
Local Group Invariant Representations via Orbit Embeddings

Anant Raj Abhishek Kumar Youssef Mroueh P. Thomas Fletcher Bernhard Schölkopf
MPI IBM Research IBM Research University of Utah MPI

Abstract

Invariance to nuisance transformations is one of the desirable properties of effective representations. We consider transformations that form a *group* and propose an approach based on kernel methods to derive local group invariant representations. Locality is achieved by defining a suitable probability distribution over the group which in turn induces distributions in the input feature space. We learn a decision function over these distributions by appealing to the powerful framework of kernel methods and generate local invariant random feature maps via kernel approximations. We show uniform convergence bounds for kernel approximation and provide generalization bounds for learning with these features. We evaluate our method on three real datasets, including Rotated MNIST and CIFAR-10, and observe that it outperforms competing kernel based approaches. The proposed method also outperforms deep CNN on Rotated-MNIST and performs comparably to the recently proposed group-equivariant CNN.

1 Introduction

Effective representation of data plays a key role in the success of learning algorithms. One of the most desirable properties of effective representations is being invariant to nuisance transformations. For instance, convolutional neural networks (CNNs) owe much of their empirical success to their ability in capturing local translation invariance through convolutional weight sharing and pooling which turns out to be a useful model prior for images. Capturing

class sensitive invariance can also result in reduction in sample complexity [1] which is particularly useful in label scarce applications. We approach the problem of learning with invariant representations from a group theoretical perspective and propose a scalable framework for incorporating invariance to nuisance group actions via kernel methods.

At an abstract level, a *group* is defined as a set G endowed with a notion of *product* on its elements that satisfies certain axioms of (i) *closure*: $a, b \in G \implies ab \in G$, (ii) *associativity*: $(ab)c = a(bc)$, and (iii) *inverse element*: for each $g \in G$, $\exists g^{-1} \in G$ such that $gg^{-1} = g^{-1}g = e \in G$, where e is the identity element satisfying $ge = eg = g, \forall g \in G$. A group is *abelian* if the group product is commutative ($gh = hg, \forall g, h \in G$). For most practical applications each element $g \in G$ can be seen as a transformation acting on an input space X , $T_g : X \mapsto X$. The *orbit* of an element $x \in X$ under the action of the group G is defined as the set $O_x = \{T_g(x) \mid g \in G\}$. The set of all rotations in a fixed 2-D plane is an example of an infinite group where the product is defined as the consecutive application of two rotations. The orbit of an image under this rotation group is the infinite set consisting of all rotated versions of the image. The closure property of the group implies that the orbit of a point x is invariant under a group action on x , i.e., $O_{x|G} = O_{T_g(x)}, \forall g \in G$. The reader is referred to [35] for a more detailed introduction to group theory.

For unimodular groups, which include compact groups and abelian groups, there exists a so called unique (up to scaling) *Haar measure* ν that is invariant to both left and right group products, i.e., $\nu(S) = \nu(gS) = \nu(Sg)$ for all measurable subsets $S \subset G$ and all $g \in G$, essentially generalizing the notion of Lebesgue measure to groups. For a compact group G , Haar measure can be normalized by $\nu(G)$ (since $\nu(G) < \infty$) to obtain the normalized Haar measure which assigns a probability mass to all measurable subsets of G . Normalized Haar measure can be seen as inducing a uniform

probability distribution on the group. Recently, Anselmi et al. [1] used the normalized Haar measure $\tilde{\nu}$ on the group to map each orbit $(O_{x|G} \forall x)$ to a probability distribution P_x on the input space, i.e., $P_{x|G}(A) = \tilde{\nu}(\{g \mid T_g(x) \in A\}), \forall A \subset X$. The distribution $P_{x|G}$ induced by each point x can be taken as its invariant representation. However, estimating this distribution directly can be challenging due to its potentially high dimensional support. Anselmi et al. [1] propose to capture histogram statistics of 1-dimensional projections of $P_{x|G}$ to generate an invariant representation that can be used for learning, i.e., $\phi_n^k(x) = 1/|G| \sum_{g \in G} \eta_n(\langle T_g(x), t_k \rangle)$ for a finite group G , where t_k are the projection directions (termed as *templates*), $\eta_n(\cdot)$ are some nonlinear functions that are expected to capture the histogram statistics. More recently, Mroueh et al. [29] analyzed the concentration properties of the linear kernel defined over these features and provided generalization bounds for learning with this linear kernel.

Our point of departure from [1, 29] is the observation that histogram based features may not be the optimal way to characterize the probability distributions P_x induced by the group on the input space and their approach has its limitations. First, there is no principled guidance provided regarding the choice of nonlinearities η_n . Second, the inner-product of histogram based features $(\{\phi_n^k(x)\})$ approximately induces a Euclidean distance (group-averaged) in the input space [29] which may render them unsuitable for learning complex nonlinear decision boundaries in the input space. Further, *locality* is achieved by restricting the uniform distribution to a chosen subset of the group (i.e. elements within the subset are allowed to transform the input with equal probability and elements outside the subset are prohibited) which can be limiting.

