

# GEOMETRIC DEEP LEARNING (L65)

Pietro Liò *University of Cambridge*

Petar Veličković *Google DeepMind / University of Cambridge*

Lent Term 2024

*CST Part III / MPhil ACS / MPhil MLMI*


# 1. INTRODUCTION TO GROUPS AND REPRESENTATIONS

*The fundamentals of capturing the regularity in nature*

Petar Veličković

# *Learning in high dimensions*

In general, learning functions in high dimensions is **intractable**  
Number of samples required grows *exponentially* with dimension

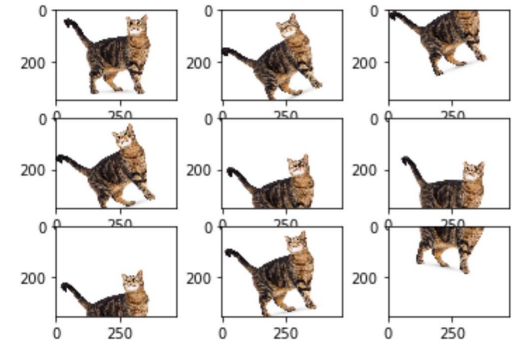


# Geometry to the rescue!

We can inject *assumptions* about **geometry** through *inductive biases*  
Restrict the functions to ones that *respect* the geometry.  
This can make the high-dimensional problem more tractable!

Some popular examples:

- **Image** data should be processed independently of **shifts**



# *Geometry to the rescue!*


We can inject *assumptions* about **geometry** through *inductive biases*

Restrict the functions to ones that *respect* the geometry.

This can make the high-dimensional problem more tractable!

Some popular examples:

- **Image** data should be processed independently of **shifts**
- **Spherical** data should be processed independently of **rotations**




# Geometry to the rescue!

We can inject *assumptions* about **geometry** through *inductive biases*  
Restrict the functions to ones that *respect* the geometry.  
This can make the high-dimensional problem more tractable!

Some popular examples:

- **Image** data should be processed independently of **shifts**
- **Spherical** data should be processed independently of **rotations**
- **Graph** data should be processed independently of **isomorphism**



# *Geometry to the rescue!*

We can inject *assumptions* about **geometry** through *inductive biases*

Restrict the functions to ones that *respect* the geometry.

This can make the high-dimensional problem more tractable!

Some popular examples:

- **Image** data should be processed independently of **shifts**
- **Spherical** data should be processed independently of **rotations**
- **Graph** data should be processed independently of **isomorphism**

We will now attempt to **formalise** this!

## *A roadmap for our formalisation*

To be able to talk about geometry of *data*, we need to formalise *where* the data lives (*domain*) and how to *featurise* it (*signal*)

Once we understand data domains, we can then formalise *symmetries* of those domains (*groups*)

Equipped with groups, we need to formalise *how* they *transform* the data domains (*group actions*)

Deep learning concerns itself with *linear algebra*; we need to be able to talk about group actions as *matrix operations* (*representations*)

Using representations, we can formalise what it means for a deep learning model to *respect symmetries* (*invariance & equivariance*)



# The space of signals on a geometric domain

A *signal* on  $\Omega$  is a function  $x : \Omega \rightarrow \mathcal{C}$ , where:

- $\Omega$  is the domain (e.g. set of pixels/nodes/...)
- $\mathcal{C}$  is a vector space, whose dimensions are called *channels*

The space of  $\mathcal{C}$ -valued signals on  $\Omega$  is defined as


$$\mathcal{X}(\Omega, \mathcal{C}) = \{x : \Omega \rightarrow \mathcal{C}\}$$

We will often omit  $\mathcal{C}$ , and just write  $\mathcal{X}(\Omega)$

Signals on *discrete, finite*  $\Omega$  expressible as *matrices*

$$\mathbf{X} \in \mathbb{R}^{|\Omega| \times \dim \mathcal{C}}$$

where the *ith row* gives *features* of the *ith node*



*Example:  $n \times n$  RGB image*

# Vector space structure of signals

We can add signals and multiply by scalars:

$$(\alpha x + \beta y)(u) = \alpha x(u) + \beta y(u), \quad \text{where } \alpha, \beta \in \mathbb{R} \text{ and } u \in \Omega$$



$\Rightarrow$  The space of signals is a *vector space!* (possibly *infinite dimensional*)

