

MLSS 2021 Summer School

GEOMETRIC DEEP LEARNING

Michael Bronstein

Imperial College London / IDSIA / Twitter

Based on joint work with J. Bruna, T. Cohen, P. Veličković

Slides adapted from AMMI Course

Fundamental Principles underlying Deep Representation Learning architectures

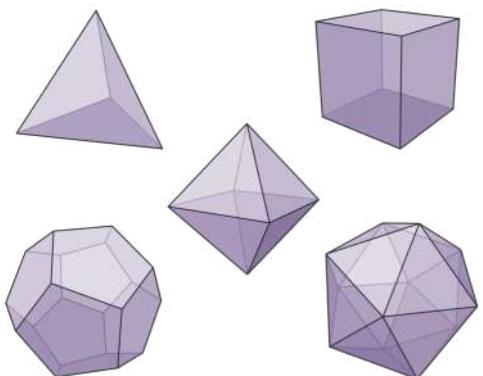
Symmetry

“Symmetry, as wide or as narrow as you may define its meaning, is one idea by which man through the ages has tried to comprehend and create order, beauty, and perfection”

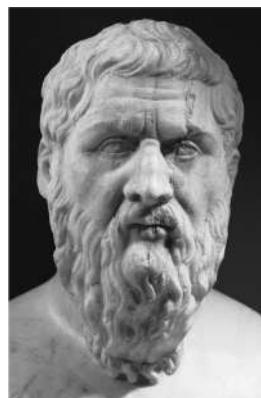


H. Weyl

συμμετρία

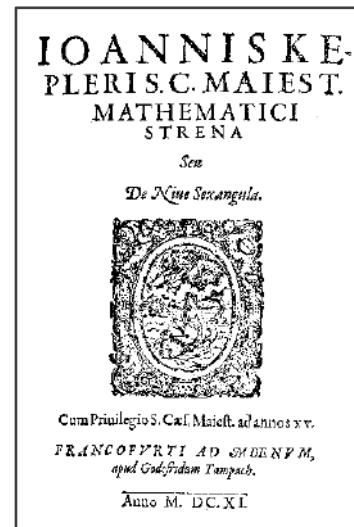


Platonic solids



Plato

~370 BC



“On the six-cornered
snowflake”

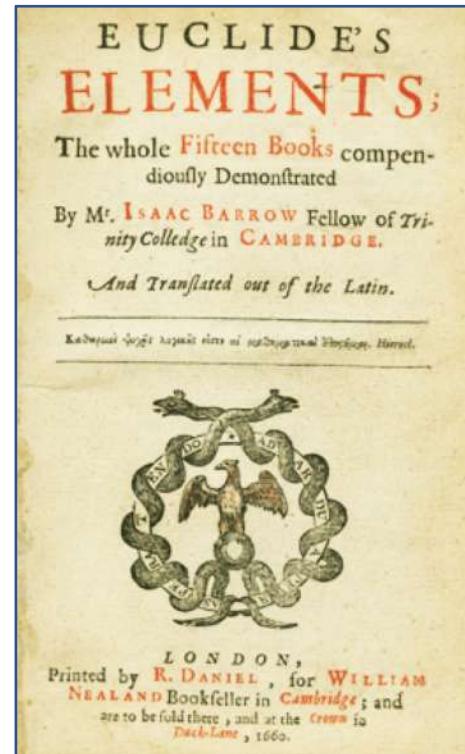


J. Kepler

1611

Euclidean geometry

“In a plane, given a line and a point not on it, at most one line parallel to the given line can be drawn through the point”



Euclid

~300 BC

End of Euclid's Monopoly



J. V. Poncelet



N. Lobachevsky



C. F. Gauss



J. Bolyai



B. Riemann

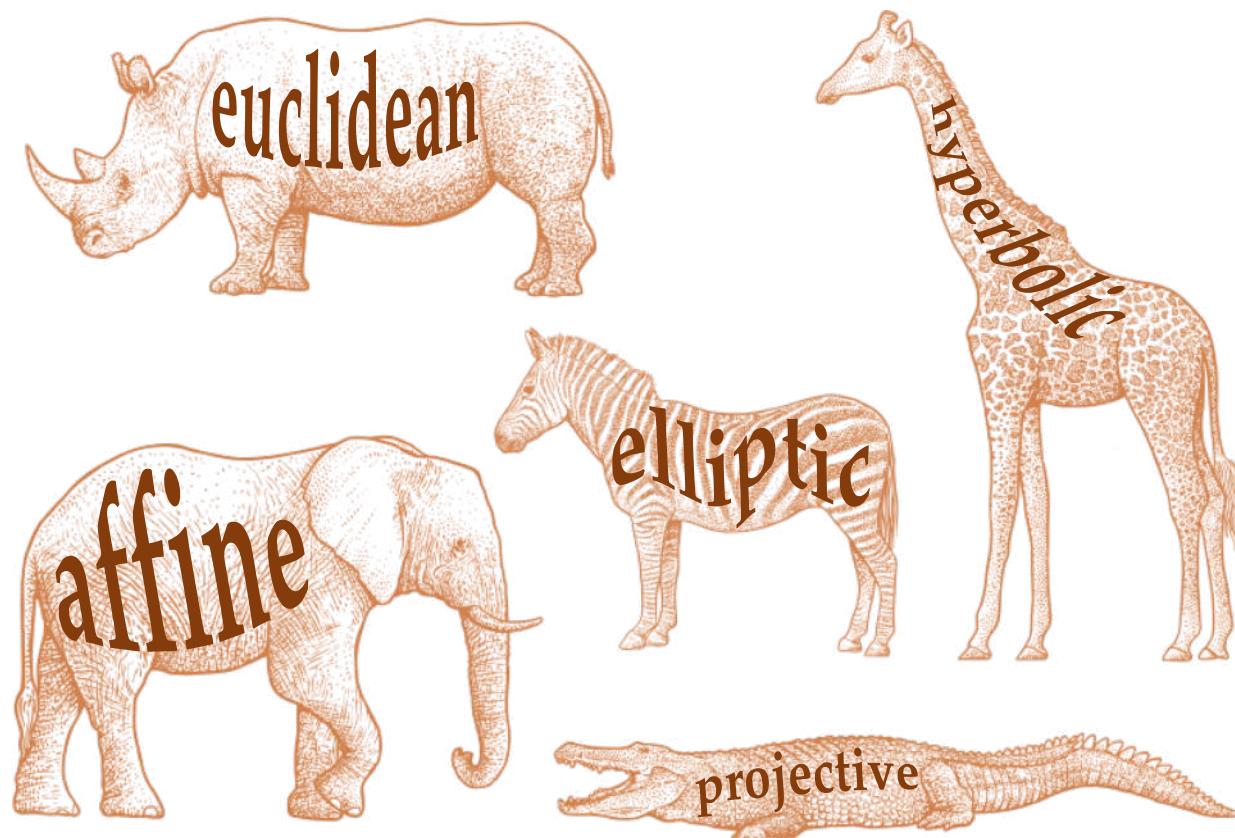
1822

1826

1832

1856

19th Century Zoo of Geometries



The Erlangen Programme

“Given a [homogeneous] manifold and a transformation group acting [transitively] on it, to investigate those properties of figures on that manifold which are invariant under transformations of that group”



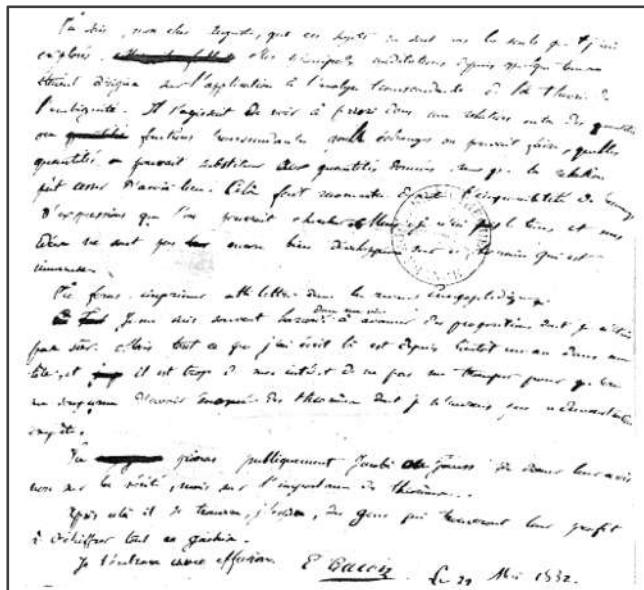
F. Klein

1872

The Erlangen Programme

	Euclidean	Affine	Projective
<i>angle</i>	+	—	—
<i>distance</i>	+	—	—
<i>area</i>	+	—	—
<i>parallelism</i>	+	+	—
<i>intersection</i>	+	+	+

Group Theory



Mon cher ami — Galois' last letter
written on the night before his duel



E. Galois



F. Klein



S. Lie

1832

Noether's Theorem

“Every [differentiable] symmetry of the action of a physical system [with conservative forces] has a corresponding conservation law”

Invariante Variationsprobleme.
(F. Klein zum fünfzigjährigen Doktorjubiläum.)
Von
Emmy Noether in Göttingen.
Vorgelegt von F. Klein in der Sitzung vom 26. Juli 1918¹⁾.

Es handelt sich um Variationsprobleme, die eine kontinuierliche Gruppe (im Lieschen Sinne) gestatten; die daraus sich ergebenden Folgerungen für die zugehörigen Differentialgleichungen finden ihren allgemeinsten Ausdruck in den in § 1 formulierten, in den folgenden Paragraphen bewiesenen Sätzen. Über diese aus Variationsproblemen entspringenden Differentialgleichungen lassen sich viel präzisere Aussagen machen als über beliebige, eine Gruppe gestattende Differentialgleichungen, die den Gegenstand der Lieschen Untersuchungen bilden. Das folgende beruht also auf einer Verbindung der Methoden der formalen Variationsrechnung mit denen der Lieschen Gruppentheorie. Für spezielle Gruppen und Variationsprobleme ist diese Verbindung der Methoden nicht neu; ich erwähne Hamel und Herglotz für spezielle endliche, Lorentz und seine Schüler (z. B. Fokker), Weyl und Klein für spezielle unendliche Gruppen²⁾. Insbesondere sind die zweite Kleinsche Note und die vorliegenden Ausführungen gegenseitig durch einander beein-

1) Die endgültige Fassung des Manuskriptes wurde erst Ende September eingesicht.

2) Hamel: Math. Ann. Bd. 59 und Zeitschrift f. Math. u. Phys. Bd. 50. Herglotz: Ann. d. Phys. (4) Bd. 36, bes. § 9, S. 511. Fokker, Verslag d. Amsterdamer Akad., 27/I. 1917. Für die weitere Literatur vergl. die zweite Note von Klein: Göttinger Nachrichten 19. Juli 1918.

In einer eben erschienenen Arbeit von Kneser (Math. Zeitschrift Bd. 2) handelt es sich um Aufstellung von Invarianten nach ähnlicher Method.

