

Deep Learning on Lie Groups for Skeleton-based Action Recognition

Zhiwu Huang[†], Chengde Wan[†], Thomas Probst[†], Luc Van Gool^{†‡}

[†]Computer Vision Lab, ETH Zurich, Switzerland [‡]VISICS, KU Leuven, Belgium

{zhiwu.huang, wanc, probstt, vangool}@vision.ee.ethz.ch

Abstract

In recent years, skeleton-based action recognition has become a popular 3D classification problem. State-of-the-art methods typically first represent each motion sequence as a high-dimensional trajectory on a Lie group with an additional dynamic time warping, and then shallowly learn favorable Lie group features. In this paper we incorporate the Lie group structure into a deep network architecture to learn more appropriate Lie group features for 3D action recognition. Within the network structure, we design rotation mapping layers to transform the input Lie group features into desirable ones, which are aligned better in the temporal domain. To reduce the high feature dimensionality, the architecture is equipped with rotation pooling layers for the elements on the Lie group. Furthermore, we propose a logarithm mapping layer to map the resulting manifold data into a tangent space that facilitates the application of regular output layers for the final classification. Evaluations of the proposed network for standard 3D human action recognition datasets clearly demonstrate its superiority over existing shallow Lie group feature learning methods as well as most conventional deep learning methods.

1. Introduction

Due to the development of depth sensors, 3D human activity analysis [27, 45, 23, 43, 41, 3, 42, 37, 44, 26, 35, 17] has attracted more interest than ever before. Recent manifold-based approaches are quite successful at 3D human action recognition thanks to their view-invariant manifold-based representations for skeletal data. Typical examples include shape silhouettes in the Kendall’s shape space [40, 3], linear dynamical systems on the Grassmann manifold [39], histograms of oriented optical flow on a hyper-sphere [11], and pairwise transformations of skeletal joints on a Lie group [41, 3, 42]. In this paper, we focus on studying manifold-based approaches [41, 3, 42] to learn more appropriate Lie group representations of skeletal action data, that have achieved state-of-the-art performances for some 3D human action recognition benchmarks.

As studied in [41, 3, 42], Lie group feature learning methods often suffer from speed variations (i.e., temporal misalignment), which tend to deteriorate classification accuracy. To handle this issue, they typically employ dynamic time warping (DTW), as originally used in speech processing [30]. Unfortunately, such process costs additional time, and also results in a two-step system that typically performs worse than an end-to-end learning scheme. Moreover, such Lie group representations for action recognition tend to be extremely high-dimensional, in part because the features are extracted per skeletal segment and then stacked. As a result, any computation on such nonlinear trajectories is expensive and complicated. To address this problem, [41, 3, 42] attempt to first flatten the underlying manifold via tangent approximation or rolling maps, and then exploit SVM or PCA-like method to learn features in the resulting flattened space. Although these methods achieve some success, they merely adopt shallow linear learning schemes, yielding sub-optimal solutions on the specific nonlinear manifolds.

Deep neural networks have shown their great power in learning compact and discriminative representations for images and videos, thanks to their ability to perform nonlinear computations and the effectiveness of gradient descent training with backpropagation. This has motivated us to build a deep neural network architecture for representation learning on Lie groups. In particular, inspired by the classical manifold learning theory [38, 36, 4, 12, 20, 19], we equip the new network structure with rotation mapping layers, with which the input Lie group features are transformed to new ones with better alignment. As a result, the effect of speed variations can be appropriately mitigated. In order to reduce the high dimensionality of the Lie group features, we design special pooling layers to compose them in terms of spatial and temporal levels, respectively. As the output data reside on nonlinear manifolds, we also propose a Riemannian computing layer, whose outputs could be fed into any regular output layers such as a softmax layer. In short, our main contributions are:

- A novel neural network architecture is introduced to deeply learn more desirable Lie group representations for the problem of skeleton-based action recognition.

- The proposed network provides a paradigm to incorporate the Lie group structure into deep learning, which generalizes the traditional neural network model to non-Euclidean Lie groups.
- To train the network within the backpropagation framework, a variant of stochastic gradient descent optimization is exploited in the context of Lie groups.

2. Relevant Work

Already quite some works [46, 34, 2, 29, 33, 14, 15] have applied aspects of Lie group theory to deep neural networks. For example, [33] investigated how stability properties of a continuous recursive neural network can be altered within neighbourhoods of equilibrium points by the use of Lie group projections operating on the synaptic weight matrix. [14] studied the behavior of unsupervised neural networks with orthonormality constraints, by exploiting the differential geometry of Lie groups. In particular, two sub-classes of the general Lie group learning theories were studied in detail, tackling first-order (gradient-based) and second-order (non-gradient-based) learning. [15] introduced deep symmetry networks (symnets), a generalization of convolutional networks that forms feature maps over arbitrary symmetry groups that are basically Lie groups. The symnets utilize kernel-based interpolation to tractably tie parameters and pool over symmetry spaces of any dimension.

Moreover, recently some deep learning models have emerged [10, 7, 28, 25, 18, 21] that deal with data in a non-Euclidean domain. For instance, [10] proposed a spectral version of convolutional networks to handle graphs. It exploits the notion of non shift-invariant convolution, relying on the analogy between the classical Fourier transform and the Laplace-Beltrami eigenbasis. [25] developed a scalable method for treating an arbitrary spatio-temporal graph as a rich recurrent neural network mixture, which can be used to transform any spatio-temporal graph by employing a certain set of well-defined steps. For shape analysis, [28] proposed a ‘geodesic convolution’ on local geodesic coordinate systems to extract local patches on the shape manifold. This approach performs convolutions by sliding a window over the manifold, and local geodesic coordinates are used instead of image patches. To deeply learn symmetric positive definite (SPD) matrices - used in many tasks - [18] developed a Riemannian network on the manifolds of SPD matrices, with some layers specially designed to deal with such structured matrices.