Contributions: In this paper, we address aforementioned points and propose a framework to generate invariant representations by embedding the orbit distributions $P_{x|G}$ into a reproducing kernel Hilbert space (RKHS) [31, 39]. We propose to use *characteristic* kernels [41] so that the resulting map from the distributions to the RKHS is injective (one-to-one), preserving all the moments of the distribution. Our use of kernel methods to embed orbit distributions also renders a large body of work on kernel approximation methods at our disposal, which enable us to scale our proposed method. In particular, we derive invariant features by approximating the kernel using Nyström method [16, 45] and random Fourier features (for shift invariant kernels) [32]. The nonlinearities in the features ($\eta_n(\cdot)$) emerge in a principled

manner as a by-product of the kernel approximation. The RKHS embedding framework also naturally allows us to use more general probability distributions on the group, apart from the uniform distribution. This allows us to have better control over selectivity of the derived features and also becomes a technical necessity when the group is non-compact. We experiment with three real datasets and observe consistent accuracy improvements over baseline random Fourier [32] and Nyström features [16] as well as over [29]. Further, on Rotated MNIST dataset [22] we outperform recent invariant deep CNN and RBM based architectures [37, 40], and perform comparably to the more recently proposed group equivariant deep convolutional nets [11].

2 Formulation

Let the input features belong to a set $X \subset \mathbb{R}^d$. A group element $g \in G$ acts on points from \mathbb{R}^d through a map $T_g : \mathbb{R}^d \mapsto \mathbb{R}^d$, and we use a shorthand notation of gx to denote $T_g(x)$. We use gS to denote the action of a group element g on the set S , i.e., $gS = \{T_g(x) \mid x \in S \subseteq X\}$. We take liberty in using the same notation to denote the product of a group element with a subset of the group, i.e., $gS = \{gh \mid h \in S \subset G\}$ and $Sg = \{hg \mid h \in S \subset G\}$.

2.1 RKHS embedding of Orbit distributions

As introduced in the previous section, the *orbit* of an element $x \in X$ under the action of the group G is defined as the set $O_{x|G} = \{gx \mid g \in G\}$. For all unimodular groups there exists a Haar measure $\nu : S \mapsto \mathbb{R}_+$ which is invariant under left and right group product i.e., $\nu(S) = \nu(gS) = \nu(Sg)$ for all measurable subsets $S \subset G$ and all $g \in G$. Let $q(\cdot)$ be the probability density function of a distribution defined over G . This probability distribution over the group can be used to map each orbit $O_{x|G}$ to a probability distribution $P_{x|G}$ on the input space, i.e., $P_{x|G}(A) = \int_{g:gx \in A} q(g) d\nu(g) \forall A \subset X$. Note that $P_{x|G}(O_{x|G}) = 1$ (for an appropriately normalized measure ν), and $P_{x|G}(A) = 0 \forall A$ for which $A \cap O_{x|G} = \emptyset$.

Let \mathcal{H} be a reproducing kernel Hilbert space (RKHS) of functions $f : X \mapsto \mathbb{R}$ induced by kernel $k : X \times X \mapsto \mathbb{R}$, with the inner-product satisfying the reproducing property, i.e., $\langle f, k(x, \cdot) \rangle = f(x), \forall f \in \mathcal{H}$ and $\langle k(x, \cdot), k(x', \cdot) \rangle = k(x, x')$. The RKHS embedding of the distribution $P_{x|G}$ is given as [39]

$$\mu[P_{x|G}] := E_{z \sim P_{x|G}} k(z, \cdot). \quad (1)$$

This expectation is well-defined under the probability measure $P_{x|G}$, which is in turn induced by the measure ν over the group. The support of $P_{x|G}$ is $O_{x|G}$ and sampling a point $z \sim P_{x|G}$ is equivalent to sampling the corresponding group element g and setting $z = gx$. Thus we can rewrite the RKHS embedding of Eq. 1 as

$$\mu[P_{x|G}] = \int_G k(gx, \cdot) q(g) d\nu(g). \quad (2)$$

If the kernel is characteristic this map from distributions to the RKHS is injective, preserving all the information about the distribution [41]. All universal kernels [43] are characteristic when the support set of the distribution is compact [39]. In addition, many shift invariant kernels (e.g., Gaussian and Laplacian kernels) are characteristic on all of \mathbb{R}^d [17]. For precise characterization of characteristic shift invariant kernels, please refer to [42].

For a characteristic kernel the embedding $\mu[P_{x|G}]$ can be used as a proxy for $P_{x|G}$ in learning problems. To this end, we introduce a hyperkernel $h : \mathcal{H} \times \mathcal{H} \mapsto \mathbb{R}$ that defines the similarity between the RKHS embeddings corresponding to two points x and x' as $k_{q,G}(x, x') := h(\mu[P_{x|G}], \mu[P_{x'|G}])$. If we take h to be the linear kernel which is the regular inner-product in \mathcal{H} , we obtain

$$\begin{aligned} k_{q,G}(x, x') &:= \langle \mu[P_{x|G}], \mu[P_{x'|G}] \rangle_{\mathcal{H}} \\ &= \int_G \int_G k(gx, g'x') q(g) q(g') d\nu(g) d\nu(g') \end{aligned} \quad (3)$$

The kernel $k_{q,G} : X \times X \mapsto \mathbb{R}$ turns out to be the expectation of the *base kernel* $k(\cdot, \cdot)$ under the predefined probability distribution on the group G . It trades off locality and group invariance through appropriately selecting the probability density $q(\cdot)$. Taking q to be a delta function over the Identity group element gives back the original base kernel $k(\cdot, \cdot)$ which does not capture any invariance. On the other hand, if we take q to be the uniform probability density, we get the global group invariant kernel (also termed as *Haar integration kernel* [20, 29])

$$k_G(x, x') = \int_G \int_G k(gx, g'x') d\nu(g) d\nu(g'), \quad (4)$$

satisfying the property $k_G(gx, g'x') = k_G(x, x')$ for any $g, g' \in G$ and any $x, x' \in X$. Haar integral kernel does not preserve any locality information (e.g., images of digits 6 and 9 will be placed under same equivalence class). Strictly speaking, we only need ν to be the normalized *right Haar measure* satisfying $\nu(S) = \nu(Sg), \forall S \subset G, \forall g \in G$ for the global

group invariance property to hold. A unique (up to scaling) right Haar measure exists for all locally compact groups and for all unimodular groups (for which left and right Haar measures coincide) [35]. All Lie groups (e.g., rotation, translation, scaling, affine) are locally compact. Additionally, all compact groups (e.g., rotation), abelian groups (e.g., translation, scaling), and discrete groups (e.g., permutation) are unimodular. However, the Haar integration kernel $k_G(x, x')$ of Eq. 4 can only be defined for compact groups since we need $\nu(G) < \infty$ to keep the integral finite. Indeed, earlier work has used Haar integration kernel for compact groups [20, 29] (however, without the RKHS embedding perspective provided in our work which motivates the use of a *characteristic* base kernel $k(\cdot, \cdot)$).

