
Supplementary Material: Quantifying and Learning Linear Symmetry-Based Disentanglement

Anonymous Author(s)

Affiliation

Address

email

1 A Preliminaries: group theory

2 In this appendix, we summarize some concepts from group theory that are important to understand the
3 main text of the paper. Group theory provides a useful language to formalize the notion of symmetry
4 transformations and their effects. For a more elaborate discussion we refer the reader to the book
5 from Hall (2015) on group theory.

6 **Group** A group is a non-empty set G together with a binary operation $\circ : G \times G \rightarrow G$ that satisfies
7 three properties:

- 8 1. *Associativity*: For all $f, g, h \in G$, it holds that $f \circ (g \circ h) = (f \circ g) \circ h$.
- 9 2. *Identity*: There exists a unique element $e \in G$ such that for all $g \in G$ it holds that
10 $e \circ g = g \circ e = g$.
- 11 3. *Inverse*: For all $g \in G$ there exists an element $g^{-1} \in G$ such that $g^{-1} \circ g = g \circ g^{-1} = e$.

12 **Direct product** Let G and G' be two groups. The *direct product*, denoted by $G \times G'$, is the group
13 with elements $(g, g') \in G \times G'$ with $g \in G$ and $g' \in G'$, and the binary operation $\circ : G \times G' \rightarrow G \times G'$
14 such that $(g, g') \circ (h, h') = (g \circ h, g' \circ h')$.

15 **Lie group** A Lie group is a group where G is a smooth manifold, this means it can be described in
16 a local scale with a set of continuous parameters and that one can interpolate continuously between
17 elements of G .

18 **Group action** Let A be a set and G a group. The *group action* of G on A is a function
19 $G_A : G \times A \rightarrow A$ that has the properties ¹

- 20 1. $G_A(e, x) = x$ for all $a \in A$
- 21 2. $G_A(g, (G_A(g', a))) = G_A(g \circ g', a)$ for all $g, g' \in G$ and $a \in A$

22 **Regular action** The action of G on A is regular if for every pair of elements $a, a' \in A$ there exists
23 a unique $g \in G$ such that $g \cdot a = a'$.

24 **Group representation** A *group representation* of G in the vector space V is a function $\rho : G \rightarrow$
25 $GL(V)$ (where $GL(V)$ is the general linear group on V) such that for all $g, g' \in G$ $\rho(g \circ g') =$
26 $\rho(g) \circ \rho(g')$ and $\rho(e) = \mathbb{I}_V$, where \mathbb{I}_V is the identity matrix.

¹To avoid notational clutter, we write $G_A(g, a) = g \cdot a$ where the set A on which $g \in G$ acts can be inferred from the context.

27 **Direct sum of representations** The direct sum of two representations $\rho_1 : G \rightarrow GL(V)$ in V and
 28 $\rho_2 : G \rightarrow GL(V')$ in V' is a group representation $\rho_1 \oplus \rho_2 : G \rightarrow GL(V \oplus V')$ over the direct sum
 29 $V \oplus V'$, defined for $v \in V$ and $v' \in V'$ as:

$$(\rho_1 \oplus \rho_2)(g) \cdot (v, v') = (\rho_1(g) \cdot v, \rho_2(g) \cdot v') \quad (1)$$

30 B Linear Symmetry-Based Disentanglement: Detailed Description

31 Higgins et al. (2018) provide a formal definition of linear disentanglement that connects symmetry
 32 transformations affecting the real world (from which data is generated) to the internal representations
 33 of a model. In the main text, we provide a definition from the perspective of a group action on the
 34 data directly, but the original definition considers an extra conceptual world state as well. Here, we
 35 describe the original setting in more detail, and explain why we choose a more direct and practical
 36 version of the definition.

37 The definition assumes the following setting. W is the set of possible world states, with underlying
 38 symmetry transformations that are described by a group G and its action $\cdot : G \times W \rightarrow W$ on W . In
 39 particular, G can be decomposed as the direct product of K groups $G = G_1 \times \dots \times G_K$. Data is
 40 obtained via an *observation* function $b : W \rightarrow X$ that maps world states to observations in a *data*
 41 *space* X . A model’s internal representation of data is modeled with the *encoding* function $h : X \rightarrow Z$
 42 that maps data to the *embedding space* Z . Together, the observation and the encoding constitute the
 43 model’s internal representation of the real world $f : W \rightarrow Z$ with $f(w) = h \circ b(w)$. The definition
 44 for Linearly Symmetry-Based Disentangled (LSBD) representations then formalizes the requirement
 45 that a model’s internal representation f should reflect and disentangle the transformation properties
 46 of the real world, and that the transformation properties of the model’s internal representations should
 47 be linear.

48 The original definition considers G acting on W and involves the model’s internal representation
 49 $f : W \rightarrow Z$, but since we do not directly observe W it is more practical to evaluate LSBD with
 50 respect to the encoding map $h : X \rightarrow Z$ instead. If the action of G on W is *regular*² and the
 51 observation map $b : W \rightarrow X$ is *injective*³ though, we can instead define LSBD with respect to the
 52 action of G on X and the encoding map h , as shown in the main text.

