINITION

1071
1072
1073
1074

To quantitatively assess the quality of the computed equivalence classes $[x]$, we introduce a *hit rate* metric that, for each $x \in \mathcal{X}$, measures the proportion of k -nearest neighbors in the computation of $[x]$ that belong to the same molecular class (SMILES) as x . A high hit rate indicates that our definition of equivalence class captures pose-invariant similarity successfully.

1075
1076
1077
1078
1079

Figure 13a shows the distribution of hit rates across all 174,481 conformers in our GEOM-QM9 dataset. The average hit rate is 0.95, with 76.8% of distinct molecules achieving an average hit rate above 0.9. This confirms that the proposed equivalence class definition behaves as expected, abstracting molecular identity across conformational variations and group transformations, while maintaining discriminative power between different molecular structures. Figure 13b shows that molecules with more conformers in the dataset tend to achieve slightly higher hit rates (correlation

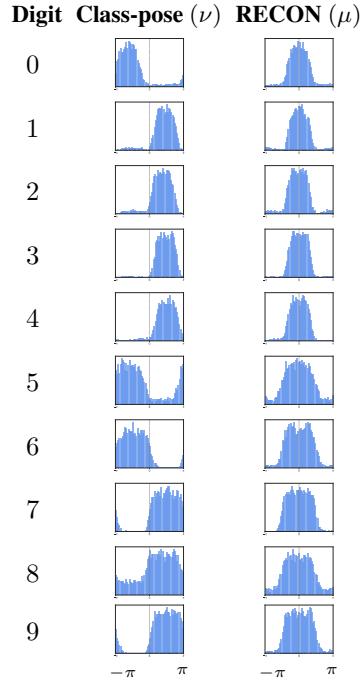


Figure 9: Probability density of recovered $SO(2)$ distributions on MNIST experiment: IE-AE, vs RECON.

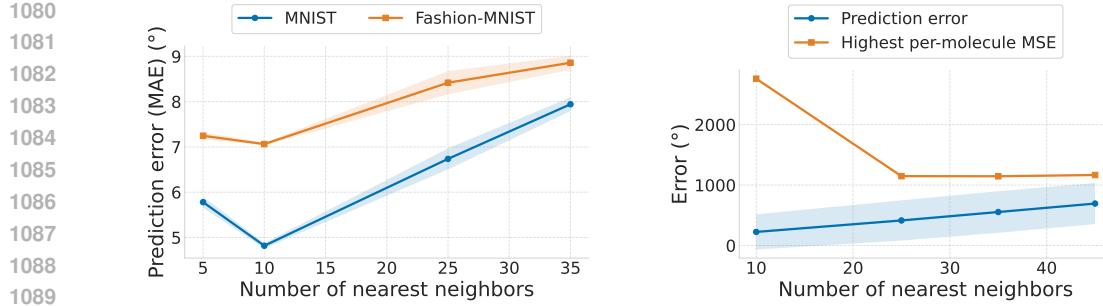


Figure 10: **Prediction error in symmetry discovery task vs number of neighbors in class computation** for MNIST and FashionMNIST experiments (left) and GEOM-QM9 (right). On the left, we plot average prediction error, which is defined as the MAE (in degrees) between the ground truth symmetry parameter and the per-class average predicted parameter, averaged across classes; shaded bands show the corresponding standard deviation across multiple runs with different seeds. On the right, we plot average and standard deviation of the prediction error, which is defined by the MSE between predicted and true matrix-Fisher parameters, averaged over all conformers. The orange curve reports the highest per-molecule mean MSE across conformers, illustrating the trade-off between average accuracy and worst case error as the number of neighbors varies.

coefficient 0.149). This is expected; molecules with a small number of conformers are more likely to include neighbors that belong to other molecules. This supports the insights about how the number of neighbors affects the quality of the symmetry discovery (Fig. 10), since a greater number of neighbors increases the probability of a reduced hit rate on low-conformer molecules.

We also report hit rate metrics for varying number of neighbors on the imaging experiments in Fig. 13c. Note that configurations with the highest hit rate ($k = 5$) do not necessarily correspond to configurations with lowest prediction error ($k = 10$) in the symmetry discovery task (Fig. 10, left). In effect, lower k generally yields higher hit rates, but using too few neighbors provides insufficient samples for accurate estimation. This explains the observed trade-off in Fig. 10, where $k = 10$ achieves the lowest prediction error despite $k = 5$ having the highest hit rate.

Limitations of the proposed equivalence class

We discuss the most prominent failure modes of our equivalence class definition by examining lowest-hit-rate cases in FashionMNIST, the dataset which exhibited the lowest hit rate score in our experiments (cf. Figure 13c). The heatmap in Figure 11 shows that most inputs with equivalence class construction imprecisions concentrate on a few semantically similar labels, particularly the footwear block (*Sneaker* \leftrightarrow *Ankle boot*) and tops (*Coat* \leftrightarrow *Shirt*). For a qualitative analysis, we visualize at random some equivalence classes with low and zero hit rate in Figure 12. We observe that neighbors of each input are visually nearly indistinguishable from the input, up to small cues (e.g., ankle height), partly due to small resolution of the data. As a result, the encoder maps them together in the latent space, despite having different semantic labels, which generates the feature overlap.

Overall, these figures indicate that the most challenging scenarios for our method involve inputs with *high inter-class overlap in the latent space* (and not large intra-class deformation). In effect, our method thrives in latent spaces with clear separation between classes, and is not required to have small within-class variations in the data.

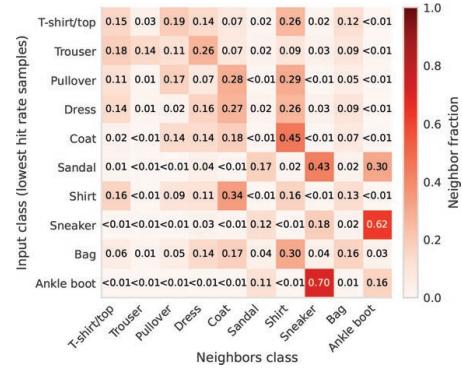


Figure 11: **Neighbor confusion matrix (FashionMNIST)** inputs with the lowest neighbor label hit rate (worst 5k samples). Each row corresponds to an input class and each column to a neighbor class; entry (i, j) gives, averaged over low-hit inputs of class i , the fraction of their k -NN that belong to class j .

