

# Introduction to Geometric Deep Learning



Yan Hu

# Background



In very broad terms, the data we use to train deep learning models belongs to two main domains:

1. **Euclidean data:** data represented in multidimensional linear spaces, it obeys Euclidean postulates, e.g, text data or tabular data.
2. **Non-Euclidean data:** in even broader terms, data that doesn't obey Euclidean postulates, e.g, molecular structures, social network interactions, or meshed 3D surface.

# Data from nature is often geometric





**Data from nature is often geometric**



# Models of nature are often geometric



# Models of nature are often geometric



# **Non-geometric models have geometric constraints**

NLPers may say “no such geometry in language”. But “geometry” is not just about spatial arrangement!

# **Non-geometric models have geometric constraints**

NLPers may say “no such geometry in language”. But “geometry” is not just about spatial arrangement!

It is about constraints: design model to “respect” regularity in data. Models like Transformers touted as “generic”, but significantly constrained.



# Non-geometric models have geometric constraints

NLPers may say “no such geometry in language”. But “geometry” is not just about spatial arrangement!

It is about constraints: design model to “respect” regularity in data. Models like Transformers touted as “generic”, but significantly constrained.

Generally, GDL offers us a perspective to categorise existing architectures. Based on which data regularity constraints they satisfy.

# Non-geometric models have geometric constraints

NLPers may say “no such geometry in language”. But “geometry” is not just about spatial arrangement!

It is about constraints: design model to “respect” regularity in data. Models like Transformers touted as “generic”, but significantly constrained.


Generally, GDL offers us a perspective to categorise existing architectures. Based on which data regularity constraints they satisfy.

This is a useful perspective even if you never encounter “geometric” data.




# Learning in High Dimensions

In general, learning functions in high dimensions is intractable.

Number of samples required grows exponentially with dimensions.



# Curse of Dimensionality



# Curse of Dimensionality

“[dimensionality is] a **curse** which has hung over the head of the physicist and astronomer for many a year. ”

