# Symmetry-adapted representation learning 

<br>${ }^{\text {a }}$ The Center for Brains, Minds, and Machines, MIT and McGovern Institute for Brain Research at MIT, Cambridge, MA, USA<br>${ }^{\mathrm{b}}$ Laboratory for Computational and Statistical Learning (LCSL), Istituto Italiano di Tecnologia, Genova, Italy

Keywords:
Representation learning
Equivariant representations
Invariant representations
Dictionary learning
Convolutional neural networks
Regularization
Data transformations


#### Abstract

In this paper, we propose the use of data symmetries, in the sense of equivalences under signal transformations, as priors for learning symmetry-adapted data representations, i.e., representations that are equivariant to these transformations. We rely on a group-theoretic definition of equivariance and provide conditions for enforcing a learned representation, for example the weights in a neural network layer or the atoms in a dictionary, to have the structure of a group and specifically the group structure in the distribution of the input. By reducing the analysis of generic group symmetries to permutation symmetries, we devise a regularization scheme for representation learning algorithm, using an unlabeled training set. The proposed regularization is aimed to be a conceptual, theoretical and computational proof of concept for symmetry-adapted representation learning, where the learned data representations are equivariant or invariant to transformations, without explicit knowledge of the underlying symmetries in the data.


## 1. Introduction

Symmetry is ubiquitous from subatomic particles to natural patterns, man-made design, art and mathematics. Invariance to symmetries in pattern recognition and computational neuroscience is an old challenging problem [1-12]. More recently, in the context of machine learning, data symmetries have been used to derive data representations with the properties of equivariance and invariance [13-17] to unknown, symmetry-generating transformations, for example geometric transformations. These properties are reflected as structure in the representation atoms and can be explicitly used for reducing the complexity of downstream supervised learning. This is achieved, for example, by constructing representations that are invariant to transformations irrelevant for the learning task, that preserve the data distribution and prediction function [18-20]. For example, for image classification, object position or scale are data symmetries that are irrelevant.

Representations that reflect symmetries inherent in the data distribution define a quotient representation space where points are equivalent up to transformations [21]. In this space, the sample complexity of learning (the size of the labeled training set) [ $16,22,23$ ] can be reduced by pooling on the representation coef-

[^0]ficients. Indeed, the pooling operation has a crucial role in Convolutional Neural Networks (CNNs) for enforcing stability to small, local perturbations [24,25]. On the other hand, learning symmetryadapted representations is in the direction of (a) generalizing CNNs to arbitrary weight sharing schemes and invariances by learning the symmetry group from the data and (b) learning, as opposed to informed designing, network architectures and feature map properties, such as locality, connectivity patterns and weight-sharing topologies.

CNNs and Convolutional Sparse Coding schemes [26] have an explicit parameterization for equivariance and robustness to shifts in the input (translations) through convolutions and pooling respectively (see also [27]). However, data symmetries extend to more general transformations depending on the data domain, for example geometric changes such as scaling, rotation or affine maps for the case of images, which, in general are unknown or complex to model. A symmetry-blind data representation will have to compensate for transformation variability using more parameters, and as a result an increased demand for labeled examples or data augmentation and adaptation assuming simple and known transformation models [28,29]. Extensions of CNNs to known transformations, beyond translations, were explored with scale-space pooling [30], convolutional maxout networks [31], pooling over neighboring values and similar filters [32], tiled CNNs [33], cyclic weight sharing and pooling [34], and wavelet scattering networks [35,36], (but see also [37]). In particular symmetry networks [23] and groupequivariant networks $[17,38]$ highlighted the complexity gains of incorporating other symmetries in the representations at each layer
of CNNs. In [39] the parameters of a neural network are tied to achieve equivariance with respect to a known group. Weight sharing over transformations capturing perceptual changes, such as speaker characteristics, were used for speech representations through group-CNNs [40].

The more general problem of learning symmetries has been previously approached as estimating the infinitesimal generators of Lie groups generating data transformations [41-44]. Symmetries in learning have been used in the context of categorizing symmetry groups (mirror, roto-translation) in random patterns [45]. The relations between typical, e.g., $l_{1}, l_{2}, l_{\infty}$, and group-based regularization schemes, for known groups, have been explored in [46].

The contributions of this work are: (1) Outlining principles for learning symmetries in data, and learning equivariant representations, without explicit knowledge of the symmetry group. As opposed to learning with known symmetries, like many existing methods, we propose learning the symmetries; (2) Reducing the analysis of generic group symmetries to permutation symmetries; (3) Deriving an analytic expression for a regularization term acting on the representation matrix that promotes a group structure and specifically the group structure in an unlabeled observation set.

The rest of the paper is organized as follows: In Section 2, we briefly recall the setting of representations with equivariance or invariance to transformations captured by group symmetries [14,16,17,21,47]. In Section 3 we formulate the problem of symmetry-adapted representation learning and provide a general principle (Section 4) for designing the regularization term. The main theoretical contributions are stated in Section 5 along with a computable, analytic form for a regularization term. Section 6 provides proof of concept results on learning exact, analytic, grouptransformations.

## 2. Equivariant and invariant representations

We briefly recall the setting of constructing data representations with equivariance or invariance to transformations using group symmetries [14,16,17,21,47]. Intuitively, a representation is equivariant with respect to a transformation on the input space, if it can be equivalently expressed as some transformation on the representation space. If this is the identity, the representation is invariant to the transformation.

Let the input space be a vector space endowed with dot product $\langle\cdot, \cdot\rangle, \mathcal{X}=\mathbb{R}^{d}$. We denote the transformations of a point $x \in \mathcal{X}$ through a representation of a group $\mathcal{G}$, of order $|\mathcal{G}|<\infty$ in $d$ dimensions. In the following, we identify $\mathcal{G}$ with the set of $d \times d$ matrices $\{g\}$, (e.g. rotation or reflection matrices), and each transformation by $g x$. The group orbit of $x \in \mathcal{X}$ is the set of transformed signals

$$
\begin{equation*}
O_{x}=\{g x \in \mathcal{X} \mid g \in \mathcal{G}\}, x \in \mathcal{X} \tag{1}
\end{equation*}
$$

Thus the action of the group, and the associated orbit definition, induces a partition of $\mathcal{X}$ into orbit sets. More precisely, it defines an equivalence relation

$$
\begin{equation*}
x \sim x^{\prime} \Leftrightarrow \exists g \in \mathcal{G}: x^{\prime}=g x, \forall x, x^{\prime} \in \mathcal{X} \tag{2}
\end{equation*}
$$

that separates points in $\mathcal{X}$ on the basis of $x$ being a transformation of $x^{\prime}$.

Orbits can be used to derive a representation $\Phi: \mathcal{X} \rightarrow$ $\mathcal{F}$ (e.g. $\mathbb{R}^{k}$ ) selected from some hypothesis space of maps with a specific parametrization. A parametrization for an equivariant, or invariant as a special case, representation is a nonlinear measurements of the projections on the orbit elements [16,21]. More specifically, given an orbit of group $\mathcal{G}$ of some vector $t_{j} \in \mathcal{X}$ written in a matrix form as

$$
\begin{equation*}
W_{j}=\left[g_{1} t_{j}, \ldots, g_{|\mathcal{G}|} t_{j}\right], t_{j} \in \mathcal{X} \tag{3}
\end{equation*}
$$

the nonlinear projection of $x$ on $W_{j}$ is

$$
\begin{equation*}
\Phi_{j, \alpha}(x)=\sigma_{\alpha}\left(W_{j}^{T} x\right)=\left(\sigma_{\alpha}\left(\left\langle g_{1} t_{j}, x\right\rangle\right), \ldots, \sigma_{\alpha}\left(\left\langle g_{|\mathcal{G}|} t_{j}, x\right\rangle\right)\right)^{T} \tag{4}
\end{equation*}
$$

where $\left\{\sigma_{\alpha}: \mathbb{R} \rightarrow \mathbb{R} \mid \alpha \in \mathbb{R}\right\}$ is a set of nonlinear functions with a scalar parameter $\alpha$, e.g. sigmoids $\sigma_{\alpha}(\cdot)=\left(1+e^{-\alpha}\right)^{-1}$ or rectifier functions $\sigma_{\alpha}(\cdot)=\max (0, \cdot-\alpha)$.