A framework allowing more general (non-uniform) probability distribution on the group serves two purposes: (i) It enables us to operate with non-compact groups in a principled manner since we only need $\int_G q(g) d\nu(g) < \infty$ to enable construction of kernels such that Eq. 3 is finite; (ii) It allows for a better control over locality of the kernel $k_{q,G}(\cdot, \cdot)$. Earlier work [1, 29] achieves locality by taking a subset $G_0 \subset G$ and restricting the domain of the Haar integration kernel to be G_0 which amounts to having a uniform distribution over G_0 . A more general non-uniform distribution (e.g., a unimodal distribution with mode at the Identity element of the group) allows us to smoothly decrease the probability of sampling more extreme group transformations rather than abruptly prohibiting group transforms falling outside a preselected subset.

2.2 Feature generation via kernel approximation

The kernel $k_{q,G}$ of Eq. 3 can be used for learning with kernel machines [38], probabilistically trading off locality and group invariance through appropriately selecting $q(\cdot)$. However, kernel based learning algorithms suffer from scalability issues due to the need to compute kernel values for all pairs of data points. In this section, we describe our approach to obtain local invariant features via approximating $k_{q,G}$.

2.2.1 Features using random Fourier approximation

We first consider the case of shift-invariant base kernel satisfying $k(x, x') = \tilde{k}(x - x')$ which is a commonly used class of kernels that includes Gaussian and Laplacian kernels. Many shift-invariant kernels

are characteristic on \mathbb{R}^d as mentioned in the previous section. We use the random Fourier features proposed in [32] that are based on the characterization of positive definite functions by Bochner [6, 36]. Bochner’s theorem establishes Fourier transform as a *bijective map* from finite non-negative Borel measures on \mathbb{R}^d to positive definite functions on \mathbb{R}^d . Applying it to shift-invariant positive definite kernels one gets

$$k(x, x') = \tilde{k}(x - x') = \int_{\mathbb{R}^d} e^{-i(x-x')^\top \omega} p(\omega) d\omega, \forall x, x', \quad (5)$$

where $p(\cdot)$ is the unique probability distribution corresponding to the kernel $k(\cdot, \cdot)$, assuming the kernel is properly scaled. We use this characterization to obtain local group invariant features as follows:

$$\begin{aligned} & k_{q,G}(x, x') \\ &= \int_G \int_G E_{\omega \sim p} \left[e^{-i(gx-g'x')^\top \omega} \right] q(g)q(g') d\nu(g)d\nu(g') \\ &= E_{\omega \sim p} \int_G \int_G e^{-i(gx-g'x')^\top \omega} q(g)q(g') d\nu(g)d\nu(g') \\ &= E_{\omega \sim p} \int_G e^{-i\langle \omega, gx \rangle} q(g) d\nu(g) \int_G e^{i\langle \omega, g'x' \rangle} q(g') d\nu(g') \\ &\approx E_{\omega \sim p} \frac{1}{r^2} \sum_{k=1}^r e^{-i\langle \omega, g_k x \rangle} \sum_{k=1}^r e^{i\langle \omega, g_k x' \rangle}, \quad (g_k \sim q) \\ &\approx \frac{1}{sr^2} \sum_{j=1}^s \sum_{k=1}^r e^{-i\langle \omega_j, g_k x \rangle} \sum_{k=1}^r e^{i\langle \omega_j, g_k x' \rangle}, \quad (g_k \sim q, \omega_j \sim p) \\ &:= \langle \psi_{RF}(x), \psi_{RF}(x') \rangle_{\mathbb{C}^s}, \end{aligned} \quad (6)$$

where

$$\begin{aligned} \psi_{RF}(x) &= \frac{1}{r\sqrt{s}} \left[\sum_{k=1}^r e^{-i\langle \omega_1, g_k x \rangle} \dots \sum_{k=1}^r e^{-i\langle \omega_s, g_k x \rangle} \right] \\ &\in \mathbb{C}^s. \end{aligned} \quad (7)$$

We use standard Monte Carlo to approximate both inner integral over the group and the outer expectation over ω . It is also possible to use quasi Monte Carlo approximation for the expectation over ω , which has been carefully studied for random Fourier features [46]. We provide uniform convergence bounds and excess risk bounds for these features in Section 3.

The feature map $\psi_{RF}(\cdot)$ requires us to apply r group actions to every data point which can be expensive in large data regime. If the group action is unitary transformation preserving norms and distances between points (i.e., $\|gx\|_2 = \|x\|_2$), the inner product satisfies $\langle x, x' \rangle = \langle gx, gx' \rangle$. This can be

used to transfer the group action from the data to the sampled template as $\langle \omega, gx \rangle = \langle g^{-1}\omega, g^{-1}gx \rangle = \langle g^{-1}\omega, x \rangle$ [1] without affecting the approximation of kernel $k_{q,G}$, as long as the pdf q is symmetric around the identity element ($q(g) = q(g^{-1}) \forall g \in G$). For instance, in the case of images which can be viewed as a function $I : \mathbb{R}^2 \mapsto \mathbb{R}$, one can show the following result¹ regarding group actions (e.g., rotation, translation, scaling, affine transformation).