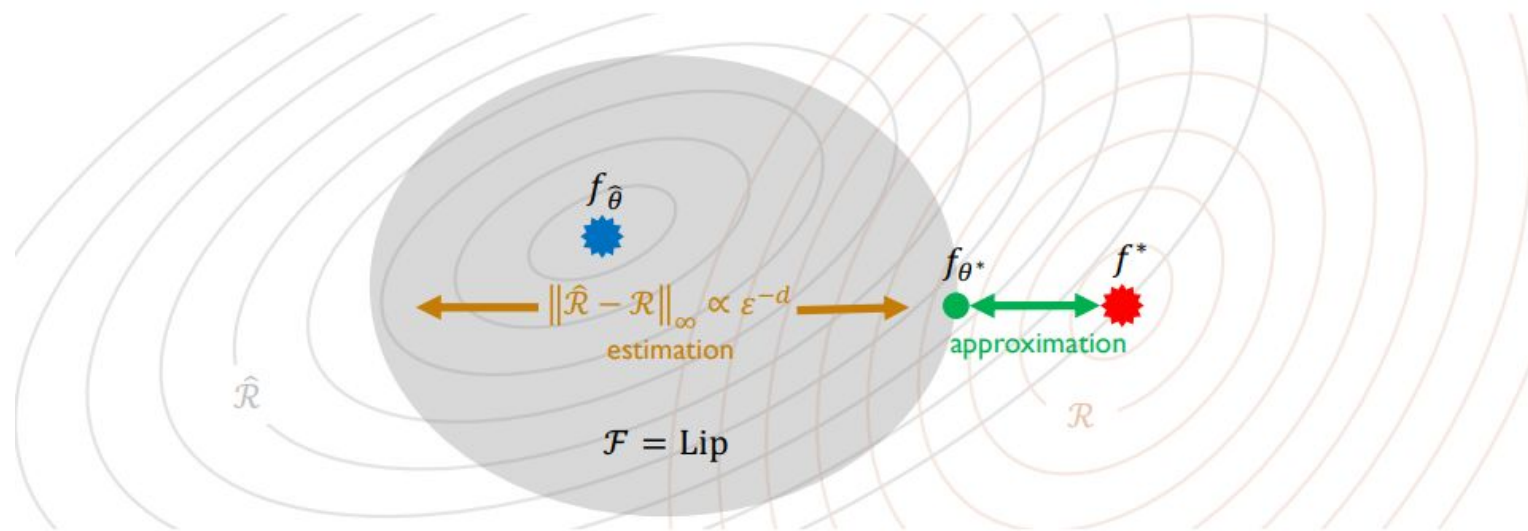
— *Dynamic Programming*



**R. Bellman**

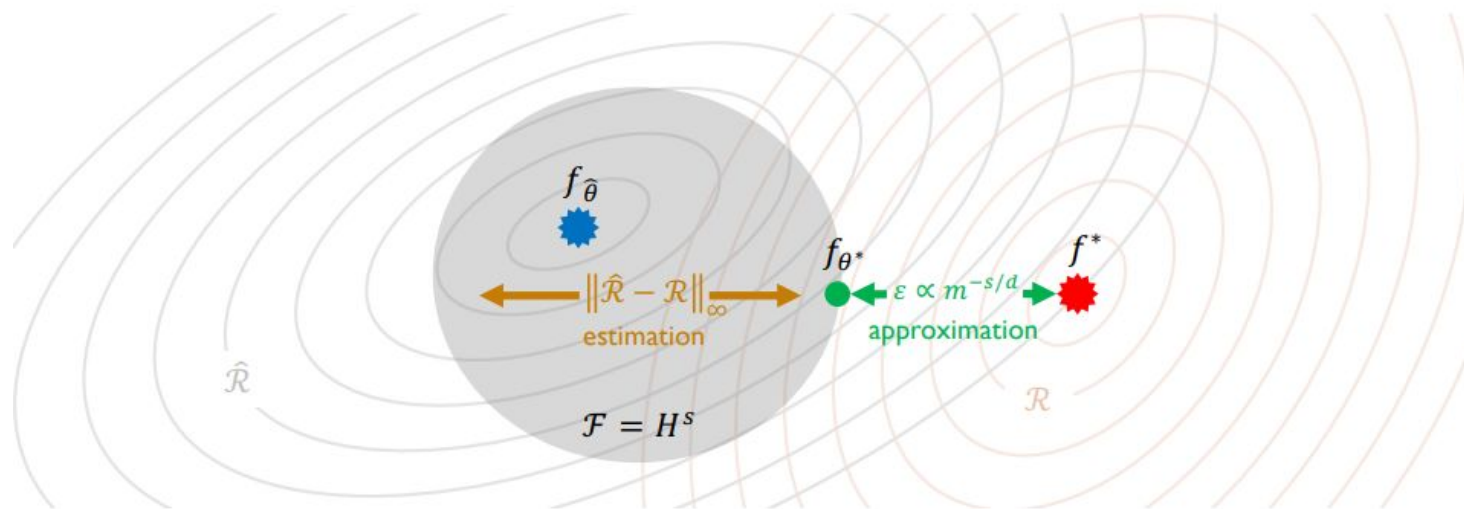
# Classical notions of regularity are of little use

Lipschitz class is too large: estimation error is dimensionality-cursed



# Classical notions of regularity are of little use

Sobolev class is too small: approximation error is dimensionality-cursed



# Takeaways

1. Learning in high dimensions is plagued by the curse of dimensionality
2. Impossible without assumptions (“priors”)
3. Classical assumptions of regularity (from low-dimensional analysis) are not appropriate priors
4. Geometric priors: inputs are signals defined over low-dimensional geometric domains

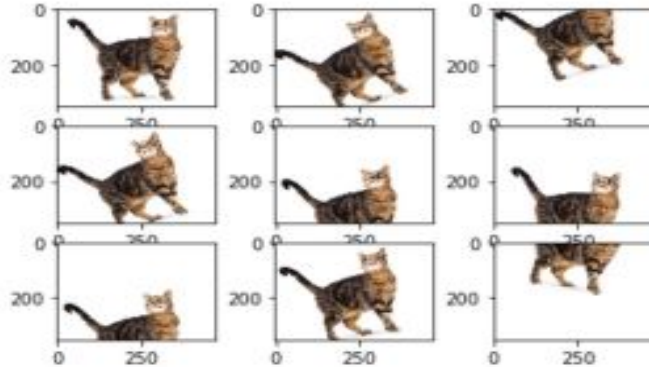


# Geometry to the rescue!

We can inject assumptions about geometry through inductive biases.