It is easy to see, using the closure property of the group composition, that $\Phi_{j, \alpha}: \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{G}|}$ is a permutation-equivariant representation w.r.t. the transformations in $\mathcal{G}$ i.e.:

$$
\begin{equation*}
\Phi_{j, \alpha}(g x)=P_{g} \Phi_{j, \alpha}(x), g \in \mathcal{G}, \forall j, \alpha \tag{5}
\end{equation*}
$$

where $P_{g} \in \mathbb{R}^{|\mathcal{G}| \times|\mathcal{G}|}$ is a permutation matrix that depends on the transformation $g$. Indeed, transforming $x$ amounts to simply reordering the entries in vector (4). This can be easily seen noting that each entry $i$ in the vector in (4) is

$$
\left(\Phi_{j, \alpha}(g x)\right)_{i}=\sigma_{\alpha}\left(\left\langle g_{i} t, g x\right\rangle\right)=\sigma_{\alpha}\left(\left\langle g^{-1} g_{i} t, x\right\rangle\right)
$$

and due to the closure property of the group composition, i.e. $g_{i} g_{k}=g_{l}, \forall i, k \exists l$, we have that setting $x \rightarrow g x$ in $\Phi_{j, \alpha}(x)$ amounts to a permutation of the vector entries.

Further, by summing the components of $\Phi_{j, \alpha}(x)$ (or any pointwise function of them), we obtain a representation $\bar{\Phi}_{j, \alpha}: \mathcal{X} \rightarrow \mathbb{R}$ invariant to transformations in $\mathcal{G}$ :

$$
\begin{equation*}
\bar{\Phi}_{j, \alpha}(x)=\sum_{i=1}^{|\mathcal{G}|} \sigma_{\alpha}\left(\left\langle g_{i} t_{j}, x\right\rangle\right)=\mathbf{1}_{|\mathcal{G}|}^{T} \Phi_{j, \alpha}(x) \tag{6}
\end{equation*}
$$

where $\alpha \in \mathbb{R}$ and $\mathbf{1}_{|\mathcal{G}|}$ is the all-ones vector of dimension $|\mathcal{G}|$. Indeed, using the permutation equivariance in (5) we have for all $g \in \mathcal{G}$ :

$$
\begin{align*}
\bar{\Phi}_{j, \alpha}(g x) & =\mathbf{1}_{|\mathcal{G}|}^{T} \Phi_{j, \alpha}(g x)=\mathbf{1}_{|\mathcal{G}|}^{T} P_{g} \Phi_{j, \alpha}(x) \\
& =\mathbf{1}_{|\mathcal{G}|}^{T} \Phi_{j, \alpha}(x)=\bar{\Phi}_{j, \alpha}(x) \tag{7}
\end{align*}
$$

since $\mathbf{1}_{|\mathcal{G}|}^{T} P_{g}=\mathbf{1}_{|\mathcal{G}|}^{T}$. Another example of invariant quantity, being $P_{g}$ unitary, is the $\ell^{2}$ norm of the representation vector:

$$
\begin{equation*}
\left\|\Phi_{j, \alpha}(g x)\right\|_{2}^{2}=\left\|P_{g} \Phi_{j, \alpha}(x)\right\|_{2}^{2}=\left\|\Phi_{j, \alpha}(x)\right\|_{2}^{2} \tag{8}
\end{equation*}
$$

Furthermore, given a set of $Q$ such orbits $\left\{W_{j}\right\}_{j=1}^{Q}$ of the same group $\mathcal{G}$, the map $\Phi: \mathcal{X} \rightarrow \mathbb{R}^{Q \times\left|\left\{\sigma_{\alpha}\right\}\right|}$ obtained by concatenating $\left\{\bar{\Phi}_{j, \alpha}\right\}$ can be shown to be selective to the partition of $\mathcal{X}$ by $\mathcal{G}$, or equivalently, sufficient for separating the equivalence classes induced by the group action [16].

## 3. Problem formulation

In the above sense, deriving equivariant and invariant representations is conditioned upon having access to an orbit set $W$, or learning a set $W$ such that it reflects the same generating symmetries of the target group. In this paper, we focus on learning such a symmetry set, without supervision, from an observation set for which we make the following simplifying assumption:
Assumption 1. The observation set $S_{N}=\left\{x_{i}\right\}_{i=1}^{N} \subset \mathcal{X}$ is a finite collection of $Q$ orbits in $\mathcal{X}=\mathbb{R}^{d}$ w.r.t. a finite group $\mathcal{G}$

$$
\begin{equation*}
S_{N}=\left\{x_{i}\right\}_{i=1}^{N}=\left\{g x_{j}, \forall g \in \mathcal{G}\right\}_{j=1}^{Q}, x_{i}, x_{j} \in \mathcal{X} \tag{9}
\end{equation*}
$$

where $N=|\mathcal{G}| Q$ and $|\mathcal{G}|$ is the group cardinality.
A symmetry-adapted representation learning problem will be formulated as learning $W$ from data $S_{N}$ such that

$$
\begin{equation*}
\underset{W \in \mathbb{R}^{d \times|G|}}{\arg \min }\left(\mathcal{L}\left(W, S_{N}\right)+\gamma \mathcal{L}^{\prime}\left(W, S_{N}\right)+\beta \mathcal{R}(W)\right) \tag{10}
\end{equation*}
$$

with $\beta, \gamma \in \mathbb{R}_{+}$where $\mathcal{L}\left(W, S_{N}\right)$ is the representation loss selected to satisfy some objective criterion, e.g. reconstruction, similarity or
empirical error (supervised) [24]; $\mathcal{R}: \mathbb{R}^{d \times|\mathcal{G}|} \rightarrow \mathbb{R}_{+}$is a regularization term, controlled by $\beta$, that enforces the columns of $W$ to be an orbit of a finite group; $\mathcal{L}^{\prime}\left(W, S_{N}\right)$ is a data-dependent regularization term to further restrict the group in $W$ to be the same as the latent group in the observation set.

The contribution of this paper is to develop analytic formulations for $\mathcal{L}^{\prime}$ and $\mathcal{R}$ assuming $S_{N}$ as above, given $|\mathcal{G}|$ and no form of supervision on the partition of $S_{N}$ or the group identity.

## 4. The Gram matrix of orbits

The simple observation for designing $\mathcal{R}(W)$ comes from the inspection of the matrix of all inner products of the vectors of an orbit (the so called Gram matrix). If the columns of $W$ correspond to an orbit of a vector $t \in \mathbb{R}^{d}$ :

$$
\begin{equation*}
W=\left[g_{1} t, \ldots, g_{|\mathcal{G}|} t\right] \tag{11}
\end{equation*}
$$

then the associated Gram matrix $G=W^{T} W \in \mathbb{R}^{|\mathcal{G}| \times|\mathcal{G}|}$ has entries of the form:

$$
\begin{equation*}
(G)_{i j}=\left\langle g_{i} t, g_{j} t\right\rangle=\left\langle t, g_{i}^{*} g_{j} t\right\rangle=v_{t}\left(g_{i}^{-1} g_{j}\right) \tag{12}
\end{equation*}
$$

where $g_{i}^{*}$ the conjugate transpose of $g_{i}$ and $\nu_{t}: \mathcal{G} \times \mathcal{G} \rightarrow \mathbb{R}$ is an injective function that depends on the vector $t$. Assuming a unitary group, i.e. $g^{*}=g^{-1}$, a hypothesis that can be relaxed, and using the closure property of group composition, we have the following proposition (see also [48], Ch5, Th. 2 and [49]).

Proposition 1. If a set of vectors forms an orbit w.r.t. a finite, unitary group, then their Gramian matrix $G$ is a permuted matrix i.e. a matrix whose columns are permutations of a single vector.

Proof. Each entry of $G$ is a pointwise function, $v_{t}$, of the multiplication table of $\mathcal{G}$. By the Cayley theorem [50], the columns of the multiplication table of a group are permuted versions one of the other. In other words the statement that $G=W^{T} W$ is a permuted matrix follows from the property that the group is closed under composition.

As an intuition for the proposition, consider the simple example of the cyclic group of order 3 . Let $\mathbb{Z}_{3}=\{0,1,2\}$ with law $i \circ j=$ $(i+j(\bmod 3))$. Each column in the multiplication table associated to the elements of $\mathbb{Z}_{3}$ is a permutation of the others.

| $\circ$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{2}$ |
| :--- | :--- | :--- | :--- |
| $\mathbf{0}$ | 0 | 1 | 2 |
| $\mathbf{1}$ | 1 | 2 | 0 |
| $\mathbf{2}$ | 2 | 0 | 1 |

Therefore a function of the multiplication table entries will give a permuted matrix.