Lemma 2.1. *Let g be a group element acting on an image $I : \mathbb{R}^2 \mapsto \mathbb{R}$. The group action defined as $T_g[I(x)] = |J_g|^{-1/2} I(g^{-1}x)$, $\forall x$, where J_g is the Jacobian of the transformation, is a unitary transformation and satisfies $\langle T_g(I), T_g(I') \rangle = \langle I, I' \rangle$.*

Proof. See appendix. \square

The lemma suggests scaling the pixel intensities of the image by a factor $|J_g|^{-1/2}$ to make the group action unitary. The Jacobian for rotating or translating an image has determinant 1 obviating the need for scaling. For general affine transformation, we need to scale the pixel intensities accordingly to keep it unitary².

2.2.2 Features using Nyström approximation

Here we consider the case of a general base kernel and derive local group invariant features using Nyström approximation [16, 45]. Nyström method starts with identifying a set of *landmark points* (also referred as *templates*) $Z = \{z_1, \dots, z_s\}$ and approximates each function $f \in \mathcal{H}$ by its orthogonal projection onto the subspace spanned by $\{k(\cdot, z_i)\}_{i=1}^s$. Several schemes for identifying the landmark points have been studied in the literature, including random sampling, sampling based on leverage scores, and clustering based landmark selection [19, 21]. We can choose landmarks from the original set X or from the orbit gX . Nyström method approximates the kernel as $k(x, x') \approx K_{Z,x}^\top K_{Z,Z}^+ K_{Z,x'}$, where $K_{Z,x} = [k(x, z_1), \dots, k(x, z_s)]^\top$ and $K_{Z,Z}$ is square kernel matrix for the landmark points with $K_{Z,Z}^+$ denoting the pseudo-inverse.

Since $K_{Z,Z}$ is a positive semi-definite matrix, let

¹This is mentioned in [1] as a remark without a formal proof. We provide a proof in the appendix for completeness.

²The Jacobian for affine transformation $T(x) = Ax + b$ is its linear component A .

$K_{Z,Z}^+ = L^\top L$, where $L \in \mathbb{R}^{\text{rank}(K_{Z,Z}) \times s}$. We have

$$\begin{aligned} & k_{q,G}(x, x') \\ & \approx \int_G \int_G K_{g_x, Z} K_{Z, Z}^+ K_{Z, g'x'} q(g) q(g') d\nu(g) d\nu(g') \\ & = \int_G \int_G K_{g_x, Z} L^\top L K_{Z, g'x'} q(g) q(g') d\nu(g) d\nu(g') \\ & = \left\langle \int_G L K_{Z, g_x} q(g) d\nu(g), \int_G L K_{Z, g'x'} q(g) d\nu(g) \right\rangle \\ & \approx \left\langle L \frac{1}{r} \sum_{k=1}^r K_{Z, g_k x}, L \frac{1}{r} \sum_{k=1}^r K_{Z, g_k x'} \right\rangle, \quad (g_k \sim q), \end{aligned}$$

where the features are given by

$$\psi_{Nys}(x) = \frac{1}{r} L \sum_{k=1}^r K_{Z, g_k x} \in \mathbb{R}^{\text{rank}(K_{Z,Z})}. \quad (8)$$

If the base kernel satisfies $k(gx, gx') = k(x, x')$, $\forall g, x, x'$, we can transfer the group action from the data points to the landmark points as $k(gx, z) = k(g^{-1}gx, g^{-1}z) = k(x, g^{-1}z)$ without affecting the Nyström approximation of $k_{q,G}$, as long as the pdf q is symmetric around the identity element ($q(g) = q(g^{-1}) \forall g \in G$). This becomes essential in large data regime where the number of data points is much larger than the number of landmarks. For the group action defined in Lemma 2.1, all dot product kernels ($\tilde{k}(\langle x, x' \rangle)$) and shift invariant kernels ($\tilde{k}(\|x - x'\|_2)$) satisfy this property.

Remarks:

(1) Earlier work [1, 29] has proposed features of the form $\phi_n^k(x) = 1/r \sum_{j=1}^r \eta_n(\langle g_j x, \omega_k \rangle)$ where $\eta_n(\cdot)$ were taken to be step functions $\eta_n(a) = 1(a < h_n)$ with preselected thresholds h_n . Nonlinearities in our proposed local invariant features emerge naturally as a result of kernel approximation, with $\eta(x, \omega) = e^{-i\langle x, \omega \rangle}$ for ψ_{RF} and $\eta(x, \omega) = k(x, \omega)$ for ψ_{Nys} .

(2) Our work can also be viewed as incorporating local group invariance in widely used random Fourier and Nyström approximation methods, however this viewpoint overlooks the Hilbert space embedding perspective motivated in this work.

(3) The kernel $k_{q,G}$ defined in Eq. (3) assumes a linear hyperkernel $h : \mathcal{H} \times \mathcal{H} \mapsto \mathbb{R}$ over RKHS embeddings of orbit distributions. It is also possible to use a nonlinear hyperkernel along the lines of [10] and [30], and approximate it using a second layer of random Fourier (RF) or Nyström features. We show empirical results for both linear and Gaussian hyperkernel (approximated using RF features) in

Sec. 4.