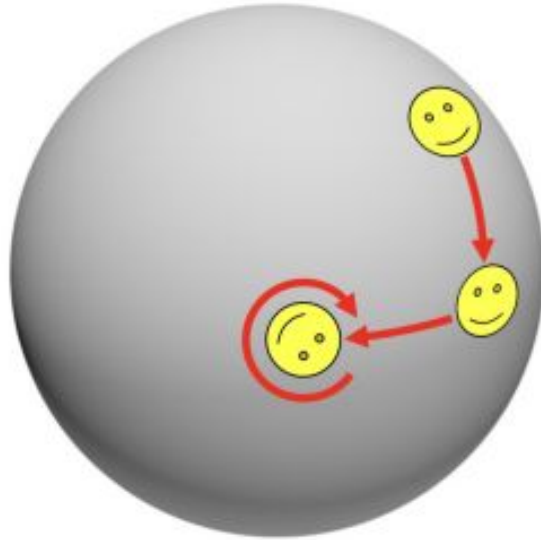
Restrict the functions to ones that respect the geometry. This can make the high-dimensional problem more tractable!

Examples: **Image** data should be processed independently of shifts





# Continued

Examples: **Spherical data** should be processed independently of rotations



# Continued

Examples: **Graph data** should be processed independently of isomorphism



# A Roadmap for Formalisation


1. To handle geometry of data, we need to formalise where the data lives (domain) and how to featurise it (signal)
2. Once we understand data domains, we can then formalise symmetries of those domains (groups)
3. Equipped with groups, we need to formalise how they transform the data domains (group actions)
4. Deep learning concerns itself with linear algebra; we need to be able to talk about group actions as matrix operations (representations)
5. Using representations, we can formalise what it means for a deep learning model to respect symmetries (invariance & equivariance)

# Geometric Domains


- **Domain**  $\Omega$  = set + some structure



**Point cloud**  
(bare set)



**Graph**  
(local neighbourhood)




**Mesh**  
(local metric)

# Signals on Geometric Domains


Signal  $x \in \mathcal{M}(\Omega, \mathbb{C}) = \{x: \Omega \rightarrow \mathbb{C}\}$ ,  $\mathbb{C}$ -valued function on  $\Omega$

- Domain  $\Omega$  (often no vector space structure, i.e., we cannot add points)
- Vector space  $\mathbb{C}$  (dimensions referred to as “channels”)



$$\Omega = \mathbb{Z}_n \times \mathbb{Z}_n$$

Image



$$\mathbb{C} = \mathbb{R}^3$$




$$\Omega = \mathcal{M}$$

Textured surface




$$\mathbb{C} = \mathbb{R}^3$$



$$\Omega = (V, E)$$

Molecular graph



$$\mathbb{C} = \mathbb{R}^m$$

# Signals on Geometric domains

The space of signals  $\mathcal{A}(\Omega, \mathbb{C})$  is a vector space (possibly infinite-dimensional)

- We can add signals and multiply them by a scalar



# Signals on Geometric Domains

The space of signals  $\mathcal{A}(\Omega, \mathbb{C})$  is a vector space (possibly infinite-dimensional)

- Given an inner product  $\langle \cdot, \cdot \rangle$  on  $\mathbb{C}$ , and a measure  $\mu$  on  $\Omega$ , we can define an inner product on  $\mathcal{A}(\Omega, \mathbb{C})$  as:

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



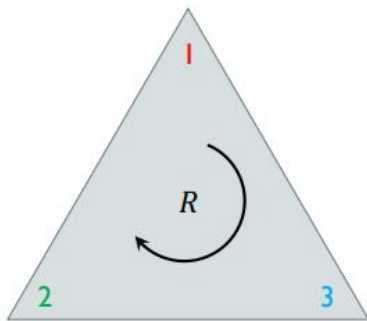
# A Roadmap for Formalisation

1. To handle geometry of data, we need to formalise where the data lives (domain) and how to featurise it (signal)
2. Once we understand data domains, we can then formalise symmetries of those domains (groups)
3. Equipped with groups, we need to formalise how they transform the data domains (group actions)
4. Deep learning concerns itself with linear algebra; we need to be able to talk about group actions as matrix operations (representations)
5. Using representations, we can formalise what it means for a deep learning model to respect symmetries (invariance & equivariance)


