

Abstract

Developing Deep Learning architectures for tensor fields on manifolds is not a trivial task. Conventional CNNs fail to learn on this type of data due to the non-Euclidean nature of the domain and the dependence of the data from a specific choice of basis. We develop an $SO(3)$ -equivariant convolutional layer for vector fields on the sphere and present an application in weather prediction. We claim that building intrinsically equivariant networks not only guarantees compliance to the symmetry group of interest, but also requires fewer data points without the need to rely on data augmentation.

Vector fields on manifolds

When dealing with data in a digital setting, we must choose how to encode it in a computer. In the specific case of vector fields on a sphere we need to choose a coordinate system for the domain as well as a basis for the vectors. Any change in the atlas (coordinate system) or in the frame (choice of basis) does not affect the underlying data but changes the way a data point is encoded and processed.

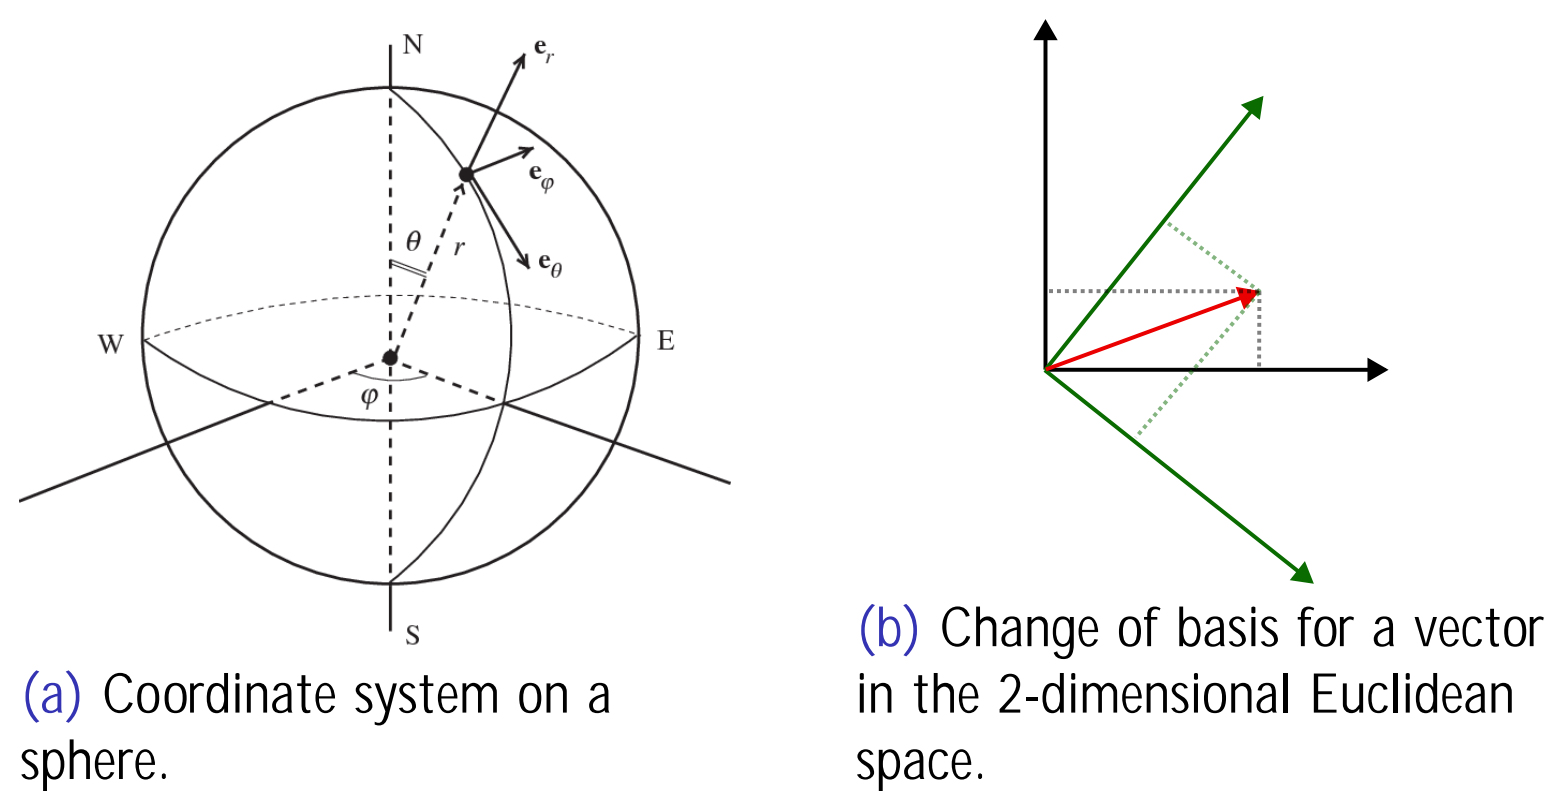
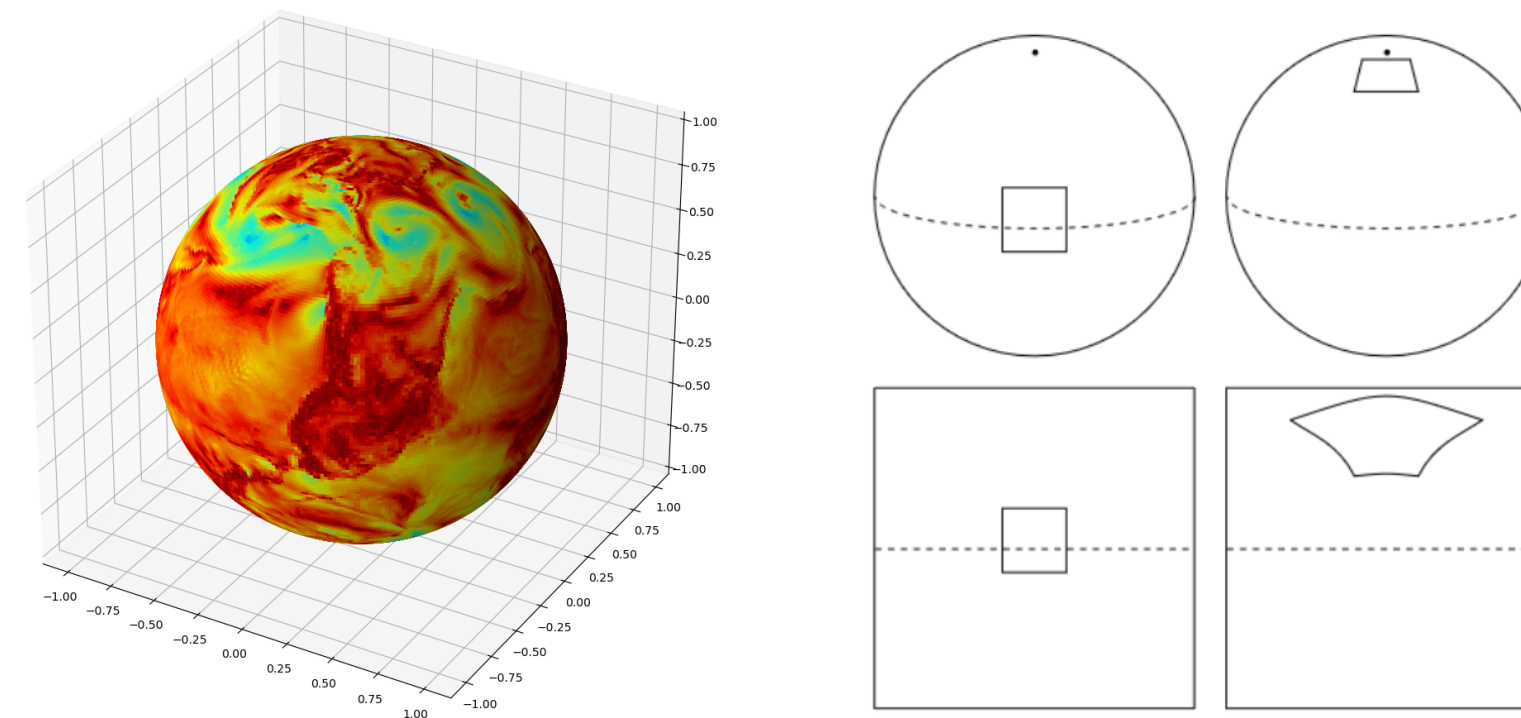


Figure: A choice of atlas or frame changes the way the data is encoded but does not affect the underlying data.

