SO(3)-Equivariant Neural Networks for Learning Vector Fields on Spheres

Francesco Ballerin

Joint work with **Erlend Grong** and **Nello Blaser**

G-StAI
@ Metochi Study Center

UNIVERSITY OF BERGEN

Fully Connected Neural Networks

$$F(x) = \hat{y}$$

$$a^{(l+1)} = \sigma(W^{(l)} \cdot a^{(l)} + b^{(l)})$$

$$a^{(0)} = x, \quad \hat{y} = a^{(L)}$$

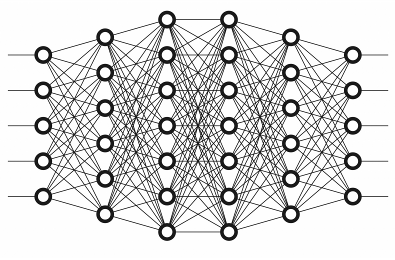

During training, the parameters ($W^{(l)}$ and $b^{(l)}$) are updated through an optimization process according to a specific loss function.

No information on the symmetry of the data is taken into consideration in the architecture.

Symmetries

Definition

A *group* is a set G equipped with a binary operation $\cdot$ that satisfies the following properties:

1. **Closure:** For all $a, b \in G$, the element $a \cdot b$ is also in G.
2. **Associativity:** For all $a, b, c \in G$, $(a \cdot b) \cdot c = a \cdot (b \cdot c)$.
3. **Identity Element:** There exists an element $e \in G$, called the identity element, such that for all $a \in G$, $a \cdot e = e \cdot a = a$.
4. **Inverse Element:** For each $a \in G$, there exists an element $a^{-1} \in G$, called the inverse of a, such that $a \cdot a^{-1} = a^{-1} \cdot a = e$.

We can identify symmetries of a certain space as group operation acting on the domain.

Symmetries in the data: 2D images

cat (0)

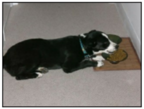

dog (1)

dog (1)

cat (0)

cat (0)

dog (1)

Example ($\mathbb{R}^2$ - unbounded or periodic 2D images)

The symmetry group of 2D images is $(\mathbb{R}^2, +)$, i.e. the group of vertical and horizontal translations.

Symmetries in the data: 2D images with rotations

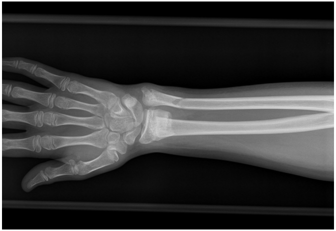

Example ($\mathbb{R}^2$ - unbounded or periodic 2D images)

The symmetry group of 2D images with rotation is $\mathbb{R}^2 \rtimes SO(2)$.

Symmetries in the data: graphs

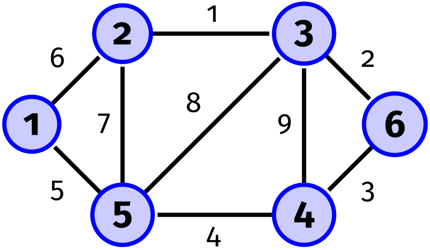

Example (Graphs)

The symmetry group for graphs with n nodes is the permutation group S_n. We can indeed list all nodes in a vector V and edges as an $n \times n$ adjacency matrix E. Then a permutation $P \in S_n$ acts on the graph (V, E) as $V' = PV$ and $E' = PEP^T$.

Symmetries in the data: S^n

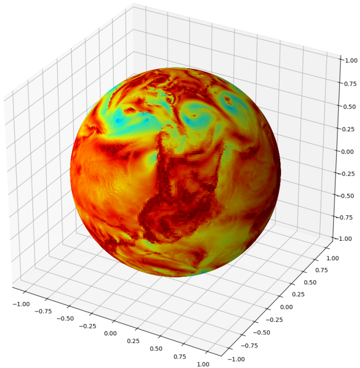

Example (Sphere)

The orientation-preserving symmetry group for S^n is $\mathrm{SO}(n+1)$.