53 C Inner Product

54 To describe the norm $\|\cdot\|_{\rho, h, \mu}$ used in the definition of $\mathcal{D}_{\text{LSBD}}$ we start with an arbitrary inner
 55 product (\cdot, \cdot) on the linear latent space Z . Assume that ρ is linearly disentangled and accordingly
 56 splits in irreducible representations $\rho_k : G \rightarrow Z_k$ where $Z = Z_1 \oplus \dots \oplus Z_K$ for some $K \in \mathbb{N}$. We
 57 will define a new inner product $\langle \cdot, \cdot \rangle_{\rho, h, \mu}$ on Z as follows. First of all we declare Z_k and Z_m to be
 58 orthogonal with respect to $\langle \cdot, \cdot \rangle_{\rho, h, \mu}$ if $k \neq m$. We denote by π_k the orthogonal projection on Z_k .

59 For $z, z' \in Z_i$, we set

$$\langle z, z' \rangle_{\rho, h, \mu} := \lambda_{k, h, \mu}^{-1} \int_{g \in G} (\rho(g) \cdot z, \rho(g) \cdot z') d\mathbf{m}(g) \quad (2)$$

60 where \mathbf{m} is the (bi-invariant) Haar measure normalized such that $\mathbf{m}(G) = 1$ and set

$$\lambda_{k, h, \mu} := \int_X \int_G \|\pi_k(h(x))\|^2 d\mathbf{m}(g) d\mu(x) \quad (3)$$

61 if the integral on the right-hand side is strictly positive and otherwise we set $\lambda_k := 1$. This construction
 62 completely specifies the new inner product, and it has the following properties:

- 63 • the subspaces Z_k are mutually orthogonal,
- 64 • $\rho_k(g)$ is orthogonal on Z_k for every $g \in G$, in other words ρ_k maps to the orthogonal group
 65 on Z_k . Moreover, ρ maps to the orthogonal group on Z . This follows directly from the
 66 bi-invariance of the Haar measure and the definition of $\langle \cdot, \cdot \rangle_{\rho, h, \mu}$.

²This assumption holds in most practical cases with a suitable description of G .

³This is typically the case, but if not it can be solved through active sensing, see Soatto (2011).

67 • If π_k is the orthogonal projection to Z_k , then

$$\int_X \|\pi_k(h(x))\|_{\rho, h, \mu}^2 d\mu(x) = 1 \quad (4)$$

68 if the integral on the left is strictly positive.

69 For an arbitrary pair $z, z' \in Z$ the inner product $\langle \cdot, \cdot \rangle_{\rho, h, \mu}$ is given by

$$\langle z, z' \rangle_{\rho, h, \mu} = \sum_{k=1}^K \lambda_{k, h, \mu}^{-1} \int_{g \in G} (\rho(g) \cdot \pi_k(z), \rho(g) \cdot \pi_k(z')) dm(g) \quad (5)$$

70 D Evaluation of Equivariance by $\mathcal{D}_{\text{LSBD}}$

71 We will now give an alternative expression for the disentanglement metric $\mathcal{D}_{\text{LSBD}}$, since it will more
 72 visibly relate to the definition of equivariance. To avoid notational cluttering, in this section we will
 73 denote the norm $\|\cdot\|_{\rho, h, \mu}$ as $\|\cdot\|_*$. Let $\rho \in \mathcal{P}(G, Z)$ be a linear disentangled representation of G
 74 in Z . By expanding the inner product (or by using usual computation rules for expectations and
 75 variances), we first find that

$$\begin{aligned} & \int_G \left\| \rho(g)^{-1} \cdot h(g \cdot x_0) - \int_G \rho(g')^{-1} \cdot h(g' \cdot x_0) d\nu(g') \right\|_*^2 d\nu(g) \\ &= \int_G \|\rho(g)^{-1} \cdot h(g \cdot x_0)\|_*^2 d\nu(g) - \left\| \int_G \rho(g)^{-1} \cdot h(g \cdot x_0) d\nu(g) \right\|_*^2 \\ &= \frac{1}{2} \int_G \int_G \|\rho(g)^{-1} \cdot h(g \cdot x_0) - \rho(g')^{-1} \cdot h(g' \cdot x_0)\|_*^2 d\nu(g) d\nu(g'). \end{aligned} \quad (6)$$

76 We now use that ρ maps to the orthogonal group for $(\cdot, \cdot)_*$, so that we can write the same expression
 77 as

$$\frac{1}{2} \int_G \int_G \|\rho(g \circ g'^{-1})^{-1} \cdot h((g \circ g'^{-1}) \cdot g') \cdot x_0 - h(g' \cdot x_0)\|_*^2 d\nu(g) d\nu(g'). \quad (7)$$

78 This brings us to the alternative characterization of $\mathcal{D}_{\text{LSBD}}$ as

$$\mathcal{D}_{\text{LSBD}} = \inf_{\rho \in \mathcal{P}(G, Z)} \frac{1}{2} \int_G \int_G \|\rho(g \circ g'^{-1})^{-1} h((g \circ g'^{-1}) \cdot g') \cdot x_0 - h(g' \cdot x_0)\|_*^2 d\nu(g) d\nu(g'). \quad (8)$$