In summary, such works have applied some theories of Lie groups to regular networks, and even generalized the common networks to non-Euclidean domains. Nevertheless, to the best of our knowledge, this is the first work that studies a deep learning architecture on Lie groups to handle the problem of skeleton-based action recognition.

3. Lie Group Representation for Skeletal Data

Let $S = (V, E)$ be a body skeleton, where $V = \{v_1, \dots, v_N\}$ denotes the set of body joints, and $E = \{e_1, \dots, e_M\}$ indicates the set of edges, i.e. oriented rigid body bones. As studied in [41, 3, 42], the relative geometry of a pair of body parts e_m and e_n can be represented in a local coordinate system attached to the other. The local coordinate system of body part e_n is calculated by rotating with minimum rotation so that its starting joint becomes the origin and it coincides with the x -axis. With the process, we consequently get the transformed 3D vectors \hat{e}_m, \hat{e}_n for the two edges e_m, e_n respectively. Then we can compute the rotation matrix $R_{m,n}$ ($R_{m,n}^T R_{m,n} = R_{m,n} R_{m,n}^T = I_n, |R_{m,n}| = 1$) from e_m to the local coordinate system of e_n . Specifically, we can firstly calculate the axis-angle representation (ω, θ) for the rotation matrix $R_{m,n}$ by

$$\omega = \frac{\hat{e}_m \otimes \hat{e}_n}{\|\hat{e}_m \otimes \hat{e}_n\|}, \quad (1)$$

$$\theta = \arccos(\hat{e}_m \cdot \hat{e}_n). \quad (2)$$

where \otimes, \cdot are outer and inner products respectively. Then, the axis-angle representation can be easily transformed to a rotation matrix $R_{m,n}$. In the same way, the rotation matrix $R_{n,m}$ from e_n to the local coordinate system of e_m can be computed. To fully encode the relative geometry between e_m and e_n , $R_{m,n}$ and $R_{n,m}$ are both used. As a result, a skeleton S at the time instance t is represented by the form $(R_{1,2}(t), R_{2,1}(t), \dots, R_{M-1,M}(t), R_{M,M-1}(t))$, where M is the number of body parts, and the number of rotation matrices is $2C_M^2$ (C_M^2 is the combination formula).

The set of $n \times n$ rotation matrices in \mathbb{R}^n forms the special orthogonal group SO_n which is actually a matrix Lie group [22, 9, 16]. Accordingly, each motion sequence of a moving skeleton is represented with a curve on the Lie group $SO_3 \times \dots \times SO_3$. It is known that the matrix Lie group is endowed with a Riemannian manifold structure that is differentiable. Hence, at each point R_0 on SO_n , one can derive the tangent space $T_{R_0}SO_n$ that is a vector space spanned by the set of skew-symmetric matrices. When the anchor point is the identity matrix $I_n \in SO_n$, the resulting tangent space is known as the Lie algebra so_n . As the tangent spaces are equipped with the inner product, the Riemannian metric on SO_n can be defined by the Frobenius inner product:

$$\langle A_1, A_2 \rangle = \text{trace}(A_1^T A_2), A_1, A_2 \in T_{R_0}SO_n. \quad (3)$$

The logarithm map \log_{R_0} and exponential map \exp_{R_0} at R_0 on SO_n associated with the Riemannian metric can be expressed in terms of the usual matrix logarithm log and exponential exp as

$$\log_{R_0}(R_1) = \log(R_1 R_0^T) \text{ with } R_0, R_1 \in SO_n, \quad (4)$$

$$\exp_{R_0}(A_1) = \exp^{A_1 R_0^T} \text{ with } A_1 \in T_{R_0}SO_n. \quad (5)$$




Figure 1. Conceptual illustration of the proposed Lie group Network (LieNet) architecture. In the network structure, the data space of each RotMap/RotPooling layer corresponds to a Lie group, while the weight spaces of the RotMap layers are Lie groups as well.

4. Lie Group Network for Skeleton-based Action Recognition

For the problem of skeleton-based action recognition, we build a deep network architecture to learn the Lie group representations of skeletal data. The network structure is dubbed as LieNet, where each input is an element on the Lie Group. Like convolutional networks (ConvNets), the LieNet also exhibits fully connected convolution-like layers and pooling layers, named rotation mapping (RotMap) layers and rotation pooling (RotPooling) layers respectively. In particular, the proposed RotMap layers perform transformations on input rotation matrices to generate new rotation matrices, which have the same manifold property, and are expected to be aligned more accurately for more reliable matching. The RotPooling layers aim to pool the resulting rotation matrices at both spatial and temporal levels such that the Lie group feature dimensionality can be reduced. Since the rotation matrices reside on non-Euclidean manifolds, we have to design a layer named logarithm mapping (LogMap) layer, to perform the Riemannian computations on. This transforms the rotation matrices into the usual skew-symmetric matrices, which lie in Euclidean space and hence can be fed into any regular output layers. The architecture of the proposed LieNet is shown in Fig. 1.

