

Learning a Lie Algebra from Unlabeled Data Pairs

Chris Ick and Vincent Lostanlen

Music and Audio Research Lab, New York University, New York, NY, USA

I. INTRODUCTION

Disentangling factors of variability from natural data is a core principle of modern representation learning [1]. For example, in computer vision, small changes in illumination and viewpoint for a given object may incur large pixel-wise perturbations. Therefore, robust models for object recognition must operate in a feature space of reduced dimensionality in which these perturbations are systematically canceled. Deep convolutional networks (convnets) show a remarkable ability to learn disentangled representations [2]. Despite this accomplishment, the mathematical foundations of representation learning in convnets is still in infancy [3].

In recent years, the generalization of deep learning to Lie groups beyond rigid motion in \mathbb{R}^n has allowed us to build convnets over datasets with non-trivial symmetries [4], such as patterns over the surface of a sphere [5]. However, one limitation of this approach is the need to explicitly define the Lie group underlying the desired invariance property before training the convnet. Whereas rotations on the sphere have a well-known symmetry group (SO(3)), the same cannot be said of many real-world factors of variability. For example, in machine listening, pitch shifts induce a complex effect in the time–frequency domain by transposing the harmonic source while leaving the spectral envelope unchanged [6]. More generally, the disentanglement of pitch, intensity dynamics, and playing technique remains a challenging task in music information retrieval [7].

This article proposes a machine learning method to discover a nonlinear transformation of the space \mathbb{R}^n which maps a collection of n -dimensional vectors $(\mathbf{x}_i)_i$ onto a collection of target vectors $(\mathbf{y}_i)_i$. The key idea is to approximate every target \mathbf{y}_i by a matrix–vector product of the form $\tilde{\mathbf{y}}_i = \phi(t_i)\mathbf{x}_i$, where the matrix $\phi(t_i)$ belongs to a one-parameter subgroup of $\text{GL}_n(\mathbb{R})$, denoted by G in the following. Crucially, the value of the parameter $t_i \in \mathbb{R}$ may change between data pairs $(\mathbf{x}_i, \mathbf{y}_i)$ and does not need to be known in advance.

II. THEORY

By Stone’s theorem [8, chapter 10], there exists a matrix \mathbf{A} such that G can be described by matrix exponentiations of the form:

$$G = \left\{ \phi(t) = e^{t\mathbf{A}} = \sum_{k=0}^{+\infty} \frac{t^k}{k!} \mathbf{A}^k \mid t \in \mathbb{R} \right\}. \quad (1)$$

The matrix \mathbf{A} is called an infinitesimal generator of the Lie group G , and its linear span $\mathfrak{g} = \{t\mathbf{A} \mid t \in \mathbb{R}\}$ is called the Lie algebra associated to G . Furthermore, the application ϕ defines a group homomorphism from \mathbb{R} to G . Evaluating the first derivative of this application near the identity ($t = 0$) yields the matrix \mathbf{A} . Hence:

$$\forall \mathbf{x} \in \mathbb{R}^n, \mathbf{Ax} = \lim_{t \rightarrow 0} \frac{\phi(t)\mathbf{x} - \mathbf{x}}{t} = \left(\frac{d\phi}{dt} \Big|_{t=0} \right) \mathbf{x}. \quad (2)$$

The optimal parameter t_i corresponding to each pair $(\mathbf{x}_i, \mathbf{y}_i)$ can be assumed to be small ($|t_i| \ll 1$) but not infinitesimal. We circumvent this issue by linearizing the orbit of G near \mathbf{x}_i :

$$\mathbf{y}_i \approx \tilde{\mathbf{y}}_i = \phi(t_i)\mathbf{x}_i = e^{t_i\mathbf{A}}\mathbf{x}_i = \mathbf{x}_i + t_i\mathbf{Ax}_i + O(t_i^2). \quad (3)$$

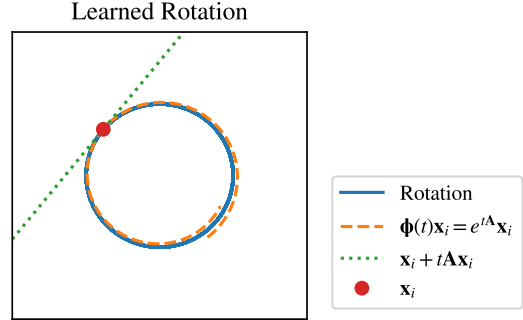


Fig. 1. Reconstructing rotation using a learned Lie algebra from data pairs.