79 In particular, if for every data point x there is a unique group element g_x such that $x = g_x \cdot x_0$, the
 80 disentanglement metric $\mathcal{D}_{\text{LSBD}}$ can also be written as

$$\inf_{\rho \in \mathcal{P}(G, Z)} \frac{1}{2} \int_G \int_X \|\rho(g \circ g_x^{-1})^{-1} h((g \circ g_x^{-1}) \cdot x) - h(x)\|_*^2 d\nu(g) d\mu(x), \quad (9)$$

81 in which the equivariance condition appears prominently. The condition becomes even more apparent
 82 if ν is in fact the Haar measure itself, in which case the metric equals

$$\inf_{\rho \in \mathcal{P}(G, Z)} \frac{1}{2} \int_G \int_X \|\rho(g)^{-1} \circ h(g \cdot x) - h(x)\|_*^2 dm(g) d\mu(x). \quad (10)$$

83 E Datasets

84 All datasets contain 64×64 pixel images. The Square, Arrow and Airplane datasets have a known
 85 group decomposition $G = \text{SO}(2) \times \text{SO}(2)$ describing the underlying transformations. In these
 86 three datasets, for each subgroup a fixed number of $|\mathcal{G}_k| = 64$ with $k \in \{1, 2\}$ transformations is
 87 selected. Each image is generated from a single initial data point upon which all possible group
 88 actions are applied, resulting in datasets with $|\mathcal{G}_1| \cdot |\mathcal{G}_2| = 4096$ images. The datasets exemplify
 89 different group actions of $\text{SO}(2)$: periodic translations, in-plane rotations, out-of-plane rotations, and
 90 periodic hue-shifts, see Figure 1.

91 The ModelNet40 and the COIL-100 datasets consist of different objects rotating with respect to a
 92 vertical axis (out-of-plane rotation). For these datasets the group $G = \text{SO}(2)$ describes the underlying
 93 transformations that each object undergoes, see Figure 1. The different objects can be seen as non-
 94 symmetric variability in the data. In this particular case, each object has its own base-point x_0 from
 95 which data is generated. The metric $\mathcal{D}_{\text{LSBD}}$ is then evaluated per object instance for the group
 96 $G = \text{SO}(2)$, the value of $\mathcal{D}_{\text{LSBD}}$ is calculated and averaged across all available objects.

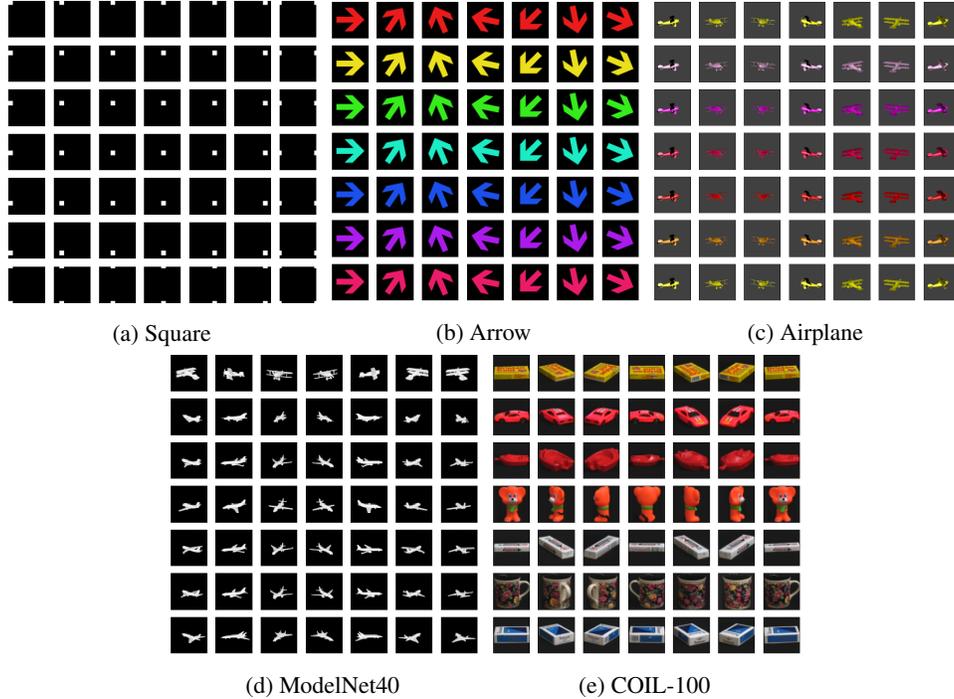


Figure 1: Example images from each of the datasets used. Each image corresponds to an example data point for a combination of two factors, e.g. color and orientation. The factors change horizontally and vertically. The boundaries for the Square, Arrow and Airplane dataset are periodic. For the ModelNet40 and COIL-100 dataset, the vertical direction represents different object instances and the horizontal direction represents the rotation of the corresponding object.

97 **Square** This dataset consists of a set of images of a black background with a square of 4×4 white
 98 pixels. The dataset is generated applying vertical and horizontal translations of the white square
 99 considering periodic boundaries.

100 **Arrow** This dataset consists of a set of images depicting a colored arrow at a given orientation. The
 101 dataset is generated by applying cyclic shifts of its color and in-plane rotations. The cyclic color
 102 shifts were obtained by preselecting a fixed set of 64 colors from a circular hue axis. The in-plane
 103 rotations were obtained by rotating the arrow along an axis perpendicular to the picture plane over 64
 104 predefined positions.