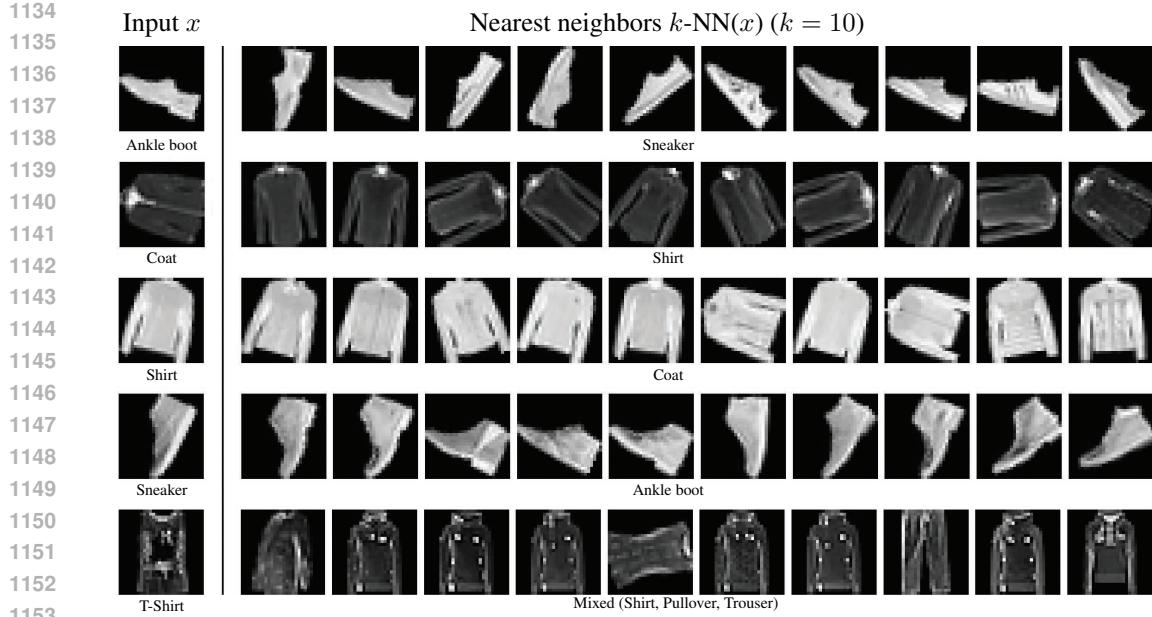


Figure 12: Qualitative analysis of lowest hit rate equivalence classes in FashionMNIST. Each row shows a training sample (left) with low hit rate (in particular, zero for rows 1-4) and its 10 nearest neighbors in the learned latent space (right), ordered by cosine similarity. Text underneath denotes labels. This illustrates inter-class feature overlap in the latent space: some inputs from different semantic classes exhibit very similar features (e.g., some *Sneaker* models look very similar to some *Ankle boot* models), therefore they are clustered together despite having different labels. This overlap is the primary cause of reduced hit rates in FashionMNIST for certain inputs, and highlights the regime where our equivalence class definition is most challenged.

As an empirical example, GEOM-QM9 exhibits non-trivial intra-class conformational variations ($\text{RMSD} \leq 1.5$ Angstrom) yet achieves a 0.95 hit rate (cf. Figure 13a). This happens because molecular identity has a very strong and clear signal from compositional and topological invariants, e.g. number of atoms of a given element, that remain perfectly constant across conformers, which allows the encoder to more easily separate conformers of different molecules. Our equivalence class definition behaves well as a result.

Mitigation strategies for inter-class overlap Mitigation strategies for these cases should therefore target inter-class overlap, i.e., increasing class separation in the invariant latent space. This challenge is well-studied, and established approaches with contrastive or self-supervised objectives (e.g., DINO (Caron et al., 2021), BYOL (Grill et al., 2020)) are known to sharpen class boundaries without supervision. Such objectives are perfectly compatible with our framework and directly address its main failure mode. Additional mitigation strategies, as discussed in our limitations section, include richer and more robust neighborhood sampling to further stabilize equivalence class construction.

D.4 PICKING UP ON PARTIAL SYMMETRIES

A notable strength of RECON is its ability to distinguish between *distinct* classes that are related by a group transformation – such as digits ‘6’ and ‘9’, which are related by a 180° rotation. Fully $SO(2)$ -equivariant methods map these inputs to an equivalent representation, and therefore, downstream tasks struggle distinguishing between them. This is a well-known example that has motivated partially equivariant methods in the past (Romero & Lohit, 2022). RECON leverages the clustering of the input’s *invariant features* and normalizes their pose distributions separately, which addresses this problem. Other class-pose decomposition methods, much like classical equivariant networks, can not pick up on partial symmetries, and therefore collapse both ‘6’s and ‘9’s into the same canonical

1188
1189 Table 1: RECON pseudo-label generation (Algorithm 1) wall-clock runtime for different datasets
1190 (naive implementation).

1191 Dataset	1192 Number of samples	1193 Nearest neighbors	1194 Runtime (seconds)
1193 MNIST	1194 12,000	1195 10	1196 15
1194 Fashion-MNIST	1195 60,000	1196 10	1197 173
1195 GEOM-QM9	1196 174,481	1197 25	1198 1033

1198 reconstruction (Figure 5b Col. 2).⁶ Consequently, the distributions $\nu_{[6]}$ and $\nu_{[9]}$ of relative poses are
1199 different with opposing peaks (Figure 9 Col. 2, Digits 6 and 9). On the contrary, RECON estimates
1200 distinct offsets ($\hat{\Gamma}_{[6]}$ and $\hat{\Gamma}_{[9]}$) based on how each digit class typically appears relative to the arbitrary
1201 canonical. As a result, normalization yields not only a *shared* symmetry pattern ($\mu_{[6]} \approx \mu_{[9]}$, Figure 9
1202 Col. 3) but also correctly associates them with *distinct* natural poses (Figure 5b Col. 3). This
1203 sensitivity to the data’s contextual orientation allows RECON to handle cases beyond perfect group
1204 orbits where full equivariance is inappropriate.

1205
1206 **D.5 COMPUTATIONAL ANALYSIS**

1207 Our method introduces three computational components on top of the class-pose backbone: (i) a
1208 (one time) pseudo-label generation step based on k -nearest neighbors (Algorithm 1), (ii) training
1209 of Θ and Λ (computation and implementation details in Appendix E) and (iii) the downstream
1210 symmetry/OOD/canonicalization inferences with the learned mappings Θ and Λ . We analyze (i) and
1211 (iii) in terms of complexity and runtime in practice.

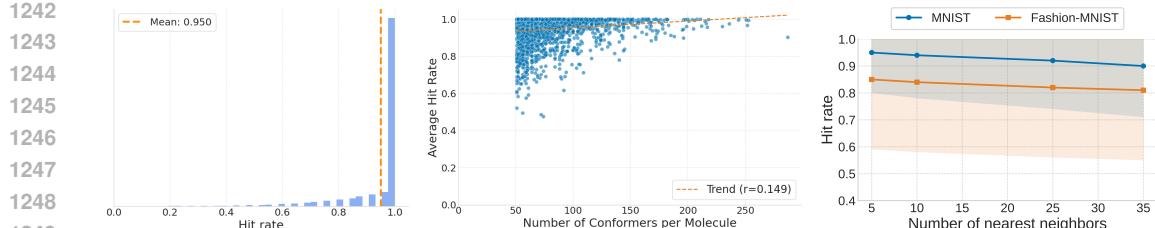
1212
1213 **Pseudo-label generation** Let N be the number of training examples, d the dimensionality of
1214 the invariant embedding z , and k the number of neighbors for Algorithm 1. The dominant cost of
1215 pseudo-label generation is the k -NN computation, with complexity $\mathcal{O}(N^2d)$. This step is ran once per
1216 dataset, and the resulting pseudo-labeled dataset is saved to disk and subsequently used to train Θ and
1217 Λ . We emphasize that the pseudo-label generation is a one-time computation that does not happen at
1218 training or inference time. On our hardware (see Appendix E), a naive, non-optimized implementation
1219 of the pseudo-label generation process completes in the order of few seconds / minutes per dataset;
1220 wall-clock times are reported in Table 1. For scaling to substantially larger N , one can use batch
1221 processing to reduce memory and replace the k -NN search by highly optimized neighbor search
1222 implementations (e.g., FAISS (Douze et al., 2024)), which greatly reduces the effective cost.

1223
1224 **RECON canonicalization overhead** Canonicalization at test time makes use of two forward passes:
1225 the IE-AE encoder on the input to obtain the relative pose $\psi(x)$ (which yields the IE-AE canonical-
1226 ization), and the forward pass $\Lambda(x)$ to obtain the centering transformation (which alongside $\psi(x)$
1227 yields the RECON canonicalization, $C(x) = \rho_{\mathcal{X}}(\Lambda(x) \cdot \psi(x)^{-1})x$). We quantify RECON’s canonical-
1228 ization overhead by measuring the average wall-clock inference time (in ms/sample) of the centering
1229 transformation computation. The computational overhead is of 0.351 ms/sample for MNIST, 0.113
1230 ms/sample for FashionMNIST and 0.068 ms/sample for GEOM-QM9. Our canonicalization adds
1231 only a small overhead, and the performance gains outweigh this cost.