# Symmetries

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

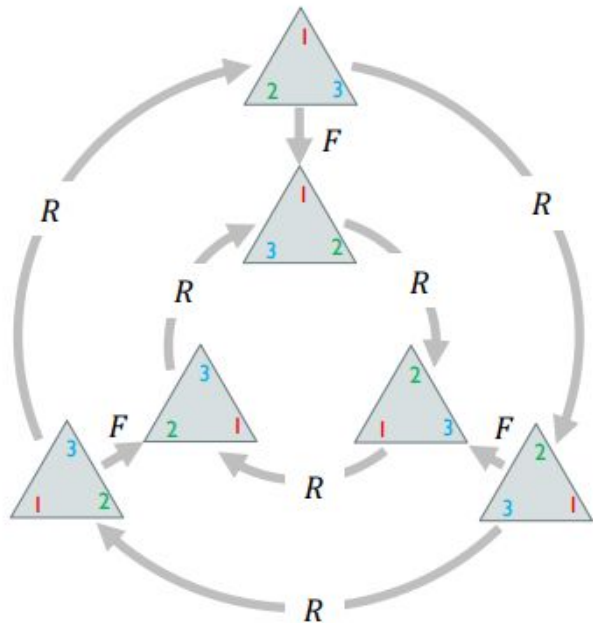
Examples: Symmetry of a triangle



rotation by  $120^\circ$



reflection



# Symmetries

- 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

# Groups


A **group**  $(G,*)$  is a set  $G$  together with binary operation  $* : G \times G \rightarrow G$  (denoted by juxtaposition  $g * h = gh$  for brevity) satisfying the following axioms:

- *Associativity:*  $(gh)k = g(hk)$  for all  $g, h, k \in G$
- *Identity:*  $\exists! e \in G$  satisfying  $eg = ge = g$  for all  $g \in G$
- *Inverse:*  $\exists! g^{-1} \in G$  for each  $g \in G$  satisfying  $g^{-1}g = gg^{-1} = e$

# A Roadmap for Formalisation

1. To handle geometry of data, we need to formalise where the data lives (domain) and how to featurise it (signal)
2. Once we understand data domains, we can then formalise symmetries of those domains (groups)
3. Equipped with groups, we need to formalise how they transform the data domains (group actions)
4. Deep learning concerns itself with linear algebra; we need to be able to talk about group actions as matrix operations (representations)
5. Using representations, we can formalise what it means for a deep learning model to respect symmetries (invariance & equivariance)


# Group actions on objects



**Point in a plane**



**Image (function)**



**Vector field**

The type of an object can be defined by the way it transforms by a group

# Group Action

Let  $G$  be a group and  $X$  a set. A **(left) group action** of  $G$  on  $X$  (often denoted  $gx = \alpha(g, x)$ ) is a mapping of the form  $\alpha : G \times X \rightarrow X$  satisfying


- *Identity:*  $\alpha(e, x) = x$  for all  $x \in X$
- *Composition:*  $\alpha(gh, x) = \alpha(g, \alpha(h, x))$  for all  $g, h \in G$  and  $x \in X$

# Group Action

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





# A Roadmap for Formalisation


1. To handle geometry of data, we need to formalise where the data lives (domain) and how to featurise it (signal)
2. Once we understand data domains, we can then formalise symmetries of those domains (groups)
3. Equipped with groups, we need to formalise how they transform the data domains (group actions)
4. Deep learning concerns itself with linear algebra; we need to be able to talk about group actions as matrix operations (representations)
5. Using representations, we can formalise what it means for a deep learning model to respect symmetries (invariance & equivariance)

# Linear Group Representation


A  $d$ -dimensional (linear) representation of  $G$  is a map  $\rho: G \rightarrow \mathbb{R}^{d \times d}$  assigning to each  $g \in G$  an invertible matrix  $\rho(g) \in \mathbb{R}^{d \times d}$  satisfying  $\rho(gh) = \rho(g)\rho(h)$  for all  $g, h \in G$ .

# Group actions on Signals on Geometric Domains

Given a group  $G$  acting on a **domain**  $\Omega$ , we automatically obtain an action of  $G$  on the space of signals  $\mathcal{X}(\Omega)$  through the **regular representation**  $(\rho(g)x)(u) = x(g^{-1}u)$



# Intuition



# Intuition

The diagram illustrates the intuition of a linear transformation. It shows two vertical bars of colored segments on the left, followed by an equals sign, a matrix of ones, and another two vertical bars on the right. The bars represent vectors in a 6-dimensional space.