Equivariance/Invariance

Let $f : A \to B$ be a function and $x \in A$. Let G be the symmetric group of interest that defines an action on both A and B. Then

Definition

f is said to be **left-equivariant** if

$$f(g \cdot x) = g \cdot f(x)$$

Definition

f is said to be **left-invariant** is

$$f(g \cdot x) = f(x)$$

Data augmentation: a naive solution to the problem of symmetries

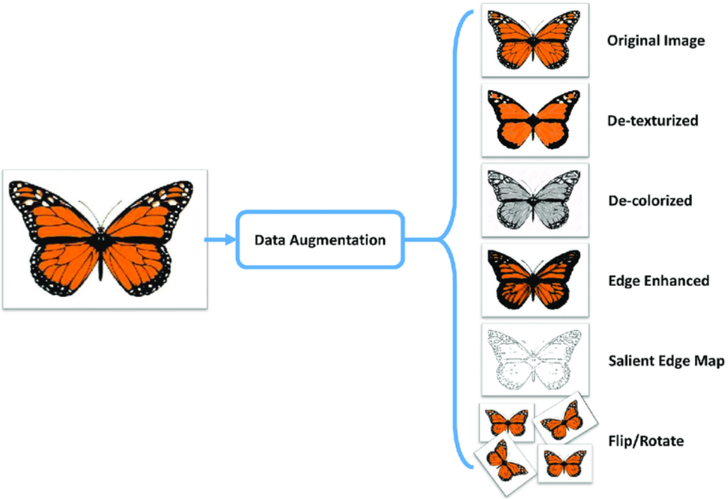

Geometric Deep Learning blueprint

The ingredients of a GDL neural network with symmetry group G are:

- G-equivariant layer
- activation function
- coarsening/pooling layer
- G-invariant layer

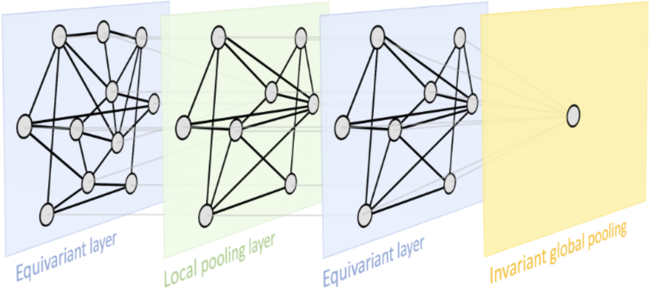

GDL blueprint in CNNs

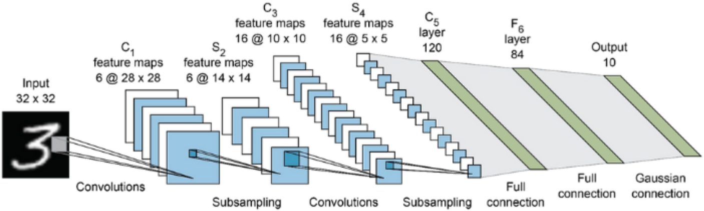

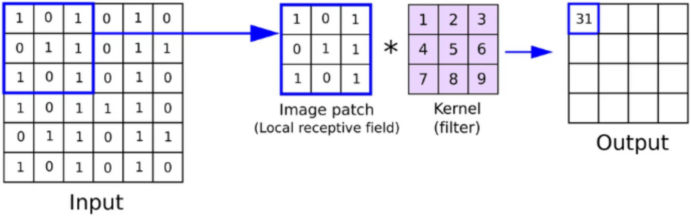

GDL blueprint in Graph Neural Networks

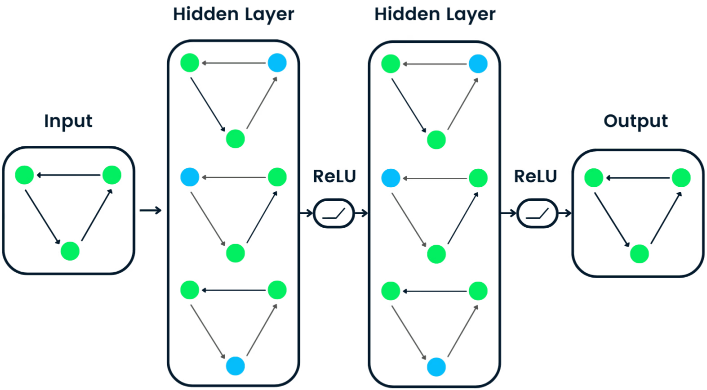

Problem description

We are given either scalar fields or vector fields on a 2-sphere, and we want to predict either a scalar field or a vector field.