A naive solution to the problem of symmetries lies within **data augmentation**: we create multiple versions of a datapoint in order to consider all (or a subset of all) possible ways the same information can be encoded. However, data augmentation is computationally expensive and does not impose any constraint on the predictive model itself. Therefore, it does not guarantee that two data points that are symmetric w.r.t. each other will yield the same result.

Our framework

In our project, we aim to predict wind speed and direction at a global level by building a Geometric Deep Learning (GDL) Neural Network based on $SO(3)$ as a symmetry group, i.e. the group of 3D rotations. A naive application of either Feedforward or Convolutional networks to a planar projection of a spherical signal is destined to fail. In particular, in Convolutional Networks the space-varying distortions introduced by a projection from the sphere to a Euclidean space will make translational weight-sharing ineffective.



(a) Wind magnitude plotted on S^2 embedded in \mathbb{R}^3 . Wind data is given as a vector field in terms of $V(\cdot)$ and $U(\cdot)$ the north-bound and east-bound components of the wind at latitude and longitude. (b) When working on a flat projection of signal on the sphere a rectangular kernel around the equator is deformed when approaching the poles. Thus weight-sharing becomes ineffective. Courtesy of [2].

We are able to treat real vector fields on S^2 as complex functions on $SO(3)$ thanks to the following proposition:

Proposition
There is a unique correspondence between real vector fields on the sphere $S^2 \rightarrow \mathbb{R}^3$ and equivariant maps $SO(3) \rightarrow \mathbb{R}^3$ such that $(A\hat{B}) = B^{-1}(A)$ with

$$\hat{B} = \begin{pmatrix} B & 0 \\ 0 & 1 \end{pmatrix}; \quad B \in SO(3); \quad A \in SO(4)$$

$L^2(SO(3); \mathbb{C})$ admits an orthonormal basis given by Wigner D-matrices $D^l := (D_{m,n}^l)$. These correspond to the usual spherical harmonics when taking $l = 0$ and a restriction to S^2 . We can then write $L^2(SO(3); \mathbb{C}) = \bigoplus_{l=0}^{\infty} D_l$ where D_l is the space of Wigner D-matrices with fixed index l . Thanks to the following propositions we are able to work in the spectral domain of Wigner D-matrices and preserve the vector field structure by restricting ourselves to D_1 .

Proposition
Equivariant complex functions on $SO(3)$ are spanned by D_1 .

Corollary
Vector fields on the sphere S^2 are in unique correspondence with complex functions on $SO(3)$ spanned by D_1 .

By introducing an $SO(3)$ -equivariant layer in the spectral domain of Wigner D-matrices (illustration below) it becomes possible to construct fully equivariant Neural Networks for vector fields on the sphere. Equivariance is obtained by considering all possible $SO(3)$ complex functions and applying a smoothing operator S corresponding to taking a restriction on $l = 1$.

We suggest global weather prediction as an application for such a network, where the task is to predict future wind directions and intensities given current wind directions and intensities. For our experiments, we have developed a UNet-type architecture operating in the spectral domain, where the coarsening operation cuts the higher orders of frequencies (illustration below).

References

- [1]: Bronstein M., Bruna J., Cohen T., Velickovic P., *Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges*, arXiv:2104.13478, (2021).
- [2]: Cohen T., Geiger M., Koehler J., Welling M., *Spherical CNNs*, arXiv:1801.10130, (2018).
- [3]: Esteves, C., Allen-Blanchette, C., Makadia, A., Daniilidis, K., *Learning $SO(3)$ Equivariant Representations with Spherical CNNs*, Int. J. Comput. Vision, 128(3), 588-600, (2019).