To further give an intuition of the idea behind the construction of the regularization term we consider the group of rotations of angles $\theta=\{0,120,240\}$. In dimension two the elements of the group are $2 \times 2$ matrices that we indicate with $R_{\theta}$. The orbit matrix of a vector $t \in \mathbb{R}^{2}$ is then

$$
W=\left(e t, R_{120} t, R_{240} t\right)
$$

with $e=R_{0}$ the matrix implementing a rotation of a angle zero i.e. the identity matrix. The associated gramian elements are given by

$$
(G)_{i j}=\left\langle t, R_{\theta_{i}}^{*} R_{\theta_{j}} t\right\rangle=\left\langle t, R_{-\theta_{i}} R_{\theta_{j}} t\right\rangle=\left\langle t, R_{\theta_{j}-\theta_{i}} t\right\rangle
$$

The explicit computation of the gramian shows:

$$
\begin{aligned}
G & =\left[\begin{array}{ccc}
\langle t, t\rangle & \left\langle t, R_{120} t\right\rangle & \left\langle t, R_{240} t\right\rangle \\
\left\langle t, R_{-120} t\right\rangle & \langle t, t\rangle & \left\langle t, R_{120} t\right\rangle \\
\left\langle t, R_{-240} t\right\rangle & \left\langle t, R_{-120} t\right\rangle & \langle t, t\rangle
\end{array}\right] \\
& =\left[\begin{array}{ccc}
\langle t, t\rangle & \left\langle t, R_{120} t\right\rangle & \left\langle t, R_{240} t\right\rangle \\
\left\langle t, R_{240} t\right\rangle & \langle t, t\rangle & \left\langle t, R_{120} t\right\rangle \\
\left\langle t, R_{120} t\right\rangle & \left\langle t, R_{240} t\right\rangle & \langle t, t\rangle
\end{array}\right]
\end{aligned}
$$

i.e. the columns of the gramian are permuted versions of each other.

## 5. Analytic expressions for regularization

In this section we provide closed-form expressions for two conditions: a) the permuted matrix condition, from Proposition 1, enforcing a group orbit structure on the representation matrix $W$ and b) the same-symmetry condition penalizing groups different than the generating group $\mathcal{G}$ of an observation set $S_{N}$.

### 5.1. Permuted matrix condition

Following Proposition 1, we provide an analytical condition for the Gram matrix $G$ to be a permuted matrix, by imposing the columns (or rows) of $G$ to be permuted versions of the same vector. A straightforward way of doing this is to impose the Euclidean distance between the distribution of column values, for all column pairs, to be null. Mathematically the condition can be written as:

$$
\begin{gather*}
\sum_{i, j=1}^{|\mathcal{G}|} \int d \lambda\left(h_{i}(\lambda)-h_{j}(\lambda)\right)^{2}=0  \tag{13}\\
h_{i}(\lambda)=\sum_{k=1}^{|\mathcal{G}|} \delta\left(G_{k i}-\lambda\right)
\end{gather*}
$$

with $G_{k i}=(G)_{k i}$ the element in row $k$ and column $i$ of $G$ and $\delta(a)=1$, if $a=0$ and 0 otherwise. This is a necessary and sufficient condition for the columns of $G$ to be permutations of a single vector, as the distribution of values is a maximal invariant [19,51] with respect to the permutation group. The following theorem provides an explicit expression for computing (13) given a matrix $G$.

Theorem 1 (Symmetry regularization). Let matrix $W \in \mathbb{R}^{d \times|\mathcal{G}|}, G=$ $W^{T} W$ and $r: \mathbb{R}^{|\mathcal{G}| \times|\mathcal{G}|} \rightarrow \mathbb{R}_{+}$such that

$$
\begin{equation*}
r(G)=\tau^{T} \delta(\operatorname{Cvec}(G)) \tag{14}
\end{equation*}
$$

where $\operatorname{vec}(G) \in \mathbb{R}^{|\mathcal{G}|^{2}}$ is the vectorization of $G, \delta$ is the elementwise function $\delta(a)=1$, if $a=0$ (otherwise 0 ), and $\tau, C$ are respectively $a$ constant weight vector and matrix that depend only on $|\mathcal{G}|$. If $r(G)=$ 0 , then $G$ is a permuted matrix and viceversa.

Proof. A direct calculation of (13) for all pairs of columns in $G$ gives:

$$
\begin{equation*}
\sum_{i, j, k, l=1}^{|\mathcal{G}|}(|\mathcal{G}| \delta(i-j)-1) \delta\left(G_{k i}-G_{l j}\right)=0 \tag{15}
\end{equation*}
$$

which can be rewritten in a compact way by vectorizing $G$ :

$$
\begin{equation*}
\tau^{T} \delta(\operatorname{Cvec}(G))=0 \tag{16}
\end{equation*}
$$

where $\tau$ is the $|\mathcal{G}|^{2}\left(|\mathcal{G}|^{2}-1\right) / 2 \times 1$ vector of weights $(|\mathcal{G}| \delta(i-j)-$ $1), i, j=1, \ldots,|\mathcal{G}|$ and $C$ is a $\left(|\mathcal{G}|^{2}-1\right)|\mathcal{G}|^{2} / 2 \times|\mathcal{G}|^{2}$ sparse, con-
stant matrix that encodes all pairwise differences in a right multiplied vector of size $|\mathcal{G}|^{2} \times 1$, namely:

$$
C=\left[\begin{array}{rrrrrrr}
1 & -1 & 0 & \ldots & 0 & 0 & 0 \\
1 & 0 & -1 & \ldots & 0 & 0 & 0 \\
\ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \\
0 & 0 & 0 & \ldots & 1 & 0 & -1 \\
0 & 0 & 0 & \ldots & 0 & 1 & -1
\end{array}\right]
$$

Since the condition is derived by a chain of equalities, the converse is also true.

The regularizer $\mathcal{R}(W)$ in (10) can then be chosen to be the function of $G$ that minimizes (14), namely $\mathcal{R}(W)=r\left(W^{T} W\right)$. The size of $W$, or equivalently the orbit size $|\mathcal{G}|$, which would correspond to the size of the representation matrix is left as a hyperparameter or can be chosen so that the size of the associated Gram matrix allows for a fast and memory-efficient computation of $\mathcal{R}(W)$. Further, using $G$ instead of the representation matrix $W$, allows us to work with the permutation group only, and thus for a uniform treatment of arbitrary, finite groups. In representations with additional structure, e.g. multilayer as in deep neural networks, the regularizer can be the sum over subsets of weights in the same layer.

We make a few additional remarks on Theorem 1: (a) Making the regularizer (14) null is a necessary and sufficient condition to have an orbit with respect to an Abelian group; see Proposition 2 in [52] and [49] where these orbits are also called geometrically uniform frames. (b) In addition, it can be used, as a pseudo-metric, to test the equivariance and invariance properties of a representation (Eq. (30) and Appendix B). (c) Choosing different nonlinear functions than $\delta$ [51] or distances in (13) can give alternative expressions for regularization. (d) For the case of permutation groups, the result can be generalized to a class of regularization terms. In this case, any elementwise function of the orbit matrix, i.e. $f\left((W)_{i j}\right), f: \mathbb{R} \rightarrow \mathbb{R}$, will induce a Gram matrix with entries $(G)_{i j}=\left\langle f\left(g_{i} t\right), f\left(g_{j} t\right)\right\rangle$ which is also a permuted matrix and the following holds

Corollary 1. For any permutation group $\mathcal{G}$ and any $f: \mathbb{R} \rightarrow \mathbb{R}$ acting pointwise on the matrix $W$, if $W^{T} W$ is a permuted matrix then $r\left(f\left(W^{T}\right) f(W)\right)=0$ and viceversa.

In fact, the action of the function on the orbit $W$ is equivalent to a change of the vector $t$ since for a permutation group $\left(f\left(g_{i} t\right)\right)_{q}=$ $\left(g_{i} f(t)\right)_{q}, \forall i, q$. This observation will be useful for having additional constraints for representation learning (see (29)).