105 **Airplane** This dataset consists of renders obtained using Blender v2.7 (Community, 2020) from a
 106 3D model of an airplane within the ModelNet40 dataset (Wu et al., 2014) (this dataset is provided for
 107 the convenience of academic research only). We created each image by varying two properties: the
 108 airplane’s color and its orientation with respect to the render camera. The orientation was changed via
 109 rotation with respect to a vertical axis (out-of-plane rotation). The colors of the model were selected
 110 from a predefined cyclic set of colors similar to the arrow rotation dataset.

111 **ModelNet40** This dataset also consists of a dataset of renders obtained using Blender v2.7 (Com-
 112 munity, 2020) from all the 3D models within the airplane category of the ModelNet40 dataset (Wu
 113 et al., 2014). We created each image by varying each airplane’s orientation with respect to the render
 114 camera, via rotation with respect to a vertical axis (out-of-plane rotation). In this case we used 64
 115 orientations for each object, i.e. $|\mathcal{G}| = 64$.

116 **COIL-100** This dataset (Nene et al., 1996) consists of images from 100 objects placed on a turntable
 117 against a black background. For each object, 72 views of the rotated object are provided. The original
 118 images have a resolution of 128×128 and were re-scaled to 64×64 to match our other datasets. In

119 this case for each object $|\mathcal{G}| = 72$. This dataset is intended for non-commercial research purposes
 120 only. This dataset was obtained by using TensorFlowDatasets (2021).

121 F Experimental Settings and Hyperparameters

122 F.1 Architectures

123 Table 1 shows the encoder and decoder architectures used for almost all methods and datasets. The
 124 encoder’s last layer depends on the method. For VAE, cc-VAE, FactorVAE, DIP-I, DIP-II, two dense
 125 layers with 4 units each were used. For LSB-D-VAE and Δ VAE two dense layers with 4 and 2 units
 126 each were used. For Quessard a single dense layer with 4 units was used. The only model that was
 127 not trained with this architectures was LSB-D-VAE/0 method for the ModelNet40 dataset the reason
 128 for this choice was that during training the loss was getting NaN values, in this case the architecture
 129 used was that of Table 2.

Table 1: Encoder and decoder architectures used in most methods.

ENCODER	
INPUT	SIZE (64,64, NUMBER CHANNELS)
CONV	FILTERS 32, KERNEL 4, STRIDE 2, RELU
CONV	FILTERS 32, KERNEL 4, STRIDE 2, RELU
CONV	FILTERS 64, KERNEL 4, STRIDE 2, RELU
CONV	FILTERS 64, KERNEL 4, STRIDE 2, RELU
DENSE	UNITS 256, RELU
DENSE(X2)	UNITS DEPEND ON METHOD
DECODER	
INPUT	SIZE (NUMBER OF LATENT DIMENSIONS)
DENSE	UNITS 256, RELU
DENSE	UNITS 4*4*64, RELU
RESHAPE	(4,4,64)
CONVT	FILTERS 64, KERNEL 4, STRIDE 2, RELU
CONVT	FILTERS 32, KERNEL 4, STRIDE 2, RELU
CONVT	FILTERS 32, KERNEL 4, STRIDE 2, RELU
CONVT	FILTERS (NUMBER CHANNELS), KERNEL 4, STRIDE 2, SIGMOID

Table 2: Encoder and decoder architecture used to train LSB-D-VAE/0 for ModelNet40 dataset.

ENCODER	
INPUT	SIZE (64, 64, NUMBER CHANNELS)
DENSE	UNITS 512, RELU, BATCH NORMALIZATION
DENSE	UNITS 256, RELU, BATCH NORMALIZATION
DENSE(X2)	UNITS DEPEND ON METHOD
DECODER	
INPUT	SIZE (NUMBER OF LATENT DIMENSIONS)
DENSE	UNITS 256, RELU, BATCH NORMALIZATION
DENSE	UNITS 512, RELU, BATCH NORMALIZATION
DENSE	UNITS 64*64*NUMBER OF CHANNELS, SIGMOID
RESHAPE	(64, 64, NUMBER OF CHANNELS)

130 F.2 Hyperparameters

131 Table 3 shows the hyperparameters used to train each model for all datasets. Table 4 shows the
 132 hyperparameters used to train the LSB-D-VAE models for each dataset. In the latter case, the number
 133 of epochs for the LSB-D-VAE model were increased. The range of values used for the scale parameter
 134 t were increased for ModelNet40 and COIL-100 datasets since it was noticed that this provided

135 better results in terms of data reconstruction and disentanglement. For the Arrow dataset, a value
 136 of $\gamma = 1$ was producing unstable results. However, the values 10, 100, 1000 or even 10000 were
 137 producing good results without significant changes among them. Therefore the value 100 was used
 138 for the datasets with the same structure (Square, Arrow and Airplane). For the ModelNet40 and
 139 COIL-100 the experiments showed that this hyperparameter for values as high as 10000 could affect
 140 the reconstructions, thus a lower value $\gamma = 1$ was chosen.