In our case we are given **wind** data as a vector field in terms of $V(\alpha, \beta)$ and $U(\alpha, \beta)$ the **north** and **east components** of the wind at latitude α and longitude β.

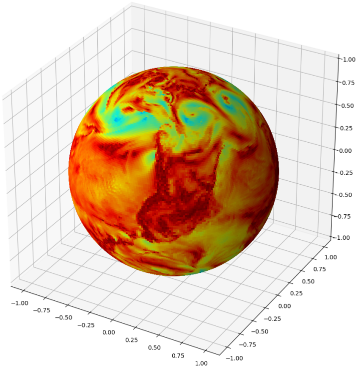

Wigner D-matrices

$L^2(SO(3), \mathbb{C})$ admits a basis given by *Wigner D-matrices* $D^l := (D^l_{m,n})$ with indices laying in the indexing set

$$I = \left\{ (l, m, n) \ : \ \begin{array}{c} l = 0, 1, 2, 3, \dots, \\ m, n \in [-l, l] \cap \mathbb{Z} \end{array} \right\}.$$

For $(l, m, n) \in I$, define function $D^l_{m,n} : \mathrm{SO}(3) \to \mathbb{C}$ by

$$D^l_{m,n}(Z(\alpha)Y(\beta)Z(\gamma)) = e^{-im\alpha} d^l_{m,n}(\beta) e^{-in\gamma},$$

$$d^l_{m,n}(\beta) = \sqrt{(l+m)!(l-m)!(l+n)!(l-n)!} \cdot$$
$$\cdot \sum_{s=s_0}^{s_1} \frac{(-1)^{m-n+s} \cos\left(\frac{\beta}{2}\right)^{2(l-s)+n-m} \sin\left(\frac{\beta}{2}\right)^{m-n+2s}}{(l+n-s)!s!(m-n+s)!(l-m-s)!}$$

with $s_0 = \max\{0, n-m\}$, $s_1 = \min\{l-m, l+n\}$.

Convolution and Wigner D-matrices

This is an orthonormal basis, so that for any $f \in L^2(SO(3), \mathbb{C})$, $A \in \mathrm{SO}(3)$

$$f(A) = \sum_{(l,m,n)\in I} \hat{f}^l_{m,n} D^l_{m,n}(A)$$

from which follows that

$$(f * \Psi)(A) = \int_{\mathrm{SO}(3)} f(B)\Psi(B^{-1}A)d\mu(B)$$

$$(f * \Psi)(A) = \sum_{l\in\mathbb{N}} \frac{1}{2l+1} \sum_{k,m,n=-l}^{l} \hat{f}^l_{m,k} \hat{\Psi}^l_{k,n} D^l_{m,n}(A)$$

$$\widehat{f * \Psi}^l_{n,m} = \frac{1}{2l+1} \sum_{k=-l}^{l} \hat{f}^l_{m,k} \hat{\Psi}^l_{k,n}$$

Wigner D-matrices and Equivariance

$$\mathcal{X}_n = span_{\mathbb{C}}\{D^l_{m,n}\}, \quad L^2(SO(3), \mathbb{C}) = \bigoplus_{n=-\infty}^{\infty} \mathcal{X}_n$$

Let $K \subseteq \mathrm{SO}(3)$ be the subgroup consisting of matrices of the form $Z(\alpha)$.