1232
1233 **E IMPLEMENTATION DETAILS OF RECON AND IE-AE BACKBONE**

1234 Our experiments are implemented using Python 3.11, primarily with Pytorch Paszke et al. (2019)
1235 and Pytorch Geometric (Fey & Lenssen, 2019) for neural network implementation and training. All
1236 experiments were seeded with a fixed random seed for reproducibility and logged using Weights &
1237 Biases Biewald (2020).

1238
1239
1240
1241
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
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691
2692
2693
2694
2695
2696
2697
2698
2699
2700
2701
2702
2703
2704
2705
2706
2707
2708
2709
2710
2711
2712
2713
2714
2715
2716
2717
2718
2719
2720
2721
2722
2723
2724
2725
2726
2727
2728
2729
2730
2731
2732
2733
2734
2735
2736
2737
2738
2739
2740
2741
2742
2743
2744
2745
2746
2747
2748
2749
2750
2751
2752
2753
2754
2755
2756
2757
2758
2759
2760
2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788
2789
2790
2791
2792
2793
2794
2795
2796
2797
2798
2799
2800
2801
2802
2803
2804<br



(a) Per-input hit rate in GEOM-QM9, measuring quality of the computed equivalence classes $[x]$. (b) Average hit rate per-molecule vs the number of conformers. Molecules with less conformers present noisier equivalence classes. (c) Average and standard deviation of hit rate across neighbors in imaging experiments.

Figure 13: Quantitative analysis of the proposed equivalence class definition for the GEOM-QM9 experiment.

E.1 IMAGES

E.1.1 DATA PREPARATION

We use augmented versions of the MNIST/FashionMNIST datasets with different true symmetry distributions for each class. We load the original training and testing sets, split images by their class label (0-9), and apply a rotation to each image sampled from a label-dependent distribution (e.g. uniformly from $[-60^\circ, 60^\circ]$ for labels 0-4, and $[-90^\circ, 90^\circ]$ for labels 5-9 in MNIST).

E.1.2 MODEL IMPLEMENTATIONS

For the equivariant architectures, we use the `esCNN` library Cesa et al. (2022); Weiler & Cesa (2019).

Encoder. The encoder processes the input image through a series of $SE(2)$ -equivariant blocks (`esCNN.nn.R2Conv` layers with kernel sizes 7, 5, 5, 5, 5, 3, 1, using `gspaces.rot2dOnR2`), with additional batch normalization and non-linearities (ReLU, `NormNonLinearity`). The final layer outputs features decomposed into invariant scalar fields (128 channels, representing the invariant component $\eta(x)$), and equivariant vector fields from which the relative pose $\psi(x)$ – parameterized by an angle – is derived.

Decoder. A standard CNN takes the 128-dimensional invariant latent vector as input. It uses a sequence of `torch.nn.Conv2d`, `BatchNorm2d`, `Dropout2d` ($p=0.2$), and ReLU layers, along with bilinear interpolation for upsampling, to reconstruct the input image. The output passes through a Sigmoid activation.

Learnable mappings Θ and Λ . First, Θ computes an invariant embedding from the input using `esCNN.nn.R2Conv` layers as before. Then, this embedding is passed through a small MLP consisting of two `torch.nn.Linear` layers with ReLU activations, predicting a single scalar, parameter of the target distribution (an angle $\hat{\theta}$ for MNIST experiment and a standard deviation angle $\hat{\sigma}$ for FashionMNIST experiment). Λ is built equivalently as the Θ predictor, but outputs a single scalar value representing the predicted transformation offset $\hat{\Gamma}$ parameterized as an angle.

E.1.3 TRAINING

- **Training the IE-AE:** Best hyperparameter configuration (as in configuration that yields the lowest reconstruction error during validation) was found by hyperparameter tuning. The IE-AE components were trained for 700 epochs using the Adam optimizer Kingma & Ba (2017) with a learning rate $\approx 8 \times 10^{-4}$ and a batch size of 128. We save model weights corresponding to the lowest validation reconstruction loss.
- **Canonical orientation normalization:** Using the frozen pre-trained encoder, we compute the invariant embedding and the relative rotation angle degrees for each training sample. We then compute the $k = 10$ nearest neighbors (for each sample) in the η space based on cosine

1296 similarity . Finally, for each sample x , we compute two pseudo-labels for training Λ and Θ
 1297 respectively:
 1298

- $\Gamma_{[x]}$: the Fréchet mean (an angle in degrees) – which is equivalent to the circular mean
 1299 in $SO(2)$ – of the set of neighbor angles $\{\psi(x)\}_{x \in \mathcal{NN}(x)}$.
- $\hat{\theta}_{[x]}$: the set of neighbor angles $\{\psi(x)\}_{x \in \mathcal{NN}(x)}$ is then normalized using the previous
 1300 Fréchet mean: $\{\psi(x)\Gamma_{[x]}^{-1}\}_{x \in \mathcal{NN}(x)}$. From this set of normalized angles, the pseudo-
 1301 label $\hat{\theta}_{[x]}$ (estimate of the parameter of the symmetry distribution) is calculated using
 1302 standard parameter estimation methods. In our case, we use methods based on moments
 1303 robust to outliers.

 1304
 • **Training the learnable mappings Θ and Λ :** The Θ predictor and Λ predictor were trained
 1305 jointly for 600 epochs by minimizing MSE between the network outputs and the pseudo-
 1306 labels. We used a combined Adam optimizer targeting the parameters of both predictors,
 1307 with a learning rate of approximately 1.35×10^{-4} . We use a batch size of 128. We weight
 1308 both losses by a weighting factor of 0.25 applied to the loss of Λ . We save model weights
 1309 corresponding to the best validation loss.

1310 For all our trainings, we employ a cosine annealing learning rate scheduler with a warm-up phase of 5
 1311 epochs and restarts. Trainings were performed using an NVIDIA A100-SXM4-80GB graphics cards,
 1312 running for approximately 6 + 1 hours for MNIST and 20 + 3 hours for FashionMNIST (IE-AE +
 1313 learnable mappings phase).

1314 E.2 MOLECULAR CONFORMATIONS

1315 E.2.1 DATA PREPARATION AND SELECTION

1316 Our starting point is the GEOM dataset (Axelrod & Gómez-Bombarelli, 2022), focusing on its
 1317 QM9 subset (Ramakrishnan et al., 2014). From this, we select molecules that possess at least
 1318 64 distinct conformers. We focus on low-energy states and only retain conformers with a Root
 1319 Mean Square Deviation (RMSD) of less than 1.5 Angstrom from their respective molecule’s
 1320 minimum-energy conformer. This threshold is applied using RDKit’s Landrum et al. (2025)
 1321 `rdkit.Chem.rdBondAlign.GetBestRMS` function, which also aligns each qualifying con-
 1322 formers to its corresponding minimum-energy reference structure. This alignment step provides an
 1323 orientation-neutral base for each set of conformers.

1324 The conformers for each selected molecule are then randomly split into training, validation, and test
 1325 sets with an 0.8, 0.1, 0.1 ratio (*per molecule*) respectively. This results in 174,481 training samples,
 1326 20,902 validation samples and 23,805 test samples across 2,221 distinct classes of molecules.