4.1. RotMap Layer

As well-known from classical manifold learning theory [38, 36, 4, 12, 20, 19], one can learn or preserve the original data structure to faithfully maintain geodesic distances for better classification. Accordingly, we design a RotMap layer to transform the input rotation matrices to new ones that are more suitable for the final classification. Formally,

the RotMap layers adopt a rotation mapping f_r as

$$\begin{aligned} f_r^{(k)}((R_1^{k-1}, R_2^{k-1}, \dots, R_M^{k-1}); W_1^k, W_2^k, \dots, W_M^k) \\ = (W_1^k R_1^{k-1}, W_2^k R_2^{k-1}, \dots, W_M^k R_M^{k-1}) \\ = (R_1^k, R_2^k, \dots, R_M^k) \end{aligned} \quad (6)$$

where $\hat{M} = 2C_M^2$ (M is the number of body bones in one skeleton, C_M^2 is the combination computation), $(R_1^{k-1}, R_2^{k-1}, \dots, R_M^{k-1}) \in SO_3 \times SO_3 \dots \times SO_3$ is the input Lie group feature (i.e., product of rotation matrices) for one skeleton in the k -th layer, $W_i^k \in \mathbb{R}^{3 \times 3}$ is the transformation matrix (connection weights), and $(R_1^k, R_2^k, \dots, R_M^k)$ is the resulting Lie group representation. Note that although there is only one transformation matrix for each rotation matrix, it would be easily extended with multiple projections for each input. To ensure the form $(R_1^k, R_2^k, \dots, R_M^k)$ becomes a valid product of rotation matrices residing on $SO_3 \times SO_3 \dots \times SO_3$, the transformation matrices $W_1^k, W_2^k, \dots, W_M^k$ are all basically required to be rotation matrices. Accordingly, both the data and the weight spaces on each RotMap layer correspond to a Lie group $SO_3 \times SO_3 \dots \times SO_3$.

Since the RotMap layers are designed to work together with the classification layer, each resulting skeleton representation is tuned for more accurate classification in an end-to-end deep learning manner. In other words, the major purpose of designing the RotMap layers is to align the Lie group representations of a moving skeleton for more faithful matching.

4.2. RotPooling Layer

In order to reduce the complexity of deep models, it is typically useful to reduce the size of the representations to

decrease the amount of parameters and computation in the network. For this purpose, it is common to insert a pooling layer in-between successive convolutional layers in a typical ConvNet architecture. The pooling layers are often designed to compute statistics in local neighborhoods, such as sum aggregation, average energy and maximum activation.

Without loss of generality, we here just introduce max pooling¹ to the LieNet setting with the equivalent notion of neighborhood. Since the input and output of the special pooling layers are both expected to be rotation matrices, we call this kind of layers as rotation pooling (RotPooling) layers. For the RotPooling, we propose two different concepts of neighborhood in this work. The first one is on the spatial level. As shown in Fig.2(a)→(b), we first pool the Lie group features on each pair of basic bones e_m, e_n in the i -th frame, which is represented by the two rotation matrices $\mathbf{R}_{m,n}^{k-1,i}, \mathbf{R}_{n,m}^{k-1,i}$ (here $k-1$ is the order of the layer) as aforementioned. Then, as depicted in Fig.2(b)→(c), we can perform pooling on the adjacent bones that belong to the same group (here, we can define five part groups, i.e., torso, two arms and two legs, of the body). However, the second step would inevitably result in a serious spatial misalignment problem, and thus lead to bad matching performances. Therefore, we finally only adopt the first step pooling. In this setting, the function of the max pooling is given by

$$f_p^{(k)}(\{\mathbf{R}_{m,n}^{k-1,i}, \mathbf{R}_{n,m}^{k-1,i}\}) = \max(\{\mathbf{R}_{m,n}^{k-1,i}, \mathbf{R}_{n,m}^{k-1,i}\}) = \begin{cases} \mathbf{R}_{m,n}^{k-1,i}, & \text{if } \Theta(\mathbf{R}_{m,n}^{k-1,i}) > \Theta(\mathbf{R}_{n,m}^{k-1,i}), \\ \mathbf{R}_{n,m}^{k-1,i}, & \text{otherwise,} \end{cases} \quad (7)$$

where $\Theta(\cdot)$ is the representation of the given rotation matrix such as quaternion, Euler angle or Euler axis-angle. For example, the Euler axis ω and angle θ representations are typically calculated by

$$\omega(\mathbf{R}_{n,m}) = \frac{1}{2 \sin(\theta(\mathbf{R}_{n,m}))} \begin{pmatrix} \mathbf{R}_{n,m}(3,2) - \mathbf{R}_{n,m}(2,3) \\ \mathbf{R}_{n,m}(1,3) - \mathbf{R}_{n,m}(3,1) \\ \mathbf{R}_{n,m}(2,1) - \mathbf{R}_{n,m}(1,2) \end{pmatrix}, \quad (8)$$

$$\theta(\mathbf{R}_{n,m}) = \arccos\left(\frac{\text{trace}(\mathbf{R}_{n,m}) - 1}{2}\right), \quad (9)$$

where $\mathbf{R}_{n,m}(i,j)$ is the i -th row, j -th column element of $\mathbf{R}_{n,m}$. Unfortunately, except the angle representation, it is non-trivial to define an ordering relation for a quaternion or an axis-angle representation. Hence, in this paper, we finally adopt the angle form Eqn.9 of rotation matrices and its simple ordering relation to calculate the function $\Theta(\cdot)$.

The other pooling scheme is on the temporal level. As shown in Fig.2 (c)→(d), the aim of the temporal pooling

¹In contrast to sum and mean poolings, max pooling can generate valid rotation matrices directly, and hence suits the proposed LieNets. On the other hand, leveraging Lie group computing to enable sum and mean pooling to work for the LieNets, however, goes beyond the scope of this paper.




Figure 2. Illustration of spatial pooling (SpaPooling) (a)→(b)→(c) and temporal pooling (TemPooling) (c)→(d) schemes.

is to obtain more compact representations for a motion sequence. This is because a sequence often contains many frames, which results in the problem of extremely high-dimensional representations. Thus, pooling in the temporal domain can reduce the model complexity as well. Formally, the function of this kind of max pooling is defined as

$$f_p^{(k)}(\{(\mathbf{R}_{1,2}^{k-1,1} \dots \mathbf{R}_{M-1,M}^{k-1,1}) \dots, (\mathbf{R}_{1,2}^{k-1,p} \dots, \mathbf{R}_{M-1,M}^{k-1,p})\}) = (\max(\{\mathbf{R}_{1,2}^{k-1,1} \dots, \mathbf{R}_{1,2}^{k-1,p}\}) \dots, \max(\{\mathbf{R}_{M-1,M}^{k-1,1} \dots, \mathbf{R}_{M-1,M}^{k-1,p}\})), \quad (10)$$

where M is the number of body parts in one skeleton, p is the number of skeleton frames for pooling, and the function $\max(\cdot)$ is defined in the way of Eqn.7.