We can define an action $\hat{\rho}$ of K on $\mathrm{SO}(3)$ by

$$\hat{\rho}(Z(\alpha))A = AZ(-\alpha) = AZ(\alpha)^{-1}$$

and a representation ρ_n on $\mathbb{C}$ by

$$\rho_n(Z(\alpha))z = e^{in\alpha}z.$$

Definition

We say that a function $f \in L^2(SO(3), \mathbb{C})$ is n-equivariant if

$$f(\hat{\rho}(Z(\alpha))A) = f(AZ(-\alpha)) = e^{in\alpha}f(A) = \rho_n(Z(\alpha))f(A).$$

Theorem

1. *Functions in $L^2(S^2, \mathbb{C})$ are in unique correspondence with functions on $L^2(SO(3), \mathbb{C})$ spanned by $\mathcal{X}_0$.*
2. *Real vector fields on the sphere are in unique correspondence with functions on $L^2(SO(3), \mathbb{C})$ spanned by $\mathcal{X}_1$.*

Proof of [1].

For complex functions on the sphere one can notice that

$$Y_\ell^m(\theta, \phi) = \sqrt{\frac{2\ell + 1}{4\pi}}\, D_{m0}^{\ell *}(\phi, \theta, 0)$$

where Y_ℓ^m are spherical harmonics, and use the fact that these are an orthogonal basis for $L^2(S^2, \mathbb{C})$. ∎

Proof of [2].

Consider the principle bundle

$$\mathrm{SO}(n) \to \mathrm{SO}(n+1) \xrightarrow{\pi} S^n$$

with $\pi : \mathrm{SO}(n+1) \to S^n$ as the map $\pi(A) = Ae_{n+1}$.

Let $\xi : S^n \to \mathbb{R}^{n+1}$ be a vector field $A = (A_1, \dots, A_{n+1}) \in \mathrm{SO}(n+1)$. Then $\pi(A) = A_{n+1}$, while $\xi(A_{n+1})$ will be in the span of $A_1, \dots, A_n$. We can then introduce a 1-equivariant map $\bar{\xi}$ by

$$\bar{\xi}(A) = \begin{pmatrix} \langle \xi(A_{n+1}), A_1 \rangle \\ \vdots \\ \langle \xi(A_{n+1}), A_n \rangle \end{pmatrix}.$$

[Continues...]

Proof of [2] (continuation).

Conversely, if we have a 1-equivariant map $\bar{\xi} : \mathrm{SO}(n+1) \to \mathbb{R}^n$, then we can define a vector field by

$$\xi(\pi(A)) = A \begin{pmatrix} \bar{\xi}(A) \\ 0 \end{pmatrix}.$$

Wigner D-matrices $D^l_{m,n}$ are n-equivariant functions, and by using their orthogonality, it follows that $f \in L^2(SO(3), \mathbb{C})$ is n-equivariant if and only if $f \in \mathcal{X}_n$. ■

G-equivariant layer in spherical CNNs

For $f \in \mathcal{X}_p$ introduce weights Ψ^l for order $l \in \mathbb{N}$. An equivariant layer $f \mapsto f * \Psi$ can be defined as

$$\widehat{f * \Psi}^{l}_{m,p} = \frac{1}{2l+1} \hat{f}^{l}_{m,p} \hat{\Psi}^{l}$$

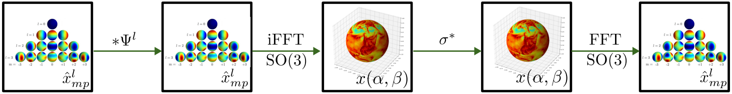

Advantages:

- Output is guaranteed to be in $\mathcal{X}_p$

Disadvantages:

- Low expressivity (weights are only order-rescalings)
- Nonlinearieties need to be equivariance-preserving