1327 To introduce controlled and diverse controlled global orientations to serve as ground truths, we
 1328 augment the data by applying random rotations to the aligned conformers. These rotations are
 1329 sampled from matrix-Fisher distributions (Mardia & Jupp, 2009), a unimodal distribution on $SO(3)$
 1330 suitable for modeling varied directional concentrations.

1331 Specifically, we utilize three distinct matrix-Fisher parameter matrices F_{true} that simulate rotations
 1332 around the standard e_1, e_2, e_3 axes in the 3D space:

- 1333 • $F_{true}^1 = \text{diag}(100, 0.001, 0.001)$
- 1334 • $F_{true}^2 = \text{diag}(0.001, 100, 0.001)$
- 1335 • $F_{true}^3 = \text{diag}(0.001, 0.001, 100)$

1336 Each molecule, along with its conformers, is randomly assigned one of these F_{true} matrices. Rotations
 1337 are then sampled for each conformer from its assigned F_{true} matrix. This process allows us to simulate
 1338 distinct, realistic and parametrically defined orientation preferences across different molecules in the
 1339 dataset.

1340 E.2.2 DATA PRE-PROCESSING

1341 In our graph-based framework, molecules are represented as graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where nodes \mathcal{V}
 1342 correspond to atoms and edges \mathcal{E} represent connections between them. The Steerable E(3)-Equivariant

1350
1351 Graph Neural Network (SEGN) (Brandstetter et al., 2022) processes these graphs, leveraging their
1352 geometric and chemical information.
1353

1354 **Initial node features.** Each atom $i \in \mathcal{V}$ is ini-
1355 tially characterized by a feature vector \mathbf{x}_i . This
1356 vector is a one-hot encoding of the atom type
1357 (e.g., Hydrogen, Carbon, Nitrogen, Oxygen, Flu-
1358 orine). For our SEGN model, this corresponds
1359 to an input irreducible representation (irrep) of
1360 $5 \times 0e$, representing five distinct scalar atom
1361 types.

1362 **Graph connectivity and edge definition.** The
1363 graph’s edges are determined using a radius
1364 graph approach. An edge (i, j) exists between
1365 atom i and atom j if their Euclidean distance
1366 is within a predefined cutoff radius r_{cut} . In our
1367 experiments, we use $r_{cut} = 2.0$. The connec-
1368 tivity is stored in `edge_index` as per PyTorch
1369 Geometric conventions.

1370 **Edge attributes.** To incorporate 3D geometry in an equivariant manner, edges are augmented
1371 with attributes derived from the relative positions of the connected atoms (Brandstetter et al., 2022).
1372 Specifically, for an edge (i, j) connecting atom i at position $\mathbf{p}_i \in \mathbb{R}^3$ to atom j at position $\mathbf{p}_j \in \mathbb{R}^3$,
1373 we compute the relative position vector $\mathbf{r}_{ij} = \mathbf{p}_i - \mathbf{p}_j$. These relative position vectors are then
1374 transformed into spherical harmonics up to a maximum degree l_{max}^{edge} (we use $l_{max}^{edge} = 3$). We use the
1375 e3nn (Weiler et al., 2018; Thomas et al., 2018; Kondor et al., 2018) library to compute spherical
1376 harmonics $\mathbf{Y}(\mathbf{r}_{ij})$, which serve as edge attributes for the SEGN. The raw relative squared Euclidean
1377 distance $d_{ij}^2 = \|\mathbf{r}_{ij}\|^2$ is also added and passed to the SEGN as an additional scalar $1 \times 0e$ type
1378 feature for each edge.

1379 **Node attributes.** In addition to the initial atom type features, nodes are also assigned geometric
1380 attributes based on the mean of the spherical harmonic attributes of their incoming edges. These node
1381 attributes are processed by the SEGN.

1383 E.2.3 MODEL IMPLEMENTATION

1385 **SEGN-based encoder.** The SEGN architecture learns equivariant node representations through
1386 message passing. Our implementation consists of 6 message passing blocks. Each layer operates
1387 on node features represented by a combination of irreducible representations (irreps): we use 100
1388 scalar channels (type 0e irreps), 64 vector channels (type 1o irreps), 16 type 2e tensor channels,
1389 and 8 type 3o tensor channels. Equivariant tensor products and gated Sigmoid Linear Units (SiLU)
1390 non-linearities are used throughout these layers to ensure equivariance. Instance normalization is
1391 applied to the features within the SEGN layers.

1392 After the main message passing sequence, a final equivariant projection layer maps the processed
1393 node features to a target set of $d_{inv} = 512$ scalar channels and $d_{eq} = 512$ vector (type 1o) channels
1394 per node. These structured node-level features are then processed as:

- 1396 **1. Pooling:** The scalar (0e) and vector (1o) components of the node features are independently
1397 pooled across all nodes to obtain graph-level invariant and equivariant features (per-graph:
1398 an invariant embedding $\mathbf{s}_{pool} \in \mathbb{R}^{512}$ and an equivariant feature set \mathbf{v}_{pool} consisting of
1399 512 3D vectors). Additionally, the invariant embedding is passed through a Multi-Layer
1400 Perceptron (linear layers with ReLU activations) to produce a final graph-level invariant
1401 latent encoding $\mathbf{z}_{inv} \in \mathbb{R}^{512}$.
- 1402 **2. SE(3) pose prediction:** The equivariant embedding must be processed to obtain group
1403 transformations $(R, t) \in SE(3)$. We pass the pooled equivariant features \mathbf{v}_{pool} through an
1404 equivariant MLP to obtain three equivariant 3D output vectors: $\mathbf{y}_1, \mathbf{y}_2$ and \mathbf{y}_t . A rotation

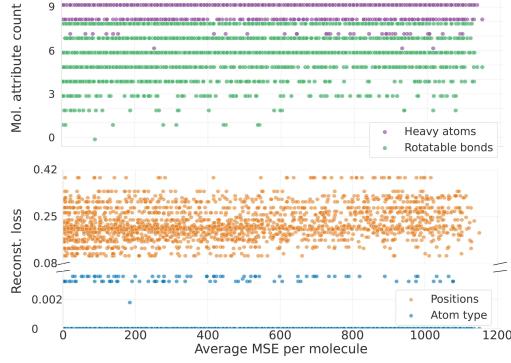


Figure 14: MSE vs reconstruction error and molecule attributes.

1404 matrix $R \in SO(3)$ is then derived from \mathbf{y}_1 and \mathbf{y}_2 using a Gram-Schmidt orthogonalization
 1405 procedure. The translation vector $\mathbf{t} \in \mathbb{R}^3$ is simply given by \mathbf{y}_t .
 1406

1407 This construction leads to an embedding \mathbf{z}_{inv} invariant to $SE(3)$ transformations of the input graph,
 1408 and to a predicted pose $(R, \mathbf{t}) \in SE(3)$ that transforms in an equivariant manner.
 1409

1410 **Decoder.** The decoder reconstructs the molecule from the invariant latent code \mathbf{z}_{inv} . It comprises
 1411 two main components: a position decoder and an atom-type decoder. Both are implemented as MLPs
 1412 with residual blocks and ReLU activations.

1413

- 1414 1. The **Position decoder** takes \mathbf{z}_{inv} as input and outputs a set of 3D coordinates for a maximum
 1415 number of 29 atoms in our dataset. This set of 3D coordinates represents the molecule in a
 1416 learned canonical orientation.
- 1417 2. The **Atom-type decoder** also takes \mathbf{z}_{inv} and predicts the logits for atom types for each of
 1418 the 29 positions.