4.3. LogMap Layer

Classification of curves on the Lie group $SO_3 \times \dots \times SO_3$ is a complicated task due to the non-Euclidean nature of the underlying space. To address the problem as in [42], we design the logarithm map (LogMap) layer to flatten the Lie group $SO_3 \times \dots \times SO_3$ to its Lie algebra $so_3 \times \dots \times so_3$. Accordingly, by using the logarithm map Eqn.4, the function of this layer can be defined as

$$f_l^{(k)}((\mathbf{R}_1^{k-1}, \mathbf{R}_2^{k-1} \dots, \mathbf{R}_M^{k-1})) = (\log(\mathbf{R}_1^{k-1}), \log(\mathbf{R}_2^{k-1}) \dots, \log(\mathbf{R}_M^{k-1})). \quad (11)$$

One typical approach to calculate the logarithm map is to use the approach $\log(\mathbf{R}) = \mathbf{U} \log(\mathbf{\Sigma}) \mathbf{U}^T$, where $\mathbf{R} = \mathbf{U} \mathbf{\Sigma} \mathbf{U}^T$, $\log(\mathbf{\Sigma})$ is the diagonal matrix of the eigenvalue logarithms. However, the spectral operation not only suffers from the problem of zeroes occurring in $\log(\mathbf{\Sigma})$ due to the property of the rotation matrix \mathbf{R} , but also consumes too much time for matrix gradient computation [24]. Therefore, we resort to other approaches to perform the function of this

layer. Fortunately, we can explore the relationship between the logarithm map and the axis-angle representation as:

$$\log(\mathbf{R}) = \begin{cases} 0, & \text{if } \theta(\mathbf{R}) = 0, \\ \frac{\theta(\mathbf{R})}{2\sin(\theta(\mathbf{R}))}(\mathbf{R} - \mathbf{R}^T), & \text{otherwise,} \end{cases} \quad (12)$$

where $\theta(\mathbf{R})$ is the angle Eqn.9 of \mathbf{R} . With this equation, the corresponding matrix gradient can be easily derived by traditional element-wise matrix calculation.

4.4. Output Layers

After performing the LogMap layers, the outputs can be transformed into vector form and concatenated directly frame by frame within one sequence due to their Euclidean nature. Then, we can add any regular network layers such as rectified linear unit (ReLU) layers and regular fully connected (FC) layers. In particular for the ReLU layer, we can simply set relatively small elements to zero as done in classical ReLU. In the FC layer, the dimensionality of the weight is set to $d_k \times d_{k-1}$, where d_k and d_{k-1} are the class number and the vector dimensionalities, respectively. For skeleton-based action recognition, we employ a common softmax layer as the final output layer. Besides, as studied in [37, 26], learning temporal dependencies over the sequential data can improve human action recognition. Hence, we can also feed the outputs into Long Short-Term Memory (LSTM) unit to learn useful temporal features. Because of the space limitation, we do not study this any further.

5. Training Procedure

In order to train the proposed LieNets, we exploit the Stochastic gradient descent (SGD) algorithm that is one of the most popular network training tools. To begin with, let the LieNet model be represented as a sequence of function compositions $f = f^{(l)} \circ f^{(l-1)} \dots \circ f^{(1)}$ with a parameter tuple $\mathbf{W} = (\mathbf{W}_l, \mathbf{W}_{l-1} \dots, \mathbf{W}_1)$, where $f^{(k)}$ is the function for the k -th layer, \mathbf{W}_k (dropping the sample index for simplicity) represents the weight parameters of the k -th layer, and l is the number of layers. The loss of the k -th layer is defined by $L^{(k)} = \ell \circ f^{(l)} \dots \circ f^{(k)}$, where ℓ is the loss function for the final output layer.

To optimize the deep model, one classical SGD algorithm needs to compute the gradient of the objective function, which is typically achieved by the backpropagation chain rule. In particular, the gradients of the weight \mathbf{W}_k and the data \mathbf{R}_{k-1} (dropping the sample index for simplicity) for the k -th layer can be respectively computed by the chain rule:

$$\frac{\partial L^{(k)}(\mathbf{R}_{k-1}, y)}{\partial \mathbf{W}_k} = \frac{\partial L^{(k+1)}(\mathbf{R}_k, y)}{\partial \mathbf{R}_k} \frac{\partial f^{(k)}(\mathbf{R}_{k-1})}{\partial \mathbf{W}_k}, \quad (13)$$

$$\frac{\partial L^{(k)}(\mathbf{R}_{k-1}, y)}{\partial \mathbf{R}_{k-1}} = \frac{\partial L^{(k+1)}(\mathbf{R}_k, y)}{\partial \mathbf{R}_k} \frac{\partial f^{(k)}(\mathbf{R}_{k-1})}{\partial \mathbf{R}_{k-1}}, \quad (14)$$

where y is the class label, $\mathbf{R}_k = f^{(k)}(\mathbf{R}_{k-1})$. Eqn.13 is the gradient for updating \mathbf{W}_k , while Eqn.14 computes the gradients in the layers below to update \mathbf{R}_{k-1} .

The gradients of the data involved in RotPooling, LogMap and regular output layers can be calculated by Eqn.14 as usual. Particularly, the gradient for the data in RotPooling can be computed with the same gradient computing approach used in a regular max pooling layer in the context of traditional ConvNets. For the data in the LogMap layer, the gradient can be obtained by the element-wise gradient computation on the involved rotation matrices.