Kgl. Ges. d. Wiss. Nachrichten, Math.-phys. Klasse, 1918, Heft 2. 17



E. Noether

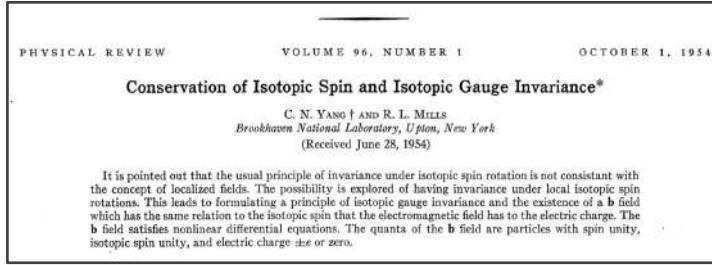
1918

Gauge Invariance



H. Weyl

1929

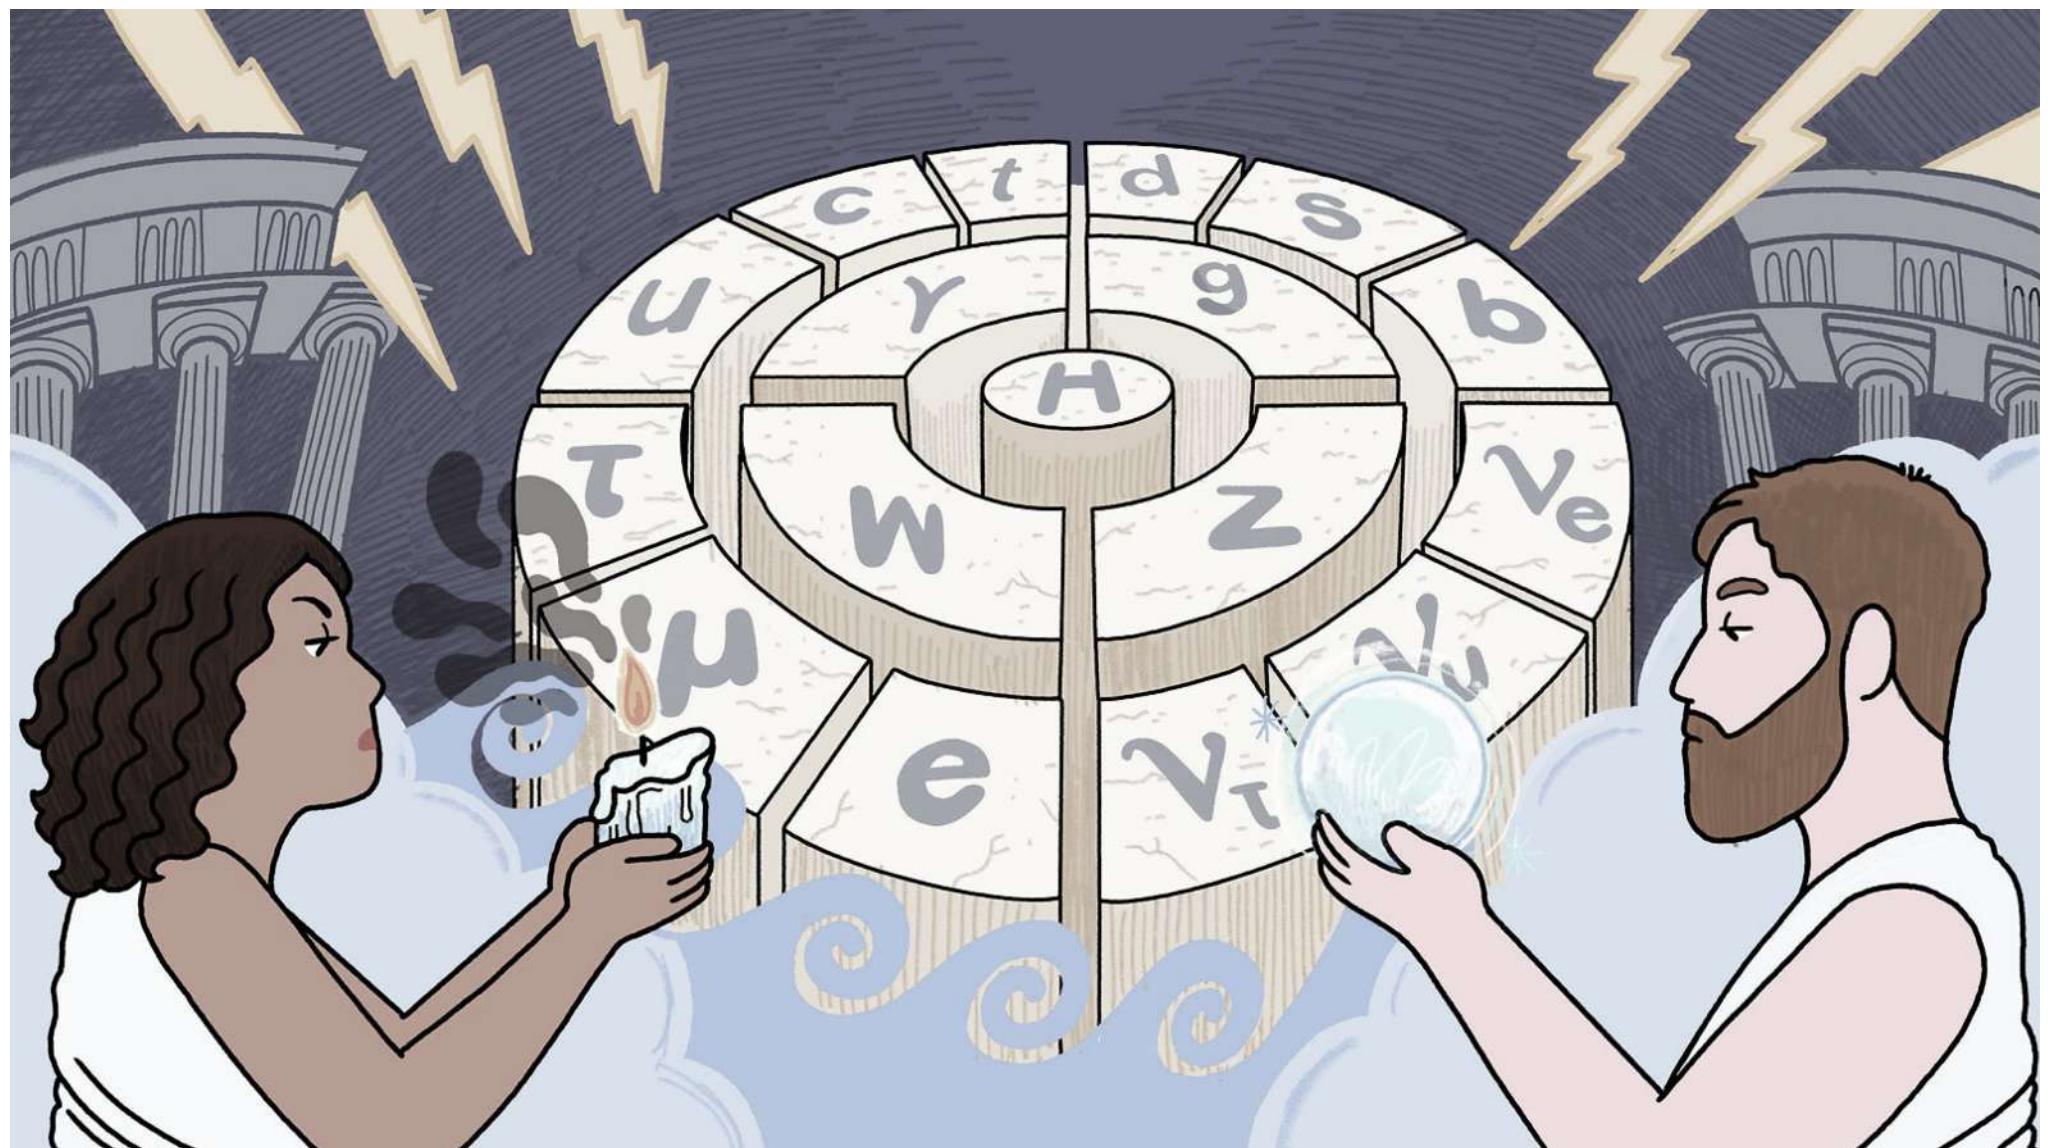


C. N. Yang



R. L. Mills

1954



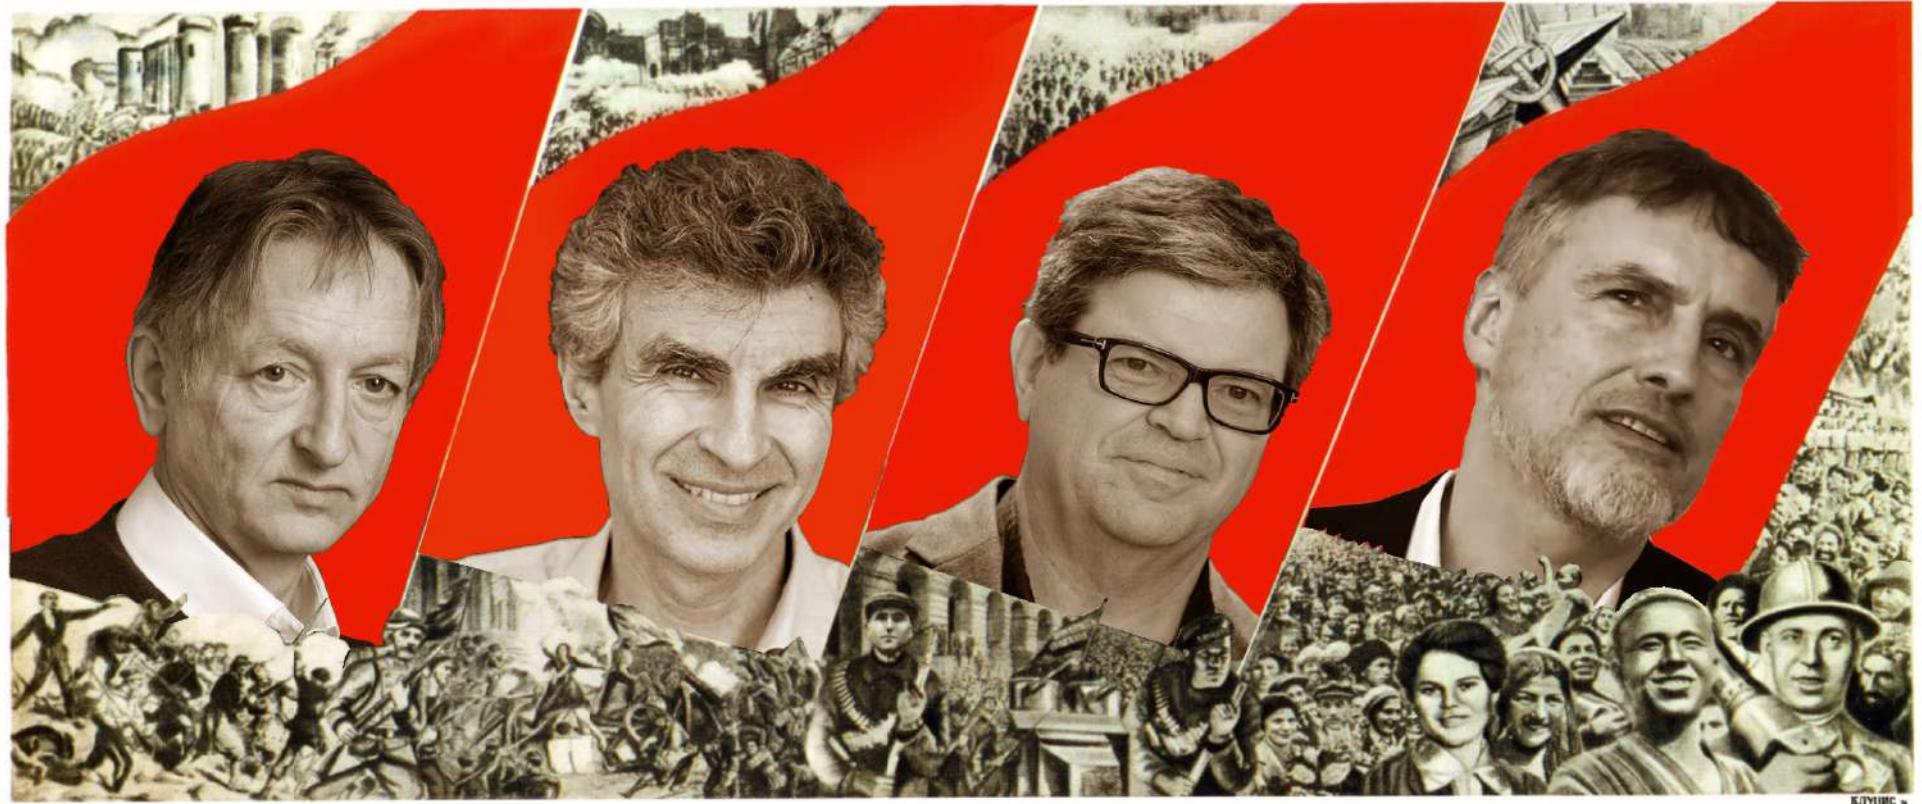
“It is only slightly overstating the case to say that
Physics is the study of symmetry”



P. Anderson

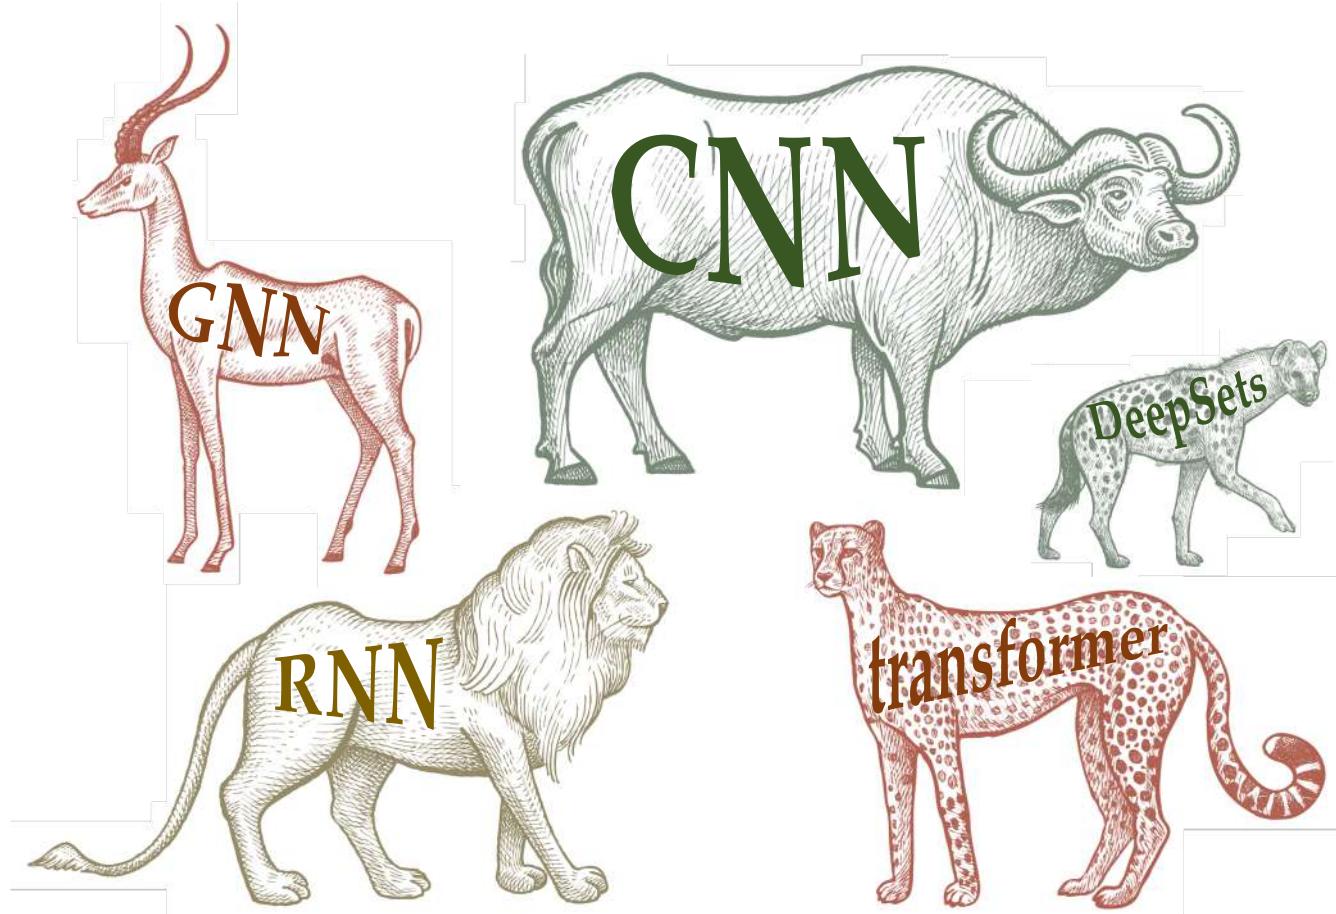
Anderson 1972

?



REVOLUTION IN DATA SCIENCE

20th Century Zoo of Neural Network Architectures



The Erlangen Programme of ML
Geometric Deep Learning

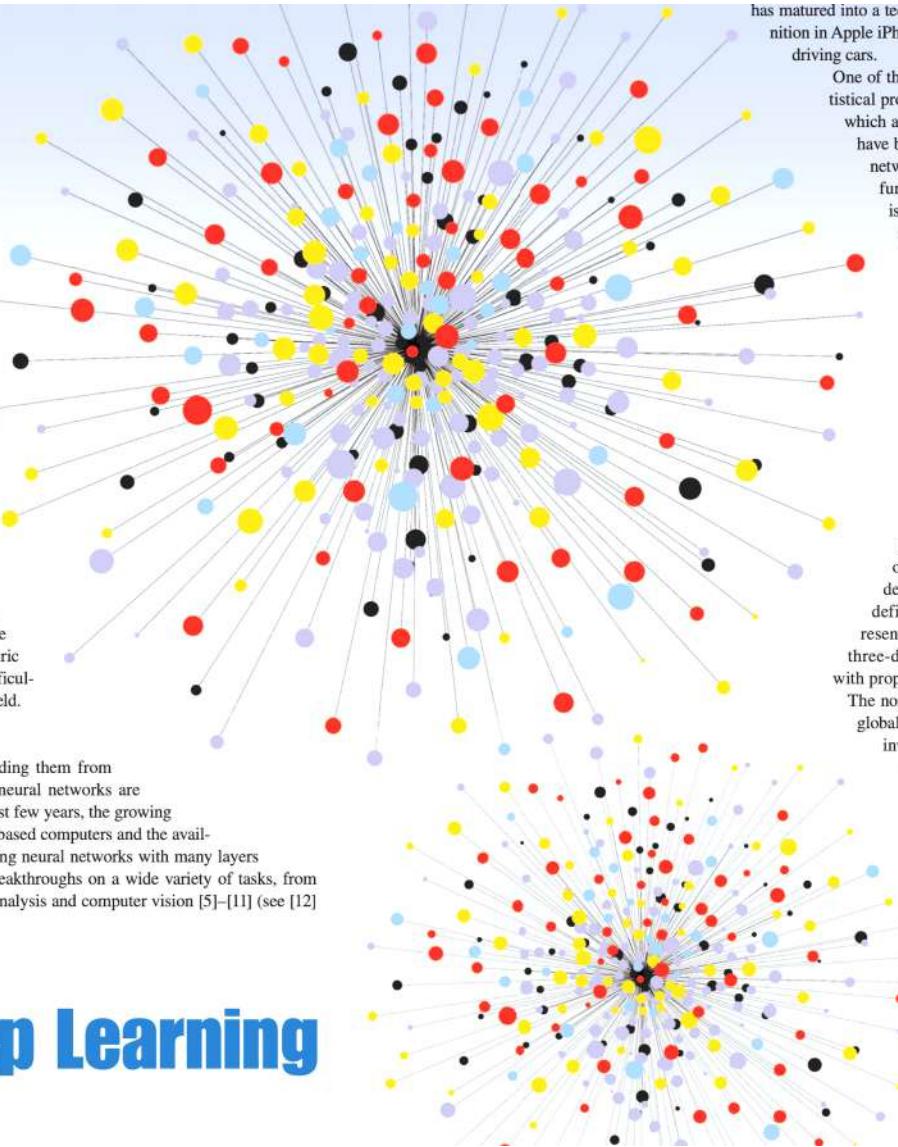
Michael M. Bronstein, Joan Bruna, Yann LeCun,
Arthur Szlam, and Pierre Vandergheynst

Many scientific fields study data with an underlying structure that is non-Euclidean. Some examples include social networks in computational social sciences, sensor networks in communications, functional networks in brain imaging, regulatory networks in genetics, and meshed surfaces in computer graphics. In many applications, such geometric data are large and complex (in the case of social networks, on the scale of billions) and are natural targets for machine-learning techniques. In particular, we would like to use deep neural networks, which have recently proven to be powerful tools for a broad range of problems from computer vision, natural-language processing, and audio analysis. However, these tools have been most successful on data with an underlying Euclidean or grid-like structure and in cases where the invariances of these structures are built into networks used to model them.

Geometric deep learning is an umbrella term for emerging techniques attempting to generalize (structured) deep neural models to non-Euclidean domains, such as graphs and manifolds. The purpose of this article is to overview different examples of geometric deep-learning problems and present available solutions, key difficulties, applications, and future research directions in this nascent field.

Overview of deep learning

Deep learning refers to learning complicated concepts by building them from simpler ones in a hierarchical or multilayer manner. Artificial neural networks are popular realizations of such deep multilayer hierarchies. In the past few years, the growing computational power of modern graphics processing unit (GPU)-based computers and the availability of large training data sets have allowed successfully training neural networks with many layers and degrees of freedom (DoF) [1]. This has led to qualitative breakthroughs on a wide variety of tasks, from speech recognition [2], [3] and machine translation [4] to image analysis and computer vision [5]–[11] (see [12]



Geometric Deep Learning

Going beyond Euclidean data

has matured into a technology that is widely used in commercial applications, including Siri speech recognition in Apple iPhone, Google text translation, and Mobileye vision-based technology for autonomously driving cars.

One of the key reasons for the success of deep neural networks is their ability to leverage statistical properties of the data, such as stationarity and compositionality through local statistics, which are present in natural images, video, and speech [14], [15]. These statistical properties have been related to physics [16] and formalized in specific classes of convolutional neural networks (CNNs) [17]–[19]. In image analysis applications, one can consider images as functions on the Euclidean space (plane), sampled on a grid. In this setting, stationarity is owed to shift invariance, locality is due to the local connectivity, and compositionality stems from the multiresolution structure of the grid. These properties are exploited by convolutional architectures [20], which are built of alternating convolutional and downsampling (pooling) layers. The use of convolutions has a twofold effect. First, it allows extracting local features that are shared across the image domain and greatly reduces the number of parameters in the network with respect to generic deep architectures (and thus also the risk of overfitting), without sacrificing the expressive capacity of the network. Second, the convolutional architecture itself imposes some priors about the data, which appear very suitable especially for natural images [17]–[19], [21].

While deep-learning models have been particularly successful when dealing with speech, image, and video signals, in which there are an underlying Euclidean structure, recently there has been a growing interest in trying to apply learning on non-Euclidean geometric data. Such kinds of data arise in numerous applications. For instance, in social networks, the characteristics of users can be modeled as signals on the vertices of the social graph [22]. Sensor networks are graph models of distributed interconnected sensors, whose readings are modeled as time-dependent signals on the vertices. In genetics, gene expression data are modeled as signals defined on the regulatory network [23]. In neuroscience, graph models are used to represent anatomical and functional structures of the brain. In computer graphics and vision, three-dimensional (3-D) objects are modeled as Riemannian manifolds (surfaces) endowed with properties such as color texture.

The non-Euclidean nature of such data implies that there are no such familiar properties as global parameterization, common system of coordinates, vector space structure, or shift invariance. Consequently, basic operations like convolution that are taken for granted in the Euclidean case are even not well defined on non-Euclidean domains. The purpose of this article is to show different methods of translating the key ingredients of successful deep-learning methods, such as CNNs, to non-Euclidean data.

Geometric learning problems

Broadly speaking, we can distinguish between two classes of geometric learning problems. In the first class of problems, the goal is to characterize the structure of the data. The second class of problems deals with analyzing functions defined on a given non-Euclidean domain. These two classes are related, because understanding the properties of functions defined on a domain conveys certain information about the domain, and vice versa, the structure of the domain imposes certain properties on the functions on it.

Structure of the domain

As an example of the first class of problems, assume to be given a set of data points with some underlying low-dimensional structure embedded into a high-dimensional Euclidean space. Recovering that low-dimensional structure

Geometric Deep Learning Grids, Groups, Graphs, Geodesics, and Gauges

Michael M. Bronstein¹, Joan Bruna², Taco Cohen³, Petar Veličković⁴

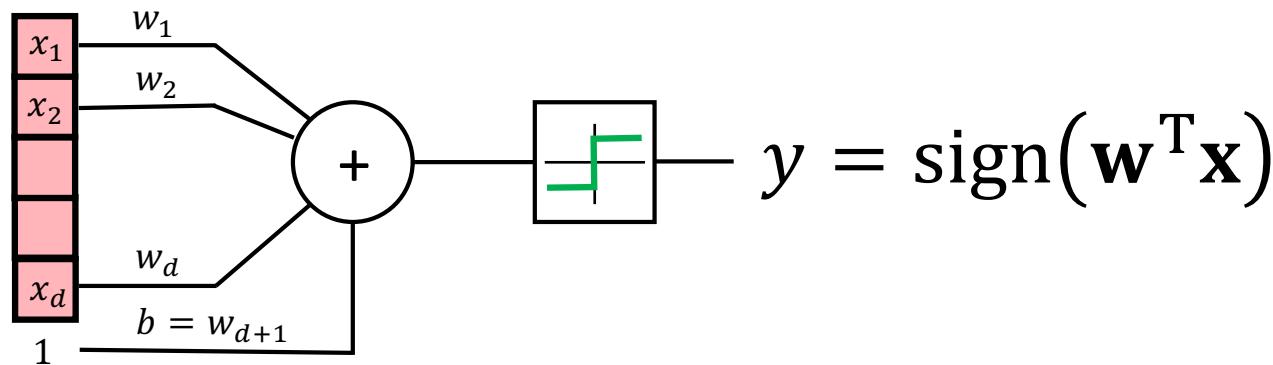
May 4, 2021

www.geometricdeeplearning.com

Supervised ML = Function Approximation



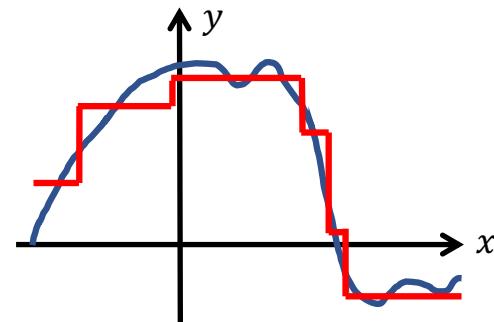
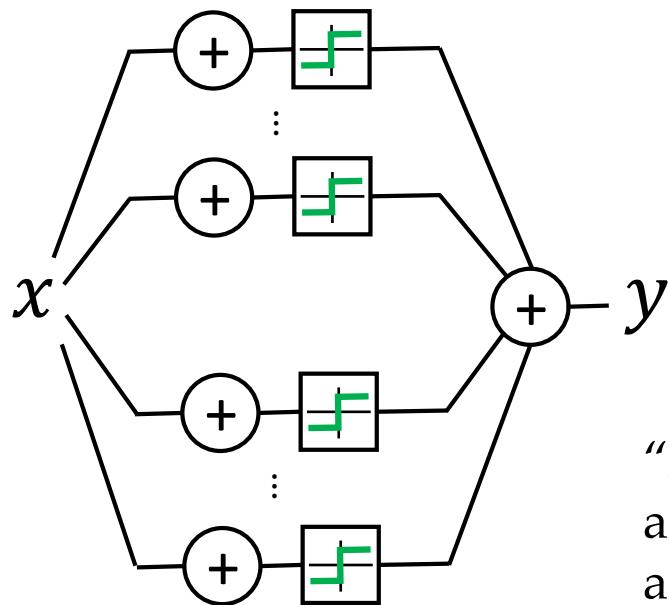
Perceptrons