1419 The final reconstructed atom positions are obtained by applying the predicted equivariant transforma-
 1420 tion (R, \mathbf{t}) to the canonical positions obtained by the decoder.
 1421

1422 **Learnable mappings Θ and Λ .** During this self-supervised learning phase we train two separate
 1423 networks using the computed pseudo-labels as outlined in Section 3.3. Both networks use the
 1424 same architectural pattern as the encoder: an SEGNN backbone to generate an invariant graph-level
 1425 embedding, followed by an MLP head.

1426 The Θ network outputs 9 parameters to form the 3×3 matrix-Fisher parameter matrix F_{pred} . The Λ
 1427 network outputs a rotation matrix representing the predicted offset $\hat{\Gamma}$.
 1428

1429 E.2.4 TRAINING

1430

- 1431 • **Training the IE-AE:** The encoder and a decoder based on $SE(3)$ -equivariant graph neural
 1432 networks as defined previously were trained for 600 epochs using the Adam optimizer with
 1433 a learning rate of $\approx 8.89 \times 10^{-5}$ and a batch size of 128. We compute two loss functions for
 1434 each of the decoder outputs – molecule’s positions and molecule’s atom types respectively.
 - 1435 – A positional reconstruction loss, computed as an L1 loss (mean absolute error) between
 1436 the true node coordinates and the coordinates obtained by applying the predicted
 1437 transformation (R, \mathbf{t}) to the decoder’s output positions.
 - 1438 – An atom-type reconstruction loss, which is a cross-entropy loss between the true atom
 1439 types and the atom types predicted by the decoder.

1440 The loss contribution by the atom-type loss is weighted by a factor of ≈ 3.614 . We save
 1441 model weights (both encoder and decoder) corresponding to the lowest validation positional
 1442 reconstruction loss.
 1443

- 1444 • **Canonical orientation normalization:** Using the frozen pre-trained encoder, we compute
 1445 the invariant embedding and the observed relative rotation matrix for each training sample
 1446 (molecule). We then compute the $k = 25$ nearest neighbors for each sample in the η space
 1447 based on cosine similarity. Finally, for each sample x we compute two pseudo-labels for
 1448 training the Λ and Θ predictors respectively:

- 1449 – $\hat{\theta}$: The set of neighbor rotation matrices $\{\psi(x)\}_{x \in \mathcal{NN}(x)}$ is then normalized using the
 1450 inverse of the centering rotation as $\{\psi(x)\hat{\Gamma}_{[x]}\}_{x \in \mathcal{NN}(x)}$. From this set of normalized
 1451 rotation matrices, we estimate the parameters of the matrix-Fisher distribution ($\hat{\theta} = \hat{F}$)
 1452 on these aligned rotations via the moment-matching approach (inverting $A(s) =$
 1453 $\coth(s) - 1/s$ by Newton’s method and reconstructing \hat{F} via SVD) (Wood, 2008;
 1454 Mardia & Jupp, 2009). This $\hat{\theta}$ is the pseudo-label for the Θ network.
 1455
- 1456 – $\hat{\Gamma}_{[x]}$: This is the Fréchet mean on $SO(3)$, estimated as the mode of the matrix-Fisher
 1457 distribution fitted to the set of observed neighbor rotations $\{\psi(x)\}_{x \in \mathcal{NN}(x)}$. We
 1458 compute it using the SVD-based moment matching estimator for the mode standard

1458 in literature (Wood, 2008; Mardia & Jupp, 2009). This mode $\hat{\Gamma}_{[x]}$ then serves as the
 1459 pseudo-label for the Λ network.

1460 • **Training the learnable mappings Θ and Λ :** The Θ and Λ predictor were trained jointly for
 1461 150 epochs. The SEGNNS for these predictors use 4 layers, with $50 \times 0e + 32 \times 1o + 8 \times 2e$
 1462 $+ 4 \times 3o$ hidden irreps, and instance normalization. We minimized the MSE between the
 1463 network outputs and their respective pseudo-labels. A combined Adam optimizer was used
 1464 for the parameters of both predictors, with a learning rate of approximately 4.83×10^{-4}
 1465 and a batch size of 128. The loss contribution from the Λ predictor was weighted by a factor
 1466 of 500. We save model weights to best validation loss.

1467 For all our trainings, we employ a cosine annealing learning rate scheduler with a warm-up phase
 1468 of 5 epochs. Gradient clipping with a maximum norm of 1.0 is applied during the training of the
 1469 learnable mappings phase. Training was performed using an NVIDIA A100-SXM4-80GB graphics
 1470 cards, running for approximately 6 + 1.5 days (IE-AE + learnable mappings phase).

1472 F IMPLEMENTATION DETAILS OF DOWNSTREAM APPLICATIONS

1473 F.1 OOD DETECTION

1474 The distributions recovered by RECON can be used to identify objects in unnatural (out-of-distribution)
 1475 poses relative to their learned symmetry profile. Given predictors $\Theta(x)$ and $\Lambda(x)$ (Sec. 3.3), consider
 1476 the absolute pose as $g_{abs} = \psi(x)\Lambda(x)^{-1}$ (normalized relative pose). The likelihood of g_{abs} under
 1477 the distribution parameterized by $\Theta(x)$ serves as an anomaly score,

$$1478 s(x) := -\log p_{\Theta(x)}(g_{abs}). \quad (23)$$

1479 Low $s(x)$ indicates in-distribution, while high $s(x)$ indicates OOD. We empirically validate this by
 1480 classifying randomly oriented $SO(2)$ (images) and $SO(3)$ (GEOM) test instances.

1481 Theoretically, one could skip RECON centering and score with an uncentered class-pose density
 1482 using arbitrary canonicals, i.e., using the score $s_{rel}(x) := -\log p_{\tilde{\Theta}(x)}(\psi(x))$ where $\tilde{\Theta}(x)$ is trained
 1483 on pseudo-labels derived from raw relatives $\{\psi(x)\}$. In effect, we show (Proposition B.2) that the
 1484 log density (and therefore the proposed score $s(x)$) is invariant to the choice of reference frame.
 1485 However, in practice, centered distributions (RECON) perform better. Centered distributions facilitate
 1486 optimization: they concentrate probability mass near the identity (reducing variance and dynamic
 1487 range of the loss), decouple the center $\hat{\Gamma}_x$ from the shape parameters (which are learned by Θ),
 1488 and overall result in more stable training. On the contrary, uncentered distributions introduce class-
 1489 specific, arbitrary shifts $\hat{\Gamma}_{[x]}$ that must be jointly learned by Θ , increasing complexity. Empirically
 1490 this results in consistently higher AUC-ROC across datasets for RECON (Table 5; cf. IE-AE vs.
 1491 RECON, ROC curve plots in Fig. 15).

1492 We now offer a detailed implementation and discussion of the OOD experiment, including dataset
 1493 preparation, metrics and uncentered (IE-AE) ablation.

1494 F.1.1 IMAGING

1495 **Pre-processing and metrics.** Each MNIST test image is rotated by a random angle $\theta \sim$
 1496 $\mathcal{U}(-180^\circ, 180^\circ]$. An input is labeled in-distribution (ID) if its applied rotation lies within the
 1497 class-specific training symmetry range (digits 0-4: $\pm 60^\circ$; digits 5-9: $\pm 90^\circ$), and out-of-distribution
 1498 (OOD) otherwise. All angular differences are computed on $(-180^\circ, 180^\circ]$; we write $\text{angdiff}(\alpha, \beta) =$
 1499 $\text{wrap}(\alpha - \beta)$.