On the other hand, the computation of the gradients of the parameter weights defined in the RotMap layers is non-trivial. This is because the weight matrices are enforced to be on the Riemannian manifold SO_3 of the rotation matrices, i.e. the Lie group. As a consequence, merely using Eqn.13 to compute their Euclidean gradients rather than Riemannian gradients in the procedure of backpropagation would not generate valid rotation weights. To handle this problem, we propose a new approach of updating the weights used in Eqn.6 for the RotMap layers. As studied in [1], the steepest descent direction for the used loss function $L^{(k)}(\mathbf{R}_{k-1}, y)$ with respect to \mathbf{W}_k on the manifold SO_3 is the Riemannian gradient $\tilde{\nabla} L_{\mathbf{W}_k}^{(k)}$, which can be obtained by parallel transporting the Euclidean gradients onto the corresponding tangent space. In particular, transporting the gradient from a point \mathbf{W}_k^t to another point \mathbf{W}_k^{t+1} requires subtracting the normal component $\bar{\nabla} L_{\mathbf{W}_k}^{(k)}$, at \mathbf{W}_k^{t+1} , which can be obtained as follows:

$$\bar{\nabla} L_{\mathbf{W}_k}^{(k)} = \nabla L_{\mathbf{W}_k}^{(k)} \mathbf{W}_k^T \mathbf{W}_k, \quad (15)$$

where the Euclidean gradient $\nabla L_{\mathbf{W}_k}^{(k)}$ is computed by using Eqn.13 as

$$\nabla L_{\mathbf{W}_k}^{(k)} = \frac{\partial L^{(k+1)}(\mathbf{R}_k, y)}{\partial \mathbf{R}_k} \mathbf{R}_{k-1}^T. \quad (16)$$

Thanks to the parallel transport, the Riemannian gradient can be calculated by

$$\tilde{\nabla} L_{\mathbf{W}_k}^{(k)} = \nabla L_{\mathbf{W}_k}^{(k)} - \bar{\nabla} L_{\mathbf{W}_k}^{(k)}. \quad (17)$$

Searching along the tangential direction takes the update in the tangent space of the SO_3 manifold. Then, such update is mapped back to the SO_3 manifold with a retraction operation. Consequently, an update of the weight \mathbf{W}_k on the SO_3 manifold is of the following form

$$\mathbf{W}_k^{t+1} = \Gamma(\mathbf{W}_k^t - \lambda \tilde{\nabla} L_{\mathbf{W}_k}^{(k)}), \quad (18)$$

where \mathbf{W}_k^t is the current weight, Γ is the retraction operation, λ is the learning rate.

6. Experiments

We employ three standard 3D human action datasets to study the effectiveness of the proposed LieNets.

6.1. Evaluation Datasets

G3D-Gaming dataset [5] contains 663 sequences of 20 different gaming motions. Each subject performed every action more than two times. Besides, 3D locations of 20 joints (i.e., 19 bones) are provided with the dataset.

HDM05 dataset [31] consists of 2,337 sequences of 130 action classes executed by various actors. Most of the motion sequences have been performed several times by all five actors according to the guidelines in a script. As G3D-Gaming [5] dataset, 3D locations of 31 joints (i.e., 30 bones) of the subjects are provided as well with this dataset.

NTU RGB+D dataset [37] is, to the best of our knowledge, currently the largest 3D action recognition dataset, which contains more than 56,000 sequences. A total of 60 different action classes are performed by 40 subjects. 3D coordinates of 25 joints (i.e., 24 bones) are also offered. Due to its large scale, the dataset is highly suitable for deep learning.

6.2. Implementation Details

For the feature extraction, we use the code of [42] to represent each human skeleton with a point on the Lie group $SO_3 \times \dots \times SO_3$. As preprocessed in [42], we normalize any sequence of motion into a fixed N -length one. As a result, for each moving skeleton, we finally compute a Lie group curve of length 100, 16, 64 for the G3D-Gaming, HDM05 and NTU RGB-D datasets, respectively.

As the focus of this work is on skeleton-based action recognition, we mainly utilize manifold-based approaches for comparison. The two baseline approaches are the special Euclidean group (SE) [41] and the special orthogonal group (SO) [42] representations based shallow learning methods. For a fair comparison, we use the source codes from the original authors, and set the involved parameters as in the original papers. For the proposed LieNet, we build its architecture with single or multiple block(s) of RotMap/RotPooling layers illustrated in Fig. 1 before the three final layers being LogMap, FC and softmax layers. The learning rate λ is fixed to 0.01, the batch size is set to 30, the weights in the RotMap layers are initialized as random rotation matrices, the number of samples for the temporal RotPooling (TemPooling) layer is set to 4. For training the LieNet, we just use an i7-6700K (4.00GHz) PC without any GPUs. As the LieNet gets promising results on all datasets with the same configuration, this shows its insensitivity to the parameter settings. Note that, for the LieNets, we do not employ the dynamic time warping (DTW) technique [30], which has been used in the SO and SE methods to solve the problem of speed variations.

Method	G3D-Gaming
RBM+HMM [32]	86.40%
SE [41]	87.23%
SO [42]	87.95%
LieNet-0Block	84.55%
LieNet-1Block	85.16%
LieNet-2Blocks	86.67%
LieNet-3Blocks	89.10%

Table 1. Recognition accuracies on the G3D-Gaming database.

6.3. Experimental Results

G3D-Gaming dataset [5]. For the dataset, we follow a cross-subject test setting, where half the subjects are used for training and the other half are employed for testing. All the results reported for this dataset are averaged over ten different combinations of training and testing datasets.