Table 3: Model hyperparameters for all datasets

MODEL	PARAMETERS
VAE	TRAINING STEPS 30000
β -VAE	$\beta = 5$, TRAINING STEPS 30000
CC-VAE	$\beta = 5, \gamma = 1000, c_{max} = 15$, ITERATION THRESHOLD 3500, TRAINING STEPS 30000
FACTOR	$\gamma = 1$, EPOCHS 30000
DIP-I	$\lambda_{od} = 1, \lambda_d = 10$, TRAINING STEPS 30000
DIP-II	$\lambda_{od} = 1, \lambda_d = 1$, TRAINING STEPS 30000
QUESSARD	$\lambda = 0.01$, TRAJECTORIES 3000

Table 4: LSB-D-VAE hyperparameters for all datasets

DATASETS	PARAMETERS
SQUARE, ARROW, AIRPLANE	$t \in [10^{-10}, 10^{-9}]$, $\gamma = 100.0$, EPOCHS 1500
MODELNET40	$t \in [10^{-10}, 10^{-5}]$, $\gamma = 1.0$, EPOCHS 1500
COIL-100	$t \in [10^{-10}, 10^{-5}]$, $\gamma = 1.0$, EPOCHS 6000

141 F.3 Hardware & Running Time

142 The hardware used across all experiments was a DGX station with 4 NVIDIA GPUs V100 and
 143 32GB . Only one GPU was used per experiment. The running time for the LSB-D-VAE across all 9
 144 degrees of supervision $L \in \{0, 256, 768, 1024, 1280, 1536, 1792, 2048\}$ and all 10 runs (total $9 \cdot 10$
 145 repetitions) for the datasets were: Arrow 33 ± 4 minutes Airplane 29 ± 4 minutes and Square 28 ± 4
 146 minutes. The running time for the LSB-D-VAE across 2 degrees of supervision and 10 runs (total
 147 $2 \cdot 10$ repetitions) for ModelNet40 was 136 ± 10 minutes and for COIL-100 90 ± 6 minutes. For
 148 the method from (Quessard et al., 2020) the training times were approximately 30 minutes across all
 149 datasets. The training times for the methods from `disentanglement_lib` (Locatello et al., 2019)
 150 were not measured.

151 F.4 Code Licenses

152 The `disentanglement_lib` (Locatello et al., 2019) code is registered with an Apache 2.0 License
 153 while the code used to reproduce the method by Quessard et al. (2020) is registered with an MIT
 154 license.

155 G Full results

156 The full results for all experiments on all datasets are given in Tables 5, 6, 7, 8, and 9. We report the
 157 mean and standard deviation over 10 runs for each experiment.

158 G.1 Limited Supervision Suffices to Learn LSB-D Representations

159 The results obtained from Tables 5, 6, 7 show that we do not need transformation-labels for all data
 160 points, only a subset of labeled pairs is sufficient to learn LSB-D representations. To further highlight
 161 this, Figure 2 shows $\mathcal{D}_{\text{LSB-D}}$ scores for LSB-D-VAE trained on the Square, Arrow, and Airplane
 162 datasets respectively, for various values for the number of labeled pairs L . For each L and each
 163 dataset, we trained 10 models so we can report box plots of the $\mathcal{D}_{\text{LSB-D}}$ scores.

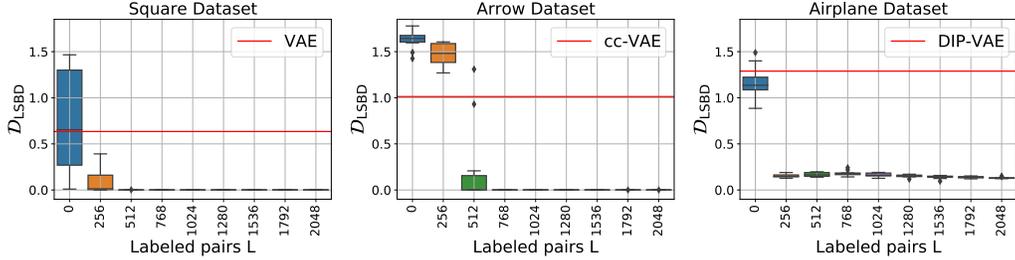


Figure 2: Box plots for $\mathcal{D}_{\text{LSBD}}$ scores over 10 training repetitions for different numbers of labeled pairs L , for all datasets. The red line indicates the best-performing traditional disentanglement method.

164 For low values of L we see worse scores and high variability. But for slightly higher L , scores
 165 are consistently good, starting already at $L = 512$ for the Square, $L = 768$ for the Arrow, and
 166 $L = 256$ for the Airplane. This corresponds to respectively 25%, 37.5%, and 12.5% of the data being
 167 involved in a labeled pair. Moreover, we see that with just a little supervision we outperform the best
 168 traditional method on $\mathcal{D}_{\text{LSBD}}$. Overall, these results suggest that with some expert knowledge (about
 169 the underlying group and a suitable representation) and limited annotation of transformations, LSBD
 170 can be achieved.

171 G.2 Quessard Arrow

172 In the main text we mentioned that we did not reproduce good results with Quessard et al. (2020)’s
 173 method on the Arrow and Square dataset. We highlight a particular case for the Arrow dataset,
 174 where the method clearly learns the rotations of the arrow but fails to learn color. Figure 3 shows
 175 reconstructed Arrow images. Since color isn’t learned well, this example doesn’t get a good $\mathcal{D}_{\text{LSBD}}$
 176 score, even though rotation is properly linearly disentangled.