### 5.2. Same-symmetry condition

The permuted-matrix condition in Proposition 1 is not conditioned on an observation set, i.e. the group generating the orbit for $W$ and the data could be different. One way to constrain the solutions to those with same symmetries as the observation set $S_{N}$ comes from looking at the data covariance matrix. In fact if $X=\left[x_{1}, \cdots, x_{N}\right] \in \mathbb{R}^{d \times N}$ is the nonzero matrix storing $S_{N}$ (as in (9)), a simple calculation shows:

$$
\begin{equation*}
X X^{T}=\sum_{i=1}^{|\mathcal{G}|} g_{i} T T^{T} g_{i}^{T} \tag{17}
\end{equation*}
$$

where $T=\left[t_{1}, \ldots, t_{Q}\right]$ is the $d \times Q$ matrix of representatives for each orbit in $S_{N}$ (e.g. an arbitrary element of each). The matrix in
(17) has a high degree of symmetry that can be used to prove the following Lemma. Let $[A, B]=A B-B A$ denote the commutator of matrices $A, B$.
Lemma 1. Let $W \in \mathbb{R}^{d \times|\tilde{G}|}$ a nonzero orbit matrix w.r.t. a finite group $\tilde{\mathcal{G}}$ and $X=\left[x_{1}, \cdots, x_{N}\right] \in \mathbb{R}^{d \times N}, x_{i} \in S_{N}$ (as in 9) with $|\tilde{\mathcal{G}}|=|\mathcal{G}|$. If $\left[X X^{T}, W W^{T}\right]=\mathbf{0}$ then $\tilde{\mathcal{G}}$ and $\mathcal{G}$ are the same up to a fixed unitary conjugation.

Towards proving this Lemma, we start with two preparatory Lemmas: the first states that the elements of a $\mathcal{G}$-orbit of an eigenvector are also eigenvectors with the same eigenvalue; the second states that the linear span of the $\mathcal{G}$-orbit of eigenvector with eigenvalue $\lambda$ coincides with the eigenspace associated to $\lambda$.

Lemma 2. If $v \in \mathbb{R}^{d}$ is an eigenvector of $X X^{T}$ with eigenvalue $\lambda$, the set $\{g v \mid g \in \mathcal{G}\}$ is a set of eigenvectors of $X X^{T}$ with the same eigenvalue.
Proof. Multiplying the right hand side of Eq. (17) by some $g_{j} \in \mathcal{G}$ we have

$$
\begin{equation*}
X X^{T} g_{j}=\sum_{i=1}^{|\mathcal{G}|} g_{i} T T^{T} g_{i}^{T} g_{j}=g_{j} \sum_{k=1}^{|\mathcal{G}|} g_{k} T T^{T} g_{k}^{T}=g_{j} X X^{T} \tag{18}
\end{equation*}
$$

where we used the change of variables $g_{k}=g_{j}^{T} g_{i}$, the unitary group assumption and the closure property of the group. Thus

$$
\begin{equation*}
\left[X X^{T}, g\right]=\mathbf{0}, \quad \forall g \in \mathcal{G} \tag{19}
\end{equation*}
$$

Let $v \in \mathbb{R}^{d}$ be an eigenvector of $X X^{T}$ with eigenvalue $\lambda$, i.e. $X X^{T} v=$ $\lambda v$. Then, using (19), for any element of the $\mathcal{G}$-orbit of $v, O_{v}=$ $\{g \nu \mid g \in \mathcal{G}\}:$

$$
\begin{equation*}
X X^{T} g v=g X X^{T} v=\lambda g v, \quad \forall g \in \mathcal{G} \tag{20}
\end{equation*}
$$

i.e. any element of the orbit $O_{v}$ is an eigenvector of $X X^{T}$ with eigenvalue $\lambda$.

Lemma 3. Let $E_{\lambda}$ the eigenspace of $X X^{T}$ associated to eigenvalue $\lambda$ and $v$ an arbitrary vector in $E_{\lambda}$. Then the eigenspace coincides with the linear span of the $\mathcal{G}$-orbit of $v$, i.e. $\operatorname{span}\left(O_{v}\right)=E_{\lambda}$. Further the representation of $\mathcal{G}$ on $E_{\lambda}$ is irreducible, i.e. $E_{\lambda}$ is the smallest subspace containing $v$ that is left invariant by the action of the group $\mathcal{G}$.
Proof. Let $B=\left\{b_{i}\right\}_{i=1}^{|B|}$ an orthogonal basis in $E_{\lambda}$ with $|B|=$ $\operatorname{dim}\left(E_{\lambda}\right)$. Any $v \in E_{\lambda}$ can be expressed as a linear combination of the basis elements

$$
\begin{equation*}
v=\sum_{i=1}^{|B|} \alpha_{i} b_{i} \tag{21}
\end{equation*}
$$

and an element of its $\mathcal{G}$-orbit $O_{v}$ as:

$$
\begin{equation*}
g v=\sum_{i=1}^{|B|} \alpha_{i} g b_{i}, \quad \forall g \in \mathcal{G} \tag{22}
\end{equation*}
$$

Since each $g b_{i}$ is an eigenvector (by Lemma 2), $g v$ is a linear combination of eigenvectors. Clearly this holds also for any linear combination of the orbit elements, i.e., we have $\operatorname{span}\left(O_{v}\right) \subseteq E_{\lambda}$. Note now that for any two $u, v \in E_{\lambda}$ :

$$
\begin{equation*}
\operatorname{span}\left(O_{u}\right)=\operatorname{span}\left(O_{v}\right), \tag{23}
\end{equation*}
$$

i.e., $\operatorname{span}\left(O_{u}\right)$ and $\operatorname{span}\left(O_{v}\right)$ coincide since they both consist of all linear combinations of the set $\left\{g b_{i}\right\}$. This implies:

$$
\begin{equation*}
\bigcup_{u \in E_{\lambda}} \operatorname{span}\left(O_{u}\right)=\operatorname{span}\left(O_{v}\right) \tag{24}
\end{equation*}
$$

However:

$$
\begin{equation*}
E_{\lambda} \subseteq \bigcup_{u \in E_{\lambda}} \operatorname{span}\left(O_{u}\right)=\operatorname{span}\left(O_{v}\right) \subseteq E_{\lambda} \tag{25}
\end{equation*}
$$

which implies that $\operatorname{span}\left(O_{v}\right)=E_{\lambda}$ for any $v \in E_{\lambda}$ and clearly is the smallest invariant space w.r.t. $\mathcal{G}$ (since any such space is a span of
an orbit of a vector in $E_{\lambda}$ ). In other words $\mathcal{G}$ acts irreducibly on $E_{\lambda}$.

We can now prove Lemma 1:
Proof. Let $E_{\lambda}$ and $\tilde{E}_{\tilde{\lambda}}$ the eigenspaces and eigenvalues respectively of the matrices $X X^{T}$ and $W W^{T}$. From Lemma 2 we have

$$
\begin{equation*}
g E_{\lambda}=E_{\lambda}, \quad \tilde{g} \tilde{E}_{\tilde{\lambda}}=\tilde{E}_{\tilde{\lambda}}, \forall g \in \mathcal{G}, \tilde{g} \in \tilde{\mathcal{G}} . \tag{26}
\end{equation*}
$$

Note that $X X^{T}, W W^{T}$ are Hermitian matrices. Now, if two Hermitian matrices commute they have the same eigenspaces and eigenvalues [53], so $E_{\lambda}=\tilde{E}_{\tilde{\lambda}}$. Thus,

$$
\begin{equation*}
\tilde{g} E_{\lambda}=g \tilde{E}_{\tilde{\lambda}}=g E_{\lambda}=E_{\lambda}, \forall \lambda, \tilde{\lambda}, \forall g \in \mathcal{G}, \forall \tilde{g} \in \tilde{\mathcal{G}} \tag{27}
\end{equation*}
$$

or in other words $\tilde{\mathcal{G}}$ and $\mathcal{G}$ preserve, respectively, the eigenspaces of $X X^{T}, W W^{T}$ i.e. $\left[g, W W^{T}\right]=\left[\tilde{g}, X X^{T}\right]=0$. Further, by Lemma 3, the action of the representation of $\mathcal{G}$ and $\tilde{\mathcal{G}}$ restricted to $E_{\lambda}$ is irreducible. Since $\mathcal{G}$ and $\tilde{\mathcal{G}}$ act irreducibly on the same subspaces $E_{\lambda}$ they are equivalent up to a unitary conjugation, i.e. $\tilde{g}=U^{T} g U, \exists U$.

### 5.3. Representation learning with symmetry conditions

The main result of this work is summarized in the following Theorem 2, which puts together Proposition 1, Theorem 1 and Lemma 1.

Theorem 2 (Symmetry-adapted regularization). Let $W \in \mathbb{R}^{d \times|\tilde{G}|} a$ nonzero matrix whose columns are the elements of an orbit of a finite group $\tilde{\mathcal{G}}$ with cardinality $|\tilde{\mathcal{G}}|$. Then $W^{T} W$ is permuted and $r\left(W^{T} W\right)=$ 0 . Further, let $X \in \mathbb{R}^{d \times Q|G|}$ a matrix whose columns are all elements from the union of $Q$ orbits of a finite group $\mathcal{G}$, with $|\mathcal{G}|=|\tilde{\mathcal{G}}|$. If $\left[X X^{T}, W W^{T}\right]=\mathbf{0}$ then $\tilde{\mathcal{G}}$ and $\mathcal{G}$ are the same up to a fixed unitary conjugation.