(4) Computational aspects. The complexity of feature computation is $rC_f + rsC_g$ where C_f is the cost of computing the vanilla random Fourier or vanilla Nyström features and C_g is the cost of computing a group action on a *template* ω . However same set of templates are used for all data points so group actions on the templates can be computed in advance. Structured random Gaussian templates can also be used in our framework to speed up the computation of random Fourier features ψ_{RF} [7, 9, 23]. Recent approaches for scaling randomized kernel machines to massive data sizes and very large number of random features can also be used [3].

3 Theory

In this section we focus on local invariance learning using the random feature map ψ_{RF} defined in Section 2.2.1 for the Gaussian base kernel $k(\cdot, \cdot)$. We first address the uniform convergence of the random feature map ψ_{RF} to the local invariant kernel $k_{q,G}$ on a set of points \mathcal{M} . In other words we show in Theorem 3.1 that for a sufficiently large number of random templates s , and group element samples r , we have $\langle \psi_{RF}(x), \psi_{RF}(y) \rangle \approx k_{q,G}(x, y)$, for all points $x, y \in \mathcal{M}$. Second we consider a supervised binary classification setting, and study generalization bounds of learning a linear classifier in the local invariant random feature space ψ_{RF} . In a nutshell Theorem 3.2 shows that linear functions in the random feature space $\langle w, \psi_{RF}(x) \rangle$, approximate functions in the RKHS induced by our local invariant kernel $k_{q,G}$.

3.1 Uniform Convergence

Theorem 3.1 provides a uniform convergence bound of our invariant random feature map ψ_{RF} for Gaussian base kernel $k(\cdot, \cdot)$.

Theorem 3.1 (Uniform convergence of Fourier Approximation). *Let X be a compact space in \mathbb{R}^d with diameter $\text{diam}(X)$. For $\varepsilon > 0, \delta_1, \delta_2 \in (0, 1)$, the following uniform convergence bound holds with probability*

$$1 - \left(\frac{64(d+1)}{\varepsilon^2 \sigma^2} \right)^{\frac{d}{d+1}} (\delta_1 + \delta_2)^{\frac{2d}{d+1}}.$$

$$\sup_{x, y \in X} \left| \langle \psi_{RF}(x), \psi_{RF}(y) \rangle - K_{q,G}(x, y) \right| \leq \varepsilon + \frac{1}{r}$$

for a number of group samples

$$r \geq C_1 \frac{d}{\varepsilon^2} \log(\text{diam}(X)/\delta_1),$$

and a number of random templates

$$s \geq C_2 \frac{d}{\epsilon^2} \log(\text{diam}(X)/\delta_2),$$

where $\sigma_p^2 \equiv E_p[\omega^\top \omega] = d/\sigma^2$ is the second moment of the Fourier transform of the Gaussian base kernel k , and C_1 and C_2 are numeric universal constants.

Proof. See Appendix. \square

3.2 Generalization Bounds

Given a labeled training set $S = \{(x_i, y_i) \mid x_i \in X, y_i \in Y = \{+1, -1\}\}$, our goal is to learn a decision function $f : X \rightarrow Y$ via empirical risk minimization (ERM)

$$\min_{f \in \mathcal{H}_K} \hat{\mathcal{E}}_V(f) = \frac{1}{N} \sum_{i=1}^N V(y_i f(x_i))$$

where V is convex and L -Lipschitz loss function. Let $\mathcal{E}_V(f) = \mathbb{E}_{x, y \sim P} V(yf(x))$ be the expected risk for $f \in \mathcal{H}_K$. According to the representer theorem, the solution of ERM is given by $f^*(\cdot) = \sum_{i=1}^N \alpha_i^* k_{q, G}(x_i, \cdot)$.

We consider linear hyperkernel h in Eq. (3) and consider \mathcal{H}_K , the RKHS induced by the kernel $k_{q, G}(x, y) = \int_G \int_G k(gx, g'x') q(g)q(g') d\nu(g)d\nu(g')$, as introduced in Sec. 2.1. Similar to [29], for $C > 0$, we define \mathcal{F}_p an infinite dimensional space to approximate \mathcal{H}_K (see [33] for a motivation for this approximation):

$$\mathcal{F}_p \equiv \left\{ f(x) = \int_{\Omega} \alpha(\omega) \int_G \phi(gx, \omega) q(g) d\nu(g) d\omega \mid |\alpha(\omega)| \leq Cp(\omega) \right\},$$

where $\phi(gx, \omega) = e^{-i\langle gx, \omega \rangle}$. Similarly define the linear space in the span of $\psi_{RF}(\cdot)$, $\hat{\mathcal{F}}_p \equiv \left\{ \hat{f}(x) = \langle \alpha, \psi_{RF}(x) \rangle = \sum_{k=1}^s \alpha_k \frac{1}{r} \sum_{i=1}^r \phi(g_i x, \omega_k) \mid |\alpha_k| \leq \frac{C}{s} \right\}$.

Theorem 3.2. *Let $\delta > 0$. Consider the training set $S = \{(x_i, y_i) \mid x_i \in X, y_i \in Y, i = 1 \dots N\}$ sampled from the input space and let f_N^* is the empirical risk minimizer such that $f_N^* = \arg \min_{f \in \hat{\mathcal{F}}} \hat{\mathcal{E}}_V(f) = \frac{1}{N} \sum_{i=1}^N V(y_i f(x_i))$, then we have with probability $1 - 3\delta$ (over the training set, random templates and group elements)*

$$\mathcal{E}_V(f_N^*) - \min_{f \in \mathcal{F}_p} \mathcal{E}_V(f) \leq$$

Method	RMSE	RMSE w/ 2nd layer RF
Original (RF)	14.01	13.78
Original (Nys)	13.97	13.81
Original (GP)	13.48	N/A
Sort-Coulomb (RF)	12.89	12.49
Sort-Coulomb (Nys)	12.83	12.51
Sort-Coulomb (GP) [28]	12.59	N/A
Rand-Coulomb [28]	11.40	N/A
GICDF [29]	12.25	N/A
LGICA(RF)	10.82	10.05
LGICA(Nys)	10.87	10.45