Table 1 compares the proposed LieNet with the state-of-the-art methods (i.e., RBM-HMM [32], SE [41] and SO [42]) reported for the G3D-Gaming dataset. To be fair, we report their results without using the Fourier Temporal Pyramid (FTP) post-processing (their accuracies are 91.09% and 90.94% after using FTP). As shown in Table 1, the LieNet shows its superiority over the two baseline methods SO and SE. Besides, our LieNet with 3 blocks of RotMap and RotPooling layers achieves the best performance. For this dataset, we also study the performances of different block numbers in the LieNet architecture. As the number of frames in each sequence was fixed to 100 as mentioned before, adding more blocks to the LieNet will finally degenerate the temporal sequence into only 1 frame. In theory, this extreme case would result in the loss of the temporal resolution and thus undermine the performance of recognizing activities. In order to keep the balance between compact spatial feature learning and temporal information encoding, we therefore equip the LieNet with a limited number of blocks in different settings. Thus, we study 4 blocks at most for our LieNets. For the case of stacking 4 blocks, we find its performance (87.28%) is lower than that of the 3-block case, which justifies the above argumentation. Nevertheless, as observed from Table 1, stacking some more RotMap/RotPooling blocks can improve the performance of the proposed LieNet.

In addition, we evaluate the performances of different LieNet configurations as shown in Fig. 3. The left of Fig. 3 verifies the necessity of using RotMap, RotPooling and LogMap layers to improve the proposed LieNet-3Blocks. In addition, we also compare the LieNet with and without DTW. On this dataset, the performance (88.89% vs. 89.10%) of these two cases is approximately equal. Therefore, the benefit of using RotMap layers somehow shows




Figure 3. (a) Comparison of different LieNet configurations: without using RotMap layers (w/o-RotMap), w/o-RotPooling layers, w/o-LogMap layers and using all (w/-All) in LieNet-3Blocks for G3D-Gaming. (b) Comparison of different pooling schemes: using 1 spatial RotPooling layer (1SpaPooling, i.e., LieNet-1Block), 2 spatial RotPooling layers (2SpaPooling), 1SpaPooling+1 temporal RotPooling layer (1Spa1TemPooling, i.e., LieNet-2Blocks), 1SpaPooling+2TemPooling (1Spa2TemPooling, i.e., LieNet-3Blocks) for G3D-Gaming.

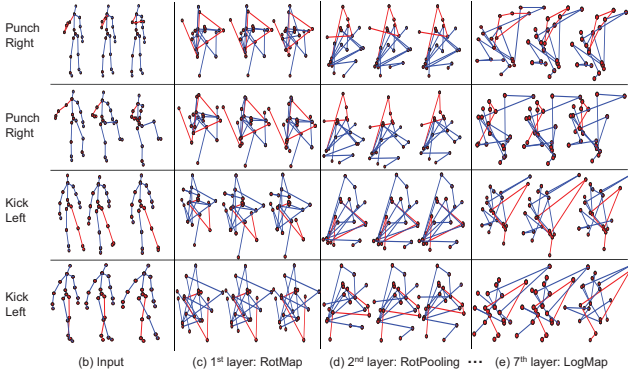


Figure 4. The example skeletons are reconstructed by the output rotation matrices of some representative LieNet layers for processing four action sequences from the G3D-Gaming dataset. The bones in red are the interesting ones for the action classes.

it can take the role of DTW that solves the problem of speed variations. The right of Fig.3 analyses the effectiveness of the 1SpaPooling case (i.e., Fig.2(a)→(b)), and shows the decreasing performance behavior of the 2SpaPooling case (i.e., Fig.2(a)→(b)→(c)). Thus, we finally utilize 1SpaPooling and 2TemPooling (i.e., Fig.2(c)→(d)) in the LieNet-3Blocks structure. Besides, we also study the behavior of adding a rectified linear unit (ReLU)-like layer (i.e., setting the matrix elements below a threshold $\epsilon = 0.1$ to zero) on the top of the LogMap layer as presented before. Yet, the performance was worse (87.58%) than without. Further, to validate the improvements are from the contribution of the RotMap and RotPooling layers rather than deeper architectures, we build a regular (LeNet-like) deep structure, i.e., LogMap→2×(FC→MaxPooling)→FC→ReLU→FC→Softmax, that applies 8 regular layers on the con-

Method	HDM05
SPDNet [18]	61.45%±1.12
SE [41]	70.26%±2.89
SO [42]	71.31%±3.21
LieNet-0Block	71.26%±2.12
LieNet-1Block	73.35%±1.14
LieNet-2Blocks	75.78%±2.26

Table 2. Recognition accuracies on the HDM05 database.

catenated output Euclidean forms of the LogMap layer. The step for MaxPooling is set to 4, and the sizes of different FC weights are set to 307800×40000 , 10000×4000 , 1000×400 and 400×20 respectively. The performance of this network is 85.49%, which supports the validation.

For a better understanding of the proposed LieNet, we also visualize the output results of some representative layers. In particular, we roughly estimate the 3D location of each body bone, given the learned rotation matrix and the 3D coordinate of the beginning edge in the torso part. In Fig.4, we present the visualization of some layers for four action sequences, that belong to the classes of ‘punch right’ and ‘kick left’. As shown in Fig.4, we observe that they yield meaningful semantic information layer by layer for specific classes. Specifically, the reconstructions from the first layer (RotMap) and the second layer (RotPooling) typically still mix some patterns specific for the action classes with some rather confusing ones. But, when arriving at the seventh layer (LogMap), the patterns for specific motion classes become more discriminative.

HDM05 dataset [31]. Following [18], we conduct 10 random evaluations, each of which randomly selects half of the sequences for training and the rest for testing.