Proof. Follows by putting together Theorem 1 and Lemma 1.
Theorem 2 provides the two regularization conditions for the symmetry-adapted representation learning problem (10):

$$
\begin{equation*}
\underset{W \in \mathbb{R}^{d \times||G|}}{\arg \min } \mathcal{L}\left(W, S_{N}\right)+\gamma\left\|\left[X X^{T}, W W^{T}\right]\right\|_{F}^{2}+\beta r\left(W^{T} W\right) \tag{28}
\end{equation*}
$$

where $\beta, \gamma \in \mathbb{R}_{+}$are regularization parameters, $\mathcal{L}^{\prime}\left(W, S_{N}\right)=$ $\left\|\left[X X^{T}, W W^{T}\right]\right\|_{F}^{2}, \mathcal{R}(W)=r\left(W^{T} W\right)$ and the loss $\mathcal{L}\left(W, S_{N}\right)$ is taskdependent. We use a smooth version of (14), substituting the $\delta$ function with a Gaussian function $g_{\sigma}: \mathbb{R} \rightarrow \mathbb{R}$ of width $\sigma \ll 1$; the regularizer is then $r\left(W^{T} W\right)=\tau^{T} g_{\sigma}\left(C \operatorname{vec}\left(W^{T} W\right)\right)$. The analytic form of the gradients, for use in gradient-based methods for minimizing (28), is provided in Appendix A.

## 6. Results on unsupervised orbit learning

As a proof of concept, we pose the following unsupervised learning problem: Given an unlabeled observation set as in (9), namely a union of orbits of the same, finite, unknown group $\mathcal{G}$, learn a single orbit $W$, of an arbitrary vector $t \in \mathbb{R}^{d}$, with respect to the (latent) group that generated the data.

### 6.1. Synthetic group data

We use synthetic data generated by known permutation groups, of different orders $|\mathcal{G}|$, acting on $\mathbb{R}^{6}$ :

- Cyclic group ( $C_{6}$ ), Abelian group, $|\mathcal{G}|=6$
- Dihedral group $\left(D_{6}\right)$, non-Abelian group, $|\mathcal{G}|=12$
- Pyritohedral group ( $T_{h}$ ), non-Abelian group, $|\mathcal{G}|=24$

The data $X$ were generated using $Q$ vectors uniformly sampled at random from the unit ball in $\mathbb{R}^{6}$, for train and validation sets of sizes $Q=1000$ and $Q=200$ orbits (or $Q|\mathcal{G}|$ observations).


Fig. 1. Permutation sensitivity of the pseudo-metric for a permutation group (Abelian group, Cyclic $C_{6}$ ) and (a) representation through a random orbit of $C_{6}$, (b) representation through (unsupervised) learned orbit. Details on the distance matrices and plots are as in Fig. 2.

### 6.2. Optimization for learning

We use (28), without a loss $\mathcal{L}\left(W, S_{N}\right)$, i.e. the commutator norm drives the representation, and $\gamma=1$. Instead of $r\left(W^{T} W\right)$, since we are testing permutation group, following Corollary 1, we use a sum of regularizers $r\left(\sigma_{j}(W)^{T} \sigma_{j}(W)\right), j=1 \ldots J$, where $\sigma_{j}(x)=$ $\max \left(0, x-\alpha_{j}\right)$ the rectifier function applied pointwise. Then, the minimization problem takes the form

$$
\begin{equation*}
\underset{W \in \mathbb{R}^{d \times|G|} \mid}{\arg \min }\left\|\left[X X^{T}, W W^{T}\right]\right\|_{F}^{2} \quad+\frac{\beta}{J} \sum_{j=1}^{J} r\left(\sigma\left(W^{T}-\alpha_{j}\right) \sigma\left(W-\alpha_{j}\right)\right), \quad \beta \in \mathbb{R}_{+} \tag{29}
\end{equation*}
$$

where $X$ is the $d \times N$ matrix storing the observation set, $\sigma\left(x-\alpha_{j}\right)=\max (x-$ $\alpha_{j}, 0$ ) is the rectifier nonlinearity with parameter $\alpha_{j}, J$ is the number of nonlinearities and $\beta$ is the regularization constant. Note that the scheme with a single $r\left(W^{T} W\right)$ term can be derived as a special case by setting $J=1$ and $\alpha_{1} \ll 0$.

Minimization was performed using quasi-Newton iterative optimization with a cubic line search, using Broyden-Fletcher-Goldfarb-Shanno (BFGS) [54] for the Hessian matrix approximation. The matrix $W$ is set to size $d \times|\mathcal{G}|$, assuming the target orbit size $|\mathcal{G}|$ is given, and initialized at random. For each $\beta$, we initialize and run the minimization process $m$ times to obtain $m$ orbit solutions, corresponding to different local minima of the loss. For the results we used schemes with $J$ set to 100 and 300 , selecting $\alpha_{j}$ values uniformly spaced in $[-1,1]$. For $\beta$, we used a range of orders of magnitude, namely $\log _{10} \beta=\{0, \ldots,-7\}$. For the learned weights (Figs. 1(b) and 2(b)) correspond to the minimum loss (commutator norm) solution for $m=50$ random initializations of $W$. In Fig. $1\left(C_{6}\right)$, we used $\beta=10^{-4}, J=100,|\mathcal{G}|=6$. In Fig. 2 we used $J=300, \beta=10^{-5},|\mathcal{G}|=12$ (for $D_{6}$ ) and $\beta=10^{-6},|\mathcal{G}|=24$ (for $T_{h}$ ).

### 6.3. Equivariance of learned data representations

To test symmetry structure in a representation $W$ for the solutions given by the algorithm we used the fact that (Section 2) a representation $\Phi$ of the form of (4) is permutation-equivariant and thus if two signals $x, x^{\prime} \in \mathbb{R}^{d}$ are part of an orbit $O_{x}$ then their representations $\Phi(x), \Phi\left(x^{\prime}\right)$, with respect to an orbit dictionary $W$, are permuted vectors if $W$ is an orbit of a vector generated by the same group. We can then define the following pseudometric

$$
\begin{equation*}
D\left(x, x^{\prime}\right)=r\left(\left[\Phi(x), \Phi\left(x^{\prime}\right)\right]\right)=r\left(W^{T}\left[x, x^{\prime}\right]\right) \tag{30}
\end{equation*}
$$



Fig. 2. Dihedral group $D_{6}$ of order 12 (a-c, left) and Pyritohedral group $T_{h}$ of order $\mathbf{2 4}$ (a-c, right): pseudo-distance $r\left(\left[\Phi\left(x_{i}\right), \Phi\left(x_{j}\right)\right]\right)(30)$, where $\Phi\left(x_{i}\right)=W^{T} x_{i}$ and $W$ a $6 \times 12$ and $6 \times 24$ dictionary, for different representations. Learning performed from a training set of 1000 orbits. Shown are the distance matrices (top) for 6 orbits (arranged in 6 blocks of 12 elements) and 3 orbits (arranged in 3 blocks of 24 elements) and inter/intra-orbit distributions on log x-axis (bottom) for 50 orbits from a validation set.
where the input to $r: \mathbb{R}^{m \times 2} \rightarrow \mathbb{R}_{+}$is now the $m \times 2$ matrix $(\Phi(x)$, $\Phi\left(x^{\prime}\right)$ ). Details are provided in Appendix B. Using $D$ we can check if $W$ induces an equivariant representation and, for the case of matrices learned with the proposed method, if the equivariance is induced by a matrix with the symmetries in the data. For evaluation purposes only, we consider a labeled validation set, i.e. the orbit assignment for each point is known.

In Figs. 1 and 2 we evaluate the following cases: a) $W$ is an orbit selected from the same distribution as the training set, i.e. a ground truth $W$; this serves as a sanity check that the same symmetry will indeed induce equivariance in $\Phi(x)$, b) $W$ is a learned representation via minimization of (29), and c) $W$ are the eigenvectors of the data covariance (PCA), to check if symmetries can be discovered through high-variance directions or the result of a sparse dictionary learning algorithm (K-SVD [55]). The last two are included as baseline comparisons to check if symmetries can be discovered through projections in max-variance directions or sparsity/reconstruction constraints. In Fig. 1 we also depict the case $W=I_{d \times|\mathcal{G}|}$, corresponding to $\Phi(x)=x$. This serves as a sanity check for the validity of the pseudodistance as, for the case of permutation groups, each orbit is composed from permuted vectors.