It follows from the equation above that, as long as t_i remains small, the vectors $(\mathbf{y}_i - \mathbf{x}_i)$ and (\mathbf{Ax}_i) are approximately colinear. Therefore, we propose to discover \mathbf{A} by minimizing the empirical risk associated to the rectified cosine distance between these vectors:

$$\mathbf{A}^* = \arg \min_{\mathbf{A} \in \mathbb{R}^{n \times n}} \sum_i \left(1 - \frac{|(\mathbf{y}_i - \mathbf{x}_i)^\top \mathbf{Ax}_i|}{\|\mathbf{y}_i - \mathbf{x}_i\|_2 \|\mathbf{Ax}_i\|_2} \right). \quad (4)$$

III. PRACTICE

We define \mathbf{A} as the weight matrix of a feedforward neural network with a single layer and a linear activation function. By gradient descent on the cosine distance loss (see Equation 4), we train \mathbf{A} to predict the azimuth of each $(\mathbf{y}_i - \mathbf{x}_i)$ from the corresponding \mathbf{x}_i .

As a toy problem, we demonstrate our approach on $G = \text{SO}(2)$, i.e., the group of rotations of the plane. In this case, the Lie algebra \mathfrak{g} associated to G has an infinitesimal generator of the form:

$$\mathbf{A}^* = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \text{ hence } \phi : t \mapsto \begin{pmatrix} \cos t & \sin t \\ -\sin t & \cos t \end{pmatrix} \quad (5)$$

We train the neural network on a synthetic dataset of 1k pairs $(\mathbf{x}_i, \mathbf{y}_i)_i$, where the source points are drawn uniformly at random on the unit circle. Each angle t_i between the source point \mathbf{x}_i and the destination point \mathbf{y}_i is drawn uniformly at random in $]0, \frac{\pi}{30}[$.

Figure 1 illustrates our result with a specific source point \mathbf{x}_i (red dot). As expected, we observe that the vector \mathbf{Ax}_i (dashed green line) is tangent to the orbit of G (blue circle). Furthermore, the locus of $\phi(t)\mathbf{x}_i$ for $t \in [0, 2\pi]$ (dashed orange curve) is almost circular.

IV. CONCLUSION

We have presented a simple method for learning the Lie algebra associated to any one-parameter subgroup of $\text{GL}_n(\mathbb{R})$. This method has the advantage of not requiring any human intervention other than collecting data samples by pairs, i.e., before and after the action of the group. Future work will strive to extend the scope of our method to real-world factors of variability, such as pitch shifts in music or viewpoints in computer vision. As a long-term goal, it would be interesting to employ the proposed method as a pre-training stage in representation learning, thus yielding convolutional operators whose invariance properties are automatically adapted to the dataset at hand.

REFERENCES

- [1] Ian Goodfellow, Aaron Courville, and Yoshua Bengio. *Deep learning*. MIT Press, 2016.
- [2] Guillaume Lample, Neil Zeghidour, Nicolas Usunier, Antoine Bordes, Ludovic Denoyer, and MarcAurelio Ranzato. *Fader Networks: Manipulating Images by Sliding Attributes*. In: Advances in Neural Information Processing Systems (NeurIPS), 2017.
- [3] Stéphane Mallat. *Understanding deep convolutional networks*. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 374(2065), 2016.
- [4] Risi Kondor, Shubhendu Trivedi. On the Generalization of Equivariance and Convolution in Neural Networks to the Action of Compact Groups. *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- [5] Taco S. Cohen, Mario Geiger, Jonas Koehler, Max Welling Spherical CNNs. *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, 2018.
- [6] Vincent Lostanlen and Carmine Emanuele Cella. Deep convolutional networks on the pitch spiral for musical instrument recognition. *Proceedings of the International Society for Music Information Retrieval (ISMIR) Conference*, 2016.
- [7] Vincent Lostanlen, Joakim Andn, and Mathieu Lagrange. Extended playing techniques: The next milestone in musical instrument recognition. *Proceedings of the International Conference on Digital Libraries for Musicology (DLfM)*, 2018.
- [8] Brian C. Hall. *Lie Groups, Lie Algebras, and Representations*. Springer, 2000.