1500 We aim to verify that our anomaly score $s(x)$ assigns larger values to OOD inputs than to ID. For this,
 1501 we report AUC-ROC, Area Under the Receiver Operating Characteristic Curve (using `sklearn`'s
 1502 `sklearn.metrics.roc_auc_score`), which summarizes performance *across all decision*
 1503 *thresholds* τ (if we classify x as OOD when $s(x) \geq \tau$). We use AUC-ROC since it is a threshold-free
 1504 measure and works well in class imbalance scenarios, a property that other metrics (e.g., accuracy)
 1505 lacks (they require a fixed threshold, whose choice depends on the use case for deployment, e.g. if
 1506 one wants to optimize threshold for minimizing false positives). We also provide AUC-ROC curve
 1507 plots in Fig. 15.

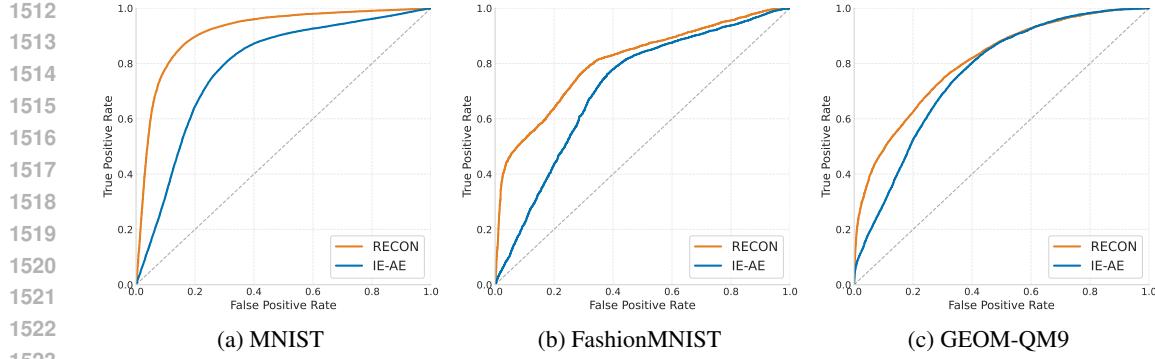


Figure 15: ROC curves for IE-AE and RECON-based anomaly scores for (a) MNIST, (b) FashionMNIST, and (c) GEOM-QM9 experiment.

Uncentered distributions (baseline, IE-AE) Given x , the model predicts an uncentered uniform support $\tilde{\Theta}(x) = (a_x, b_x)$ (in degrees). Let the input’s relative pose be $g_{\text{rel}} = \text{wrap}(\psi(x))$ and the support midpoint $\mu_x = \text{wrap}(\frac{a_x+b_x}{2})$. We use the absolute deviation

$$s_{\text{rel}}(x) = |\text{angdiff}(g_{\text{rel}}, \mu_x)|. \quad (24)$$

(Equivalently, one may use distance to the predicted support,

$$s_{\text{rel}}(x) = \max\left\{0, |\text{angdiff}(g_{\text{rel}}, \mu_x)| - w_x\right\}, \quad w_x = \frac{1}{2} \text{wrap}(b_x - a_x). \quad (25)$$

Both are surrogates of the uniform negative log likelihood, differing only by a within-support constant.)

Centered distributions (RECON) Let Γ_x be the predicted centering transformation (in degrees) from $\Lambda(x)$, and define the absolute (centered) pose $g_{\text{abs}} = \text{wrap}(g_{\text{rel}} - \Gamma_x)$. For identity-centered supports, the anomaly score reduces to

$$s(x) = |g_{\text{abs}}|. \quad (26)$$

For FashionMNIST, we mirror the MNIST setup, but assign ID/OOD labels using the FashionMNIST class-specific symmetry ranges specified in Section 5. No other changes are made.

F.1.2 GEOMETRIC GRAPHS

Pre-processing and metrics. We start from the (augmented) GEOM-QM9 test set as described in Section 5 and group conformers by molecule (SMILES). For each molecule, we randomly split its conformers into two halves: one half is kept as in-distribution (ID) and left unchanged; for the other half we generate OOD *candidates*. This split avoids a trivial task: because $SO(3)$ is large, naively applying a random rotation to every conformer would produce an overwhelming fraction of OOD examples.

For OOD candidates, we first *canonicalize* by undoing the augmentation used to create the conformer (apply the rotation used to create the augmented sample R_{aug} to bring coordinates to the original canonical frame), and then apply a fresh random rotation $R_{\text{candidate}} \sim \text{Haar}(SO(3))$ to obtain the new coordinates. ID examples are those left unmodified.

Ground-truth labels for the candidate half are decided through likelihood: if the sample rotation $R_{\text{candidate}}$ has a high likelihood of coming from F (the molecule’s matrix-Fisher parameter that was assigned during the creation of the augmented dataset), then that sample is labeled as ID. Otherwise, the sample is labeled as OOD.

Note that since the matrix-Fisher density on $SO(3)$ is proportional to $\exp(\text{tr}(F^T R))$, then the unnormalized log-likelihood of $R_{\text{candidate}}$ is $\text{tr}(F^T R_{\text{candidate}})$ up to an additive constant. Therefore, we mark a candidate as OOD if $\text{tr}(F^T R_{\text{candidate}}) < \tau$ (low likelihood), where τ is a fixed threshold controlling class balance (we use $\tau = 96$).

1566 **Uncentered distributions (baseline, IE-AE)** Let $\psi(x) \in \text{SO}(3)$ be the predicted pose and let
 1567 $F(x) \in \mathbb{R}^{3 \times 3}$ be the matrix-Fisher parameter predicted from x using the learned mapping. We score
 1568 anomalies with the negative unnormalized log likelihood

$$1569 \quad s_{\text{rel}}(x) = -\langle F(x), \psi(x) \rangle = -\text{tr}(F(x)^\top \psi(x)),$$

1571 where $\langle A, B \rangle := \text{tr}(A^\top B)$. When $F(x)$ is axis-symmetric of the form $F(x) = \kappa(x) \mu_x$ with mode
 1572 $\mu_x \in \text{SO}(3)$, $s_{\text{rel}}(x)$ is a monotone surrogate of the geodesic deviation $\phi(\mu_x^{-1} \psi(x))$.
 1573

1574 **Centered distributions (RECON)** Let $\Lambda(x) \in \text{SO}(3)$ be the predicted centering transform and
 1575 define the absolute (centered) pose

$$1576 \quad g_{\text{abs}} = \psi(x) \Lambda(x)^{-1}.$$

1577 We then score

$$1579 \quad s(x) = -\langle F(x), g_{\text{abs}} \rangle = -\text{tr}(F(x)^\top g_{\text{abs}}),$$

1580 a monotone surrogate of the geodesic angle $\phi(g_{\text{abs}})$.
 1581

1582 F.2 GRANTING GROUP INVARIANCE TO PRE-TRAINED BACKBONES

1584 F.2.1 IMAGING

1585 We evaluate inference-only canonicalization using a classifier (ResNet18 backbone) trained on the
 1586 raw (no augmentations) MNIST and FashionMNIST datasets. We then evaluate the performance
 1587 of this pre-trained model on the rotated dataset variations created for the symmetry discovery
 1588 experiment outlined in Section 5 (that is, $\pm 60^\circ / \pm 90^\circ$ rotations for MNIST, etc). For both MNIST
 1589 and FashionMNIST, we train a ResNet-18 with cross-entropy and Adam optimizer (lr=1e-3, batch
 1590 size 128, 100 epochs) and save the best classifier checkpoint based on best validation accuracy. At
 1591 test time, we compare three input pre-processing modes before feeding images into the classifier: (i)
 1592 no canonicalization (using the pre-trained classifier), (ii) arbitrary canonicalization (taken from the
 1593 IE-AE) and (iii) RECON canonicalization. As a reference value, we additionally report accuracy on
 1594 the original (un-augmented) test set (Table 2), which matches the training distribution by construction.
 1595 This corresponds to a perfect canonicalization (reversing the augmentations), and serves as an upper
 1596 bound of the best accuracy that can be obtained if we provide a perfect canonicalization function
 1597 during inference.