Table 1: RMSE on Quantum Machine data

$$\mathcal{O} \left(\left(\frac{1}{\sqrt{N}} + \frac{1}{\sqrt{s}} + \frac{1}{\sqrt{r}} \right) LC \sqrt{\log \frac{1}{\delta}} \right).$$

Proof. See Appendix. \square

4 Empirical Observations

We evaluate the proposed method (referred as LGICA here) on three real datasets. We use Gaussian kernel as the base kernel in all our experiments. For methods that produce random (unsupervised) features, which include the proposed approach as well as regular random Fourier (abbrv. as RF) [32] and Nystrom [45] method, we report performance with: (i) linear decision boundary on these features (linear SVM or linear regularized least squares (RLS)), and (ii) nonlinear decision boundary which is realized by having a Gaussian kernel on top of the features and approximating it through random Fourier features [32], followed by a linear SVM or RLS. The later can also be viewed as using a nonlinear hyperkernel over RKHS embeddings of orbit distributions (also see Remark (3) at the end of Sec. 2). Parameters for all the methods are selected using grid search on a hold-out validation set unless otherwise stated.

4.1 Quantum Machine dataset

This data consists of 7165 Coulomb matrices of size 23×23 (each matrix corresponding to a molecule) and their associated atomization energies in kcal/mol. It is a small subset of a large dataset collected by Blum and Raymond (2009) [4], and was recently used by Montavon et al. (2012) [28] for evaluation. The goal is to predict atomization energies

of molecules which is modeled as a regression task.

The atomization energy is known to be invariant to permutations of rows/columns of the Coulomb matrix which motivates the use of representations invariant to the permutation group. We follow the experimental methodology of [28] and report mean cross-validation accuracy on the five folds provided in the dataset. An inner cross-validation is used for tuning the parameters for each fold as in [28]. We compare the performance of our method with several baselines in Table 1: (i) *Original (GP/RF/Nys)*: Gaussian Process regression on original Coulomb matrices and its approximation via random Fourier (RF) [32] and Nystrom features [45], (ii) *Sort-Coulomb (GP/RF/Nys)*: GP regression on sorted Coulomb matrices (sorted according to row norms) [28] and its approximation, (iii) *Rand-Coulomb*: permutation invariant kernel proposed in [28], and (iv) *GICDF*: Group invariant CDF (histogram) based features proposed in [29]. The results for *Sort-Coulomb (GP)* and *Rand-Coulomb* are taken directly from [28]. For all RF and Nyström based features we use $10k$ random templates (ω). For GICDF and our method, we sample 70 random permutations ($r = 70$ in Eq. 7) using the same scheme as in [28]. The proposed LGIKA outperforms all these directly competing methods including *Rand-Coulomb* and *GICDF*. Neural network based features used in [28] can also be used within our framework but we stick to raw Coulomb matrices for simplicity sake.




Figure 1: Kernel approximation error (normalized) in spectral and Frobenius norms vs number of random features, for 20 (left) and 40 (right) group transformations

Kernel Approximation. We also report empirical results on approximation error for kernel matrix (in terms of spectral norm and Frobenius norm) in

Method	Accuracy	Accuracy w/ 2nd layer RF
Original (RF)	87.75	88.01
Original (Nys)	88.93	88.98
Original (RBF)	90	N/A
TI-RBM [40]	95.8	N/A
RC-RBM [37]	97.02	N/A
GICDF [29]	93.81	N/A
Z2-CNN [11]	94.97	N/A
P4-CNN [11]	97.72	N/A
LGIKA(RF)	96.83	97.18
LGIKA(Nys)	96.81	97.21

Table 2: Rotated MNIST results

Fig. 4.1. The plots show the approximation error for different number of group actions as the number of random Fourier features are increased. The kernel used is the Gaussian kernel. The true kernel has been computed using 70 group elements randomly sampled from the permutation group. The normalized error for all the cases goes down with the number of random Fourier features which is in line with our theoretical convergence results.

4.2 Rotated MNIST

Rotated MNIST dataset [22] consists of total $62k$ images of digits ($12k$ for training and $50k$ for test), obtained by rotating original MNIST images by an angle sampled uniformly between 0 and 2π . We compare the proposed method with several other approaches in Table 4.1. We use von-misses distribution ($p(\theta) = \exp(-\kappa \cos(\theta))$ with $\kappa = 0.2$, selected using cross-validation) to sample $r = 100$ rotations. We use $s = 7k$ random templates for both RF and Nystrom approximations, and use $17k$ random templates for layer-2 RF approximation. The results for the cited methods in Table 4.1 are directly taken from the respective papers, except for GICDF [29] which we implemented ourselves. The proposed LGIKA outperforms most of the competing methods including deep architectures like rotation-invariant convolutional RBM (RC-RBM) [37], transformation invariant RBM (TI-RBM) [40], and regular deep CNN (Z2-CNN) [11]. Our method also performs close to the recently proposed group-equivariant CNN (P4-CNN) [11].