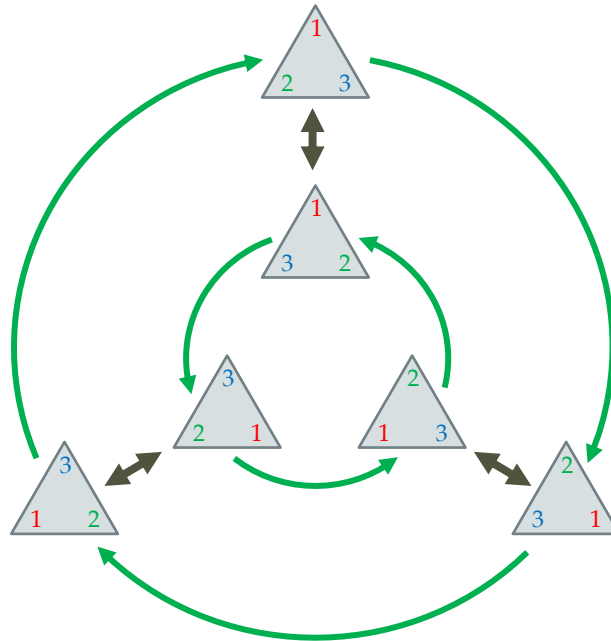
Can also define an *inner product* on signals, given inner product  $\langle \cdot, \cdot \rangle_{\mathcal{C}}$  on  $\mathcal{C}$  and a measure  $\mu$  on  $\Omega$  ( $\Rightarrow$  The space of signals is a *Hilbert space!*)

$$\langle x, y \rangle = \int_{\Omega} \langle x(u), y(u) \rangle_{\mathcal{C}} d\mu(u)$$

**Exercise:** Verify that the above satisfies the inner product axioms

# Symmetries

A **symmetry** of an object is a transformation of that object that leaves it unchanged



The symmetries of a triangle,  
as generated by 120-degree **rotations**  $R$  and **flips**  $F$ .

# *Symmetry group*

A **symmetry** of an object is a transformation of that object that leaves it **unchanged**

Observe that this immediately defines some properties:


- The **identity** transformation is always a symmetry
- Given two symmetry transformations, their **composition** (doing one after the other) is also a symmetry
- Given any symmetry, it must be **invertible**
- Moreover, its **inverse** is also a symmetry

Collecting all these *axioms* together, we recover a standard mathematical object: the **group**

# Abstract groups

A *group* is a set  $\mathfrak{G}$  with a binary operation denoted  $gh$  satisfying the following properties:

- *Associativity*:  $(gh)k = g(hk)$  for all  $g, h, k \in \mathfrak{G}$
- *Identity*: there exists a unique  $e \in \mathfrak{G}$  satisfying
$$ge = eg = g$$
- *Inverse*: for each  $g \in \mathfrak{G}$  there is a unique inverse  $g^{-1} \in \mathfrak{G}$ , such that  $gg^{-1} = g^{-1}g = e$
- *Closure*: for every  $g, h \in \mathfrak{G}$ , we have  $gh \in \mathfrak{G}$



*Rotational symmetries of the cube  
(group  $O_h$ )*

# *Symmetry groups, abstract groups & group actions*

**Symmetry group:** a group of transformations  $g : \Omega \rightarrow \Omega$

The group operation is *composition*

# *Symmetry groups, abstract groups & group actions*

**Symmetry group:** a group of transformations  $g : \Omega \rightarrow \Omega$

The group operation is *composition*

**Abstract group:** a set of elements together with a composition rule, satisfying the group axioms

*(an abstract group does not directly tell us how to transform data!)*

# Symmetry groups, abstract groups & group actions

**Symmetry group:** a group of transformations  $g : \Omega \rightarrow \Omega$

The group operation is *composition*

**Abstract group:** a set of elements together with a composition rule, satisfying the group axioms

**Group action:** a map  $\mathfrak{G} \times \Omega \rightarrow \Omega$  (denoted  $(g, u) \mapsto gu$ ) such that


$$g(hu) = (gh)u$$

$$eu = u$$

e.g.: Euclidean 2D motions  $\mathfrak{G} = \mathbb{R}^3$  (angle + translation) acting on  $\Omega = \mathbb{R}^2$ :

$$(\theta, t_x, t_y)(x, y) \mapsto (x \cos \theta + y \sin \theta + t_x, x \sin \theta + y \cos \theta + t_y)$$

**Exercise:** Verify this satisfies the group action axioms






## *Symmetries of $\Omega$ acting on signals $\mathcal{X}(\Omega)$*

Given an action of  $\mathfrak{G}$  on  $\Omega$ , we automatically obtain an action of  $\mathfrak{G}$  on the space of signals  $\mathcal{X}(\Omega)$  :