Figure: The developed $SO(3)$ -equivariant layer. A convolution is performed between an equivariant function on $SO(3)$ and a vector field \hat{x}_m^l to obtain a non-equivariant function on $SO(3)$. Consequently, a non-linearity is applied in the spatial domain. Finally, the smoothing operator S is applied to the function on $SO(3)$ to guarantee equivariance.

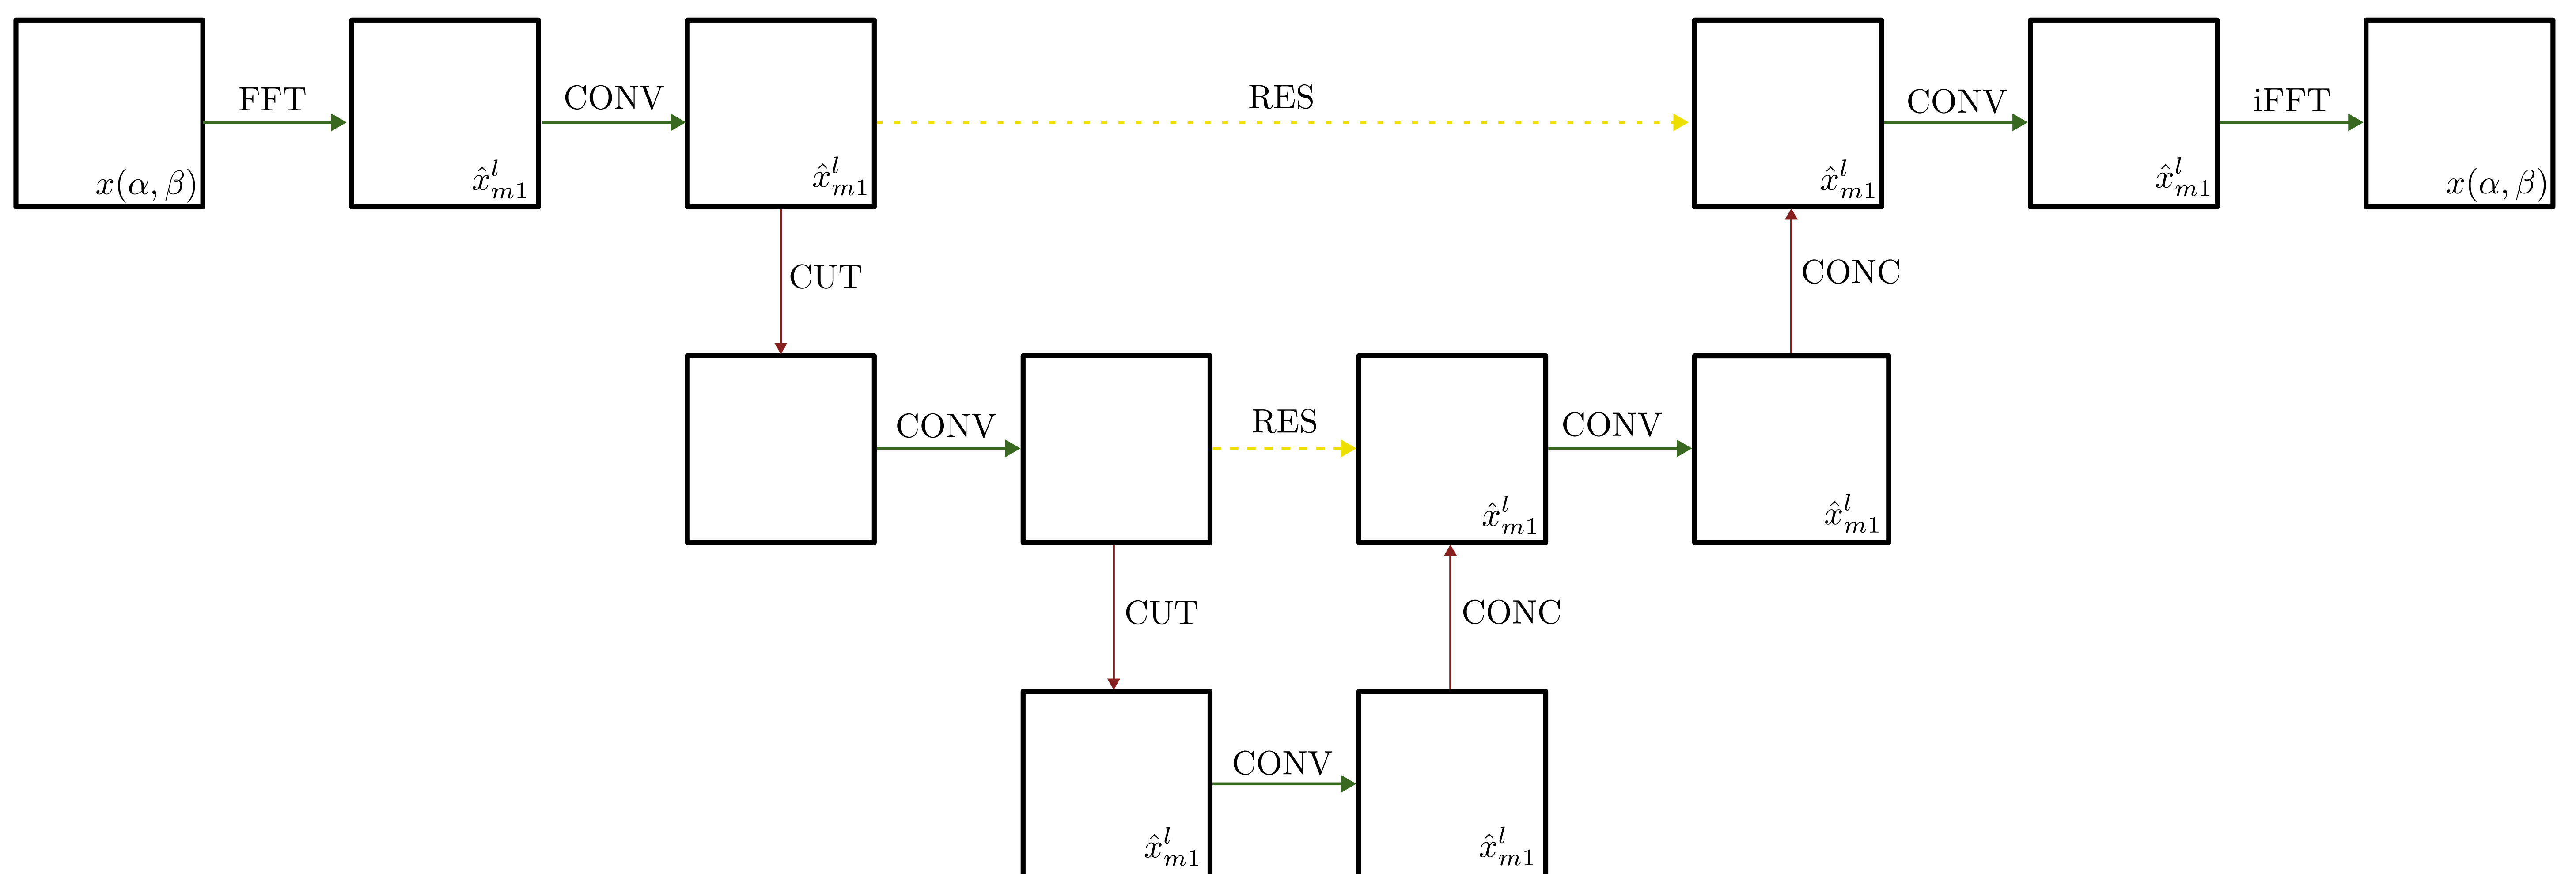


Figure: $SO(3)$ -UNet: A UNet architecture in the spectral domain of Wigner D-matrices. By truncating the higher orders of frequencies computational cost is reduced, allowing for more layers to be computed. By using residual connections high-frequency information is preserved.