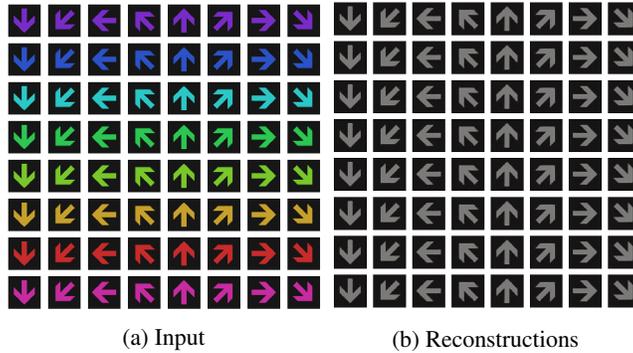


Figure 3: Results from Quessard et al. (2020)’s method on the Arrow dataset

177 References

178 Community, B. O. (2020). *Blender - a 3D modelling and rendering package*. Blender Foundation,
 179 Stichting Blender Foundation, Amsterdam.

180 Hall, B. C. (2015). *Lie Groups, Lie Algebras, and Representations*, volume 222 of *Graduate Texts in*
 181 *Mathematics*. Springer International Publishing, Cham.

182 Higgins, I., Amos, D., Pfau, D., Racaniere, S., Matthey, L., Rezende, D., and Lerchner, A. (2018).
 183 Towards a definition of disentangled representations. *arXiv preprint arXiv:1812.02230*.

184 Locatello, F., Bauer, S., Lucic, M., Gelly, S., Schölkopf, B., and Bachem, O. (2019). Challenging
 185 common assumptions in the unsupervised learning of disentangled representations. In *International*
 186 *Conference on Machine Learning*.

187 Nene, S. A., Nayar, S. K., Murase, H., et al. (1996). Columbia object image library (coil-20).

188 Quessard, R., Barrett, T. D., and Clements, W. R. (2020). Learning Group Structure and Disentangled
189 Representations of Dynamical Environments. *Advances in Neural Information Processing Systems*,
190 33.

191 Soatto, S. (2011). Steps Towards a Theory of Visual Information: Active Perception, Signal-
192 to-Symbol Conversion and the Interplay Between Sensing and Control. *arXiv preprint*
193 *arXiv:1110.2053*.

194 TensorFlowDatasets, . (2021). TensorFlow Datasets, a collection of ready-to-use datasets. <https://www.tensorflow.org/datasets>.

195

196 Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., and Xiao, J. (2014). 3D ShapeNets: A
197 Deep Representation for Volumetric Shapes.

Table 5: Scores for the Square dataset.

MODEL	BETA \uparrow	FACTOR \uparrow	SAP \uparrow	DCI \uparrow	MIG \uparrow	MOD \uparrow	$\mathcal{D}_{\text{LSBD}}$ \downarrow
VAE	.945 \pm .061	.835 \pm .140	.019 \pm .004	.009 \pm .005	.013 \pm .004	.579 \pm .202	.634 \pm .440
β -VAE	.980 \pm .033	.913 \pm .095	.021 \pm .006	.017 \pm .011	.021 \pm .014	.642 \pm .147	.732 \pm .488
CC-VAE	.508 \pm .023	.000 \pm .000	.003 \pm .002	.007 \pm .002	.014 \pm .004	.222 \pm .110	1.905 \pm .023
FACTOR	.974 \pm .048	.910 \pm .104	.020 \pm .003	.019 \pm .017	.017 \pm .010	.712 \pm .183	.667 \pm .428
DIP-I	.972 \pm .042	.861 \pm .097	.020 \pm .005	.010 \pm .002	.011 \pm .002	.618 \pm .117	1.109 \pm .312
DIP-II	.930 \pm .119	.848 \pm .137	.018 \pm .004	.010 \pm .004	.015 \pm .007	.607 \pm .207	.907 \pm .559
QUESSARD	.504 \pm .021	.000 \pm .000	.004 \pm .003	.007 \pm .004	.018 \pm .008	.354 \pm .213	1.686 \pm .294
LSBD-VAE /0	.970 \pm .079	.913 \pm .121	.018 \pm .003	.052 \pm .052	.018 \pm .004	.884 \pm .183	.749 \pm .554
LSBD-VAE /256	1.000 \pm .000	1.000 \pm .001	.021 \pm .004	.267 \pm .152	.027 \pm .007	.986 \pm .023	.104 \pm .147
LSBD-VAE /512	1.000 \pm .000	1.000 \pm .000	.021 \pm .006	.393 \pm .022	.025 \pm .005	.999 \pm .000	.000 \pm .000
LSBD-VAE /768	1.000 \pm .000	1.000 \pm .000	.019 \pm .004	.387 \pm .014	.025 \pm .004	.999 \pm .000	.000 \pm .000
LSBD-VAE /1024	1.000 \pm .000	1.000 \pm .000	.022 \pm .005	.398 \pm .020	.024 \pm .003	.999 \pm .000	.000 \pm .000
LSBD-VAE /1280	1.000 \pm .000	1.000 \pm .000	.023 \pm .003	.389 \pm .016	.023 \pm .003	.999 \pm .000	.000 \pm .000
LSBD-VAE /1536	1.000 \pm .000	1.000 \pm .000	.022 \pm .004	.398 \pm .013	.027 \pm .002	.999 \pm .000	.000 \pm .000
LSBD-VAE /1792	1.000 \pm .000	1.000 \pm .000	.020 \pm .004	.397 \pm .016	.027 \pm .005	.999 \pm .000	.000 \pm .000
LSBD-VAE /2048	1.000 \pm .000	1.000 \pm .000	.021 \pm .006	.380 \pm .027	.027 \pm .005	.999 \pm .000	.000 \pm .000

Table 6: Scores for the Arrow dataset.