Learning functions on a sphere

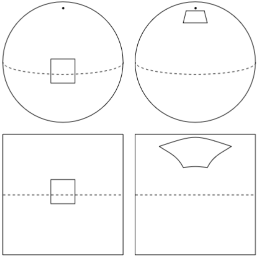

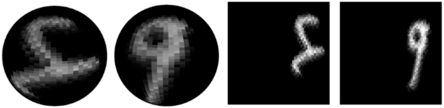

	NR / NR	R / R	NR / R
planar	0.98	0.23	0.11
spherical	0.96	0.95	0.94

The smoothing operator $\mathscr{S}_q$

Introduce the orthogonal projection $\mathscr{S}_q : \mathcal{X} \to \mathcal{X}_q$ defined by

$$\mathscr{S}_q(x) = \frac{1}{2\pi} \int_0^{2\pi} e^{iq\theta} r_{Z(-\theta)} x \, d\theta.$$

This can be rewritten in the spectral domain as

$$\widehat{\mathscr{S}_q x} = \hat{x}^l_{m,n} \delta_{n,q}.$$

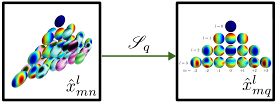

A more expressive G-equivariant layer

$$\widehat{f * \Psi}^{l}_{m,n} = \frac{1}{2l+1} \sum_{s=-l}^{l} \hat{f}^{l}_{m,s} \hat{\Psi}^{l}_{s,n}$$

For $f \in \mathcal{X}_p$ we can restrict Ψ to coefficients $\hat{\Psi}^l_n$

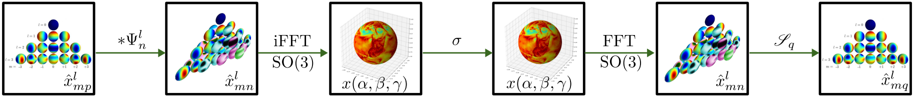

Advantages:

- More expressive
- Can use any activation function

Disadvantages:

- More weights per layer
- One extra dimension in FT

Coarsening layer

We implement coarsening layer $\mathcal{C}$ by limiting the bandwidth of the function at a certain order L.

$$\mathcal{C}(f) = \sum_{(l,m,n)\in I_L} \hat{f}^l_{m,n} D^l_{m,n}$$

We need to keep in mind that:

- the convolution layer and nonlinearity do not mix frequencies at different orders;
- in practice it cuts high-frequency information.

Therefore coarsening needs to be paired with an activation function in the spatial domain.

Our proposed architecture

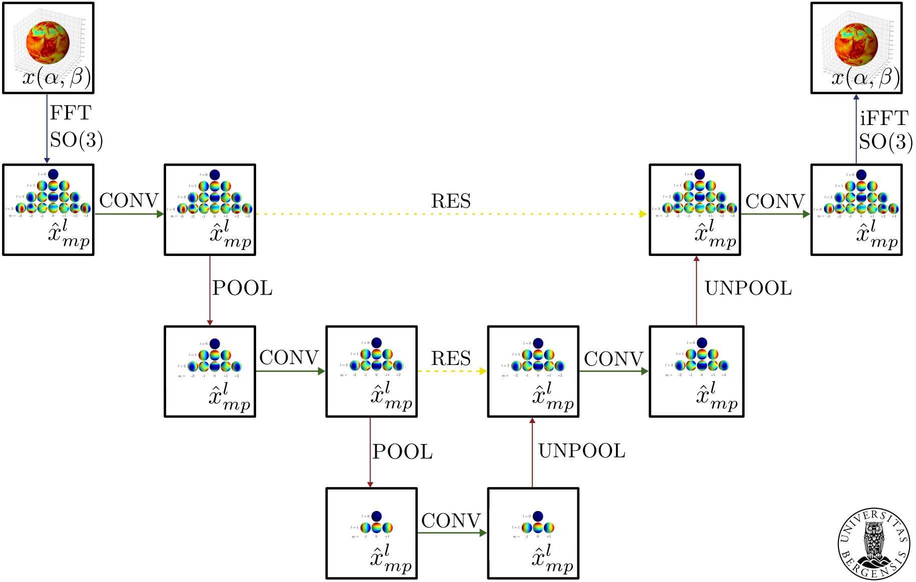

ERA5 Dataset

As a proof of concept we use use the **ERA5 meteorological dataset**.

It contains **hourly global measurements** of different quantities **from 1940 to today**.