4.3 CIFAR-10

The CIFAR-10 dataset consists of $60k$ RGB images ($50k/10k$ for train/test) of size 32×32 , divided into 10 classes. We consider a sub-group of the affine

Original (RF) [32]		LGIKA	
1-layer	2-layer	1-layer	2-layer
61.02	62.79	64.19	67.32

Table 3: CIFAR-10 results

group $Aff(2)$ consisting of rotations, translations and isotropic scaling. Instead of operating with a distribution (e.g. Gaussian) over this subgroup, we use three individual distributions to have better control over the three variations: a log-normal distribution over the scaling group ($\mu = 0, \sigma = 0.3$), a Gaussian distribution over the translation group ($\mu = 0, \sigma = 0.3$), and a von-misses distribution over the rotation group ($\kappa = 9$). We observe that working with wider distributions over these groups actually hurts the performance, highlighting the importance of locality for CIFAR-10. We use the normalized pixel intensities as our input features and use the group action defined in Lemma 2.1 to keep it unitary. We use $s = 10k$ random templates and $r = 50$ group transforms for the first layer RF features (Eq. 7), and use $30k$ random templates for second layer RF features. The proposed LGIKA outperforms vanilla RF features as shown in Table 4.3. Nyström based features gave similar results as random Fourier features in our early explorations. We were not able to scale GICDF [29] to a suitable number of random templates due to memory issues (for every random template, GICDF generates n features (number of bins, set to 25 following [29]) blowing up the overall feature dimension to $n \times 10k$). Note that the performance of LGIKA on this data is still significantly worse than deep CNNs [11] since LGIKA treats the image as a vector ignoring the spatial neighborhood structure taken into account by CNNs through translation invariance over small image patches. Incorporating orbit statistics of image patches in our framework is left for future work.

5 Related Work

Invariant Kernel Methods. [2] introduced Tomographic Probabilistic Representations (TPR) that embed orbits to probability distributions. Unlike TPR, our representation maps orbits or local portions of the orbit via kernel mean embedding to an RKHS and allows to define similarity between orbits in this space. Indeed our representation is infinite dimensional and is related to Haar Invariant Kernel [20]. As discussed earlier it can be approximated via random features or Nyström sampling. Other

approaches for building invariant kernels were defined in [44] that focuses on dilation invariances. A kernel view of histogram of gradients was introduced in [5], where finite dimensional features were defined through kernel PCA. Kernel convolutional networks introduced in [27], [26], considers the composition of multilayer kernels, where local image patches are represented as points in a reproducing kernel. However they do not consider general group invariances. The work of [12] considers the general problem of learning from conditional distributions. When applied to invariant learning, their optimization approach needs to sample a group transformed example in every SGD iteration whereas our approach allows working with group actions on the random templates.

Invariance in Neural Networks. Inducing invariances in neural networks has attracted many recent research streams. It is now well established that convolutional neural networks (CNN) [24] ensure translation invariance. [15] showed that mapping orbits of rotated and flipped images through a shared fully connected network builds some invariance in the network. Scattering networks [8] have built in invariances for the roto-translation group. [18] generalizes CNN to general group transformations. [14] exploits cyclic symmetry to have invariant prediction in the network. More recently, [11] designs a convolutional neural network that is equivariant to group transforms by introducing convolution over the group.

6 Concluding Remarks

The proposed approach can be suitable for large-scale problems, benefiting from the recent advances in scalability of randomized kernel methods [3, 13, 25]. As a future direction, we would like to extend our framework to operate at the level of image patches, enabling us to capture local spatial structure. Further, the proposed approach requires computation of all r group transformations for all the sampled random templates. Reducing the required number of group transformations is an important direction for future work. Our work also assumes that the appropriate group actions are given. Extension to the case when the group transformations are learned from the data (e.g., using local tangent space [34]) is also an important direction for future work.

Acknowledgments: We thank Dmitry Malioutov for several insightful discussions. This work was done while Anant Raj was a summer intern at IBM Research.

References

- [1] Fabio Anselmi, Joel Z Leibo, Lorenzo Rosasco, Jim Mutch, Andrea Tacchetti, and Tomaso Poggio. Unsupervised learning of invariant representations in hierarchical architectures. *arXiv preprint arXiv:1311.4158*, 2014.
- [2] Fabio Anselmi, Lorenzo Rosasco, and Tomaso Poggio. On invariance and selectivity in representation learning. *Information and Inference*, 2016.
- [3] Haim Avron and Vikas Sindhwani. High-performance kernel machines with implicit distributed optimization and randomization. *Technometrics*, 2015.
- [4] Lorenz C Blum and Jean-Louis Reymond. 970 million druglike small molecules for virtual screening in the chemical universe database gdb-13. *Journal of the American Chemical Society*, 131(25), 2009.
- [5] Liefeng Bo, Xiaofeng Ren, and Dieter Fox. Kernel descriptors for visual recognition. In *NIPS*, 2010.
- [6] S. Bochner. Monotone funktionen, stieltjes integrale und harmonische analyse. *Math. Ann.*, 108, 1933.
- [7] Mariusz Bojarski, Anna Choromanska, Krzysztof Choromanski, Francois Fagan, Cedric Gouy-Pailler, Anne Morvan, Nouri Sakr, Tamas Sarlos, and Jamal Atif. Structured adaptive and random spinners for fast machine learning computations. *arXiv preprint arXiv:1610.06209*, 2016.
- [8] Joan Bruna and Stephane Mallat. Invariant scattering convolution networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2013.
- [9] Krzysztof Choromanski and Vikas Sindhwani. Recycling randomness with structure for sub-linear time kernel expansions. *arXiv preprint arXiv:1605.09049*, 2016.
- [10] Andreas Christmann and Ingo Steinwart. Universal kernels on non-standard input spaces. In *Advances in neural information processing systems*, pages 406–414, 2010.
- [11] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *Proceedings of the 33rd International Conference on Machine Learning*, 2016.
- [12] Bo Dai, Niao He, Yunpeng Pan, Byron Boots, and Le Song. Learning from conditional distributions via dual kernel embeddings. *arXiv preprint arXiv:1607.04579*, 2016.
- [13] Bo Dai, Bo Xie, Niao He, Yingyu Liang, Anant Raj, Maria-Florina F Balcan, and Le Song. Scalable kernel methods via doubly stochastic gradients. In *Advances in Neural Information Processing Systems*, pages 3041–3049, 2014.
- [14] Sander Dieleman, Jeffrey De Fauw, and Koray Kavukcuoglu. Exploiting cyclic symmetry in convolutional neural networks. *ICML*, 2016.
- [15] Sander Dieleman, Kyle Willett, and Joni Dambre. Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Monthly Notices Of The Royal Astronomical Society*, 2015.
- [16] Petros Drineas and Michael W Mahoney. On the nyström method for approximating a gram matrix for improved kernel-based learning. *Journal of Machine Learning Research*, 6(Dec):2153–2175, 2005.
- [17] Kenji Fukumizu, Arthur Gretton, Xiaohai Sun, and Bernhard Schölkopf. Kernel measures of conditional dependence. In *NIPS*, volume 20, pages 489–496, 2007.
- [18] Robert Gens and Pedro M Domingos. Deep symmetry networks. In *NIPS*, 2014.
- [19] Alex Gittens and Michael W Mahoney. Revisiting the nystrom method for improved large-scale machine learning. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, 2013.
- [20] Bernard Haasdonk, A Vossen, and Hans Burkhardt. Invariance in kernel methods by haar-integration kernels. In *Scandinavian Conference on Image Analysis*, pages 841–851. Springer, 2005.
- [21] Sanjiv Kumar, Mehryar Mohri, and Ameet Talwalkar. Sampling methods for the nyström method. *Journal of Machine Learning Research*, 13, 2012.
- [22] Hugo Larochelle, Dumitru Erhan, Aaron Courville, James Bergstra, and Yoshua Bengio. An empirical evaluation of deep architectures on problems with many factors of variation. In *Proceedings of the 24th international conference on Machine learning*, 2007.