Perceptron: the simplest neural network

Rosenblatt 1957

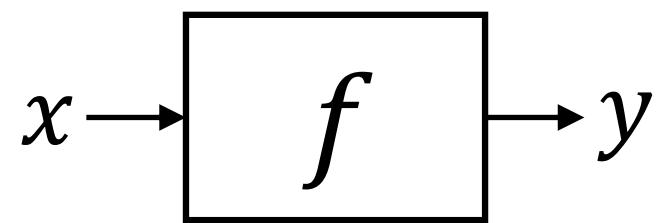
Universal Approximation



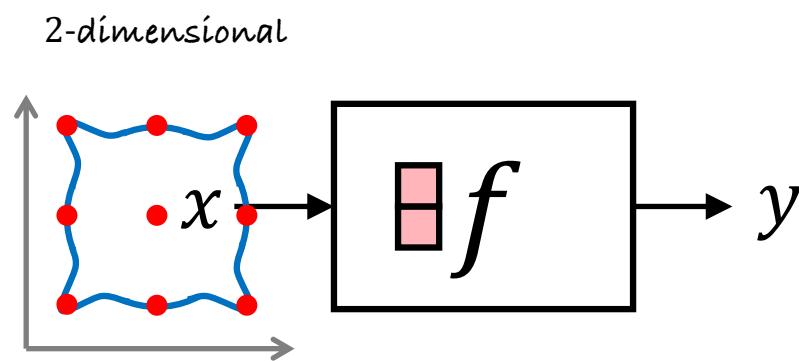
“A 2-layer perceptron can approximate a continuous function to any desired accuracy”

Cybenko 1989; Hornik 1991; Barron 1993; Leshno et al 1993; Maiorov 1999; Pinkus 1999

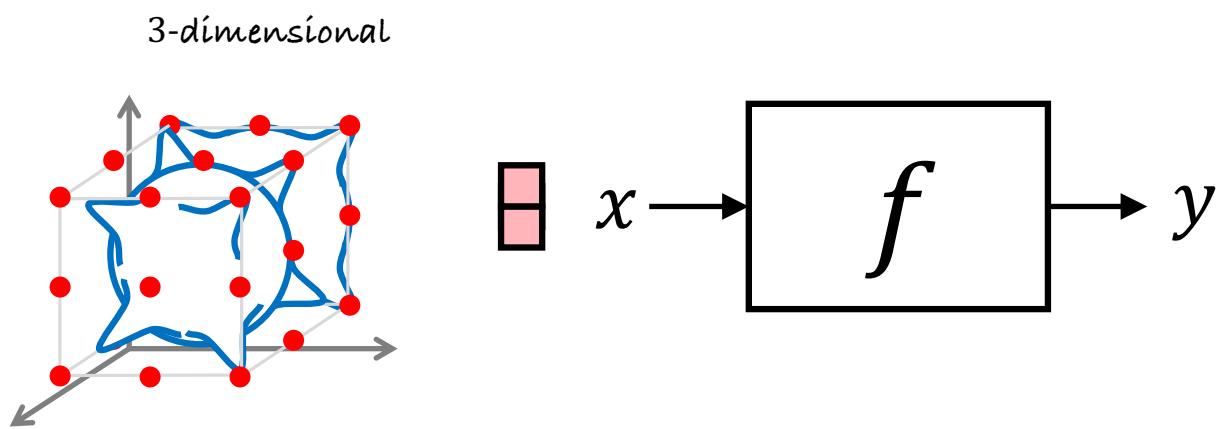
The Curse of Dimensionality



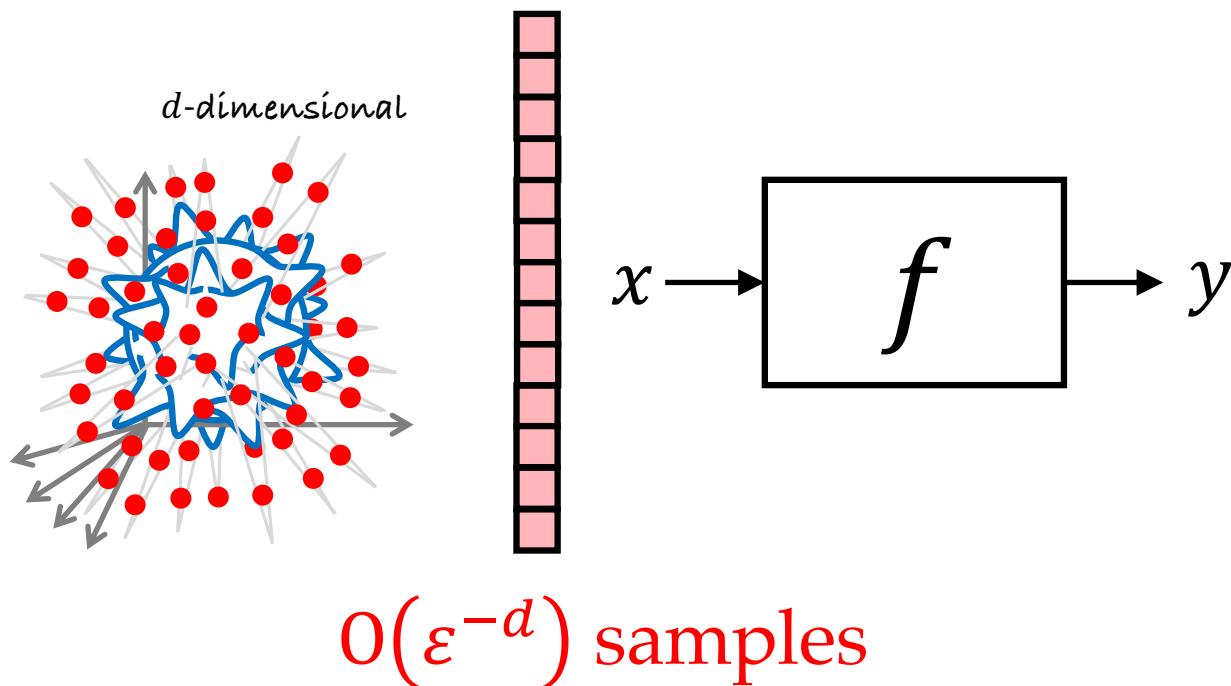
The Curse of Dimensionality



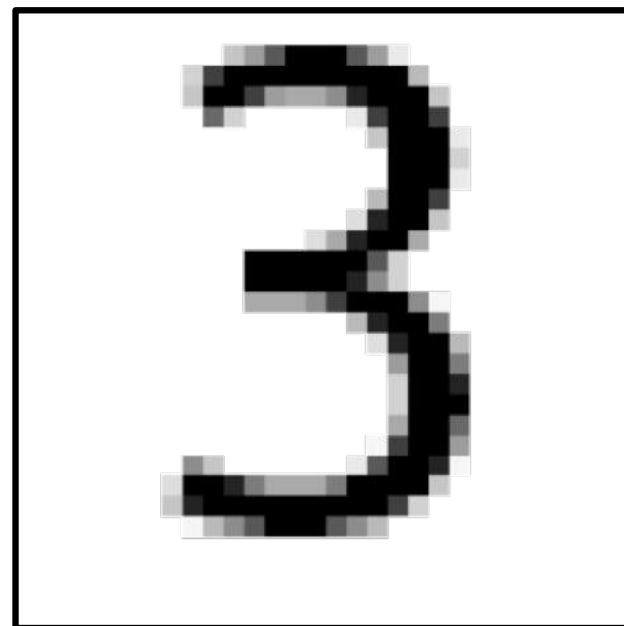
The Curse of Dimensionality



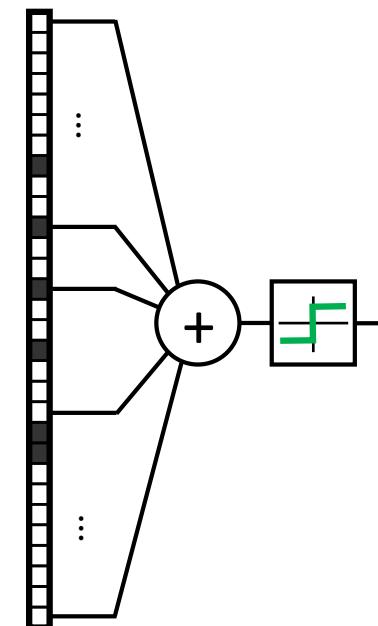
The Curse of Dimensionality



Computer Vision

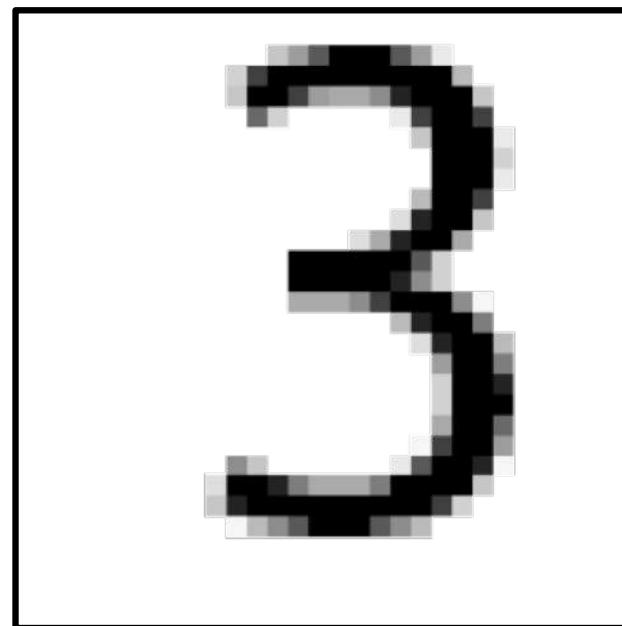


input image

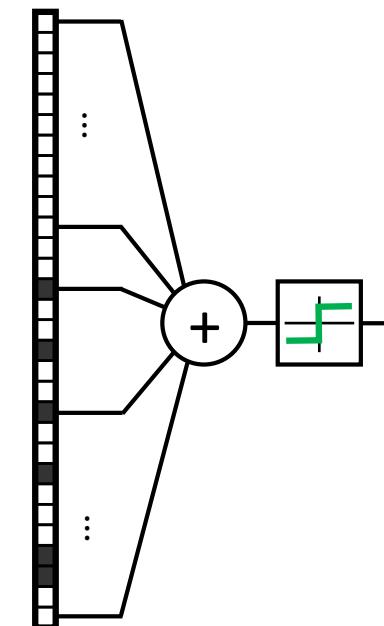


input vector

Computer Vision



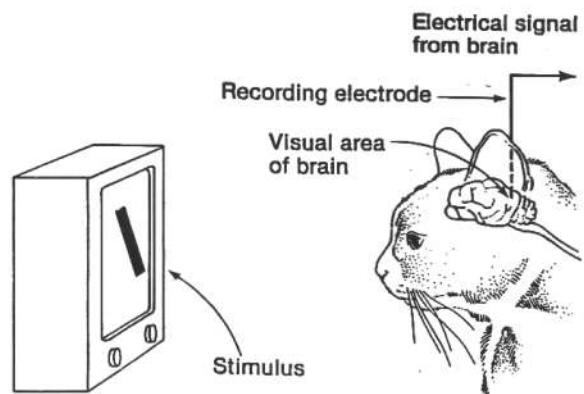
input image



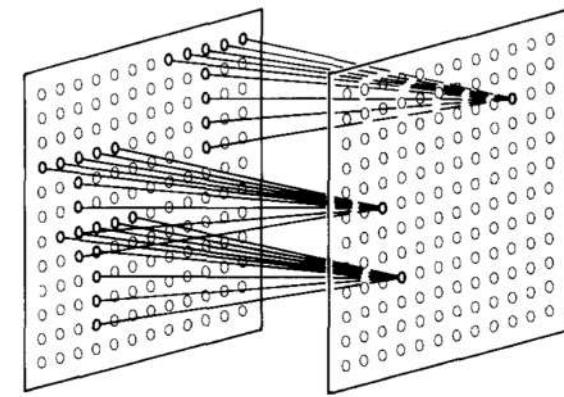
input vector

must learn shift invariance from data!