Left side:  $0.5$  [Bar 1] +  $0.5$  [Bar 2]

Matrix:  $\begin{bmatrix} 1 & & & & & \\ & 1 & & & & \\ & & 1 & & & \\ & & & 1 & & \\ & & & & 1 & \\ & & & & & 1 \end{bmatrix}$


Right side:  $0.5$  [Bar 3] +  $0.5$  [Bar 4]

The transformation maps the first bar to the third bar and the second bar to the fourth bar. The colors in the bars are: Bar 1 (orange, light blue, brown, light yellow, orange, brown), Bar 2 (dark grey, olive green, dark grey, orange, yellow, olive green), Bar 3 (light blue, brown, light yellow, orange, brown, orange), Bar 4 (olive green, dark grey, orange, yellow, olive green, dark grey).


# Symmetry in Learning

Symmetries of the Label Function:

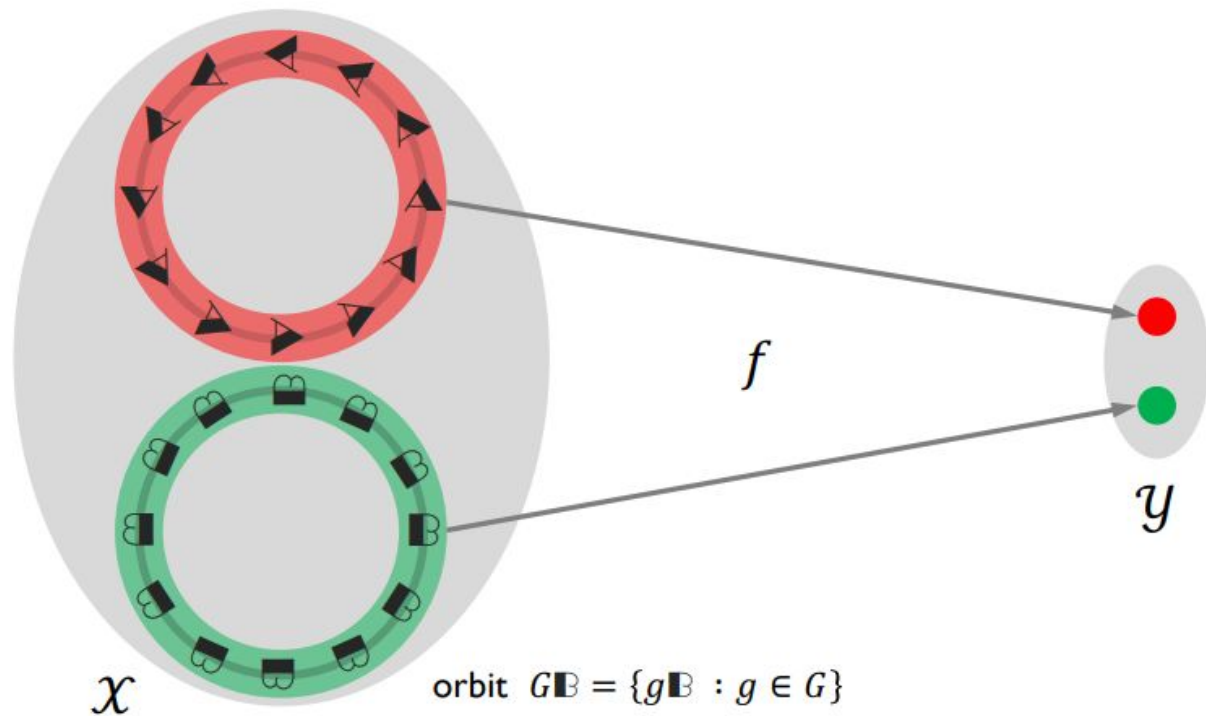
- Label function  $f: X \rightarrow Y$  e.g., classification task ( $Y=\{1,\dots,K\}$ )
- Symmetry of a label function is an invertible label-preserving map  $g: X \rightarrow X$ ,  
i.e.  $(f \circ g)(x) = f(x)$  for all  $x \in X$