$$(\mathfrak{g} x)(u) = x(\mathfrak{g}^{-1}u)$$

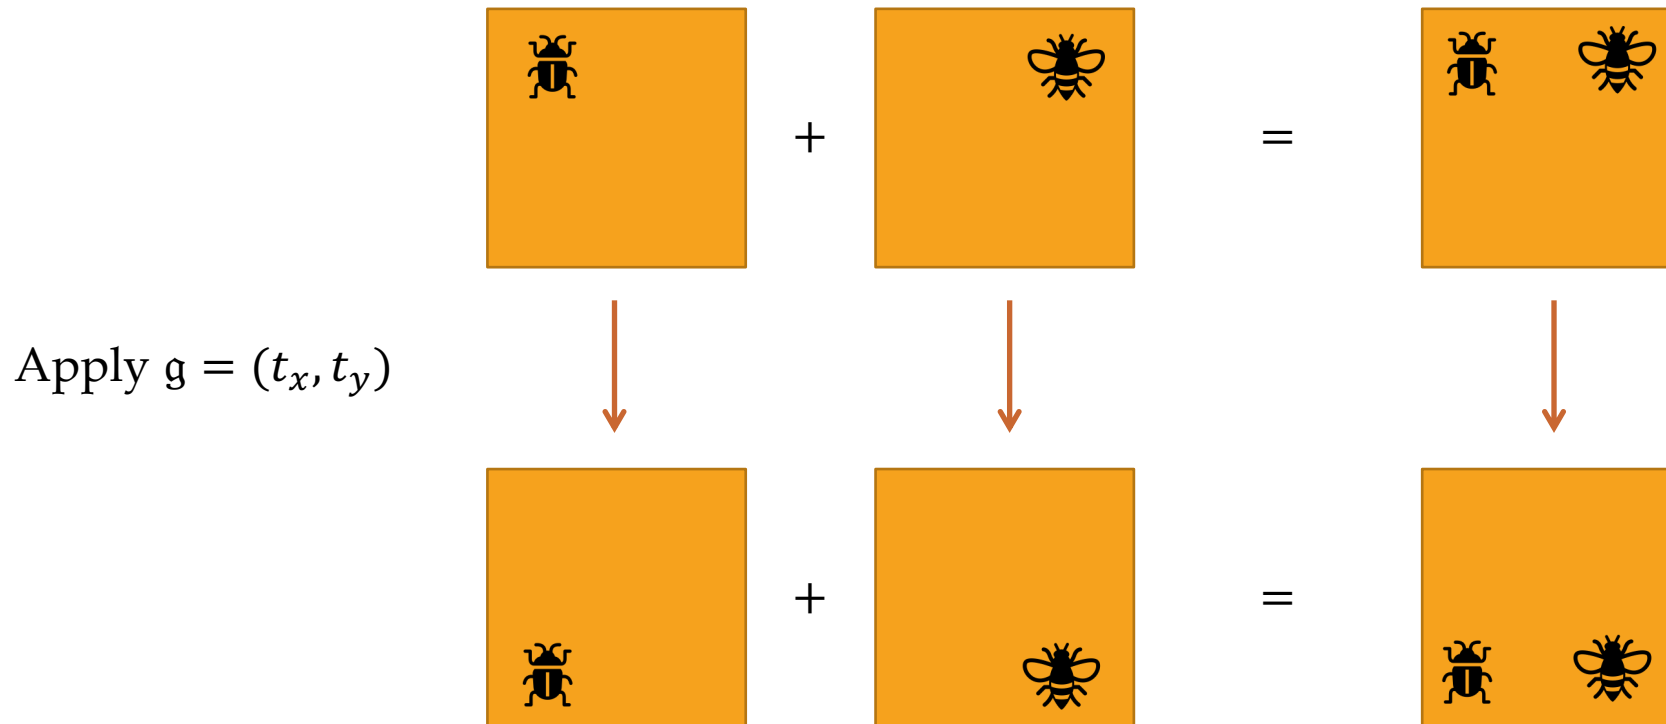


# Linearity of the group action

If the signals support a vector space, such a group action on signals

$$(g x)(u) = x(g^{-1}u)$$

is *linear*!



## *Linearity of the group action*

If the signals support a vector space, such a group action on signals

$$(g x)(u) = x(g^{-1}u)$$

is *linear*!

This is *excellent* news for us, as deep learning is basically *linear algebra*  
And we will be able to describe group actions as *matrix multiplication*!

For the time being, we will assume our domain is *discrete* and *finite*  
That is, that we can represent our domain using a matrix  $\mathbf{X} \in \mathbb{R}^{|\Omega| \times k}$

# Group representations

A real representation of  $\mathfrak{G}$  on a finite vector space  $\mathcal{X}$  is a map  $\rho_{\mathcal{X}} : \mathfrak{G} \rightarrow \mathbb{R}^{n \times n}$ , assigning to each element  $g \in \mathfrak{G}$  an invertible matrix  $\rho_{\mathcal{X}}(g)$ , and satisfying

$$\rho_{\mathcal{X}}(gh) = \rho_{\mathcal{X}}(g)\rho_{\mathcal{X}}(h), \quad \forall g, h \in \mathfrak{G}$$