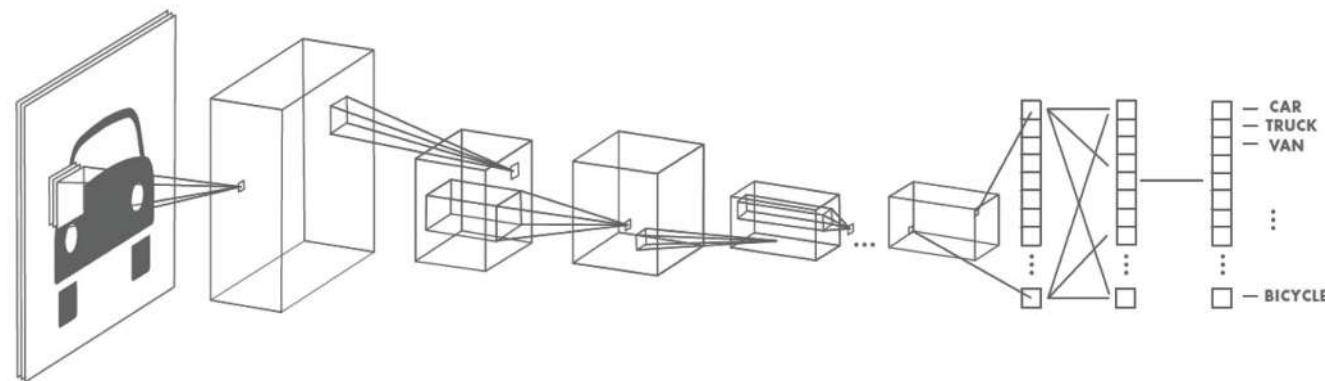
Computer Vision



Hubel, Wiesel 1962;  1981

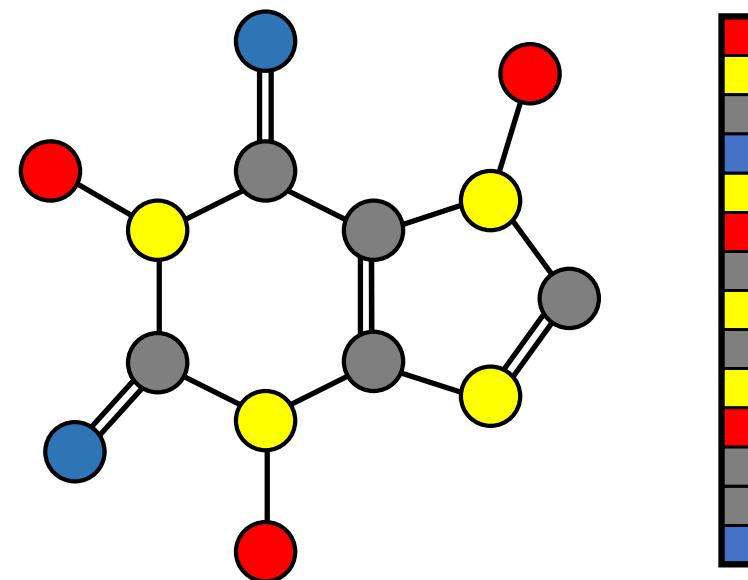


Fukushima 1980



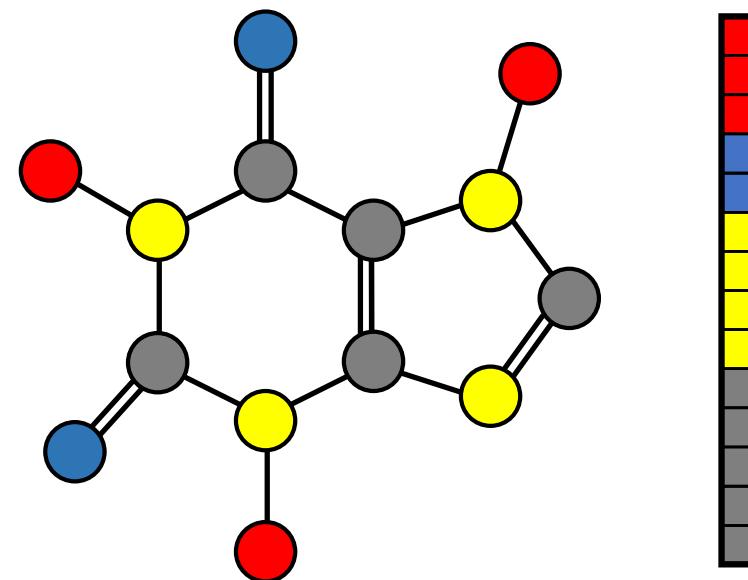
LeCun et al. 1989

Beyond Grids

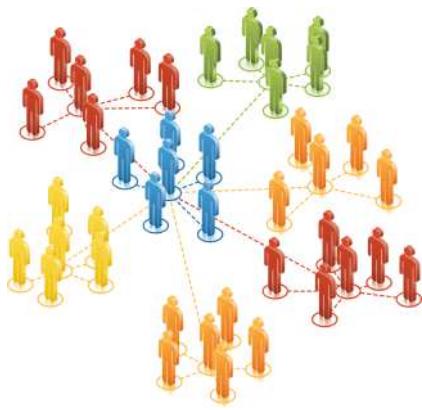


energy $U = ?$

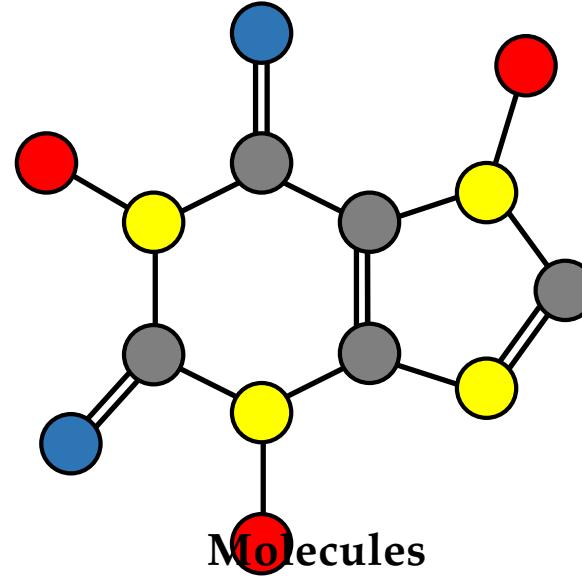
Beyond Grids



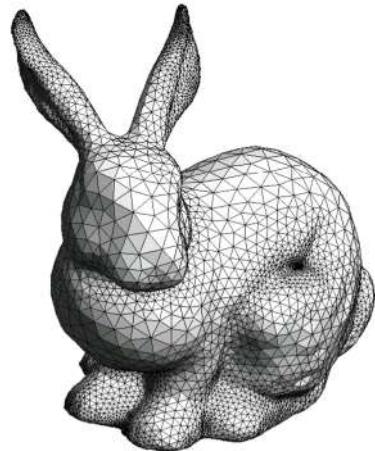
energy $U = ?$



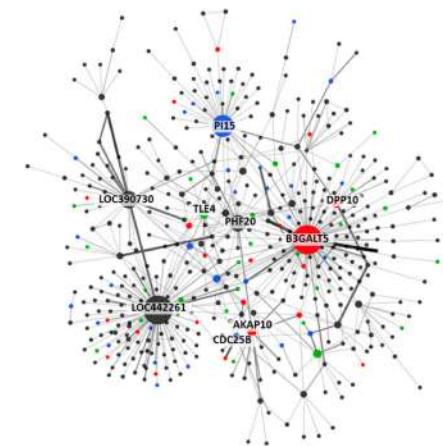
Social networks



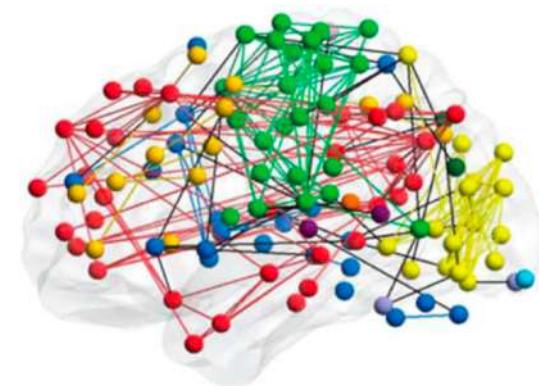
Molecules



Meshes

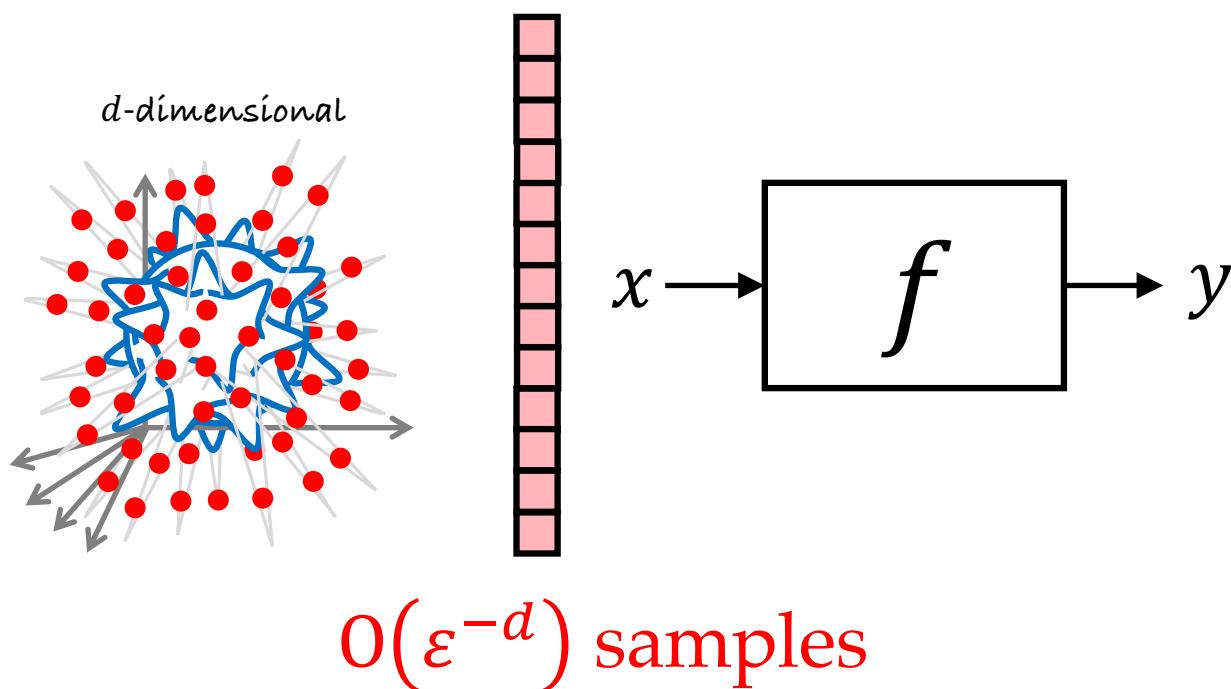


Interaction networks

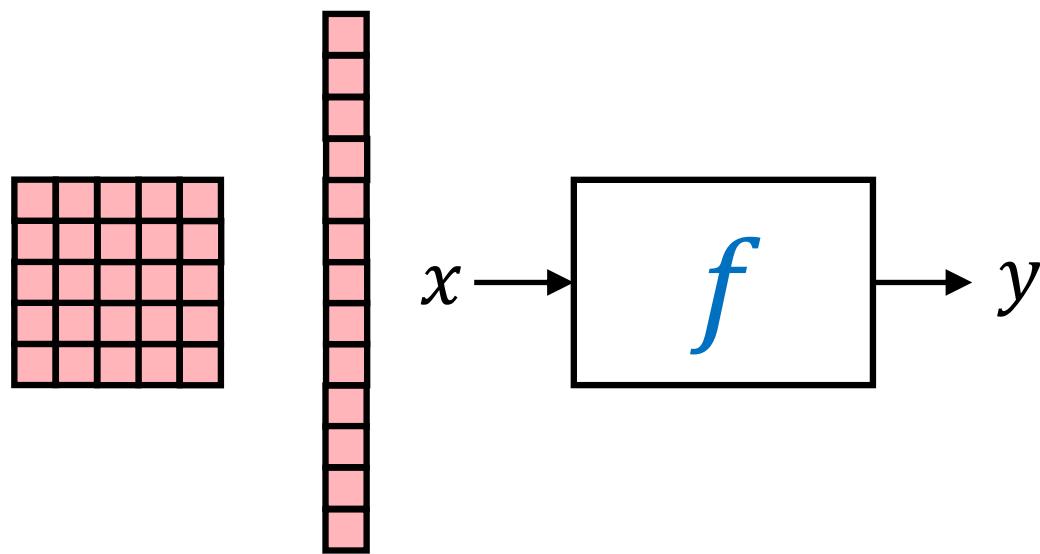


Functional networks

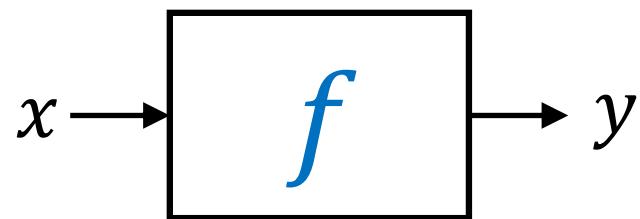
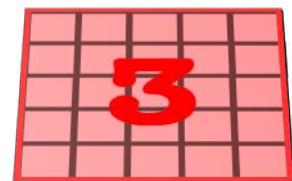
The Curse of Dimensionality



Geometric Priors

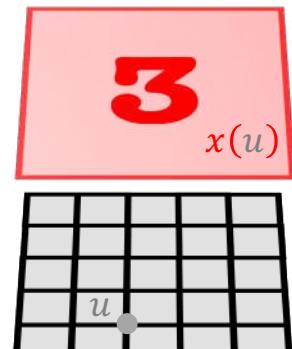


Geometric Priors

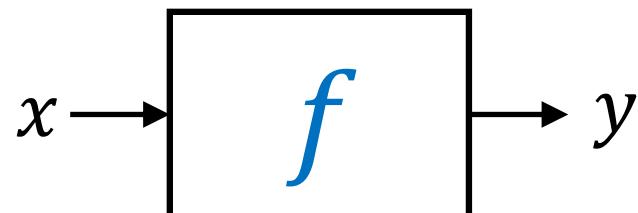


Symmetry Prior

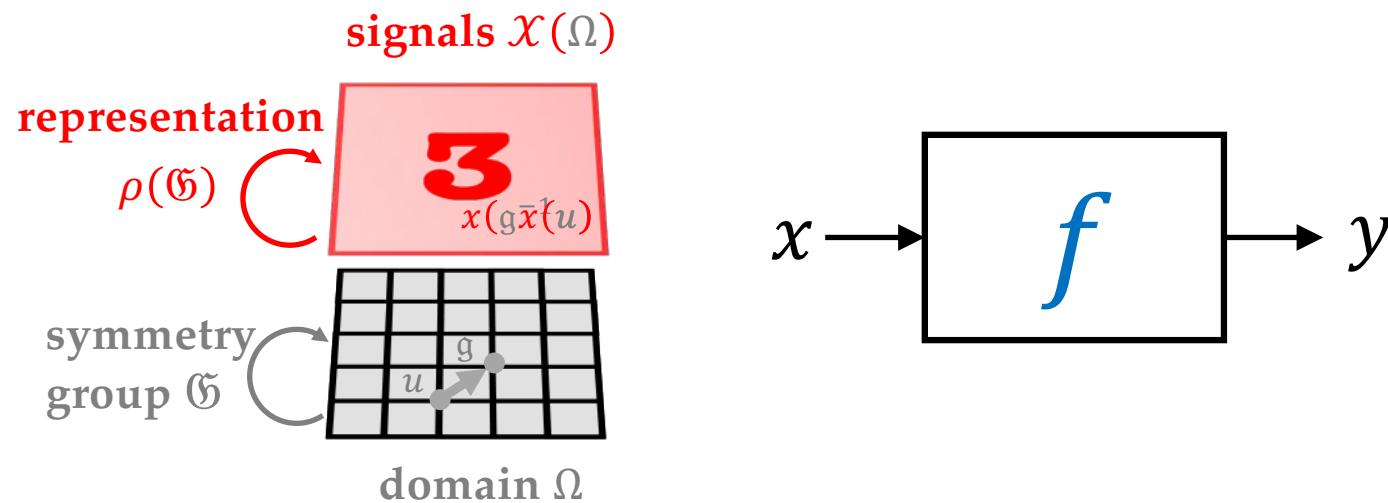
signals $\chi(\Omega)$



domain Ω



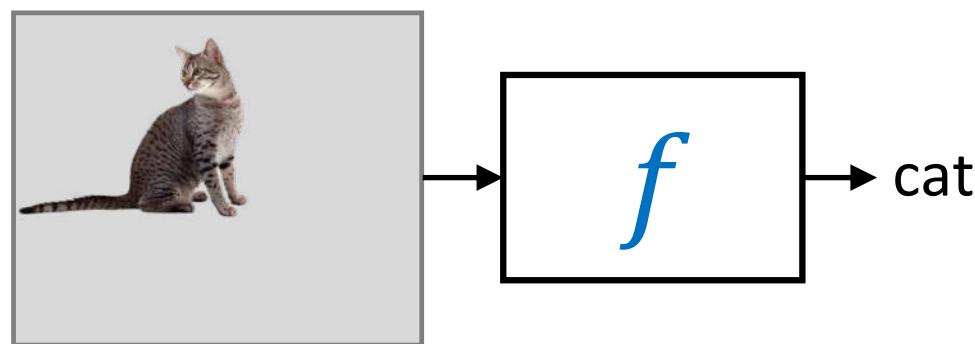
Symmetry Prior



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

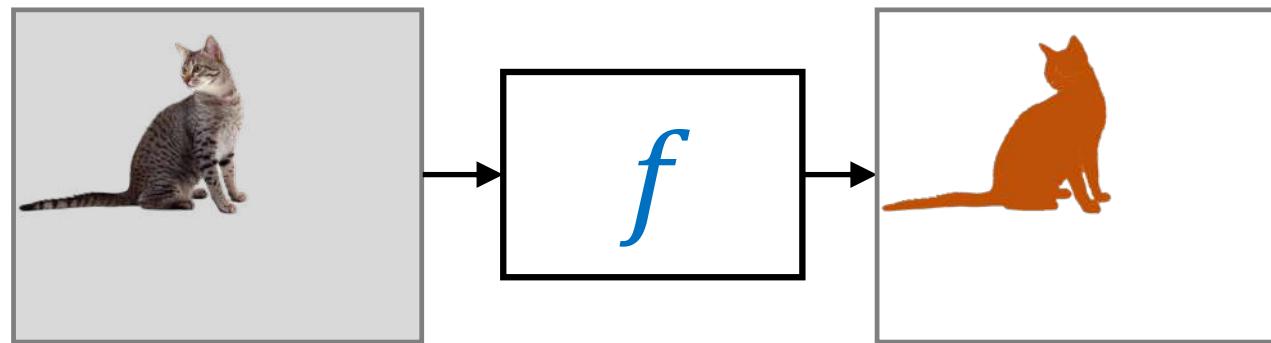
Invariant functions: Image Classification

\mathfrak{G} -invariance $f(\rho(g)x) = f(x)$

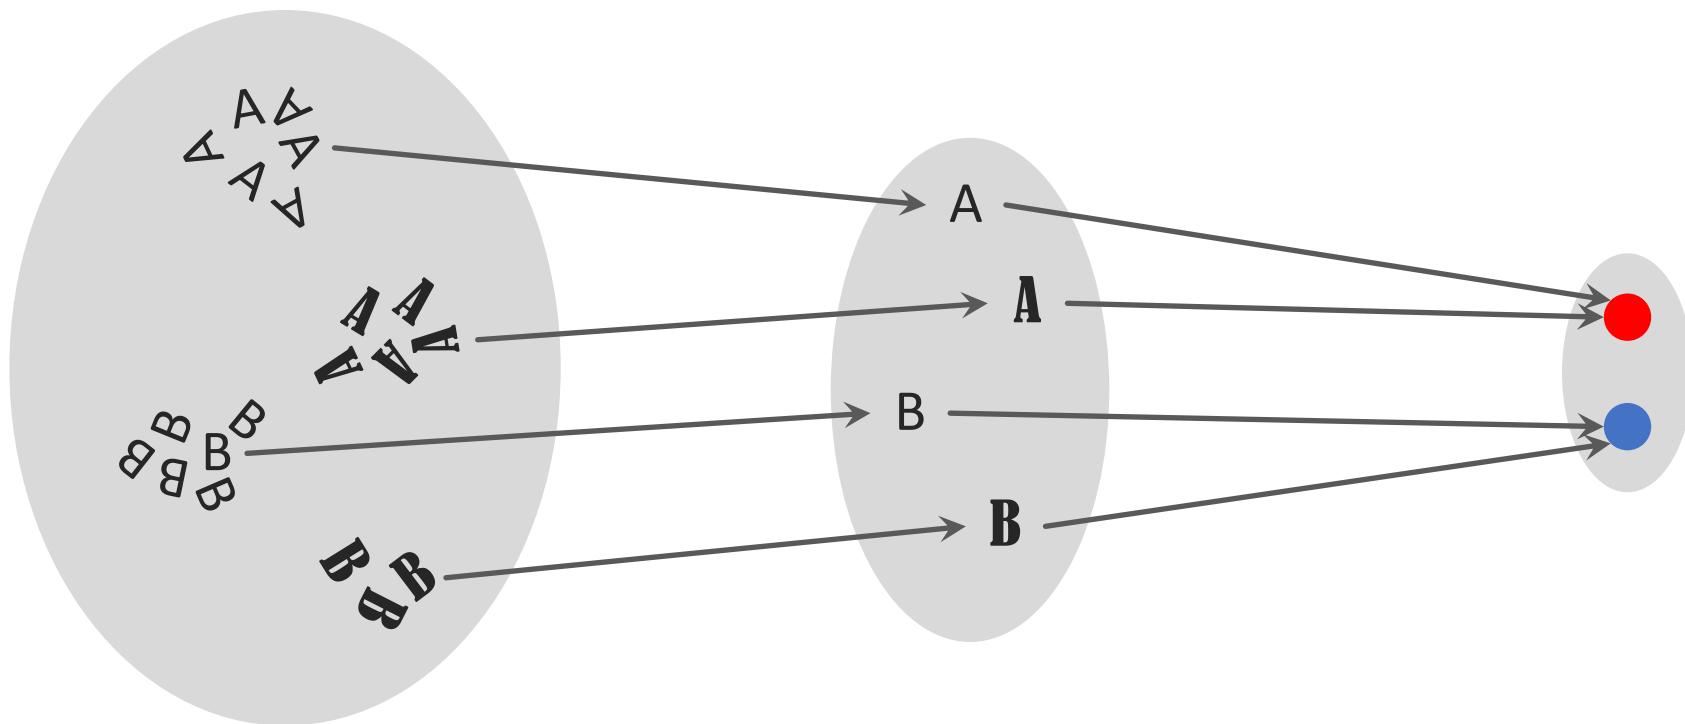


Equivariant functions: Image Segmentation

$$\mathfrak{G}\text{-equivariance } f(\rho(g)x) = \rho(g)f(x)$$



Invariant Representations



Problem with Invariant Representations

To recognise whole we must recognize parts. Relative positions of parts are important.

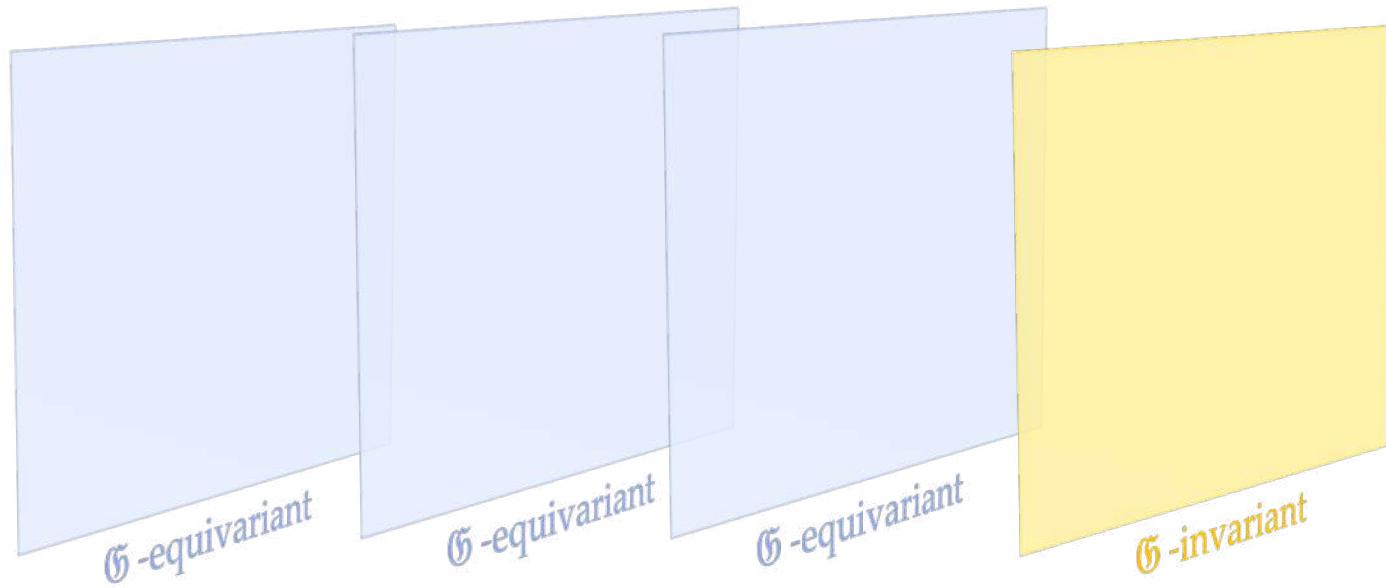
—Hinton



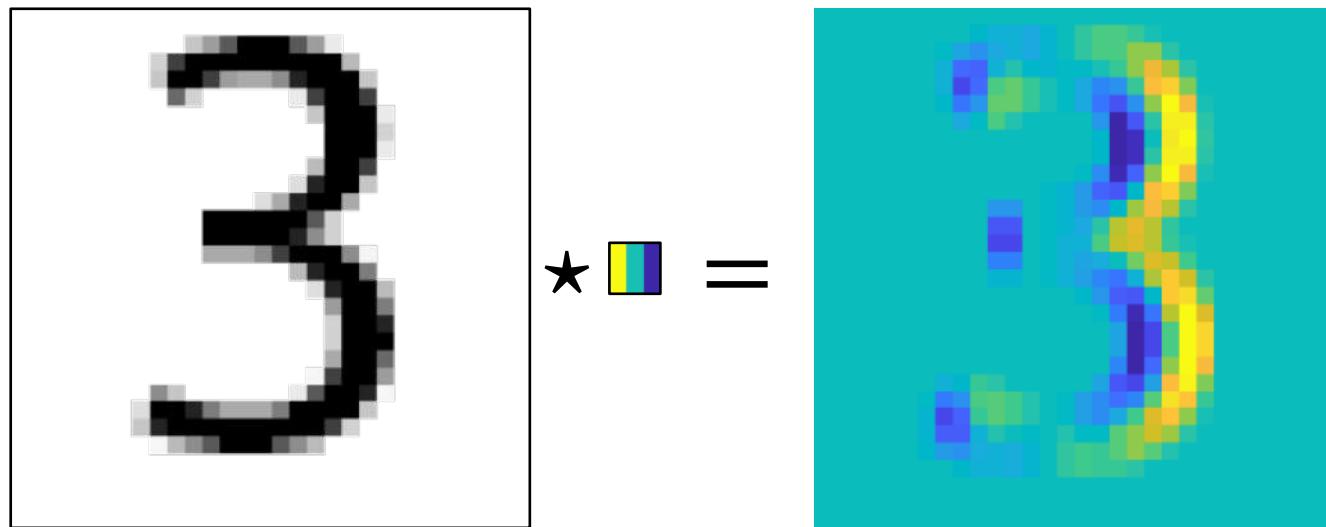
loose critical information!