The figures depict the distance matrices (of 12, 6 and 3 random orbits resp.) for $C_{6}, D_{6}$ and $T_{h}$ and the probability distributions of pairwise distances for the validation set ( 50 orbits and 300 points in total), estimated for inter- and intra-orbit pairs separately. The block diagonal structure of the distance matrices support the main claims of this work (Theorem 2), and the separability of the interand intra- distributions suggest that the learned $W$ implies a quotient representation with respect to the unknown group of transformations.

## 7. Conclusion

We studied the problem of learning data symmetries, in particular group symmetries, as a prior on structure due to transformations in the data. Our motivation was to derive representations that are adapted to symmetries and reduce the sample complexity of downstream supervised learning. In particular we explored mathematical conditions that can drive, in an unsupervised way, a learned representation to reflect the symmetries in an unlabeled training set.The approach is particularly relevant for data that have a low-dimensional intrinsic structure (e.g. transforming images or sounds) and can be applied to any finite group since it reduces their analysis to permutation groups whose size |\mathcall(G)|can
be chosen to be efficiently computable. This is in contrast with other proposed methods, $[41,42]$ where transformations matrices of size $d^{2} \gg \mid$ mathcall $(\mathrm{G})^{\wedge} 2 \mid$ are learned from data.

In this work, we focused on global group transformations as a proof of concept. However, the same theoretical framework can be applied to real data, for dealing with non-group, arbitrary transformations. The key notion for relaxing the assumptions on the observation set is locality. Arbitrary transformations can be approximated by smooth transformations of a number of instances, on local neighborhoods (of each instance). The transformation can then be described in terms of a group, i.e. the Lie group associated to the tangent space of the manifold of the signal transformations. The proposed regularization scheme can then be applied for learning localized filters, each representing a small, local signal neighborhood. This would be applicable to a general family of CNNs, where local filters are employed at each layer for local, positionrobust feature detection.

Additional directions for future work include addressing the dependency on noise (both in the unsupervised setting and supervised/semisupervised setting, [56-58]), size of the dictionary and partial orbits. For the latter, Lemma 1 can be restated in terms of empirical covariance matrices and the validity of the approximation can be measured using concentration inequalities. Interestingly, group codes (first studied by Slepian in his seminal paper [59]) or dictionaries ([60]) tend to have low self-coherence which is related to the exact recovery condition of a signal $x$ of given sparsity [61].

The long-term goal of this work is to learn data-symmetries driven group convolutions and network topologies in CNNs. Imposing structure in subsets of weights, corresponding to multiple filters transforming under one or more groups, will necessitate an extension to learning multiple orbits or multiple symmetries. Moreover, extensions to multilayer representations, e.g. deep networks, can be formulated by noting that the representation will be permutation-equivariant after the first layer, given a map as in (5). Subsequent layers will then process permutation-transformed signals.

## Acknowledgments

This material is based upon work supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216 and the Italian Institute of Technology. The authors gratefully acknowledge Alan Yuille, Maximilian Nickel, Silvia Villa and Carlo Cilliberto, for the insightful discussions and feedback on
early versions of this work, and Saverio Salzo for pointing to the proof of Lemma 3.

## Appendix

In the following we let $\mathcal{G}, \tilde{\mathcal{G}}$ two finite groups with cardinality $|\mathcal{G}|$ and $X \in \mathbb{R}^{d \times Q|\mathcal{G}|}$ a matrix whose columns are all elements from the union of orbits of $Q$ vectors in $\mathbb{R}^{d}$ w.r.t. $\mathcal{G}$.

## Appendix A. Analytic gradients

## Regularizer gradient:

We calculate the derivative of $r\left(W^{T} W\right)=\tau^{T} g_{\sigma}\left(\operatorname{Cvec}\left(W^{T} W\right)\right)$ with respect to the vectorization of $W$; using the chain rule we have:

$$
\begin{equation*}
\frac{\partial r\left(W^{T} W\right)}{\partial \operatorname{vec}(W)}=\left(\frac{\partial\langle\tau, p\rangle}{\partial p} \frac{\partial p}{\partial q} \frac{\partial q}{\partial s} \frac{\partial s}{\partial \operatorname{vec}(W)}\right)^{T} \tag{A.1}
\end{equation*}
$$

with $p=g_{\sigma}(q), q=C s$ and $s=\operatorname{vec}\left(W^{T} W\right)$. The first three factors are easy to calculate and the right side of Eq. (A.1) is:

$$
\begin{align*}
\frac{\partial r\left(W^{T} W\right)}{\partial \operatorname{vec}(W)}= & -\frac{2}{\sigma^{2}}\left(\tau^{T} \operatorname{diag}\left((\operatorname{Cvec}(G)) \odot g_{\sigma}\left(\operatorname{Cvec}\left(W^{T} W\right)\right)\right)\right. \\
& \left.\times C \frac{\partial \operatorname{vec}(G)}{\partial \operatorname{vec}(W)}\right)^{T} \tag{A.2}
\end{align*}
$$

where $\odot$ denotes the Hadamard product and $\operatorname{diag}(\cdot)$ a $\left(|\mathcal{G}|^{2}\left(|\mathcal{G}|^{2}-\right.\right.$ $1) / 2) \times\left(|\mathcal{G}|^{2}\left(|\mathcal{G}|^{2}-1\right) / 2\right)$ diagonal matrix. For the last partial derivative we have:

$$
\begin{equation*}
\frac{\partial \operatorname{vec}\left(W^{T} W\right)}{\partial \operatorname{vec}(W)}=\left(\mathbb{I}_{|\mathcal{G}|^{2}}+T\right)\left(\mathbb{I}_{|\mathcal{G}|} \otimes W^{T}\right) \tag{A.3}
\end{equation*}
$$

where the matrix $T$ is defined as $T \operatorname{vec}(v)=\operatorname{vec}\left(v^{T}\right)$ and $\mathbb{I}_{|\mathcal{G}|^{2}}, \mathbb{I}_{|\mathcal{G}|}$ are the identity matrices of dimensions $|\mathcal{G}|^{2} \times|\mathcal{G}|^{2},|\mathcal{G}| \times|\mathcal{G}|$. It can be shown that $T$ can be written explicitly as:

$$
\left\{\begin{array}{l}
T_{i j}=1, \text { if } j=1+|\mathcal{G}|(i-1)-\left(|\mathcal{G}|^{2}-1\right) \text { floor }\left(\frac{i-1}{|\mathcal{G}|}\right)  \tag{A.4}\\
T_{i j}=0, \text { otherwise. }
\end{array}\right.
$$

Putting everything together, right hand side of (A.1) becomes:

$$
\begin{align*}
\frac{\partial r\left(W^{T} W\right)}{\partial \operatorname{vec}(W)}= & -\frac{2}{\sigma^{2}}\left(\mathbb{I}_{|\mathcal{G}|} \otimes W\right) P^{T}\left(\operatorname{Cvec}\left(W^{T} W\right)\right) \odot g_{\sigma} \\
& \times\left(\operatorname{Cvec}\left(W^{T} W\right)\right) \odot \tau \tag{A.5}
\end{align*}
$$

with $P=C\left(\left.\mathbb{I}_{\mid \mathcal{G}}\right|^{2}+T\right)$ a constant matrix, which can be precalculated given the group cardinality. Correctness of the gradient has been tested with numerical simulations.

Commutator gradient: It is easy to derive the analytic gradient of the norm $\|M(X, W)\|_{F}^{2}=\operatorname{trace}\left(M(X, W)^{T} M(X, W)\right)$, where $M(X, W)=\left[X X^{T}, W W^{T}\right]$ is the commutator of the data and representation covariance matrices. A simple calculation shows that the gradient has the analytic form:

$$
\begin{equation*}
\frac{\partial\|M(X, W)\|_{F}^{2}}{\partial W}=-4\left[M(X, W), X X^{T}\right] W \tag{A.6}
\end{equation*}
$$

by expanding the gradient as follows:

$$
\begin{aligned}
& \frac{\partial}{\partial W}\left(-2 \operatorname{Tr}\left(X X^{T} W W^{T} X X^{T} W W^{T}\right)+2 \operatorname{Tr}\left(X X^{T} W W^{T} W W^{T} X X^{T}\right)\right) \\
& =-8 X X^{T} W W^{T} X X^{T}+4 X X^{T} X X^{T} W W^{T} W+4 W W^{T} X X^{T} X X^{T} W \\
& =-4\left[X X^{T}, W W^{T}\right] X X^{T} W+4 X X^{T}\left[X X^{T}, W W^{T}\right] W \\
& =-4\left[\left[X X^{T}, W W^{T}\right], X X^{T}\right] W \\
& =-4\left[M(X, W), X X^{T}\right] W .
\end{aligned}
$$

## Appendix B. Permutation invariant pseudometric

Quantifying equivalence under permutations is a way to test for equivalence under general symmetries, provided that $W$ has the same symmetries, and thus a way to quantify same-symmetry, orbit structure in a given $W$. To measure equivalence under permutations, we can use a modified version of the regularizer function (14), where the input is a 2 -column matrix instead of the Gram matrix. Indeed the representations $\Phi(x), \Phi\left(x^{\prime}\right)$ of two elements $x, x^{\prime} \in \mathbb{R}^{d}$ of the same orbit $O_{x}$, w.r.t. a same-group orbit dictionary $W$ are permuted vectors.