- [23] Quoc Le, Tamás Szepesvári, and Alex Smola. Fastfood-approximating kernel expansions in loglinear time. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.
- [24] Yann Lecun, Lon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998.
- [25] Zhiyun Lu, Avner May, Kuan Liu, Alireza Bagheri Garakani, Dong Guo, Aurélien Bellet, Linxi Fan, Michael Collins, Brian Kingsbury, Michael Picheny, et al. How to scale up kernel methods to be as good as deep neural nets. *arXiv preprint arXiv:1411.4000*, 2014.
- [26] Julien Mairal. End-to-end kernel learning with supervised convolutional kernel networks. In *NIPS*, 2016.
- [27] Julien Mairal, Piotr Koniusz, Zaid Harchaoui, and Cordelia Schmid. Convolutional kernel networks. In *NIPS*, 2014.
- [28] Grégoire Montavon, Katja Hansen, Siamac Fazli, Matthias Rupp, Franziska Biegler, Andreas Ziehe, Alexandre Tkatchenko, Anatole V Lilienfeld, and Klaus-Robert Müller. Learning invariant representations of molecules for atomization energy prediction. In *Advances in Neural Information Processing Systems*, pages 440–448, 2012.
- [29] Youssef Mroueh, Stephen Voinea, and Tomaso A Poggio. Learning with group invariant features: A kernel perspective. In *Advances in Neural Information Processing Systems*, pages 1558–1566, 2015.
- [30] Krikamol Muandet, Kenji Fukumizu, Francesco Dinuzzo, and Bernhard Schölkopf. Learning from distributions via support measure machines. In *Advances in neural information processing systems*, pages 10–18, 2012.
- [31] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyonds. *arXiv preprint arXiv:1605.09522*, 2016.
- [32] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184, 2007.
- [33] Ali Rahimi and Benjamin Recht. Uniform approximation of functions with random bases. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*. IEEE, 2008.
- [34] Salah Rifai, Yann Dauphin, Pascal Vincent, Yoshua Bengio, and Xavier Muller. The manifold tangent classifier. In *NIPS*, volume 271, page 523, 2011.
- [35] Joseph Rotman. *An introduction to the theory of groups*, volume 148. Springer Science & Business Media, 2012.
- [36] Walter Rudin. *Fourier analysis on groups*. John Wiley & Sons, 2011.
- [37] Uwe Schmidt and Stefan Roth. Learning rotation-aware features: From invariant priors to equivariant descriptors. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, 2012.
- [38] Bernhard Scholkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- [39] Alex Smola, Arthur Gretton, Le Song, and Bernhard Schölkopf. A hilbert space embedding for distributions. In *International Conference on Algorithmic Learning Theory*, pages 13–31. Springer, 2007.
- [40] Kihyuk Sohn and Honglak Lee. Learning invariant representations with local transformations. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.
- [41] Bharath K Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Gert RG Lanckriet, and Bernhard Schölkopf. Injective hilbert space embeddings of probability measures. In *COLT*, volume 21, pages 111–122, 2008.
- [42] Bharath K Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert RG Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11(Apr):1517–1561, 2010.
- [43] Ingo Steinwart. On the influence of the kernel on the consistency of support vector machines. *Journal of machine learning research*, 2(Nov):67–93, 2001.

- [44] C. Walder and O. Chapelle. Learning with transformation invariant kernels. In *NIPS*, 2007.
- [45] Christopher Williams and Matthias Seeger. Using the nyström method to speed up kernel machines. In *Proceedings of the 14th annual conference on neural information processing systems*, pages 682–688, 2001.
- [46] Jiyang Yang, Vikas Sindhwani, Haim Avron, and Michael W. Mahoney. Quasi-monte carlo feature maps for shift-invariant kernels. In *Proceedings of the 31st International Conference on Machine learning*, 2014.