# Symmetries of the Label Function

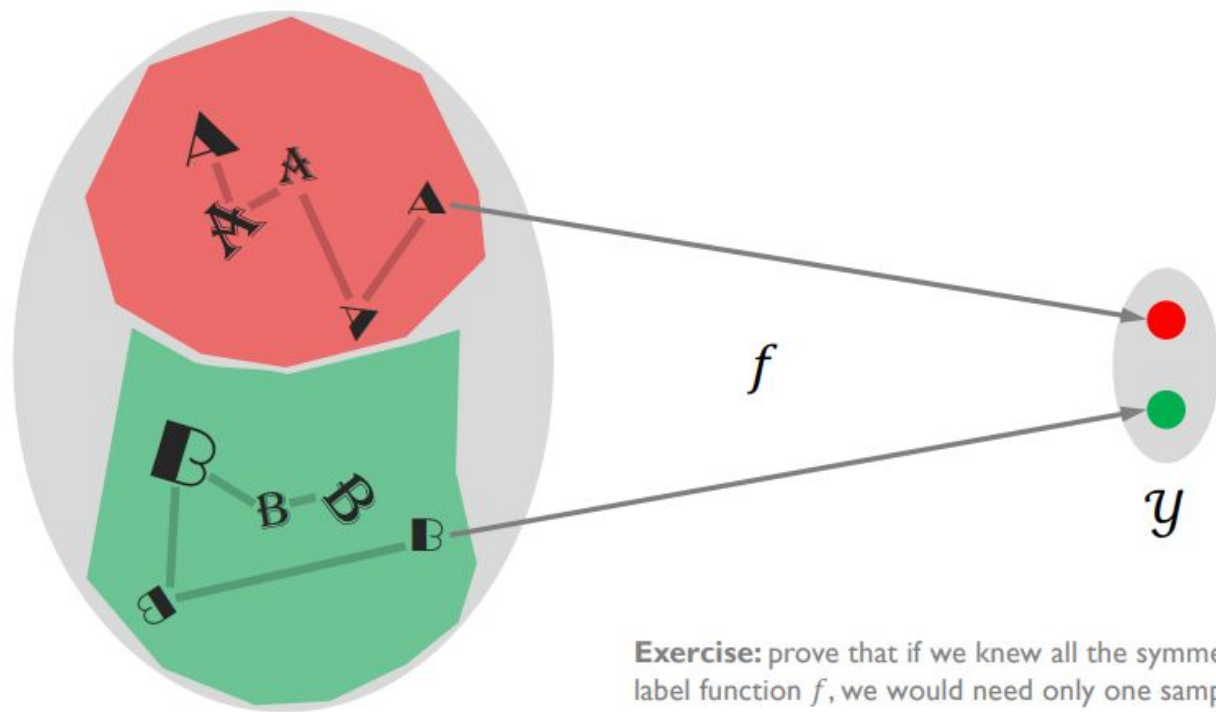


# Symmetries of the Label Function





# Symmetries of the Label Function





# Symmetries of the Weights

Let  $f_\theta: \mathcal{X} \times \Theta \rightarrow \mathcal{Y}$  be a parametric model (neural network)

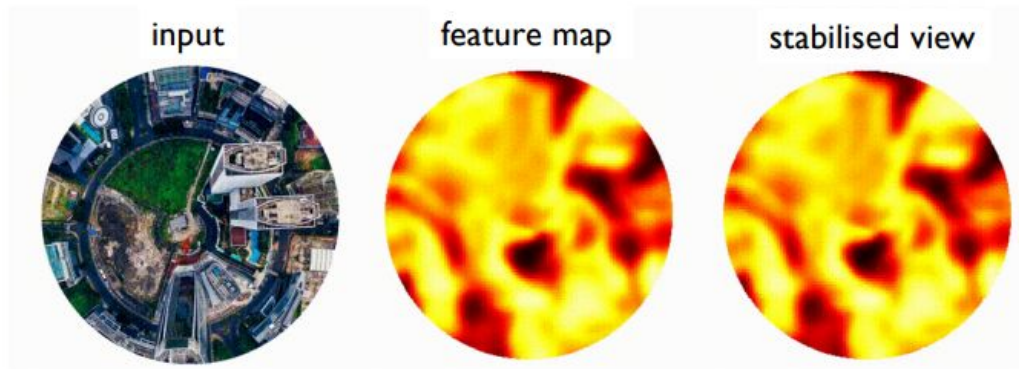
A transformation  $h: \Theta \rightarrow \Theta$  is a **symmetry of the weights** if, for all  $x \in \mathcal{X}$  and  $\theta \in \Theta$