Hinton et al. 2011

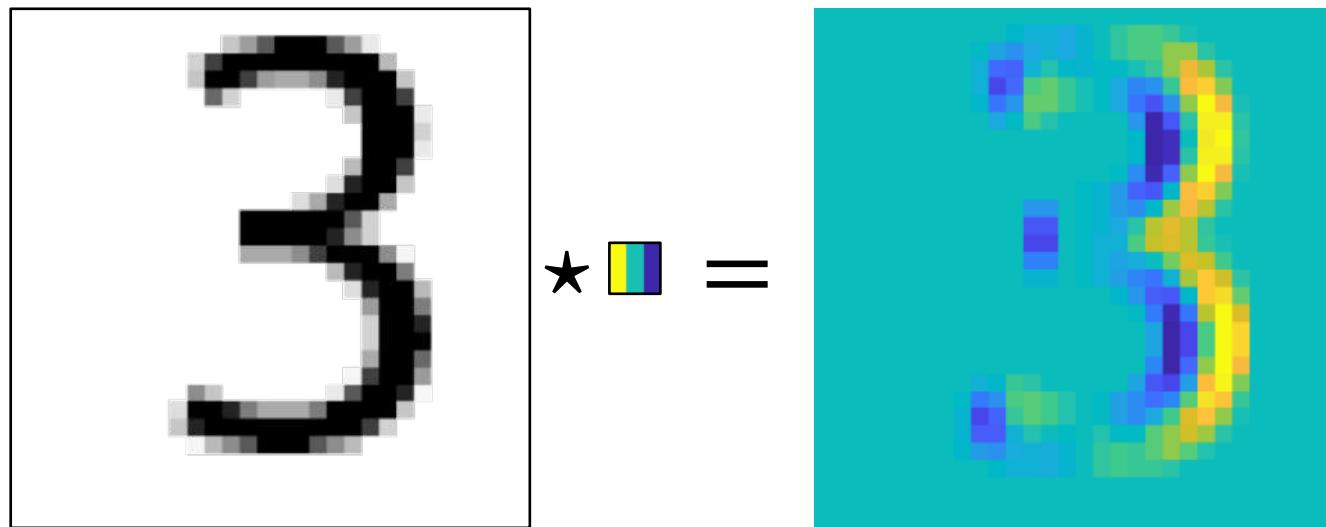
Geometric Deep Learning Blueprint



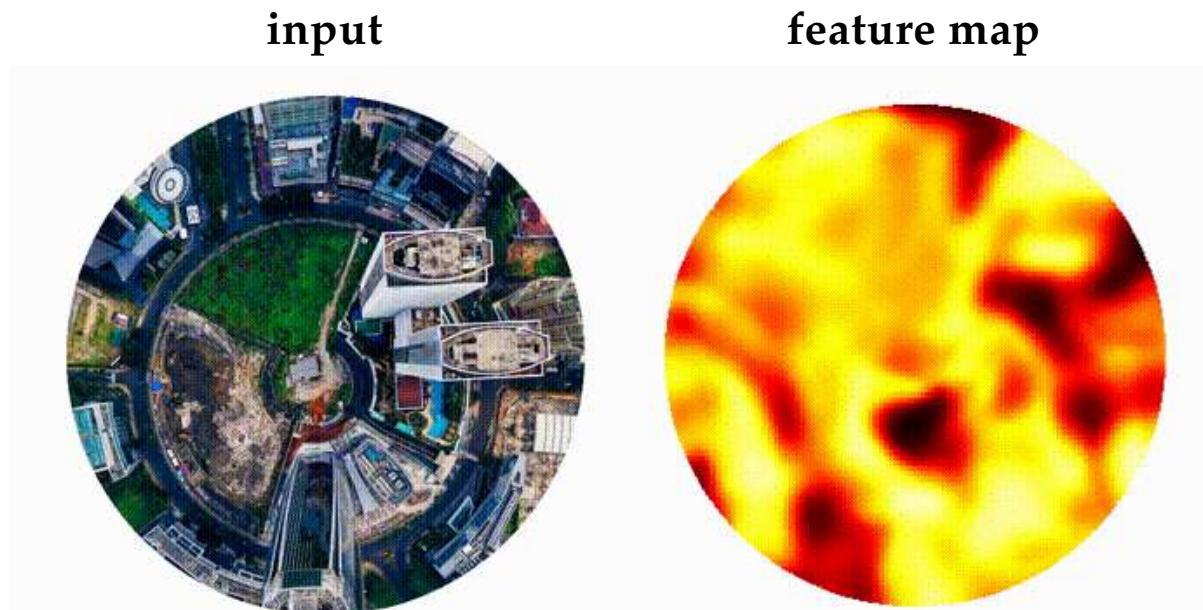
Translation Equivariance in CNNs



Translation Equivariance in CNNs

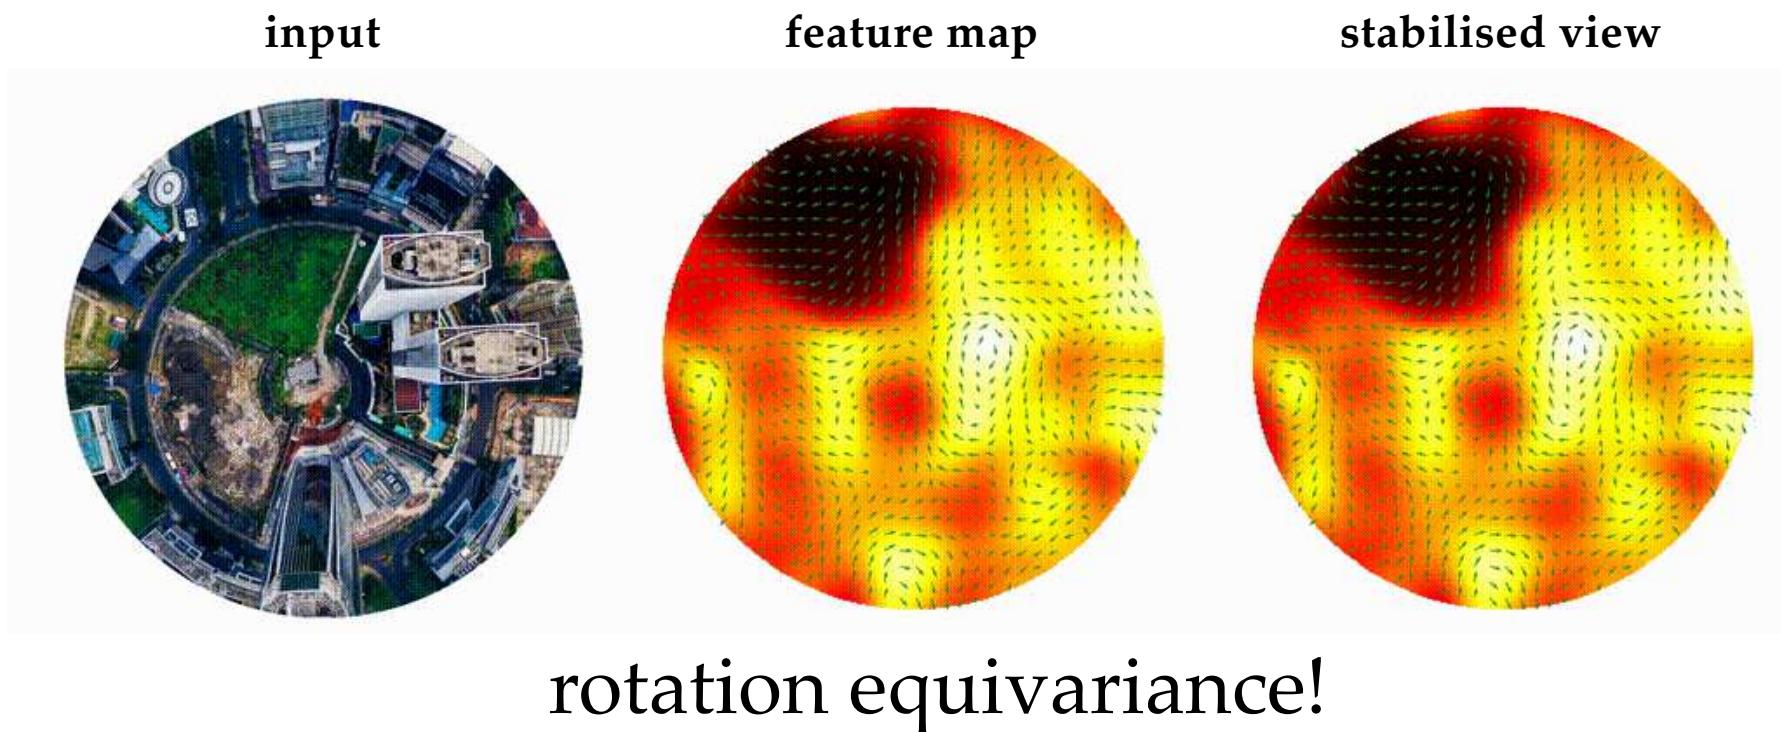


Rotation Equivariance in CNNs

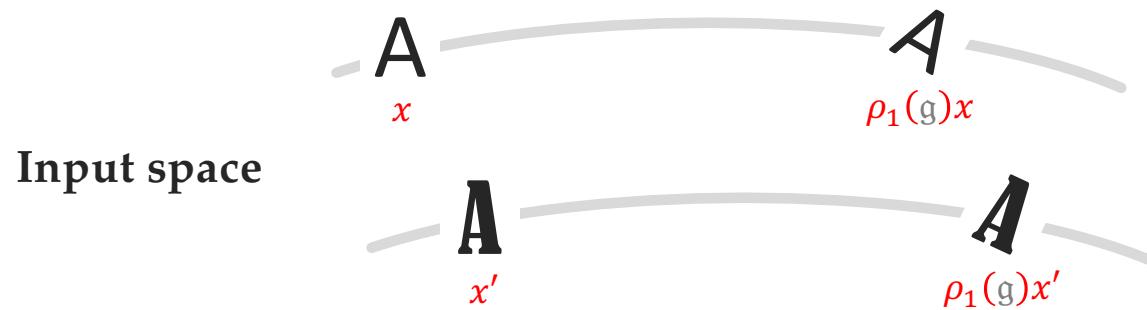


no rotation equivariance!

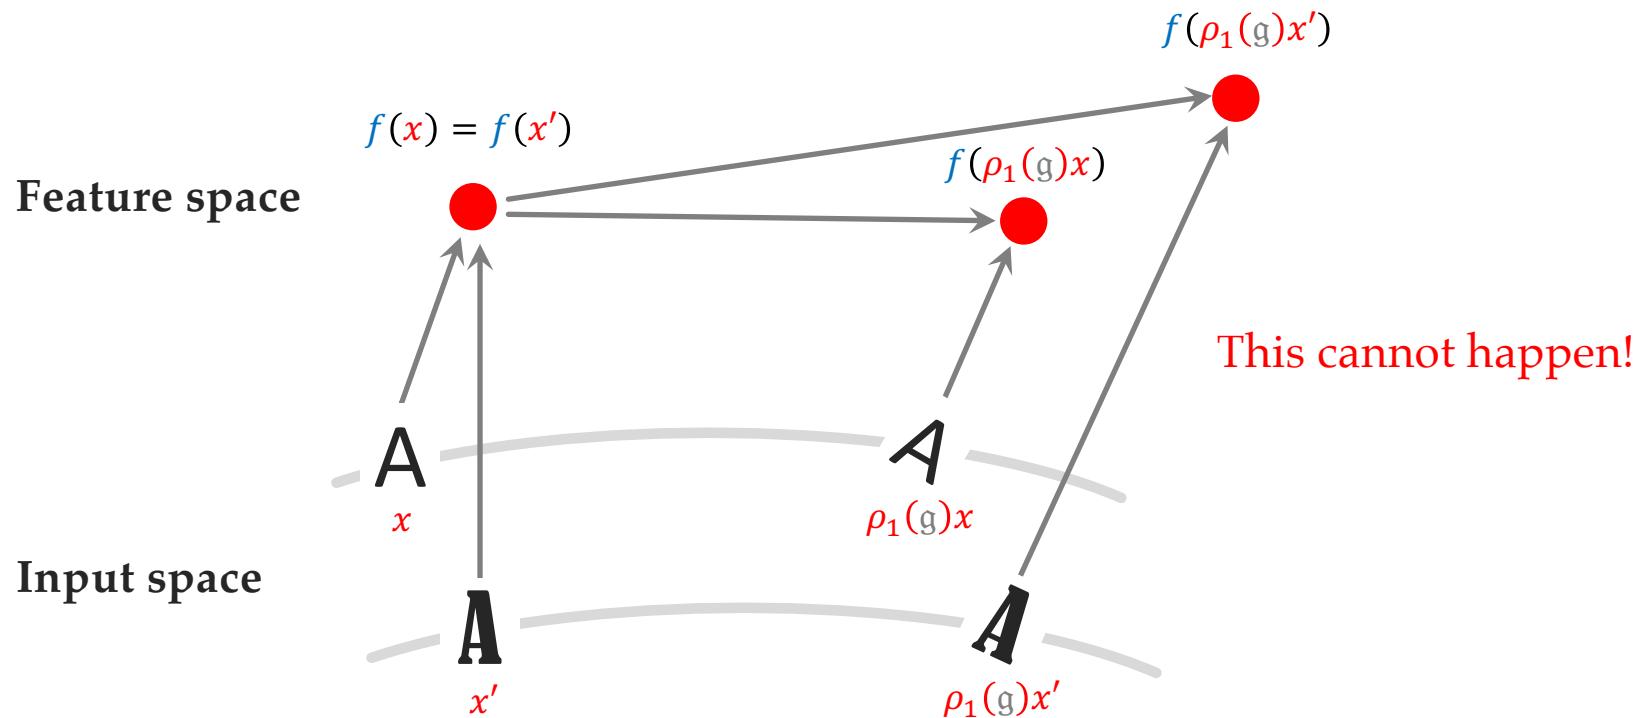
Rotation Equivariance in Group CNNs



Equivariance as Symmetry-Consistent Generalisation

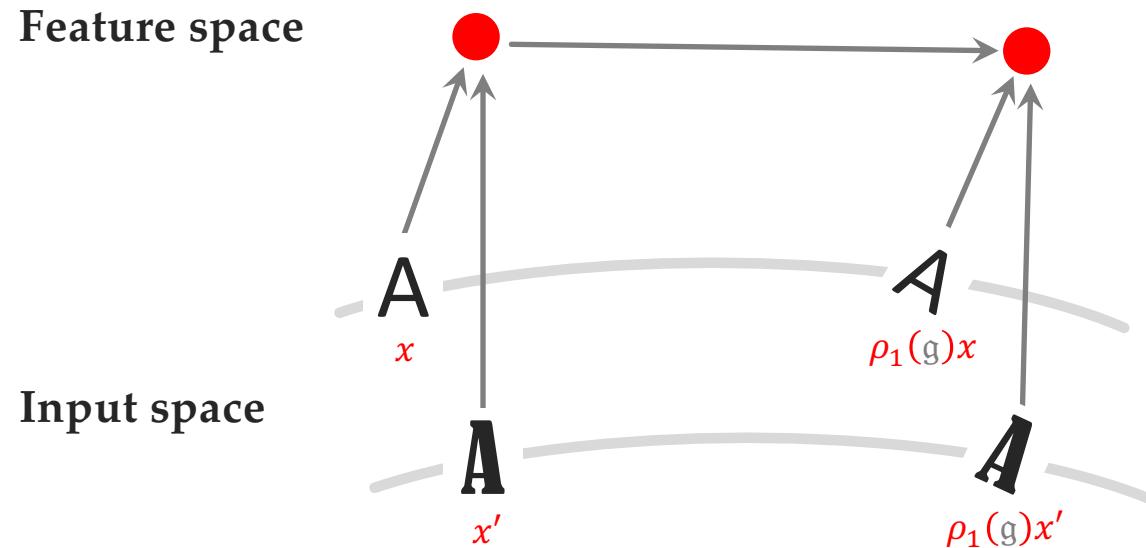


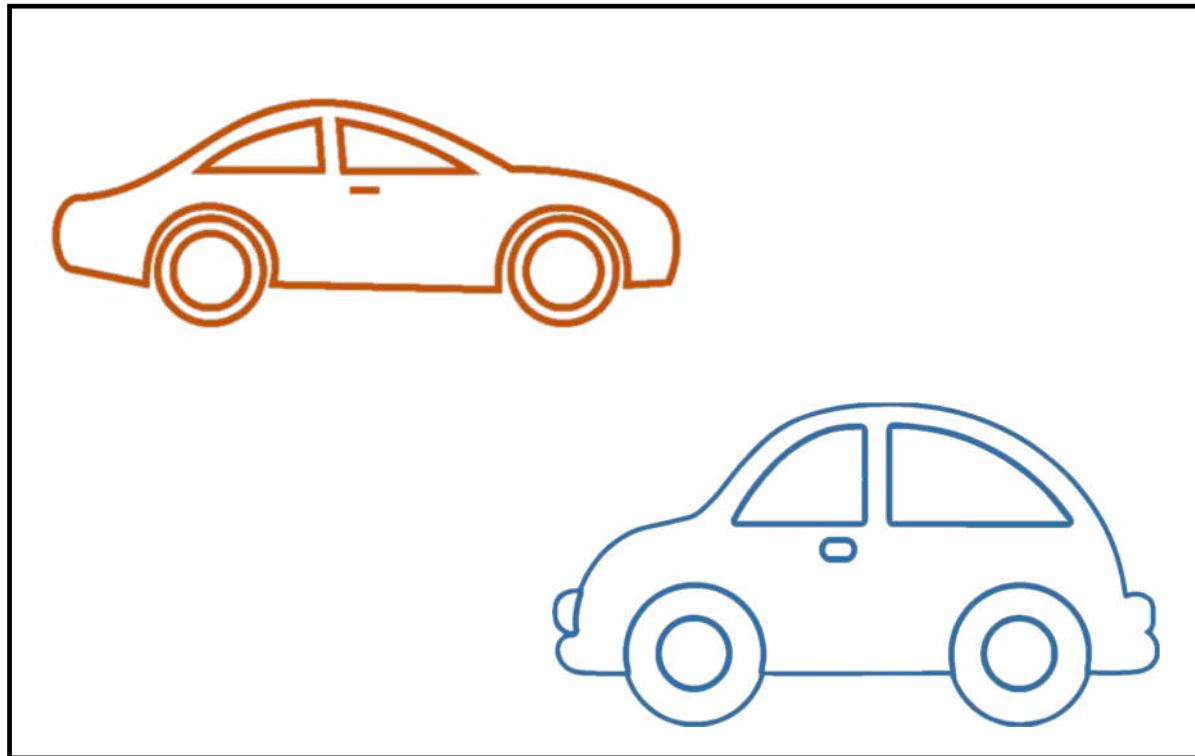
Equivariance as Symmetry-Consistent Generalisation



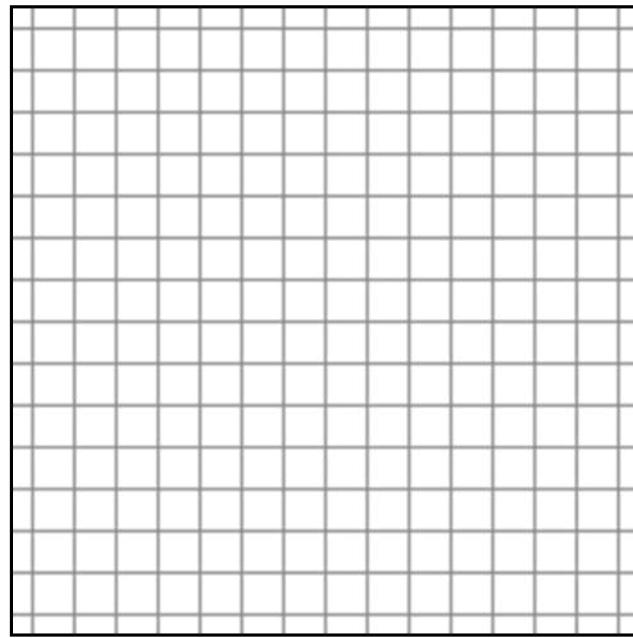
Equivariance as Symmetry-Consistent Generalisation