As listed in Table 2, besides to the two baseline methods SE and SO, we also study the SPDNet method [18] that has reached the best performance so far for this dataset. The large improvement of SE and SO over SPDNet suggests the effectiveness of the Lie group representations for the problem of skeleton-based action recognition. As the last experiment on the G3D-Gaming dataset, we also study the proposed LieNet with different numbers of blocks of RotMap and RotPooling layers. Note that since the length of each sequence in this database is fixed to 16 frames, as studied in the last evaluations, adding too much LieNet blocks will lead to the loss of the temporal resolution. Thus, we implemented the LieNet with 3 blocks at most for the dataset. As adding 3 blocks will generate 1 frame for each video, its performance (70.42%) is not as promising as other cases. In contrast, as reported in Table 2, using more blocks (below 3 blocks) improves over using less blocks, and gets the state-of-the-art on the dataset, again showing its advantages over SE and SO shallow learning methods.

Method	RGB+D-subject	RGB+D-view
HBRNN [13]	59.07%	63.97%
Deep RNN [37]	56.29%	64.09%
Deep LSTM [37]	60.69%	67.29%
PA-LSTM [37]	62.93%	70.27%
ST-LSTM [26]	69.2%	77.7%
SE [41]	50.08%	52.76%
SO [42]	52.13%	53.42%
LieNet-0Block	53.54%	54.78%
LieNet-1Block	56.35%	60.14%
LieNet-2Blocks	58.02%	62.52%
LieNet-3Blocks	61.37%	66.95%

Table 3. Recognition accuracies for the cross-subject and cross-view evaluations on the NTU RGB+D database.

NTU RGB+D dataset [37]. This dataset has two standard testing protocols. One is cross-subject test, for which half of the subjects are used for training and the rest is for testing. The other one is cross-view test, for which two views are employed for training and the rest one is utilized for testing.

Since this dataset is large enough to train deep networks, recent works [37, 26] studied typical Recurrent Neural Networks (deep RNN and deep LSTM) as well as two variants, i.e., part-aware (PA) and spatio-temporal (ST) versions of LSTM. The common advantage of these deep networks is to learn temporal information, and significantly outperform the Lie group representation learning methods SE and SO, which are good at learning spatial information, but are not deep learning models. In this paper, our LieNet fills the gap by showing the effectiveness of deep learning on the spatial representations. As shown in Table 3, our LieNet with more stacked blocks can significantly improve the two baseline methods SE and SO, which validates the effectiveness of the deep learning. By comparing with the state-of-the-art methods, our LieNet behaves better or equally well as most deep networks (e.g., deep RNN and deep LSTM) that exploit temporal information. The LieNet is still outperformed by the recently proposed PA-LSTM and ST-LSTM however, which jointly learn spatial and temporal features of moving skeletons. This is reasonable because the LieNet is mainly designed to learn the spatial features with only pooling temporal information.

Properties of LieNet training algorithm. While the convergence of the used SGD algorithm on Riemannian manifolds has been studied well in [8, 6] already, the convergence behavior (see Fig.5) of our LieNet training algorithm also demonstrates that it can converge to a stable solution after 100 epochs. In terms of the runtime, training LieNet-3Blocks takes about 6 minutes(m) per epoch on the G3D-gaming, and 514m per epoch on the NTU RGB+D. Training LieNet-2Blocks takes around 122m per epoch on the




Figure 5. The convergence behavior of the proposed LieNet for the G3D-Gaming dataset.

HDM05. In testing, the LieNet (i.e., the forward pass) takes about 3m, 4m and 86m respectively on G3D, HDM05 and NTU RGB+D. Note that, like usual network toolboxes, the current LieNet can be sped up a lot by implementing a GPU version. Regarding the memory requirement, LieNet-3Block-G3D, LieNet-2Block-HDM05 and LieNet-3Blocks-NTURGBD require around 1.2GB, 1.1GB and 1.4GB, respectively.

7. Summary and Future Work

We studied a deep network architecture in the domain of Lie group features, that is successful for skeleton-based action recognition. In order to handle the key issues of speed variation and high dimensionality of the Lie group features, we designed special mapping layers and pooling layers to process the resulting rotation matrices. In addition, we also exploited logarithm mapping layers to perform Riemannian computing on the representations, with which regular output layers are supplied in the new network structure. The final evaluations on three standard 3D action datasets not only demonstrated the effectiveness of the proposed network, but also compared its different configurations. Moreover, we also showed an interesting visualization for the network, which somewhat discloses its intrinsic mechanism.

As the proposed network is, to the best of our knowledge, the first attempt to perform deep learning on Lie groups for skeleton-based action recognition, there are quite a few open issues. For example, studying multiple rotation mappings per RotMap layer and exploiting a ReLU-like layer in the context of a Lie group network are worth paying attention to. Besides, building a deeper network, beginning from the raw 3D joint locations up to the Lie group features in an end-to-end learning manner, could be more effective. Last but not least, encouraged by the success of the deep spatio-temporal networks [37, 26], exploring the potential of the proposed network in the temporal setting would also be an interesting direction.

Acknowledgments: This work is supported by EU Framework Seven project ReMeDi (grant 610902).