MODEL	BETA \uparrow	FACTOR \uparrow	SAP \uparrow	DCI \uparrow	MIG \uparrow	MOD \uparrow	$\mathcal{D}_{\text{LSBD}}$ \downarrow
VAE	1.000 \pm .000	.646 \pm .032	.017 \pm .004	.009 \pm .003	.013 \pm .004	.961 \pm .012	1.316 \pm .193
β -VAE	.999 \pm .002	.588 \pm .045	.018 \pm .004	.008 \pm .002	.015 \pm .005	.898 \pm .032	1.178 \pm .065
CC-VAE	.982 \pm .056	.707 \pm .102	.019 \pm .004	.011 \pm .005	.016 \pm .004	.980 \pm .038	1.013 \pm .096
FACTOR	1.000 \pm .000	.659 \pm .028	.017 \pm .003	.008 \pm .003	.014 \pm .002	.935 \pm .037	1.526 \pm .125
DIP-I	1.000 \pm .000	.624 \pm .042	.020 \pm .004	.008 \pm .002	.012 \pm .003	.967 \pm .027	1.521 \pm .113
DIP-II	1.000 \pm .000	.644 \pm .064	.020 \pm .004	.009 \pm .003	.013 \pm .004	.973 \pm .011	1.616 \pm .102
QUESSARD	1.000 \pm .000	.596 \pm .032	.016 \pm .006	.008 \pm .004	.017 \pm .008	.999 \pm .000	1.183 \pm .412
LSBD-VAE /0	1.000 \pm .001	.664 \pm .105	.016 \pm .002	.009 \pm .004	.019 \pm .005	.897 \pm .108	1.627 \pm .104
LSBD-VAE /256	1.000 \pm .000	.662 \pm .046	.017 \pm .005	.009 \pm .004	.020 \pm .005	.963 \pm .010	1.475 \pm .121
LSBD-VAE /512	1.000 \pm .000	.956 \pm .119	.021 \pm .006	.297 \pm .157	.023 \pm .003	.967 \pm .092	.245 \pm .474
LSBD-VAE /768	1.000 \pm .000	1.000 \pm .000	.022 \pm .006	.390 \pm .022	.026 \pm .003	.999 \pm .000	.000 \pm .000
LSBD-VAE /1024	1.000 \pm .000	1.000 \pm .000	.022 \pm .003	.396 \pm .026	.026 \pm .006	.999 \pm .000	.000 \pm .000
LSBD-VAE /1280	1.000 \pm .000	1.000 \pm .000	.019 \pm .005	.401 \pm .018	.026 \pm .004	.999 \pm .000	.000 \pm .000
LSBD-VAE /1536	1.000 \pm .000	1.000 \pm .000	.019 \pm .005	.397 \pm .017	.026 \pm .007	.999 \pm .000	.000 \pm .000
LSBD-VAE /1792	1.000 \pm .000	1.000 \pm .000	.020 \pm .004	.399 \pm .018	.026 \pm .004	.999 \pm .000	.000 \pm .000
LSBD-VAE /2048	1.000 \pm .000	1.000 \pm .000	.020 \pm .006	.444 \pm .186	.027 \pm .004	.999 \pm .000	.000 \pm .000

Table 7: Scores for the Airplane dataset.

MODEL	BETA \uparrow	FACTOR \uparrow	SAP \uparrow	DCI \uparrow	MIG \uparrow	MOD \uparrow	$\mathcal{D}_{\text{LSBD}} \downarrow$
VAE	1.000 \pm .001	.947 \pm .054	.023 \pm .005	.013 \pm .005	.020 \pm .017	.801 \pm .045	1.342 \pm .084
β -VAE	1.000 \pm .001	.997 \pm .005	.018 \pm .005	.036 \pm .012	.028 \pm .012	.816 \pm .104	1.481 \pm .129
CC-VAE	.858 \pm .194	.646 \pm .353	.010 \pm .006	.021 \pm .011	.018 \pm .009	.969 \pm .034	1.481 \pm .174
FACTOR	1.000 \pm .000	.984 \pm .015	.020 \pm .003	.021 \pm .008	.026 \pm .013	.810 \pm .040	1.382 \pm .171
DIP-I	1.000 \pm .000	.994 \pm .008	.022 \pm .004	.029 \pm .012	.026 \pm .012	.842 \pm .073	1.289 \pm .150
DIP-II	.998 \pm .005	.972 \pm .031	.021 \pm .004	.022 \pm .013	.030 \pm .019	.780 \pm .054	1.367 \pm .129
QUESSARD	.999 \pm .003	.987 \pm .026	.018 \pm .007	.016 \pm .009	.018 \pm .005	.795 \pm .107	.558 \pm .239
LSBD-VAE /0	.536 \pm .065	.000 \pm .000	.002 \pm .001	.007 \pm .004	.005 \pm .003	.956 \pm .046	1.165 \pm .180
LSBD-VAE /256	1.000 \pm .000	1.000 \pm .000	.022 \pm .006	.144 \pm .011	.023 \pm .004	.870 \pm .039	.153 \pm .021
LSBD-VAE /512	1.000 \pm .000	1.000 \pm .000	.023 \pm .008	.151 \pm .015	.020 \pm .004	.846 \pm .032	.168 \pm .022
LSBD-VAE /768	1.000 \pm .000	1.000 \pm .000	.022 \pm .004	.140 \pm .014	.022 \pm .005	.832 \pm .034	.180 \pm .030
LSBD-VAE /1024	1.000 \pm .000	1.000 \pm .000	.020 \pm .005	.160 \pm .015	.022 \pm .005	.859 \pm .032	.165 \pm .021
LSBD-VAE /1280	1.000 \pm .000	1.000 \pm .000	.024 \pm .004	.153 \pm .013	.022 \pm .003	.876 \pm .016	.151 \pm .015
LSBD-VAE /1536	1.000 \pm .000	1.000 \pm .000	.021 \pm .005	.160 \pm .016	.022 \pm .004	.896 \pm .025	.140 \pm .018
LSBD-VAE /1792	1.000 \pm .000	1.000 \pm .000	.022 \pm .005	.163 \pm .022	.023 \pm .003	.904 \pm .016	.138 \pm .010
LSBD-VAE /2048	1.000 \pm .000	1.000 \pm .000	.016 \pm .008	.161 \pm .024	.021 \pm .006	.913 \pm .018	.132 \pm .009