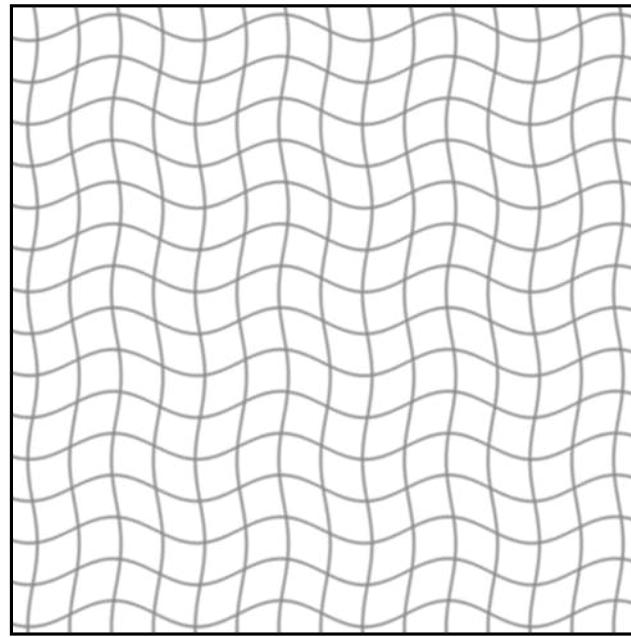
$$f(\rho_1(g)x) = \rho_2(g)f(x)$$



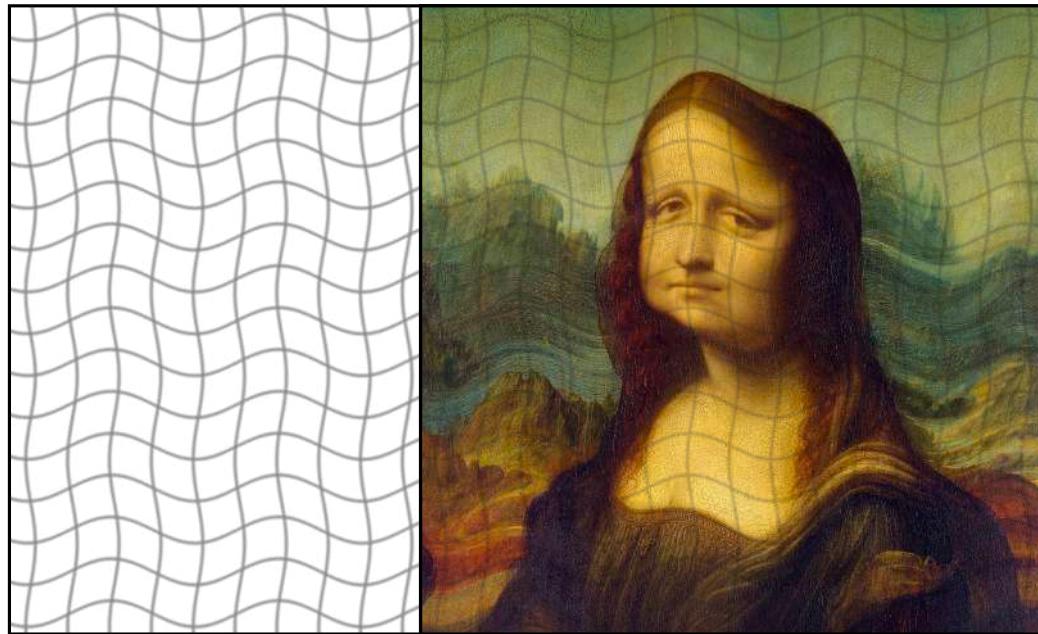


no global translation!



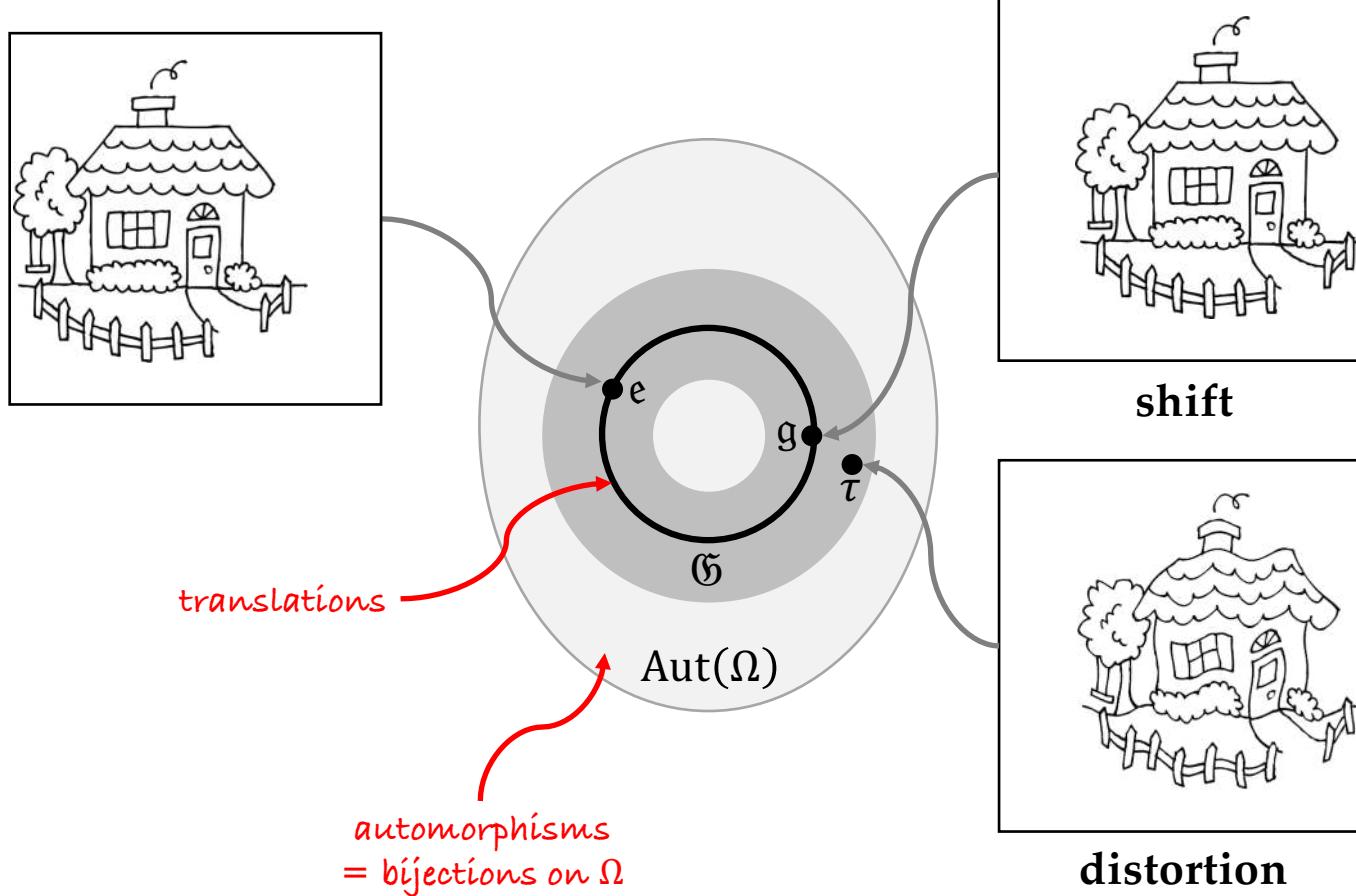


$$\tau: \Omega \rightarrow \Omega$$



$$\tau: \Omega \rightarrow (\underline{L}_\tau x)(u) = x(\tau^{-1}u)$$

Geometric Stability



Geometric Stability

“deformation insensitivity”

$$\|f(L_\tau \mathbf{x}) - f(\mathbf{x})\| \leq c \|\nabla \tau\| \cdot \|\mathbf{x}\|$$


stability under
deformation


“Dirichlet energy”

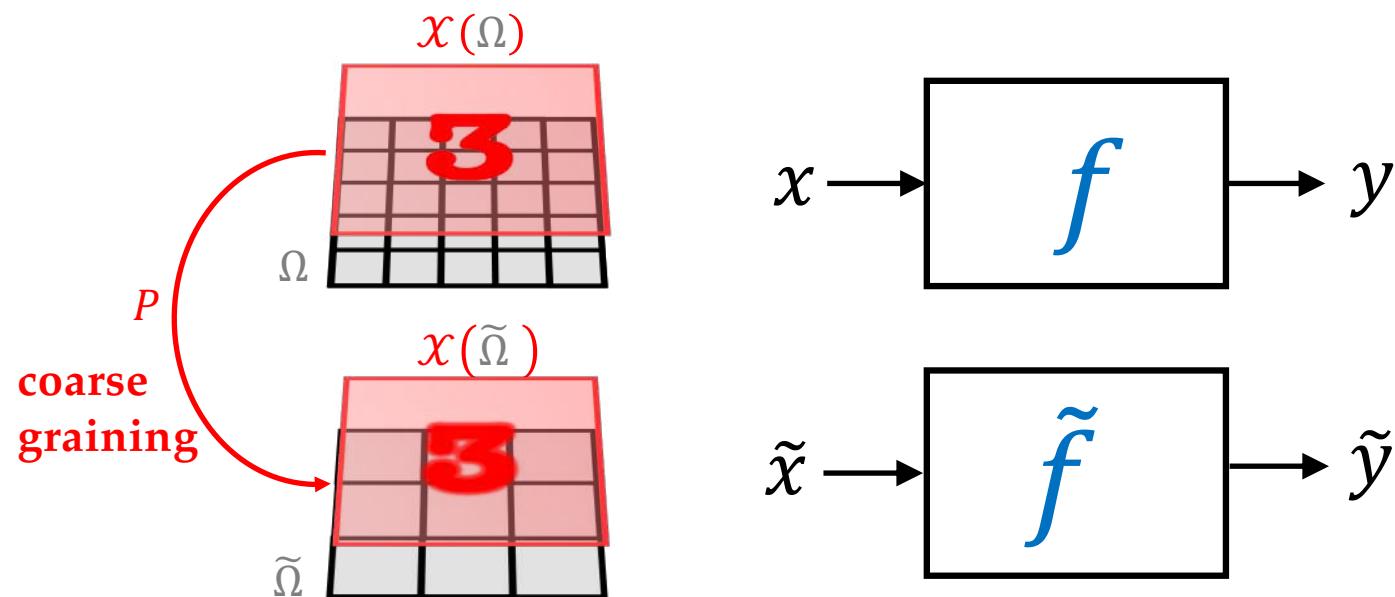
Geometric Stability

generalization of shift invariance

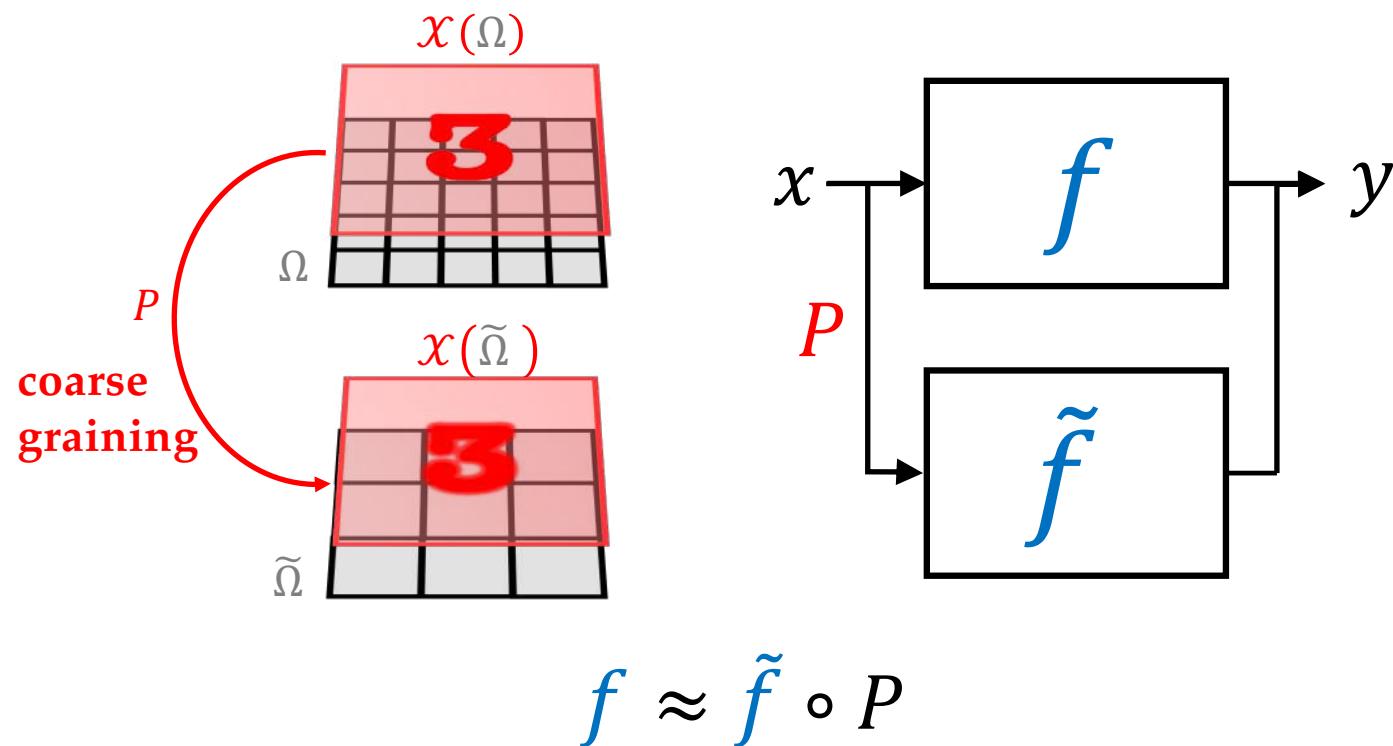
$$\|f(T_v \mathbf{x}) - f(\mathbf{x})\| = 0$$

shift $\tau(u) = u - v$
 $\nabla \tau = 0$

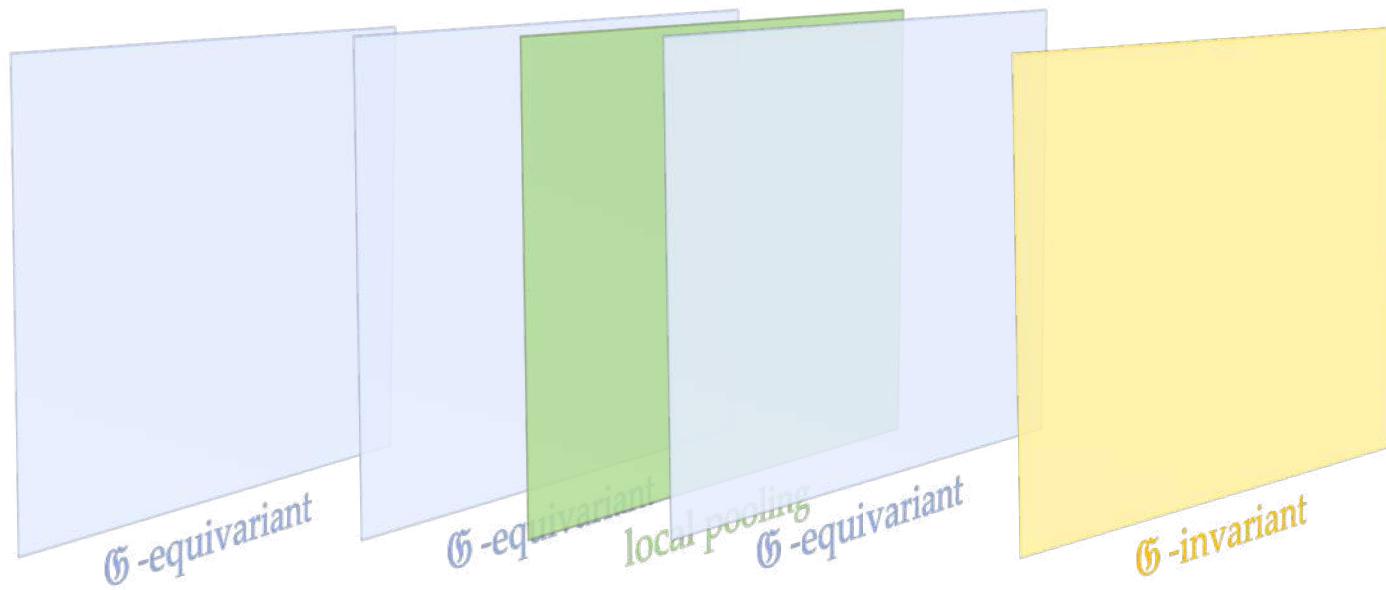
Scale Separation Prior



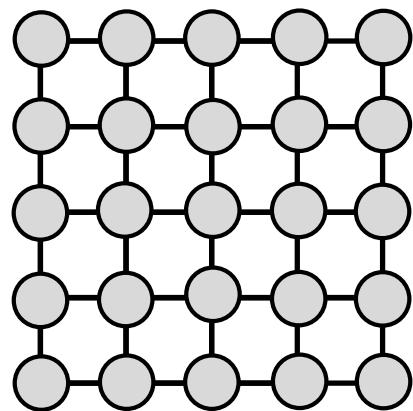
Scale Separation Prior



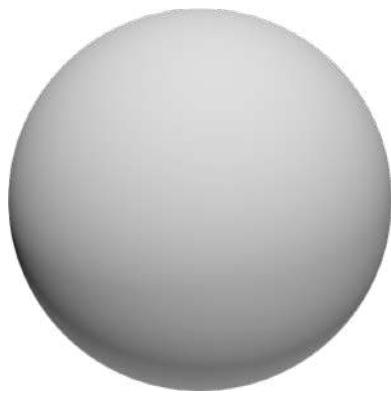
Geometric Deep Learning Blueprint



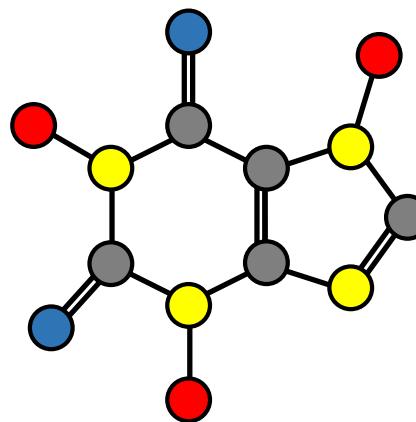
The “5G” of Geometric Deep Learning



Grids



Groups

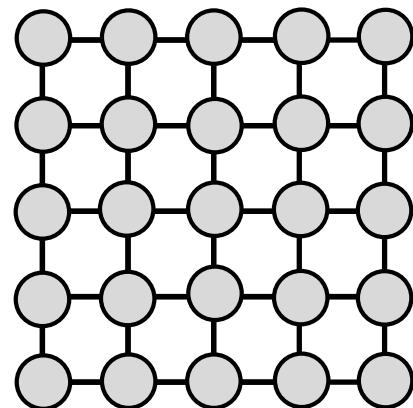


Graphs

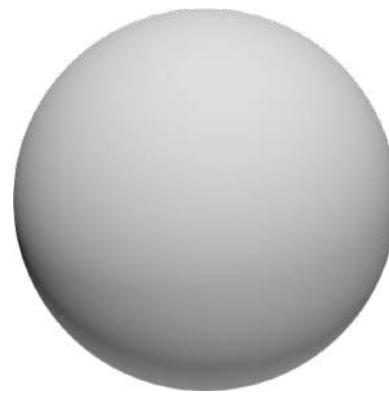


Geodesics &
Gauges

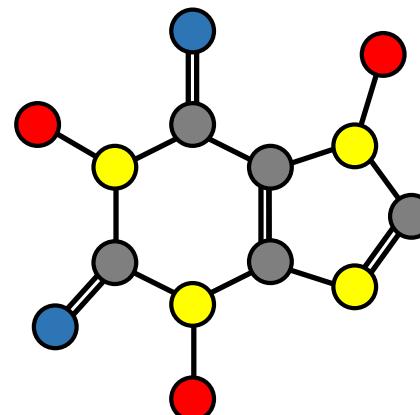
The “5G” of Geometric Deep Learning



Images &
Sequences



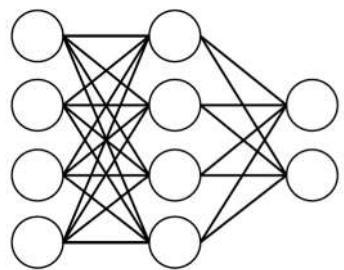
Homogeneous
spaces



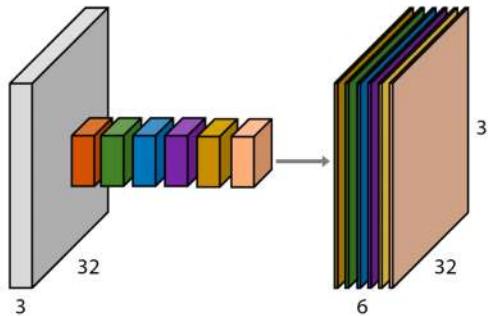
Graphs & Sets



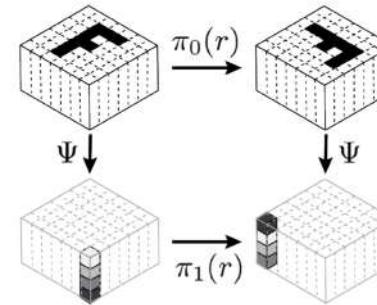
Manifolds, Meshes &
Geometric graphs



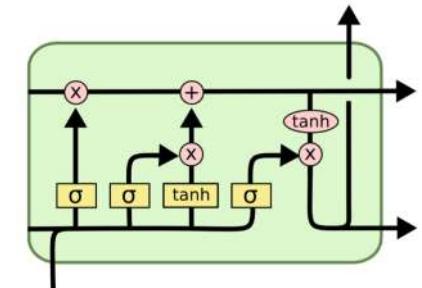
Perceptrons
Function regularity



CNNs
Translation



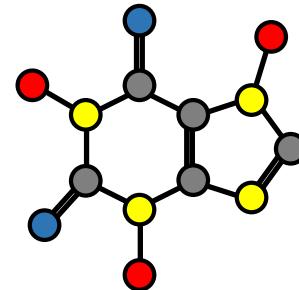
Group-CNNs
Translation+Rotation,
Global groups