# Group representations

A real representation of  $\mathfrak{G}$  on a finite vector space  $\mathcal{X}$  is a map  $\rho_{\mathcal{X}} : \mathfrak{G} \rightarrow \mathbb{R}^{n \times n}$ , assigning to each element  $g \in \mathfrak{G}$  an invertible matrix  $\rho_{\mathcal{X}}(g)$ , and satisfying

$$\rho_{\mathcal{X}}(gh) = \rho_{\mathcal{X}}(g)\rho_{\mathcal{X}}(h), \quad \forall g, h \in \mathfrak{G}$$

The dimensionality of the matrix produced by  $\rho_{\mathcal{X}}$  may depend on many factors, such as the size of the corresponding  $\Omega$

$\rho_{\mathcal{X}}$  does *not* need to assign *different* matrices to *different* elements of  $\mathfrak{G}$  (if it does, then it is a *faithful representation*)

Representations can be easily generalised to *infinite* spaces---the map  $\rho_{\mathcal{X}} : \mathfrak{G} \rightarrow (\mathcal{X}(\Omega) \rightarrow \mathcal{X}(\Omega))$  would output an *invertible linear function* (strictly speaking,  $\rho_{\mathcal{X}} : \mathfrak{G} \rightarrow \text{GL}(\mathcal{X})$ , where  $\text{GL}(\mathcal{X})$  is the *general linear group* over  $\mathcal{X}$ )

# Group representations


A real representation of  $\mathfrak{G}$  on a finite vector space  $\mathcal{X}$  is a map  $\rho_{\mathcal{X}} : \mathfrak{G} \rightarrow \mathbb{R}^{n \times n}$ , assigning to each element  $g \in \mathfrak{G}$  an invertible matrix  $\rho_{\mathcal{X}}(g)$ , and satisfying

$$\rho_{\mathcal{X}}(gh) = \rho_{\mathcal{X}}(g)\rho_{\mathcal{X}}(h), \quad \forall g, h \in \mathfrak{G}$$

Example:

- The group of 1D (circular) shifts,  $\mathfrak{G} = (\mathbb{Z}, +)$
- The domain  $\Omega = \mathbb{Z}_5 = \{0, 1, 2, 3, 4\}$  (e.g. short audio signal)
- The action of  $g = n$  on  $u \in \Omega$ :  $(n, u) \mapsto n + u \pmod{5}$
- The representation on  $\mathcal{X}(\Omega)$ :

$$\rho_{\mathcal{X}}(n) = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix}^n$$



**Exercise:** Derive the representation for the group of 90-degree rotations on 3x3 grids

## *Group invariance*

We can now *formally* describe how to *exploit* the symmetries in  $\mathfrak{G}$ !

# Group invariance

We can now *formally* describe how to *exploit* the symmetries in  $\mathfrak{G}$ !

A function  $f : \mathcal{X}(\Omega) \rightarrow \mathcal{Y}$  is  $\mathfrak{G}$ -invariant if  $f(\rho_{\mathcal{X}}(g)x) = f(x)$  for all  $g \in \mathfrak{G}$ , i.e., its output is unaffected by the group action on the input.



# Group invariance

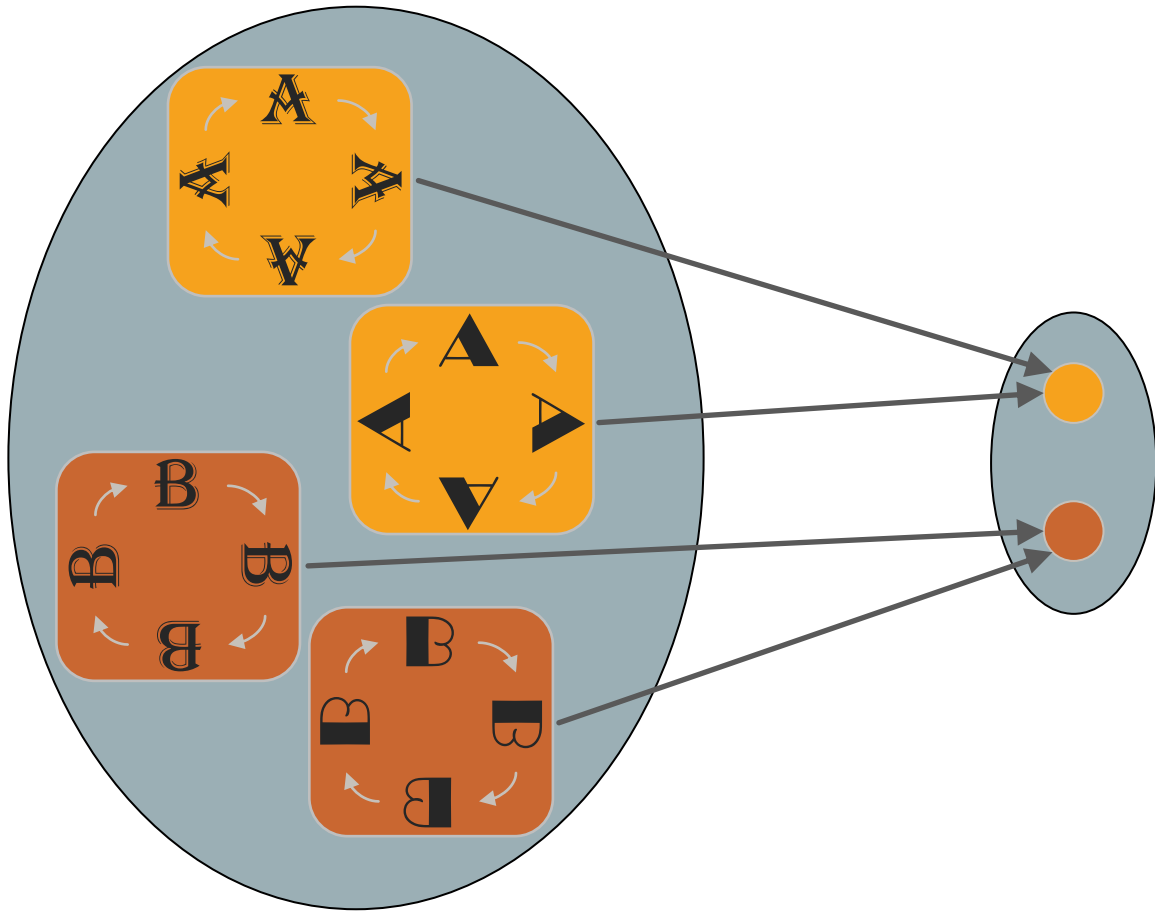
We can now *formally* describe how to *exploit* the symmetries in  $\mathfrak{G}$ !