$$f_{h\theta}(x) = f_\theta(x)$$




# Example

CNN



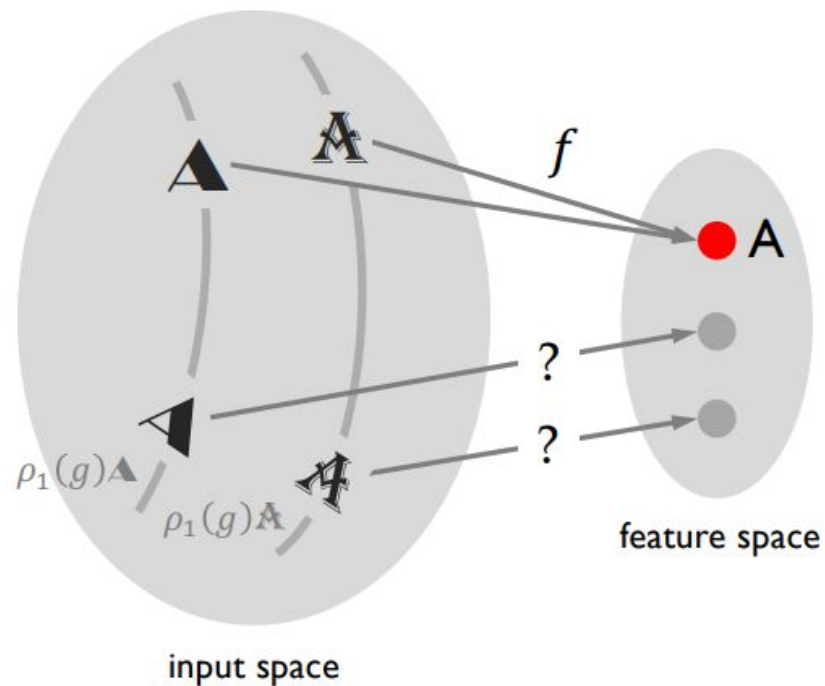
Rotation-  
equivariant  
CNN



# A Roadmap for Formalisation

1. To handle geometry of data, we need to formalise where the data lives (domain) and how to featurise it (signal)
2. Once we understand data domains, we can then formalise symmetries of those domains (groups)
3. Equipped with groups, we need to formalise how they transform the data domains (group actions)
4. Deep learning concerns itself with linear algebra; we need to be able to talk about group actions as matrix operations (representations)
5. Using representations, we can formalise what it means for a deep learning model to respect symmetries (invariance & equivariance)