LSTMs
Time warping



DeepSets / Transformers
Permutation



GNNs
Permutation



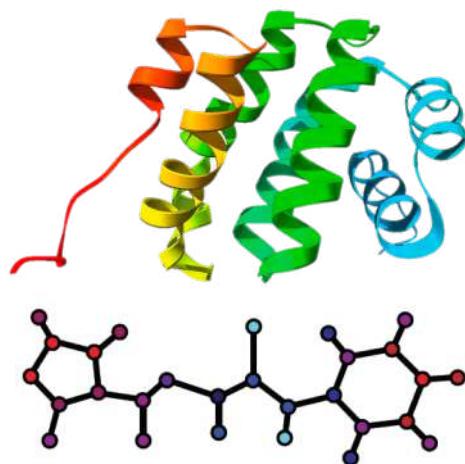
Intrinsic CNNs
Isometry / Gauge choice

Popular architectures as instances of GDL Blueprint

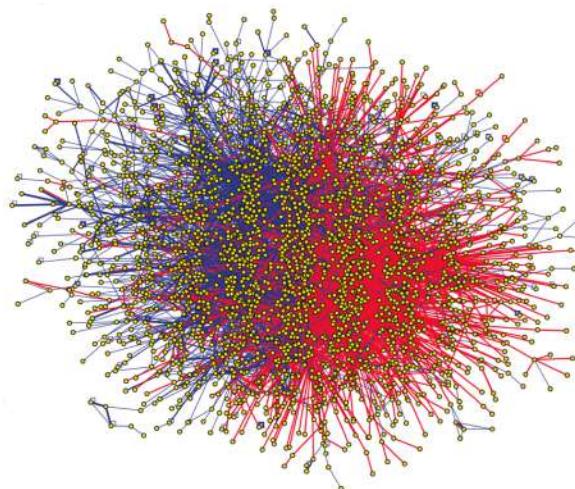
Architecture	Domain Ω	Symmetry Group \mathfrak{G}
<i>CNN</i>	Grid	Translation
<i>Spherical CNN</i>	Sphere / $SO(3)$	Rotation $SO(3)$
<i>Intrinsic / Mesh CNN</i>	Manifold	Isometry $Iso(\Omega)$ / Gauge Symmetry $SO(2)$
<i>GNN</i>	Graph	Permutation Σ_n
<i>Deep Sets</i>	Set	Permutation Σ_n
<i>Transformer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

GRAPHS

Graphs = Systems of Relations and Interactions



Molecules

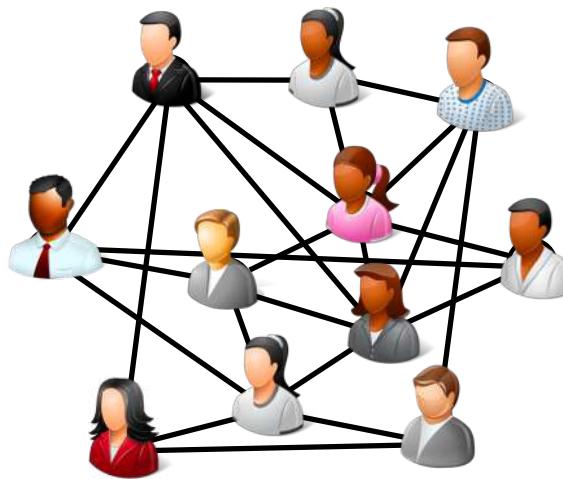


Interactomes



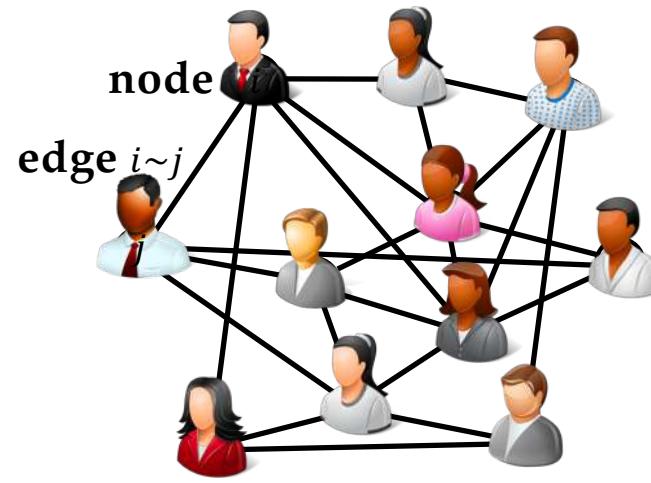
Social networks

Graphs



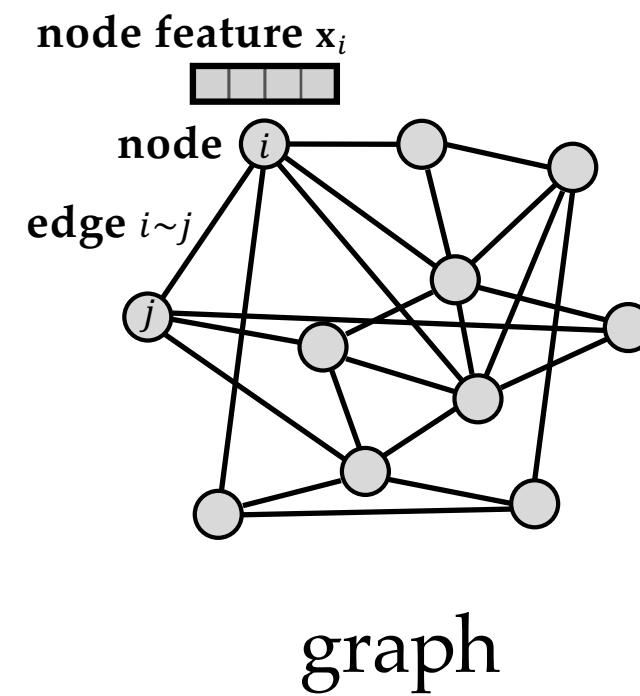
social network

Graphs

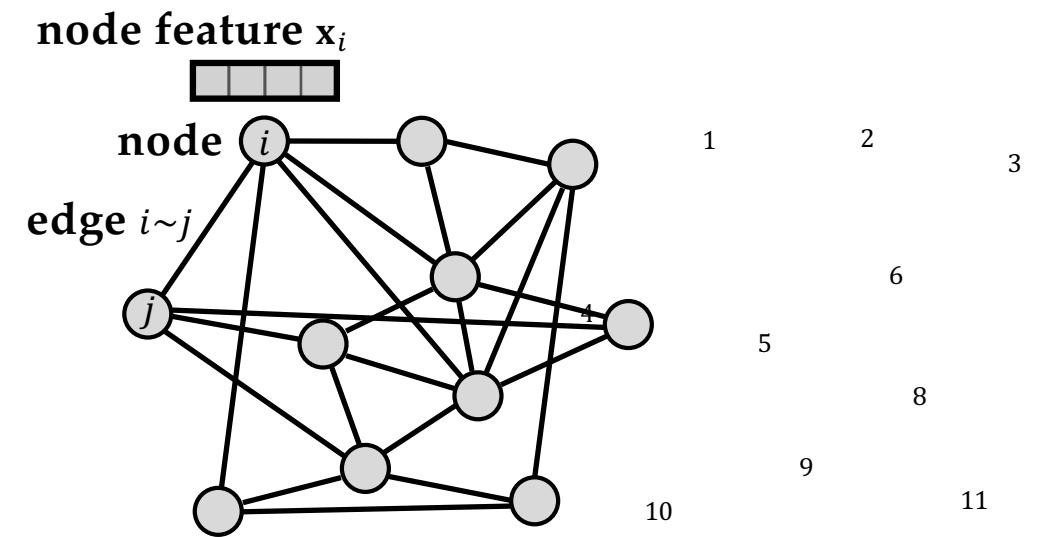


social network

Graphs



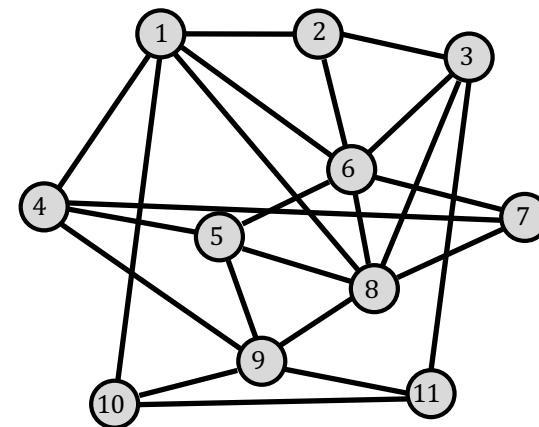
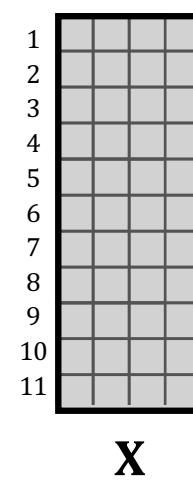
Key Structural Properties of Graphs



arbitrary graph ordering of nodes

Key Structural Properties of Graphs

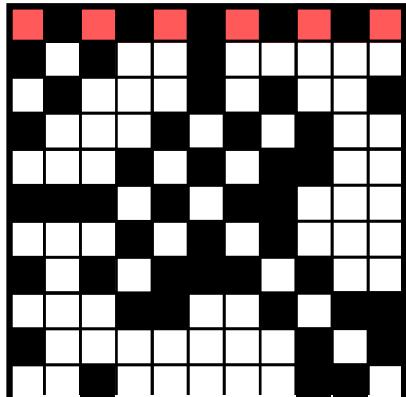
Feature
matrix $n \times d$



arbitrary ordering of nodes

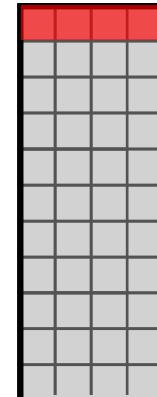
Key Structural Properties of Graphs

Adjacency
matrix $n \times n$

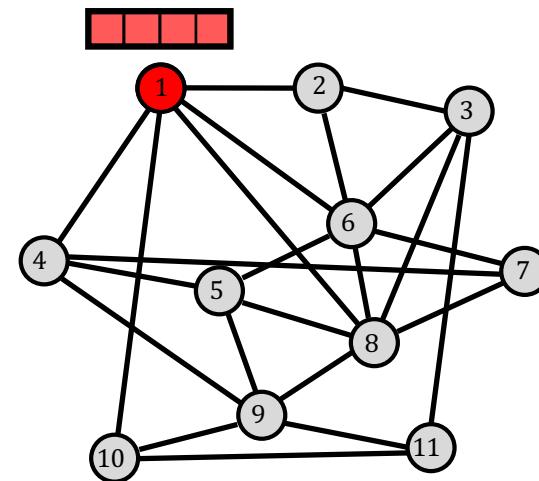


A

Feature
matrix $n \times d$



X



arbitrary ordering of nodes

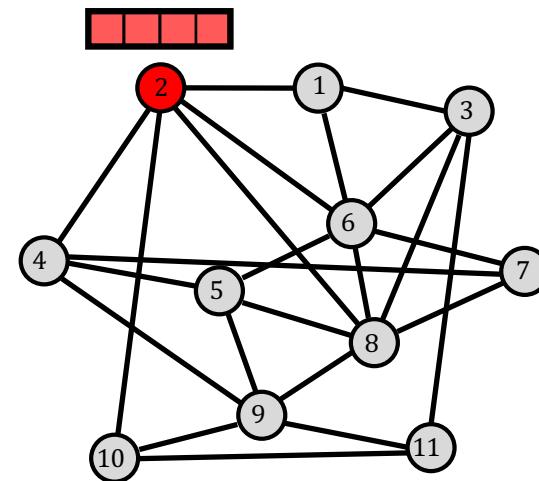
Key Structural Properties of Graphs

Adjacency
matrix $n \times n$

$\mathbf{P} \mathbf{A} \mathbf{P}^T$

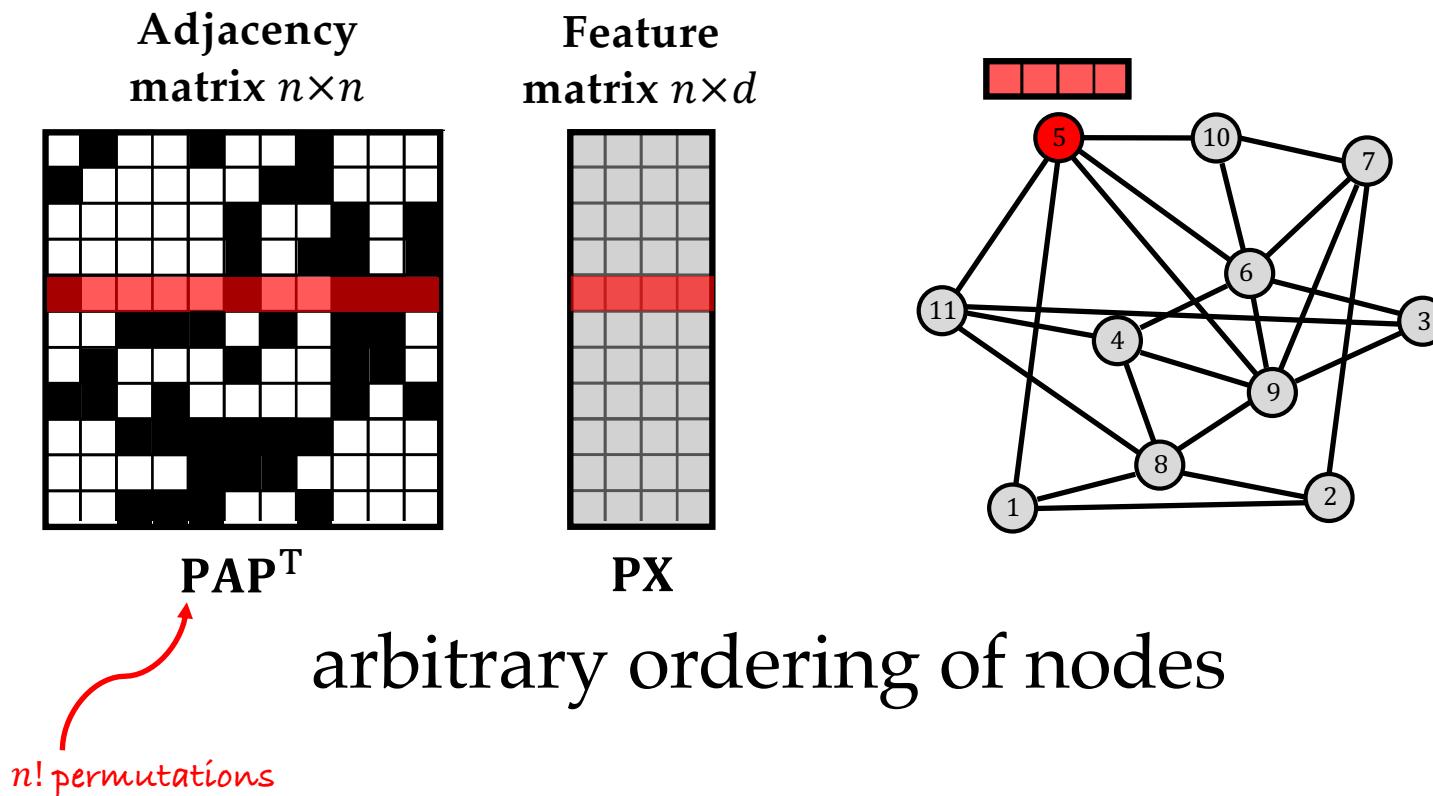
Feature
matrix $n \times d$

$\mathbf{P} \mathbf{X}$



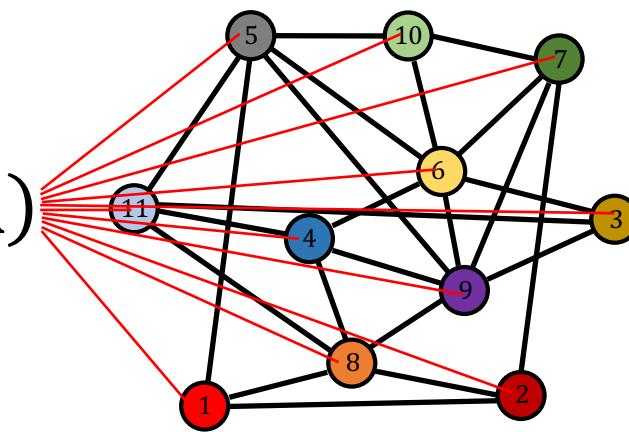
arbitrary ordering of nodes

Key Structural Properties of Graphs



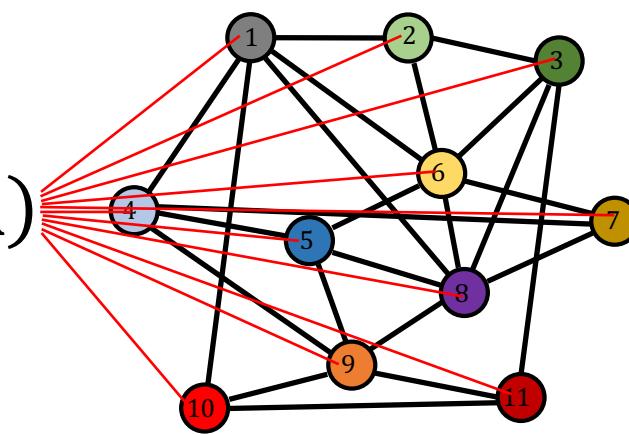
Invariant Graph Functions

graph function $f(\mathbf{X}, \mathbf{A})$



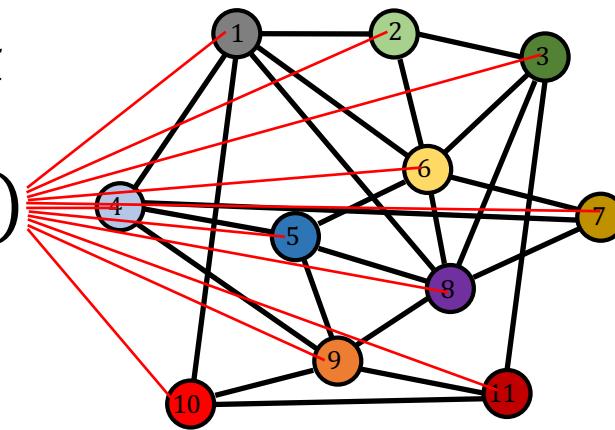
Invariant Graph Functions

graph function $f(\mathbf{X}, \mathbf{A})$



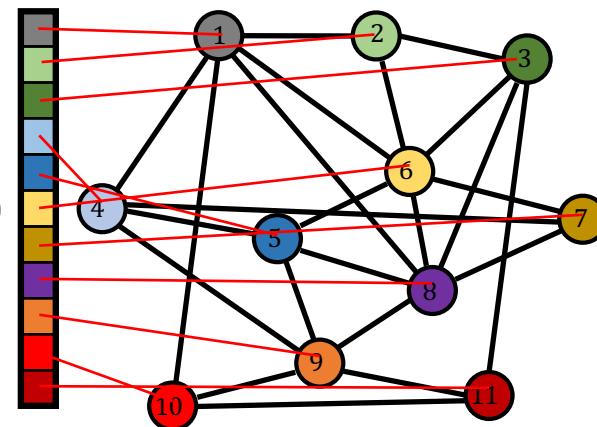
Invariant Graph Functions

permutation-invariant
 $f(\mathbf{P}\mathbf{X}, \mathbf{P}\mathbf{A}\mathbf{P}^T) = f(\mathbf{X}, \mathbf{A})$