1598 **Comparisons** We compare with EquiAdapt (Mondal et al., 2023), another test-time canonical-
 1599 ization method. We reproduce their *zero-shot setup*, that is, only the canonicalization function is
 1600 trained, and the pre-trained model weights are kept frozen (same setting as RECON). We use the
 1601 $SO(2)$ canonicalization prior and steerable network from their public `equiadapt` package for
 1602 the canonicalization function. The canonicalization function is attached on top of the pre-trained
 1603 ResNet18 classifier, and is trained for 50 epochs with a batch size of 128. We choose the best
 1604 hyperparameter configuration based on best test accuracy (see hyperparameter search in Figure 4).

1605 For reference, Table 3 reports results for specialized group equivariant (and partially equivariant)
 1606 architectures trained from scratch on the same non-rotated datasets. Equivariant models are by
 1607 construction designed to be robust to symmetry-induced test-time distribution shifts. On MNIST,
 1608 our canonicalization (RECON, 90.96%) shows a moderately small gap w.r.t. the $SE(2)$ -equivariant
 1609 ESCNN classifier (94.72%), while keeping the benefits of operating purely as an unsupervised, test-
 1610 time, plug-and-play module for arbitrary classifiers. On FashionMNIST however, RECON actually
 1611 performs better than the fully equivariant ESCNN. Lastly, we consistently outperform Partial G-CNNs
 1612 in both datasets, sometimes by a large margin (67.72% vs 81.96% in FashionMNIST).

1613 The ESCNN classifier is built using the same architecture as the ESCNN backbone for the IE-
 1614 AE described in Appendix E.1, but with a 2-layer MLP classification head attached at the end.
 1615 The invariant embedding dimension before the classification head is of size 1024 and the ESCNN
 1616 backbone has a hidden dimension of 128, totaling 3,384,463 trainable parameters in this case. We
 1617 train for 50 epochs with batch size 128 and learning rate 1e-3. Hyperparameter configuration was
 1618 chosen based on best test accuracy across different learning rates, hidden dimensions and embedding
 1619 dimensions (Table 5). The Partial G-CNNs were configured and trained similarly, aiming to keep an
 approximate same number of parameters as the ESCNN architectures.

1620 Table 2: Test accuracy on the in-distribution test set (i.e., non-augmented) obtained with the pre-
 1621 trained classifier trained on non-rotated, vanilla datasets: reference canonicalization upper bound
 1622 representing a perfect canonicalization function.

Dataset	Augmentation-reversed test set acc. (upper bound, perfect canon.)
MNIST	98.14 \pm 0.02%
FashionMNIST	91.11 \pm 0.1%
GEOM-QM9	96.40 \pm 0.0%

1630 Table 3: **Test accuracy on the per-class rotated datasets obtained with fully and partially equivari-
 1631 ant classifiers trained on non-rotated, vanilla datasets; auxiliary reference metric representing the
 1632 performance obtained with specialized equivariant architectures that are, by construction, robust to
 1633 symmetry-induced distribution shifts during inference.**

Dataset (Train / Test)	Group	Equivariance	Backbone	Test set acc.
MNIST (Orig / Rot)	$SE(2)$	Full Learned	ESCNN Partial G-CNN	94.59% 90.20%
FashionMNIST (Orig / Rot)	$SE(2)$	Full Learned	ESCNN Partial G-CNN	77.94% 67.72%
GEOM-QM9 (Orig / Rot)	$SE(3)$	Full	SEGN	98.55%

F.2.2 GEOMETRIC GRAPHS

We evaluate inference-only canonicalization on GEOM-QM9 molecular graphs. We first train a graph CNN on aligned conformers only: a 3 layer GCN (using Pytorch Geometric’s `GCNConv` layer with a 128 hidden size) followed by a two layer MLP head (64 hidden units, dropout 0.1), trained with cross-entropy and Adam (batch size 128 for 50 epochs), saving the best checkpoint by validation accuracy. Node inputs to the classifier are geometric attributes based on the mean of the spherical harmonic attributes of their incoming edges (see details in data pre-processing section of GEOM-QM9 dataset in Appendix E.2.2). At test time, we evaluate on the rotated conformer test set described in Appendix E.2.1, and compare three input pre-processing modes applied per graph before classification: (i) none (no canonicalization, using the pre-trained classifier), (ii) arbitrary canonicalization (taken from the IE-AE) and (iii) RECON canonicalization. Note that after rotation by either canonicalization method, we have to recompute the geometric features (edge/node attributes) from the updated coordinates, since those are the input to the pre-trained GCN. For a reference on the upper bound of a perfect canonicalization, we also report accuracy on the test dataset obtained after reversing the test-set augmentations (Table 2). All runs use a fixed global seed with deterministic settings.

Comparisons For comparison against EquiAdapt (Mondal et al., 2023), we attach and train their continuous $SO(3)$ point-cloud canonicalizer from the public `equiadapt` package on top of the pre-trained GCN. For each input molecular conformation, we canonicalize its atomic coordinates with the canonicalizer, and then compute the resulting conformation’s geometric features (as in Appx. E.2.2) for input into the frozen GCN (same approach as in our test-time canonicalization). We use Adam with learning rate $1e-3$, batch size 128, and train for 50 epochs. During training, the GCN weights are frozen and only the canonicalizer is optimized (coined the *zero-shot setup* in EquiAdapt (Mondal et al., 2023), equivalent to our test-time canonicalization setup), as opposed to the alternative *fine-tuning setup* in which the pre-trained model is retrained.

On GEOM-QM9, the $SO(3)$ -equivariant SEGN reaches 98.55% accuracy (Table 3), far above the pre-trained non-equivariant baseline (with or without canonicalization). This evidences that test-time canonicalization is significantly more challenging in the 3D domain than in images (at least for molecular conformations with varying per-molecule rotations). Nevertheless, RECON is the

1674 Table 4: **EquiAdapt** hyperparameter search and final test accuracy on the per-class rotated MNIST,
 1675 FashionMNIST and GEOM-QM9 dataset variations.