A function  $f : \mathcal{X}(\Omega) \rightarrow \mathcal{Y}$  is  $\mathfrak{G}$ -invariant if  $f(\rho_x(g)x) = f(x)$  for all  $g \in \mathfrak{G}$ , i.e., its output is unaffected by the group action on the input.

e.g. **image classification**: output class won't depend on image **shifts**

$$f\left(\begin{array}{c} \text{[orange square with beetle icon]} \\ \rho_x(g)x \end{array}\right) = f\left(\begin{array}{c} \text{[orange square with beetle icon]} \\ x \end{array}\right) = \text{[beetle icon]}$$

# Orbits and equivalence relations



$$O_x = \{gx \mid x \in X, g \in G\}$$


## $G$ -equivalence

$$x \sim_G y \iff \exists g \in G : gx = y$$

Satisfies the axioms of an equivalence relation:

1. Reflexivity:  $x \sim_G x$ 
  - (Because  $G$  contains the identity)
2. Transitivity:
$$x \sim_G y \wedge y \sim_G z \iff x \sim_G z$$
  - (Because  $G$  is closed under composition)
3. Symmetry:  $x \sim_G y \iff y \sim_G x$ 
  - (Because  $G$  is closed under inverses)


*$\mathcal{G}$ -invariant representations*



## *The problem with invariance*

Invariance is suitable when we need a *single* output over the *entire* domain. What if we need an output in *each* domain element?

Also, even if a single output is OK, making the intermediate representations invariant may lose *critical* information:



The *relative pose* of object parts contains critical information  
(Hinton *et al.*, ICANN'11)

# Group equivariance

We proceed to define a more *fine-grained* notion of regularity:

A function  $f : \mathcal{X}(\Omega) \rightarrow \mathcal{Z}(\Omega)$  is  $\mathfrak{G}$ -equivariant if, for all  $g \in \mathfrak{G}$ ,  
 $f(\rho_{\mathcal{X}}(g)x) = \rho_{\mathcal{Z}}(g)f(x)$ , i.e., applying a group action on the input  
affects the output in the same way.

# Group equivariance

We proceed to define a more *fine-grained* notion of regularity:


A function  $f : \mathcal{X}(\Omega) \rightarrow \mathcal{Z}(\Omega)$  is  $\mathfrak{G}$ -equivariant if, for all  $g \in \mathfrak{G}$ ,  $f(\rho_{\mathcal{X}}(g)x) = \rho_{\mathcal{Z}}(g)f(x)$ , i.e., applying a group action on the input affects the output in the same way.

e.g. **image segmentation**: segmentation mask must **follow** any shifts in the input



Note that invariance is a *special case* of equivariance (for which  $\rho_{\mathcal{Z}}$ ?)

# *Equivariance as symmetry-consistent generalisation*




This **cannot** happen!  
Equivariant net must generalize consistently across the whole orbit.

**Equivariance:**

$$f(\rho_x(\mathfrak{g}) x) = \rho_z(\mathfrak{g}) f(x)$$

# Equivariance as symmetry-consistent generalisation





# *The building blocks of Geometric Deep Learning*

Let  $\Omega$  and  $\Omega'$  be domains,  $\mathfrak{G}$  a symmetry group over  $\Omega$ .  
Write  $\Omega' \subseteq \Omega$  if  $\Omega'$  can be considered a compact version of  $\Omega$ .