In our experiments, we use **wind data** at 10m of elevation and **temperature data** at 2m of elevation.

For training and model selection we use a coarser dataset of 52 **weekly datapoints** per year for both wind and temperature. Years from 2000 to 2009 (included) have been used for training, while the years 2020 and 2021 have been used for validation and model selection. Years 2022 and 2023 have been used for testing.

Equivariance

Model	Ground truth	Pred. $\beta = 0$	Pred. $\beta = \frac{\pi}{4}$	Error
CNN	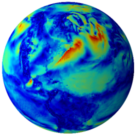	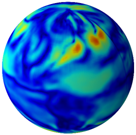	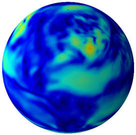	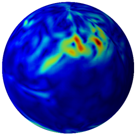
Ours	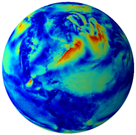	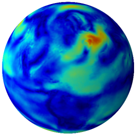	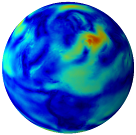	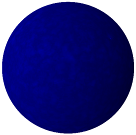

0 km/h  28 km/h

Wind to wind prediction

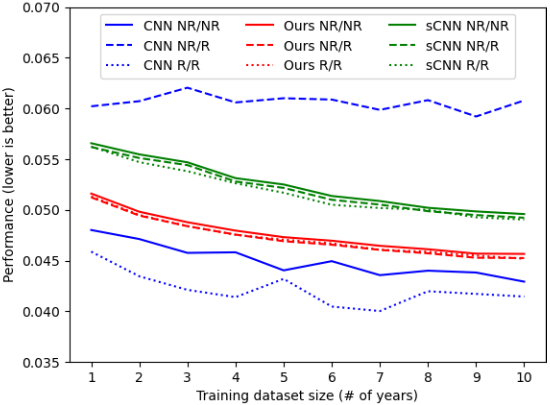

Figure: Wind to Wind t+24h prediction.

Temperature to wind estimation

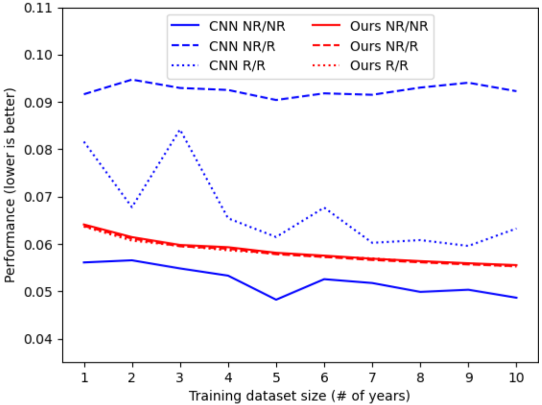

Figure: Temperature to Wind estimation.

Autoencoder compression

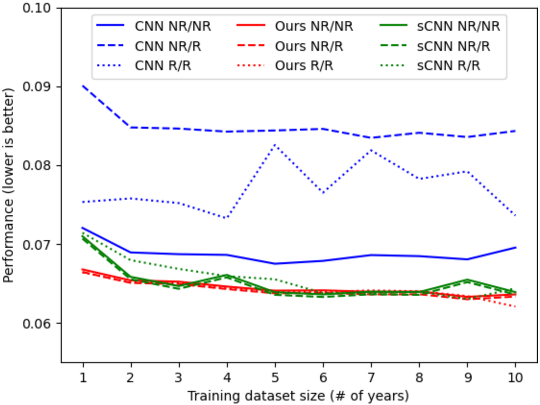

Figure: Autoencoder compression on wind data.

References

Bronstein M., Bruna J., Cohen T., Veličković P.
Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges, arXiv:2104.13478 (2021).

Cohen T., Geiger M., Koehler J., Welling M.
"Spherical CNNs", ICLR 2018.

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).

Esteves C., Slotine J., Makadia A.
Scaling Spherical CNNs, ICML2023.

Ballerin F., Blaser N., Grong E.
SO(3)-Equivariant Neural Networks for Learning Vector Fields on Spheres, arXiv:2503.09456 (2025)