1677	1678	1679	1680	1681	1682	1683	1684	1685	1686	1687	1688	1689	1690	1691	1692	1693	1694	1695	1696	1697	1698	1699	1700	1701	1702	1703	1704	1705	1706	1707	1708	1709	1710	1711	1712	1713	1714	1715	1716	1717	1718	1719	1720	1721	1722	1723	1724	1725	1726	1727	1728	1729	1730	1731	1732	1733	1734	1735	1736	1737	1738	1739	1740	1741	1742	1743	1744	1745	1746	1747	1748	1749	1750	1751	1752	1753	1754	1755	1756	1757	1758	1759	1760	1761	1762	1763	1764	1765	1766	1767	1768	1769	1770	1771	1772	1773	1774	1775	1776	1777	1778	1779	1780	1781	1782	1783	1784	1785	1786	1787	1788	1789	1790	1791	1792	1793	1794	1795	1796	1797	1798	1799	1800	1801	1802	1803	1804	1805	1806	1807	1808	1809	1810	1811	1812	1813	1814	1815	1816	1817	1818	1819	1820	1821	1822	1823	1824	1825	1826	1827	1828	1829	1830	1831	1832	1833	1834	1835	1836	1837	1838	1839	1840	1841	1842	1843	1844	1845	1846	1847	1848	1849	1850	1851	1852	1853	1854	1855	1856	1857	1858	1859	1860	1861	1862	1863	1864	1865	1866	1867	1868	1869	1870	1871	1872	1873	1874	1875	1876	1877	1878	1879	1880	1881	1882	1883	1884	1885	1886	1887	1888	1889	1890	1891	1892	1893	1894	1895	1896	1897	1898	1899	1900	1901	1902	1903	1904	1905	1906	1907	1908	1909	1910	1911	1912	1913	1914	1915	1916	1917	1918	1919	1920	1921	1922	1923	1924	1925	1926	1927	1928	1929	1930	1931	1932	1933	1934	1935	1936	1937	1938	1939	1940	1941	1942	1943	1944	1945	1946	1947	1948	1949	1950	1951	1952	1953	1954	1955	1956	1957	1958	1959	1960	1961	1962	1963	1964	1965	1966	1967	1968	1969	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100	2101	2102	2103	2104	2105	2106	2107	2108	2109	2110	2111	2112	2113	2114	2115	2116	2117	2118	2119	2120	2121	2122	2123	2124	2125	2126	2127	2128	2129	2130	2131	2132	2133	2134	2135	2136	2137	2138	2139	2140	2141	2142	2143	2144	2145	2146	2147	2148	2149	2150	2151	2152	2153	2154	2155	2156	2157	2158	2159	2160	2161	2162	2163	2164	2165	2166	2167	2168	2169	2170	2171	2172	2173	2174	2175	2176	2177	2178	2179	2180	2181	2182	2183	2184	2185	2186	2187	2188	2189	2190	2191	2192	2193	2194	2195	2196	2197	2198	2199	2200	2201	2202	2203	2204	2205	2206	2207	2208	2209	2210	2211	2212	2213	2214	2215	2216	2217	2218	2219	2220	2221	2222	2223	2224	2225	2226	2227	2228	2229	2230	2231	2232	2233	2234	2235	2236	2237	2238	2239	2240	2241	2242	2243	2244	2245	2246	2247	2248	2249	2250	2251	2252	2253	2254	2255	2256	2257	2258	2259	2260	2261	2262	2263	2264	2265	2266	2267	2268	2269	2270	2271	2272	2273	2274	2275	2276	2277	2278	2279	2280	2281	2282	2283	2284	2285	2286	2287	2288	2289	2290	2291	2292	2293	2294	2295	2296	2297	2298	2299	2300	2301	2302	2303	2304	2305	2306	2307	2308	2309	2310	2311	2312	2313	2314	2315	2316	2317	2318	2319	2320	2321	2322	2323	2324	2325	2326	2327	2328	2329	2330	2331	2332	2333	2334	2335	2336	2337	2338	2339	2340	2341	2342	2343	2344	2345	2346	2347	2348	2349	2350	2351	2352	2353	2354	2355	2356	2357	2358	2359	2360	2361	2362	2363	2364	2365	2366	2367	2368	2369	2370	2371	2372	2373	2374	2375	2376	2377	2378	2379	2380	2381	2382	2383	2384	2385	2386	2387	2388	2389	2390	2391	2392	2393	2394	2395	2396	2397	2398	2399	2400	2401	2402	2403	2404	2405	2406	2407	2408	2409	2410	2411	2412	2413	2414	2415	2416	2417	2418	2419	2420	2421	2422	2423	2424	2425	2426	2427	2428	2429	2430	2431	2432	2433	2434	2435	2436	2437	2438	2439	2440	2441	2442	2443	2444	2445	2446	2447	2448	2449	2450	2451	2452	2453	2454	2455	2456	2457	2458	2459	2460	2461	2462	2463	2464	2465	2466	2467	2468	2469	2470	2471	2472	2473	2474	2475	2476	2477	2478	2479	2480	2481	2482	2483	2484	2485	2486	2487	2488	2489	2490	2491	2492	2493	2494	2495	2496	2497	2498	2499	2500	2501	2502	2503	2504	2505	2506	2507	2508	2509	2510	2511	2512	2513	2514	2515	2516	2517	2518	2519	2520	2521	2522	2523	2524	2525	2526	2527	2528	2529	2530	2531	2532	2533	2534	2535	2536	2537	2538	2539	2540	2541	2542	2543	2544	2545	2546	2547	2548	2549	2550	2551	2552	2553	2554	2555	2556	2557	2558	2559	2560	2561	2562	2563	2564	2565	2566	2567	2568	2569	2570	2571	2572	2573	2574	2575	2576	2577	2578	2579	2580	2581	2582	2583	2584	2585	2586	2587	2588	2589	2590	2591	2592	2593	2594	2595	2596	2597	2598	2599	2600	2601	2602	2603	2604	2605	2606	2607	2608	2609	2610	2611	2612	2613	2614	2615	2616	2617	2618	2619	2620	2621	2622	2623	2624	2625	2626	2627	2628	2629	2630	2631	2632	2633	2634	2635	2636	2637	2638	2639	2640	2641	2642	2643	2644	2645	2646	2647	2648	2649	2650	2651	2652	2653	2654	2655	2656	2657	2658	2659	2660	2661	2662	2663	2664	2665	2666	2667	2668	2669	2670	2671	2672	2673	2674	2675	2676	2677	2678	2679	2680	2681	2682	2683	2684	2685	2686	2687	2688	2689	2690	2691	2692	2693	2694	2695	2696	2697	2698	2699	2700	2701	2702	2703	2704	2705	2706	2707	2708	2709	2710	2711	2712	2713	2714	2715	2716	2717	2718	2719	2720	2721	2722	2723	2724	2725	2726	2727	2728	2729	2730	2731	2732	2733	2734	2735	2736	2737	2738	2739	2740	2741	2742	2743	2744	2745	2746	2747	2748	2749	2750	2751	2752	2753	2754	2755	2756	2757	2758	2759	2760	2761	2762	2763	2764	2765	2766	2767	2768	2769	2770	2771	2772	2773	2774	2775	2776	2777	2778	2779	2780	2781	2782	2783	2784	2785	2786	2787	2788	2789	2790	2791	2792	2793	2794	2795	2796	2797	2798	2799	2800	2801	2802	2803	2804	2805	2806	2807	2808	2809	2810	2811	2812	2813	2814	2815	2816	2817	2818	2819	2820	2821	2822	2823	2824	2825	2826	2827	2828	2829	2830	2831	2832	2833	2834	2835	2836	2837	2838	2839	2840	2841	2842	2843	2844	2845	2846	2847	2848	2849	2850	2851	2852	2853	2854	2855	2856	2857	2858	2859	2860	2861	2862	2863	2864	2865	2866	2867	2868	2869	2870	2871	2872	2873	2874	2875	2876	2877	2878	2879	2880	2881	2882	2883	2884	2885	2886	2887	2888	2889	2890	2891	2892	2893	2894	2895	2896	2897	2898	2899	2900	2901	2902	2903	2904	2905	2906	2907	2908	2909	2910	2911	2912	2913	2914	2915	2916	2917	2918	2919	2920	2921	2922	2923	29

1728 G.3 ETHICAL CONCERNS
17291730 Because our method relies on a data-driven approach to identify natural poses, it may be susceptible
1731 to dataset bias. Beyond this, we do not anticipate significant ethical concerns or negative societal
1732 impacts.

1733

1734

1735

1736

1737

1738

1739

1740

1741

1742

1743

1744

1745

1746

1747

1748

1749

1750

1751

1752

1753

1754

1755

1756

1757

1758

1759

1760

1761

1762

1763

1764

1765

1766

1767

1768

1769

1770

1771

1772

1773

1774

1775

1776

1777

1778

1779

1780

1781