We define the following building blocks:

*Linear  $\mathfrak{G}$ -equivariant layer*  $B : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C}')$ ,  
satisfying  $B(\mathfrak{g}.x) = \mathfrak{g}.B(x)$  for all  $\mathfrak{g} \in \mathfrak{G}$  and  $x \in \mathcal{X}(\Omega, \mathcal{C})$ .

*Nonlinearity*  $\sigma : \mathcal{C} \rightarrow \mathcal{C}'$  applied element-wise as  $(\sigma(x))(u) = \sigma(x(u))$ .

*Local pooling (coarsening)*  $P : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C})$ , such that  $\Omega' \subseteq \Omega$ .

*$\mathfrak{G}$ -invariant layer (global pooling)*  $A : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{Y}$ ,  
satisfying  $A(\mathfrak{g}.x) = A(x)$  for all  $\mathfrak{g} \in \mathfrak{G}$  and  $x \in \mathcal{X}(\Omega, \mathcal{C})$ .

# The building blocks of Geometric Deep Learning

Let  $\Omega$  and  $\Omega'$  be domains,  $\mathfrak{G}$  a symmetry group over  $\Omega$ .  
Write  $\Omega' \subseteq \Omega$  if  $\Omega'$  can be considered a compact version of  $\Omega$ .

We define the following building blocks:

*Linear  $\mathfrak{G}$ -equivariant layer  $B : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C}')$ ,  
satisfying  $B(g \cdot x) = g \cdot B(x)$  for all  $g \in \mathfrak{G}$  and  $x \in \mathcal{X}(\Omega, \mathcal{C})$ .*

**Linear equivariant layer**

*Nonlinearity  $\sigma : \mathcal{C} \rightarrow \mathcal{C}'$  applied element-wise as  $(\sigma(x))(u) = \sigma(x(u))$ .*

*Local pooling (coarsening)  $P : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C})$ , such that  $\Omega' \subseteq \Omega$ .*

*$\mathfrak{G}$ -invariant layer (global pooling)  $A : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{Y}$ ,  
satisfying  $A(g \cdot x) = A(x)$  for all  $g \in \mathfrak{G}$  and  $x \in \mathcal{X}(\Omega, \mathcal{C})$ .*

**Invariant “tail” layer**  
(if necessary)

# The building blocks of Geometric Deep Learning

Let  $\Omega$  and  $\Omega'$  be domains,  $\mathfrak{G}$  a symmetry group over  $\Omega$ .  
Write  $\Omega' \subseteq \Omega$  if  $\Omega'$  can be considered a compact version of  $\Omega$ .

We define the following building blocks:

*Linear  $\mathfrak{G}$ -equivariant layer*  $B : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C}')$ ,  
satisfying  $B(\mathfrak{g}.x) = \mathfrak{g}.B(x)$  for all  $\mathfrak{g} \in \mathfrak{G}$  and  $x \in \mathcal{X}(\Omega, \mathcal{C})$ .

**Activation function**

*Nonlinearity*  $\sigma : \mathcal{C} \rightarrow \mathcal{C}'$  applied element-wise as  $(\sigma(x))(u) = \sigma(x(u))$ .

*Local pooling (coarsening)*  $P : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C})$ , such that  $\Omega' \subseteq \Omega$ .

*$\mathfrak{G}$ -invariant layer (global pooling)*  $A : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{Y}$ ,  
satisfying  $A(\mathfrak{g}.x) = A(x)$  for all  $\mathfrak{g} \in \mathfrak{G}$  and  $x \in \mathcal{X}(\Omega, \mathcal{C})$ .

# The building blocks of Geometric Deep Learning

Let  $\Omega$  and  $\Omega'$  be domains,  $\mathfrak{G}$  a symmetry group over  $\Omega$ .  
Write  $\Omega' \subseteq \Omega$  if  $\Omega'$  can be considered a compact version of  $\Omega$ .

We define the following building blocks:

*Linear  $\mathfrak{G}$ -equivariant layer*  $B : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C}')$ ,  
satisfying  $B(\mathfrak{g}.x) = \mathfrak{g}.B(x)$  for all  $\mathfrak{g} \in \mathfrak{G}$  and  $x \in \mathcal{X}(\Omega, \mathcal{C})$ .