# Equivariance = Symmetry-consistent Generalisation



$$f(\rho_1(g)\blacktriangle) = \rho_2(g) f(\blacktriangle)$$

$$f(\rho_1(g)\hat{\blacktriangle}) = \rho_2(g) f(\hat{\blacktriangle})$$

# Equivariance = Symmetry-consistent Generalisation



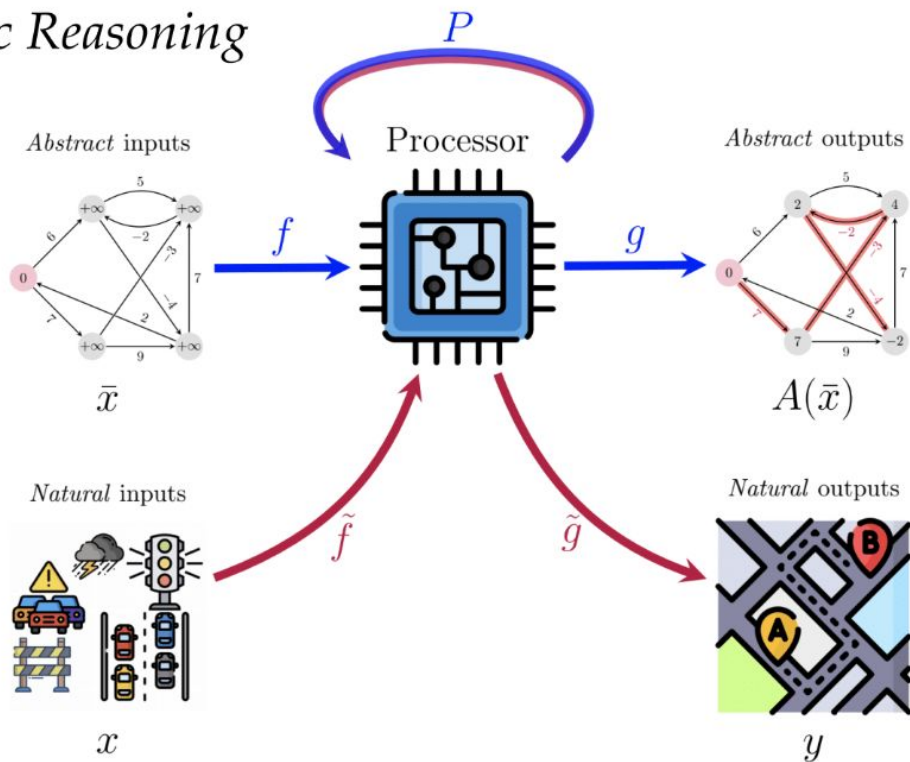
$$f(\rho_1(g)A) = \rho_2(g) f(A)$$

||

$$f(\rho_1(g)A) = \rho_2(g) f(A)$$

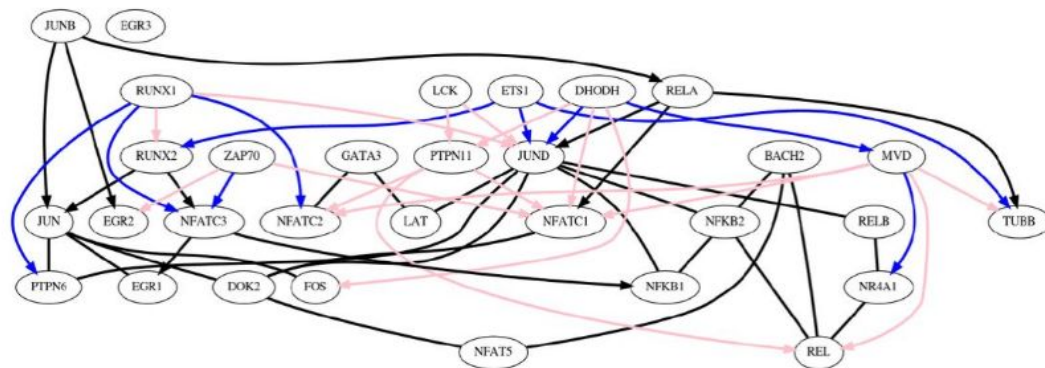

# What's next?

## Algorithmic Reasoning



# What's next?

## *Causal Inference*






# Social Networks



# 3D VISION & GRAPHICS

## *Self-Driving Cars*




# 3D VISION & GRAPHICS




# STRUCTURAL BIOLOGY

## *Protein Folding*

Protein folding



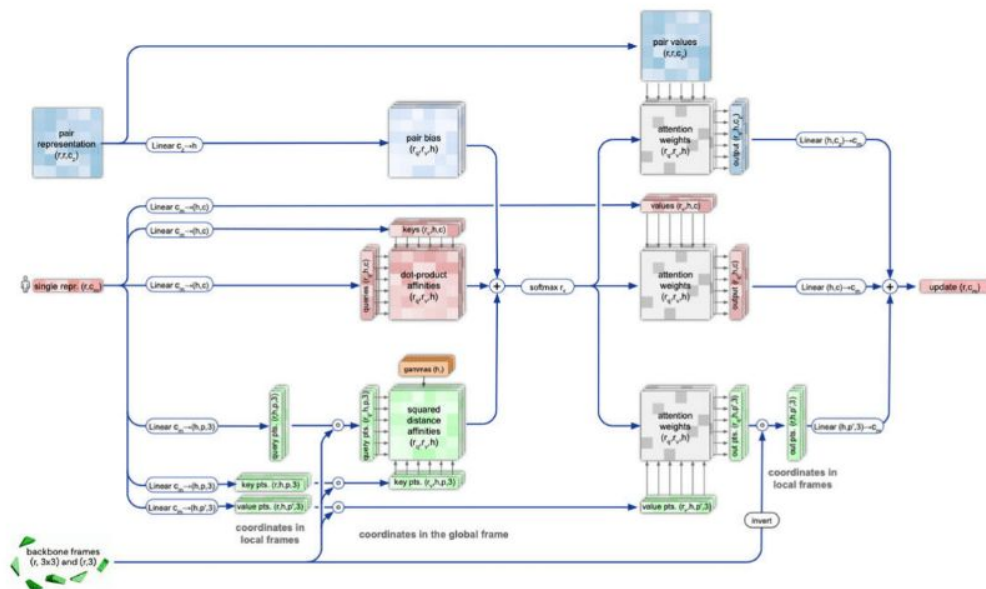
Primary protein  
structure



Tertiary protein  
structure

# STRUCTURAL BIOLOGY


## *AlphaFold 2*



Invariant Point Attention

# STRUCTURAL BIOLOGY

*Protein Design = "Inverse Folding"*



# STRUCTURAL BIOLOGY

Sequence → Structure → Function

*MaSIF: Protein Function Prediction*