Table 8: Scores for the Modelnet40 Airplanes dataset.

MODEL	BETA \uparrow	FACTOR \uparrow	SAP \uparrow	DCI \uparrow	MIG \uparrow	MOD \uparrow	$\mathcal{D}_{\text{LSBD}} \downarrow$
VAE	.995 \pm .004	.838 \pm .030	.013 \pm .002	.013 \pm .002	.009 \pm .002	.415 \pm .058	.393 \pm .110
β -VAE	.995 \pm .005	.857 \pm .045	.012 \pm .003	.015 \pm .003	.009 \pm .002	.447 \pm .067	.285 \pm .045
CC-VAE	.997 \pm .003	.818 \pm .093	.011 \pm .003	.017 \pm .004	.011 \pm .003	.567 \pm .063	.281 \pm .191
FACTOR	.996 \pm .004	.856 \pm .052	.012 \pm .002	.014 \pm .003	.010 \pm .003	.444 \pm .077	.388 \pm .096
DIP-I	.988 \pm .009	.783 \pm .070	.012 \pm .002	.013 \pm .002	.008 \pm .001	.343 \pm .082	.416 \pm .142
DIP-II	.994 \pm .006	.832 \pm .042	.013 \pm .003	.014 \pm .003	.011 \pm .002	.433 \pm .080	.379 \pm .130
QUESSARD	.907 \pm .192	.727 \pm .384	.010 \pm .005	.015 \pm .007	.009 \pm .004	.563 \pm .108	.134 \pm .294
LSBD-VAE /0	.990 \pm .009	.863 \pm .038	.011 \pm .003	.015 \pm .003	.014 \pm .003	.538 \pm .103	.731 \pm .068
LSBD-VAE /256	1.000 \pm .000	.990 \pm .004	.012 \pm .005	.052 \pm .009	.020 \pm .006	.947 \pm .007	.041 \pm .007

Table 9: Scores for COIL 100 dataset.

MODEL	BETA \uparrow	FACTOR \uparrow	SAP \uparrow	DCI \uparrow	MIG \uparrow	MOD \uparrow	$\mathcal{D}_{\text{LSBD}} \downarrow$
VAE	1.000 \pm .000	.674 \pm .049	.014 \pm .003	.016 \pm .003	.011 \pm .002	.986 \pm .001	.463 \pm .030
β -VAE	1.000 \pm .001	.740 \pm .024	.015 \pm .004	.014 \pm .004	.013 \pm .003	.982 \pm .001	.579 \pm .095
CC-VAE	.999 \pm .003	.723 \pm .026	.013 \pm .005	.014 \pm .003	.013 \pm .004	.985 \pm .001	.406 \pm .057
FACTOR	1.000 \pm .001	.684 \pm .041	.014 \pm .002	.012 \pm .002	.013 \pm .004	.984 \pm .001	.490 \pm .024
DIP-I	.999 \pm .002	.631 \pm .025	.013 \pm .004	.012 \pm .002	.010 \pm .002	.986 \pm .001	.525 \pm .109
DIP-II	1.000 \pm .001	.643 \pm .043	.013 \pm .003	.014 \pm .002	.011 \pm .002	.985 \pm .001	.568 \pm .079
QUESSARD	1.000 \pm .000	.780 \pm .044	.014 \pm .004	.014 \pm .002	.011 \pm .003	.973 \pm .004	.396 \pm .055
LSBD-VAE /0	1.000 \pm .001	.739 \pm .047	.014 \pm .003	.014 \pm .001	.011 \pm .001	.982 \pm .004	.515 \pm .099
LSBD-VAE /256	1.000 \pm .000	.655 \pm .028	.015 \pm .004	.029 \pm .003	.013 \pm .003	.802 \pm .056	.112 \pm .026