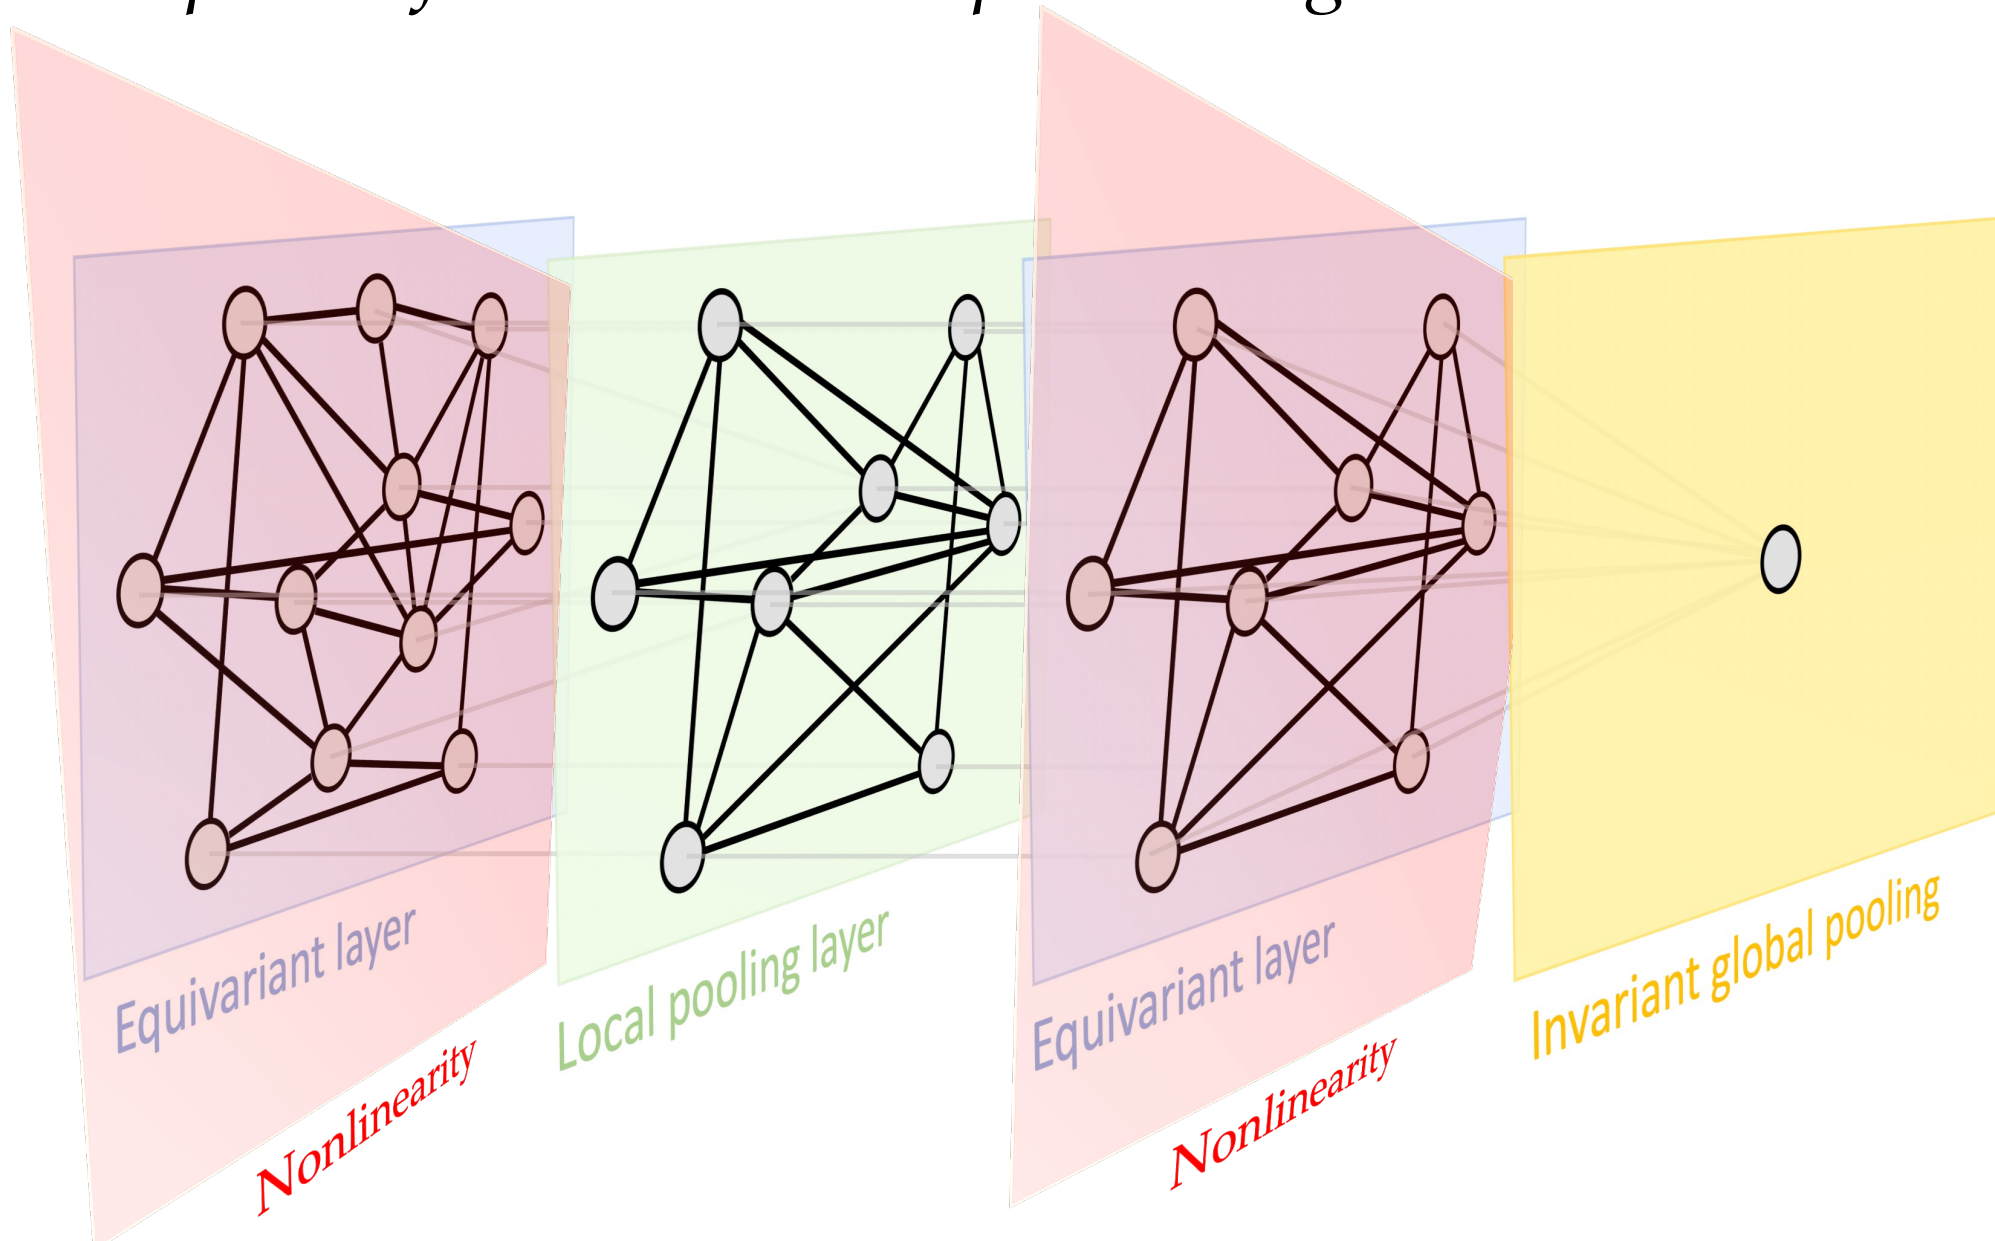
*Nonlinearity*  $\sigma : \mathcal{C} \rightarrow \mathcal{C}'$  applied element-wise as  $(\sigma(x))(u) = \sigma(x(u))$ .