Given matrix $W \in \mathbb{R}^{d \times m}$, and the corresponding representation $\Phi(x)=W^{T} x=\left(\left\langle w_{1}, x\right\rangle, \ldots\left\langle w_{m}, x\right\rangle\right), \Phi: \mathbb{R}^{d} \rightarrow \mathbb{R}^{m}$, let

$$
D\left(x, x^{\prime}\right)=r\left(\left[\Phi(x), \Phi\left(x^{\prime}\right)\right]\right)=\tau^{T} \delta\left(C\left[\Phi(x), \Phi\left(x^{\prime}\right)\right]\right)
$$

where the input to $r$ is now the $m \times 2$ matrix [ $\left.\Phi(x), \Phi\left(x^{\prime}\right)\right]$. The quantity $r: \mathbb{R}^{m \times 2} \rightarrow \mathbb{R}_{+}$is a pseudometric and

$$
\begin{equation*}
\Phi(x)=P_{g} \Phi\left(x^{\prime}\right) \Longleftrightarrow r\left(\left[\Phi(x), \Phi\left(x^{\prime}\right)\right]\right)=0 . \tag{B.1}
\end{equation*}
$$

In addition, if the conditions $\left\|\left[X X^{T}, W W^{T}\right]\right\|_{F}^{2}=0$ and $r\left(W^{T} W\right)=0$ are also satisfied, as is the case for $W$ learned by minimizing (29), we have the following.

Lemma 4. If $W$ satisfies the conditions of Theorem 2 and is invertible, then:

$$
\begin{equation*}
r\left(W^{T}\left[x, x^{\prime}\right]\right)=0 \Longleftrightarrow x=g x^{\prime}, \forall x, x^{\prime} \in \mathcal{X}, \quad \forall g \in \mathcal{G} \tag{B.2}
\end{equation*}
$$

Proof. First note that the Gram matrix $W^{T} W$ identifies $W$ up to an arbitrary (but fixed) unitary transformation $U$; together with the condition $r\left(W^{T} W\right)=0$, this implies that $W$ can be written as $W=$ $U W_{t}$, where $W_{t}$ is a group orbit matrix of some vector $t \in \mathbb{R}^{d}$ and $W_{t}^{T} W_{t}=W^{T} W$.

For the direct implication, if $r\left(W^{T}\left[x, x^{\prime}\right]\right)=0$ then there exists a permutation matrix $P_{g}$ such that

$$
\begin{equation*}
W^{T} x=P_{g} W^{T} x^{\prime} \Rightarrow W_{t}^{T} U^{T} x=P_{g} W_{t}^{T} U^{T} x^{\prime}=W_{t}^{T} g U^{T} x^{\prime} \tag{B.3}
\end{equation*}
$$

since the action of $g$ on $W_{t}$ is a permutation of its columns i.e. $g W_{t}=W_{t} P_{g}$. If $W$ is invertible from the previous equation we have

$$
\begin{equation*}
U^{T} x=g U^{T} x^{\prime} \tag{B.4}
\end{equation*}
$$

which suffices since the rotation $U^{T}$ is irrelevant from the learning point of view since $U^{T} \mathcal{X}=\mathcal{X}$. For the inverse implication we proceed by contradiction. Suppose $x=g x^{\prime}$ but $r\left(\left[\Phi(x), \Phi\left(x^{\prime}\right)\right]\right) \neq 0$. Then:

$$
\begin{equation*}
W^{T} x \neq P W^{T} x^{\prime}, \quad x=g x^{\prime} \tag{B.5}
\end{equation*}
$$

for all permutations $P$. By the assumptions

$$
\begin{equation*}
(W)_{:, i}=\left(U W_{t}\right)_{:, i}=U g_{i} t=U g_{i} U^{T} U t=\tilde{g}_{i} \tilde{t}=\left(\widetilde{W}_{\tilde{t}}\right)_{:, i}, \tag{B.6}
\end{equation*}
$$

i.e. the columns of matrix $W$ are an orbit of vector $\tilde{t}$ w.r.t. the conjugate representation $\tilde{\mathcal{G}}$ of $\mathcal{G}$ induced by $U$, namely $W_{t}=\widetilde{W}_{\tilde{t}}$. We can then write (B.5) as

$$
\begin{equation*}
\tilde{W}_{\tilde{t}}^{T} x \neq P \tilde{W}_{\tilde{t}}^{T} x^{\prime}=\tilde{W}_{\tilde{t}}^{T} \tilde{g} x^{\prime}, \quad \forall \tilde{g} \in \tilde{\mathcal{G}} \tag{B.7}
\end{equation*}
$$

As $W_{t}$ is invertible, so is $\tilde{W}_{\tilde{t}}$. Thus, from (B.7) we have that $x \neq \tilde{g} x^{\prime}$ and we can conclude as in the direct implication.

Lemma 4 states that $r$ is zero for elements of the same orbit (inter-orbit distance) and large for elements of different orbits (intra-orbit distance). Examples of the distributions of these distances are shown in Figs. 1 and 2. Concretely, the value of $r$ for pairs of vectors can be used to test the following: (a) are $\Phi(x)$ and $\Phi\left(x^{\prime}\right)$ permuted vectors (from (B.1))? and (b) for $\Phi(x)=W^{T} x$ and $W$ satisfying the symmetry conditions, are $x, x^{\prime}$ on the same orbit ? And, as an extension, does the representation matrix $W$ provide an equivariant map, i.e. are the column vectors an orbit of the same group?