References

- [1] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2008.
- [2] F. Albertini and E. D. Sontag. For neural networks, function determines form. In *Decision and Control, 1992., Proceedings of the 31st IEEE Conference on*, pages 26–31. IEEE, 1992.
- [3] R. Anirudh, P. Turaga, J. Su, and A. Srivastava. Elastic functional coding of Riemannian trajectories. *IEEE T-PAMI*, 2016.
- [4] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- [5] V. Bloom, D. Makris, and V. Argyriou. G3D: A gaming action dataset and real time action recognition evaluation framework. In *CVPR Workshops*, 2012.
- [6] S. Bonnabel. Stochastic gradient descent on Riemannian manifolds. *IEEE T-AC*, 58(9):2217–2229, 2013.
- [7] D. Boscaini, J. Masci, S. Melzi, M. Bronstein, U. Castellani, and P. Vandergheynst. Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. In *Computer Graphics Forum*, volume 34, pages 13–23. Wiley Online Library, 2015.
- [8] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *COMPSTAT*, 2010.
- [9] N. Boumal and P.-A. Absil. A discrete regression method on manifolds and its application to data on $SO(n)$. *IFAC Proceedings Volumes*, 44(1):2284–2289, 2011.
- [10] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. In *ICLR*, 2014.
- [11] R. Chaudhry, A. Ravichandran, G. Hager, and R. Vidal. Histograms of oriented optical flow and binet-cauchy kernels on nonlinear dynamical systems for the recognition of human actions. In *CVPR*, 2009.
- [12] D. L. Donoho and C. Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596, 2003.
- [13] Y. Du, W. Wang, and L. Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *CVPR*, 2015.
- [14] S. Fiori. Unsupervised neural learning on Lie group. *International Journal of Neural Systems*, 12(03n04):219–246, 2002.
- [15] R. Gens and P. M. Domingos. Deep symmetry networks. In *NIPS*, 2014.
- [16] B. C. Hall. *Lie groups, Lie algebras, and representations: an elementary introduction*. Springer, 2015.
- [17] F. Han, B. Reily, W. Hoff, and H. Zhang. Space-time representation of people based on 3D skeletal data: a review. *arXiv preprint arXiv:1601.01006*, 2016.
- [18] Z. Huang and L. Van Gool. A Riemannian network for SPD matrix learning. In *AAAI*, 2017.
- [19] Z. Huang, R. Wang, S. Shan, and X. Chen. Projection metric learning on Grassmann manifold with application to video based face recognition. In *CVPR*, 2015.
- [20] Z. Huang, R. Wang, S. Shan, X. Li, and X. Chen. Log-Euclidean metric learning on symmetric positive definite manifold with application to image set classification. In *ICML*, 2015.
- [21] Z. Huang, J. Wu, and L. Van Gool. Building deep networks on Grassmann manifolds. *arXiv preprint arXiv:1611.05742*, 2016.
- [22] K. Hüper and F. S. Leite. On the geometry of rolling and interpolation curves on $S(n)$, $SO(n)$, and Grassmann manifolds. *Journal of Dynamical and Control Systems*, 13(4):467–502, 2007.
- [23] M. E. Hussein, M. Torki, M. A. Gawayyed, and M. El-Saban. Human action recognition using a temporal hierarchy of covariance descriptors on 3D joint locations. In *IJCAI*, 2013.
- [24] C. Ionescu, O. Vantzos, and C. Sminchisescu. Matrix back-propagation for deep networks with structured layers. In *ICCV*, 2015.
- [25] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena. Structural-RNN: Deep learning on spatio-temporal graphs. In *CVPR*, 2016.
- [26] J. Liu, A. Shahroudy, D. Xu, and G. Wang. Spatio-temporal LSTM with trust gates for 3D human action recognition. In *ECCV*, 2016.
- [27] F. Lv and R. Nevatia. Recognition and segmentation of 3D human action using HMM and multiclass adaboost. In *ECCV*, 2006.
- [28] J. Masci, D. Boscaini, M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on Riemannian manifolds. In *ICCV Workshops*, 2015.
- [29] Y. Moreau and J. Vandewalle. A Lie algebraic approach to dynamical system prediction. In *Circuits and Systems, 1996. ISCAS'96., Connecting the World., 1996 IEEE International Symposium on*, volume 3, pages 182–185. IEEE, 1996.
- [30] M. Müller. *Information retrieval for music and motion*, volume 2. Springer, 2007.
- [31] M. Müller, T. Röder, M. Clausen, B. Eberhardt, B. Krüger, and A. Weber. Documentation: Mocap database HDM05. *Tech. Rep. CG-2007-2*, 2007.
- [32] S. Nie and Q. Ji. Capturing global and local dynamics for human action recognition. In *ICPR*, 2014.
- [33] D. W. Pearson. Changing network weights by Lie groups. In *Artificial Neural Nets and Genetic Algorithms*, pages 249–252. Springer, 1995.
- [34] D. W. Pearson. Hopfield networks and symmetry groups. *Neurocomputing*, 8(3):305–314, 1995.
- [35] L. L. Presti and M. La Cascia. 3D skeleton-based human action classification: A survey. *Pattern Recognition*, 53:130–147, 2016.
- [36] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.
- [37] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang. NTU RGB+D: A large scale dataset for 3D human activity analysis. In *CVPR*, 2016.

- [38] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- [39] P. Turaga and R. Chellappa. Locally time-invariant models of human activities using trajectories on the Grassmannian. In *CVPR*, 2009.
- [40] A. Veeraraghavan, A. K. Roy-Chowdhury, and R. Chellappa. Matching shape sequences in video with applications in human movement analysis. *IEEE T-PAMI*, 27(12):1896–1909, 2005.
- [41] R. Vemulapalli, F. Arrate, and R. Chellappa. Human action recognition by representing 3D skeletons as points in a Lie group. In *CVPR*, 2014.
- [42] R. Vemulapalli and R. Chellappa. Rolling rotations for recognizing human actions from 3D skeletal data. In *CVPR*, 2016.
- [43] C. Wang, Y. Wang, and A. L. Yuille. An approach to pose-based action recognition. In *CVPR*, 2013.
- [44] C. Wang, Y. Wang, and A. L. Yuille. Mining 3D key-pose-motifs for action recognition. In *CVPR*, 2016.
- [45] J. Wang, Z. Liu, Y. Wu, and J. Yuan. Mining actionlet ensemble for action recognition with depth cameras. In *CVPR*, 2012.
- [46] R. Zbikowski. Lie algebra of recurrent neural networks and identifiability. In *American Control Conference*, number 30, pages 2900–2901, 1993.