*Local pooling (coarsening)*  $P : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{X}(\Omega', \mathcal{C})$ , such that  $\Omega' \subseteq \Omega$ .

Coarsening layer

*$\mathfrak{G}$ -invariant layer (global pooling)*  $A : \mathcal{X}(\Omega, \mathcal{C}) \rightarrow \mathcal{Y}$ ,  
satisfying  $A(\mathfrak{g}.x) = A(x)$  for all  $\mathfrak{g} \in \mathfrak{G}$  and  $x \in \mathcal{X}(\Omega, \mathcal{C})$ .


# The blueprint of Geometric Deep Learning




## *Popular architectures as instances of GDL blueprint*

<b>Architecture</b>	<b>Domain <math>\Omega</math></b>	<b>Symmetry Group <math>\mathfrak{G}</math></b>
<i>CNN</i>	Grid	Translation
<i>Spherical CNN</i>	Sphere / $SO(3)$	Rotation $SO(3)$
<i>Mesh CNN</i>	Manifold	Isometry $Iso(\Omega)$ / Gauge Symmetry $SO(2)$
<i>GNN</i>	Graph	Permutation $\Sigma_n$
<i>Deep Sets</i>	Set	Permutation $\Sigma_n$
<i>Transformer</i>	Complete Graph	Permutation $\Sigma_n$
<i>E(3) GNN</i>	Geometric Graph	Permutation $\Sigma_n \times$ Euclidean $E(3)$
<i>LSTM</i>	1D Grid	Time warping


# Architectures of interest




**Perceptrons**  
Function regularity



**CNNs**  
Translation




**Group-CNNs**  
Global groups



**LSTMs**  
Time warping



**Deep Sets / Transformers**  
Permutation



**GNNs**  
Permutation



**Intrinsic CNNs**  
Isometry / Gauge choice



*...now it's our  
turn to study  
geometry! 😊*