## References

[1] W.S. McCulloch, W. Pitts, A logical calculus of the ideas immanent in nervous activity, Bull. Math. Biophys. 5 (4) (1943) 115-133.
[2] K. Fukushima, Neocognitron: a self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position, Biol. Cybern. 36 (4) (1980) 193-202.
[3] T.M. Caelli, Z.-Q. Liu, On the minimum number of templates required for shift, rotation and size invariant pattern recognition, Pattern Recognit. 21 (3) (1988) 205-216.
[4] R. Lenz, Group invariant pattern recognition, Pattern Recognit. 23 (1) (1990) 199-217.
[5] P. Foldiak, Learning invariance from transformation sequences, Neural Comput. 3 (2) (1991) 194-200.
[6] A. Grace, M. Spann, A comparison between Fourier-Mellin descriptors and moment based features for invariant object recognition using neural networks, Pattern Recognit. Lett. 12 (10) (1991) 635-643.
[7] J. Flusser, T. Suk, Pattern recognition by affine moment invariants, Pattern Recognit. 26 (1) (1993) 167-174.
[8] B.A. Olshausen, C.H. Anderson, D.C. Van Essen, A multiscale dynamic routing circuit for forming size- and position-invariant object representations, J. Comput. Neurosci. 2 (1) (1995) 45-62.
[9] L. Van Goola, T. Moons, E. Pauwels, A. Oosterlinck, Vision and Lie's approach to invariance, Image Vis. Comput. 13 (4) (1995) 259-277.
[10] M. Michaelis, G. Sommer, A Lie group approach to steerable filters, Pattern Recognit. Lett. 16 (11) (1995) 1165-1174.
[11] J. Wood, Invariant pattern recognition: a review, Pattern Recognit. 29 (1) (1996) 1-17.
[12] M. Riesenhuber, T. Poggio, Hierarchical models of object recognition in cortex, Nat. Neurosci. 2 (11) (1999) 1019-1025.
[13] M. Lessmann, R.P. Wrtz, Learning invariant object recognition from temporal correlation in a hierarchical network, Neural Netw. 54 (2014) 70-84.
[14] K. Lenc, A. Vedaldi, Understanding image representations by measuring their equivariance and equivalence, in: IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2015.
[15] Z. Shao, Y. Li, Integral invariants for space motion trajectory matching and recognition, Pattern Recognit. 48 (8) (2015) 2418-2432.
[16] F. Anselmi, L. Rosasco, T. Poggio, On invariance and selectivity in representation learning, Inf. Inference 5 (2) (2015) 134-158.
[17] T.S. Cohen, M. Welling, Group equivariant convolutional networks, in: International Conference on Machine Learning (ICML), 2016.
[18] A. Achille, S. Soatto, Emergence of invariance and disentangling in deep representations, in: Proceedings of the ICML Workshop on Principled Approaches to Deep Learning, 2017.
[19] S. Soatto, Steps towards a theory of visual information: active perception, signal-to-symbol conversion and the interplay between sensing and control, arXiv:1110.2053, 2011.
[20] B. Haasdonk, H. Burkhardt, Invariant kernel functions for pattern analysis and machine learning, Mach. Learn. 68 (1) (2007) 35-61.
[21] F. Anselmi, J.Z. Leibo, L. Rosasco, J. Mutch, A. Tacchetti, T. Poggio, Unsupervised learning of invariant representations, Theor. Comput. Sci. 633 (2016) 112-121.
[22] Y.S. Abu-Mostafa, Learning from hints in neural networks, J. Complex. 6 (2) (1990) 192-198.
[23] R. Gens, P.M. Domingos, Deep symmetry networks, in: Advances in Neural Information Processing System (NIPS), 2014, pp. 2537-2545.
[24] Y. Bengio, A. Courville, P. Vincent, Representation learning: a review and new perspectives, IEEE Trans. Pattern Anal. Mach. Intell. 35 (8) (2013) 1798-1828.
[25] J. Schmidhuber, Deep learning in neural networks: an overview, Neural Netw. 61 (2015) 85-117.
[26] V. Papyan, Y. Romano, M. Elad, Convolutional neural networks analyzed via convolutional sparse coding, J. Mach. Learn. Res. 18 (1) (2017) 2887-2938.
[27] S. Zhang, J. Wang, X. Tao, Y. Gong, N. Zheng, Constructing deep sparse coding network for image classification, Pattern Recognit. 64 (2017) 130-140.
[28] M. Jaderberg, K. Simonyan, A. Zisserman, Spatial transformer networks, Advances in Neural Information Processing Systems (NIPS), 2015.
[29] J.J. Kivinen, C.K.I. Williams, Transformation equivariant Boltzmann machines, in: International Conference on Artificial Neural Networks (ICANN), 2011.
[30] T. Serre, L. Wolf, S. Bileschi, M. Riesenhuber, T. Poggio, Robust object recognition with cortex-like mechanisms, IEEE Trans. Pattern Anal. Mach. Intell. 29 (3) (2007) 411-426.
[31] L. Toth, Combining time- and frequency-domain convolution in convolutional neural network-based phone recognition, in: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2014, pp. 190-194.
[32] K. Kavukcuoglu, M. Ranzato, R. Fergus, Y. LeCun, Learning invariant features through topographic filter maps, in: IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops, 2009, pp. 1605-1612.
[33] Q.V. Le, J. Ngiam, Z. Chen, D. Chia, P.W. Koh, A.Y. Ng, Tiled convolutional neural networks, in: Advances in Neural Information Processing Systems 23, 2010, pp. 1279-1287.
[34] S. Dieleman, J. De Fauw, K. Kavukcuoglu, Exploiting cyclic symmetry in convolutional neural networks, in: International Conference on Machine Learning (ICML), 2016, pp. 1889-1898.
[35] L. Sifre, S. Mallat, Rotation, scaling and deformation invariant scattering for texture discrimination, in: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, IEEE, 2013, pp. 1233-1240.
[36] J. Bruna, S. Mallat, Invariant scattering convolution networks, IEEE Trans. Pattern Anal. Mach. Intell. 35 (8) (2013) 1872-1886.
[37] M.I. Khalil, M.M. Bayoumi, Invariant 2d object recognition using the wavelet modulus maxima, Pattern Recognit. Lett. 21 (9) (2000) 863-872.
[38] T.S. Cohen, M. Welling, Steerable CNNs, in: International Conference on Learning Representations (ICLR), 2017.
[39] S. Ravanbakhsh, J. Schneider, B. Poczos, Equivariance through parameter-sharing, Proceedings of the 34th International Conference on Machine Learning 70 (2017) 2892-2901.
[40] C. Zhang, S. Voinea, G. Evangelopoulos, L. Rosasco, T. Poggio, Discriminative template learning in group-convolutional networks for invariant speech representations, in: INTERSPEECH 2015, Annual Conf. of the International Speech Communication Association, 2015, pp. 3229-3233.
[41] X. Miao, R.P.N. Rao, Learning the Lie groups of visual invariance, Neural Comput. 19 (10) (2007) 2665-2693.
[42] R.P.N. Rao, D.L. Ruderman, Learning Lie groups for invariant visual perception, Adv. Neural Inf. Process. Syst. 11 (1999) 810-816.
[43] J. Sohl-Dickstein, C.M. Wang, B.A. Olshausen, An unsupervised algorithm for learning lie group transformations, 2010, arXiv:1001.1027.
[44] C.M. Wang, J. Shol-Dickstein, I. Tosic, B.A. Olshausen, Lie group transformation models for predictive video coding, in: Data Compression Conference Proceedings, IEEE, 2011, pp. 83-92.
[45] T.J. Sejnowski, P.K. Kienker, G.E. Hinton, Learning symmetry groups with hidden units: beyond the perceptron, Physica D 22 (1-3) (1986) 260-275.
[46] R. Negrinho, A. Martins, Orbit regularization, in: Advances in Neural Information Processing Systems (NIPS), 2014, pp. 3221-3229.
[47] S. Mallat, Group invariant scattering, Commun. Pure Appl. Math. 65 (10) (2012) 1331-1398.
[48] P.G. Casazza, G. Kutyniok, Finite Frames: Theory and Applications, Birkhauser Basel, 2012.
[49] Y. Eldar, Least-squares inner product shaping, Linear Algebra Appl. 348 (1-3) (2002) 153-174.
[50] A. Cayley, On the theory of groups, as depending on the symbolic equation $\theta^{n}=1$, Philos. Mag. Series 4 VII (42) (1854) 40-47,408-409.
[51] J.J. Rodrigues, P.M.Q. Aguiar, J.M.F. Xavier, ANSIG - an analytic signature for permutation-invariant two-dimensional shape representation, in: IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2008.
[52] Y. Eldar, H. Bolcskei, Geometrically uniform frames, IEEE Trans. Inf. Theory 49 (4) (2003) 993-1006.
[53] A. Laub, Matrix Analysis, Cambridge University Press, 2005.
[54] J. Nocedal, S.J. Wright, Numerical Optimization, second ed., Springer, 2006.
[55] M. Aharon, M. Elad, A.M. Bruckstein, K-SVD: an algorithm for designing overcomplete dictionaries for sparse representation, IEEE Trans. Signal Process. 54 (11) (2006) 4311-4322.
[56] B. Bo Han, I. Tsang, L. Chen, C. P., C. Yu, S. Fung, Progressive stochastic learning for noisy labels, IEEE Trans. Neural Netw. Learn. Syst. (2018).
[57] B. Han, Y. Pan, I. Tsang, Robust Plackettluce model for k-ary crowdsourced preferences, Machine Learning 107 (2018) 675-702.
[58] C. Gong, T. Liu, D. Tao, K. Fu, E. Tu, J. Yang, Deformed graph Laplacian for semisupervised learning, IEEE Trans. Neural Netw. Learn. Syst. 26 (2015) 2261-2274.
[59] D. Slepian, Group codes for the Gaussian channel, Bell Syst. Tech. J. 47 (4) (1968) 575-602.
[60] M. Thill, B. Hassibi, Group frames with few distinct inner products and low coherence, IEEE Trans. Signal Process. 63 (19) (2015) 5222-5237.
[61] D.L. Donoho, X. Huo, Uncertainty principles and ideal atomic decomposition, IEEE Trans. Inf. Theory 47 (7) (2001) 2845-2862.

Anselmi Fabio is a postdoctoral fellow in the Istituto Italiano di Tecnologia and the Laboratory for Computational and Statistical Learning at MIT and part of the Center for Brains, Minds, and Machines.


[^0]:    * Corresponding author at: Department of Brain and Cognitive Science, 43 Vassar Street, Cambridge, MA 02139, United States.

    E-mail addresses: anselmi@mit.edu (F. Anselmi), gevang@mit.edu (G. Evangelopoulos), Irosasco@mit.edu (L. Rosasco), tp@mit.edu (T. Poggio).
    ${ }^{1}$ Current affiliation: X (Alphabet Inc.).