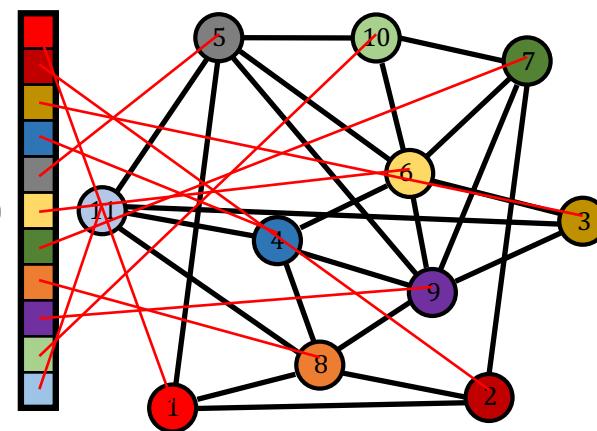
Equivariant Graph Functions

node function $F(X, A)$



Equivariant Graph Functions

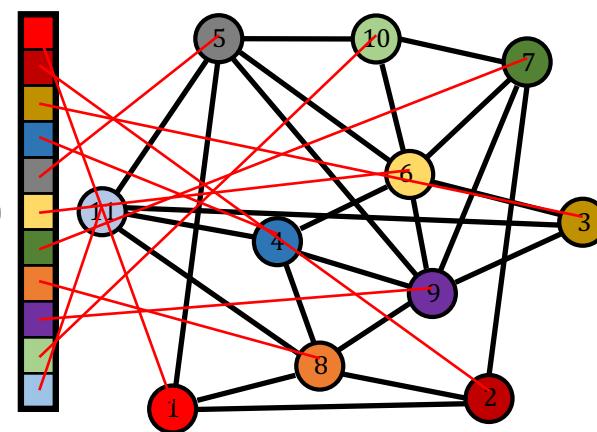
node function $F(X, A)$



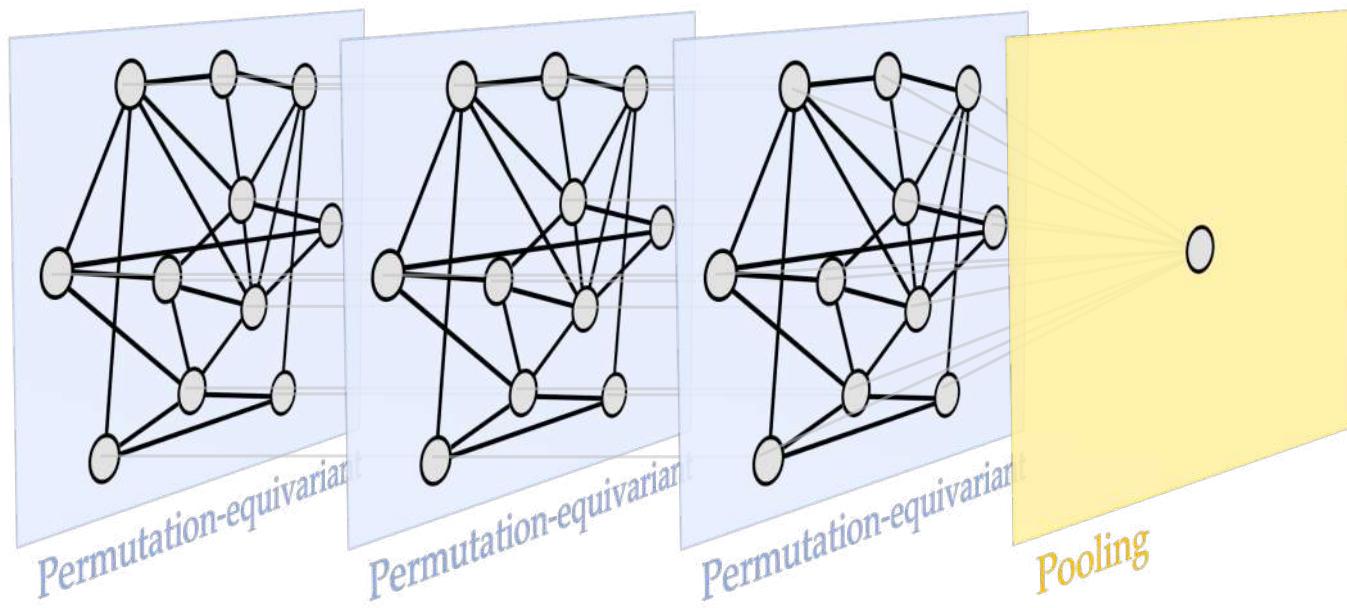
Equivariant Graph Functions

permutation-equivariant

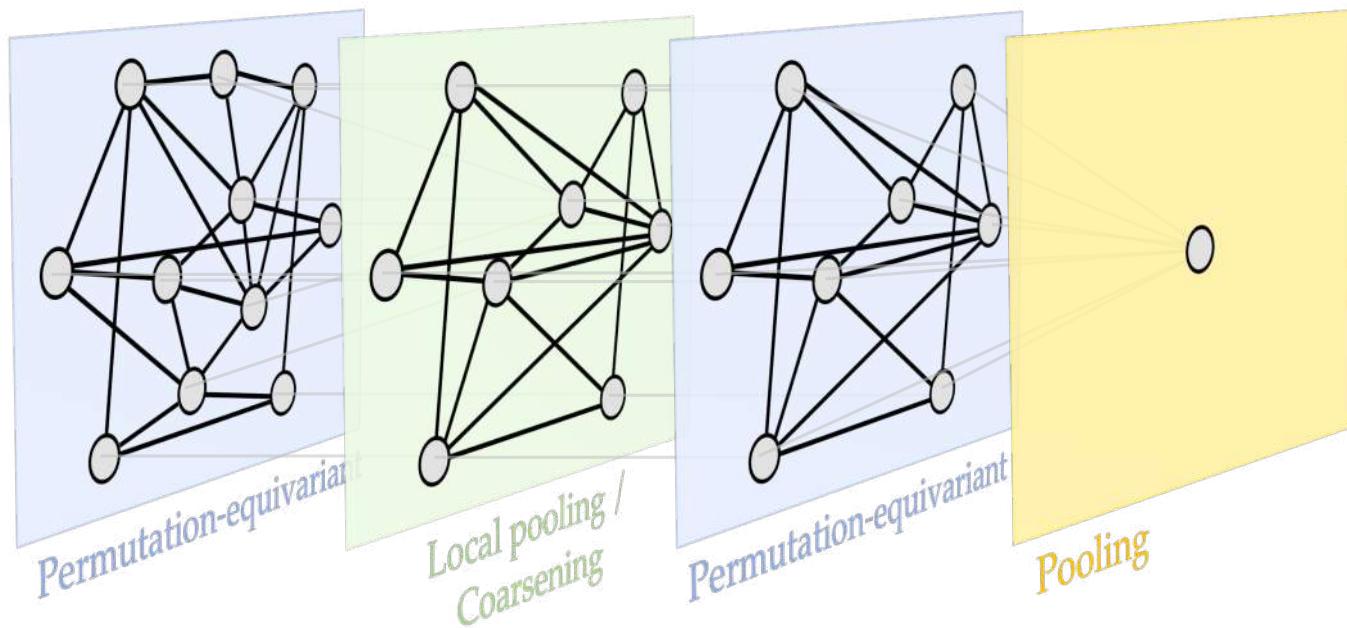
$$F(PX, PAP^\top) = PF(X, A)$$



Graph Neural Networks



Graph Neural Networks



History of Graph Neural Networks according to Machine Learners



A. Sperduti



C. Goller



A. Küchler



M. Gori



F. Scarselli



Y. Li

Labeling RAAM

1994

Backprop through structure

1996

Graph Neural Networks

2005, 2008

Gated GNN

2015

History of Graph Neural Networks according to Graph Theorists



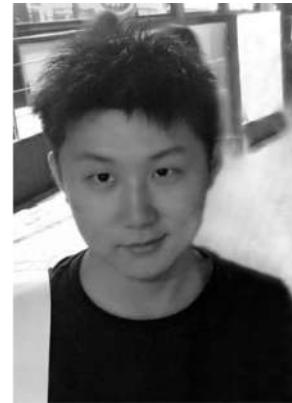
N. Shervashidze



K. Borgwardt



C. Morris



K. Xu



H. Maron

Weisfeiler-Lehman kernels

k-GNN

GIN

Provably powerful GNN

2009

2019

Weisfeiler-Lehman Test



A. Lehman B. Weisfeiler



B. Weisfeiler

Weisfeiler, Lehman 1968; Ponomarenko 2018 (history)

History of Graph Neural Networks according to Chemists



D. Kireev

ChemNet

1995



I. Baskin

Neural descriptors

1997



C. Merkwirth

Molecular graph net

2005



D. Duvenaud

Molecular fingerprints

2015

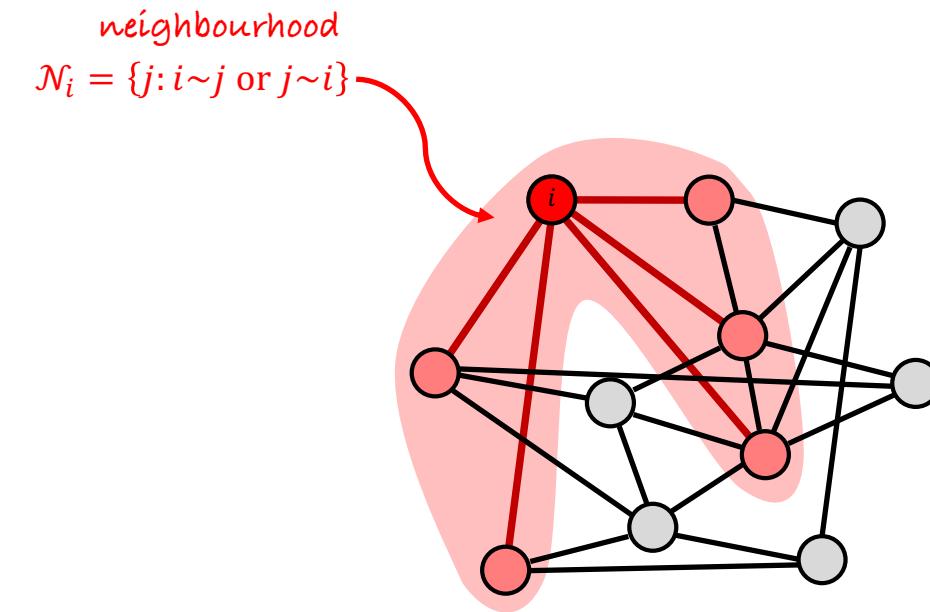


J. Gilmer

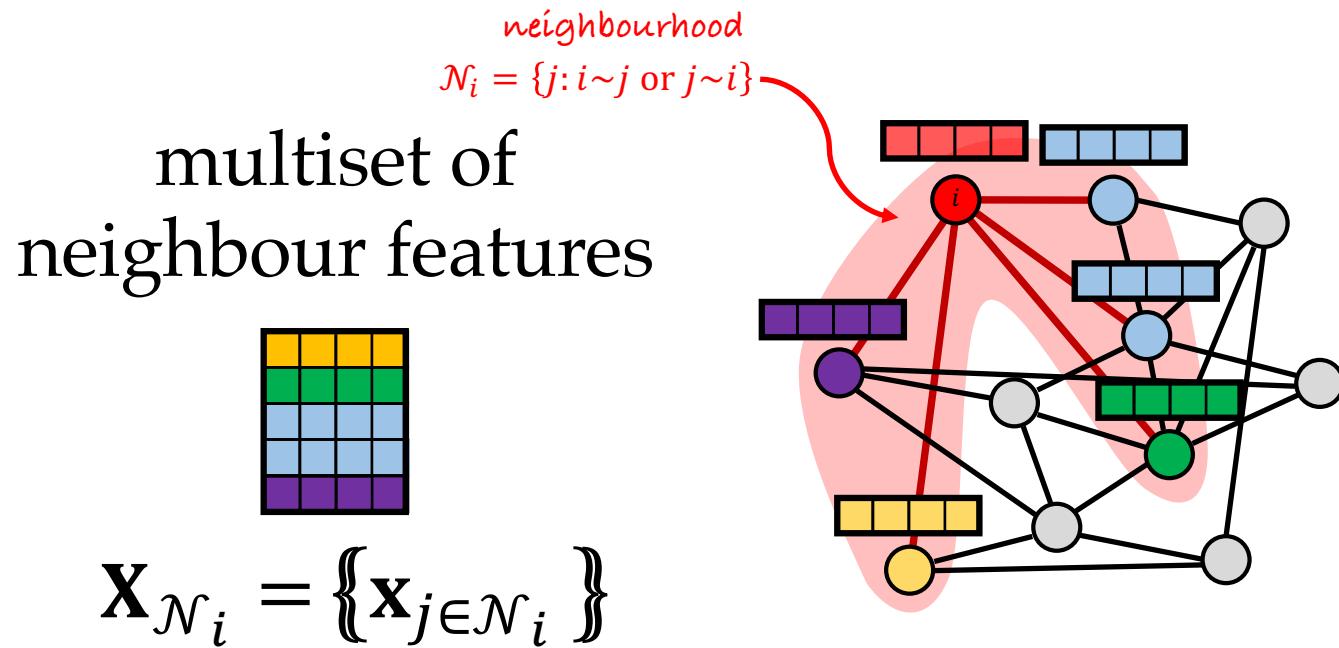
MPNNs

2017

A General Blueprint for Constructing Graph Functions

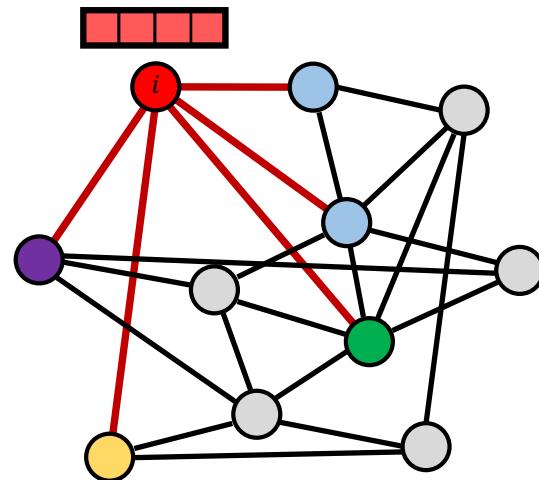


A General Blueprint for Constructing Graph Functions



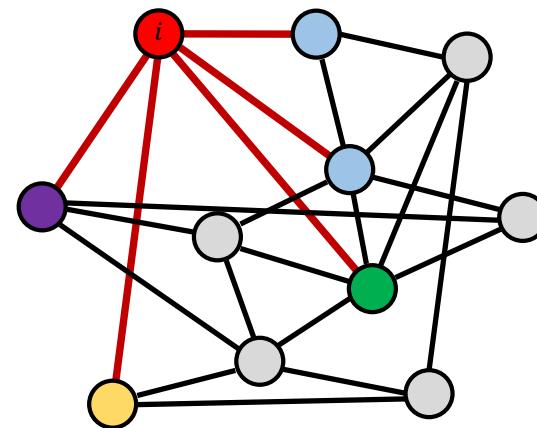
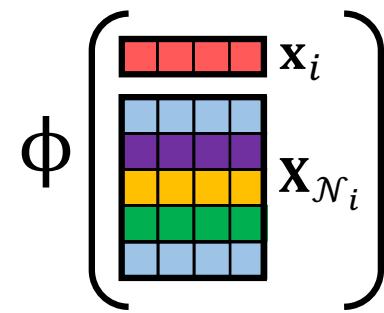
A General Blueprint for Constructing Graph Functions

multiset of
local function
neighbour features

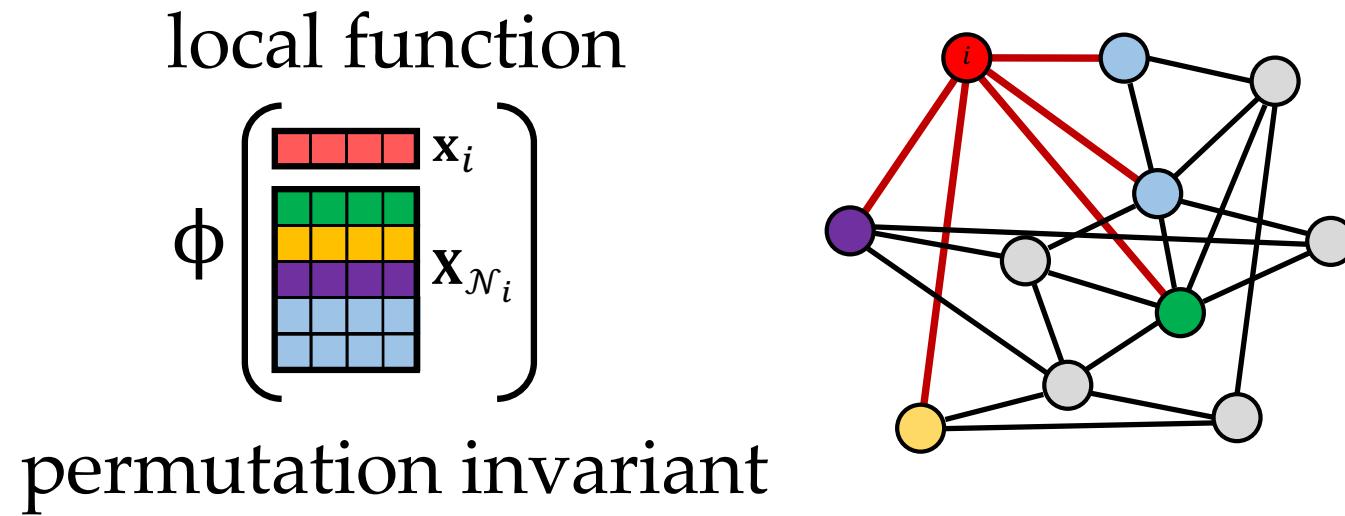
$$\phi \left(\begin{array}{c} \text{matrix of} \\ \text{neighbour features} \\ \text{for node } x_i \\ \text{in set } \mathcal{N}_i \end{array} \right) = \{x_j \in \mathcal{N}_i\}$$


A General Blueprint for Constructing Graph Functions

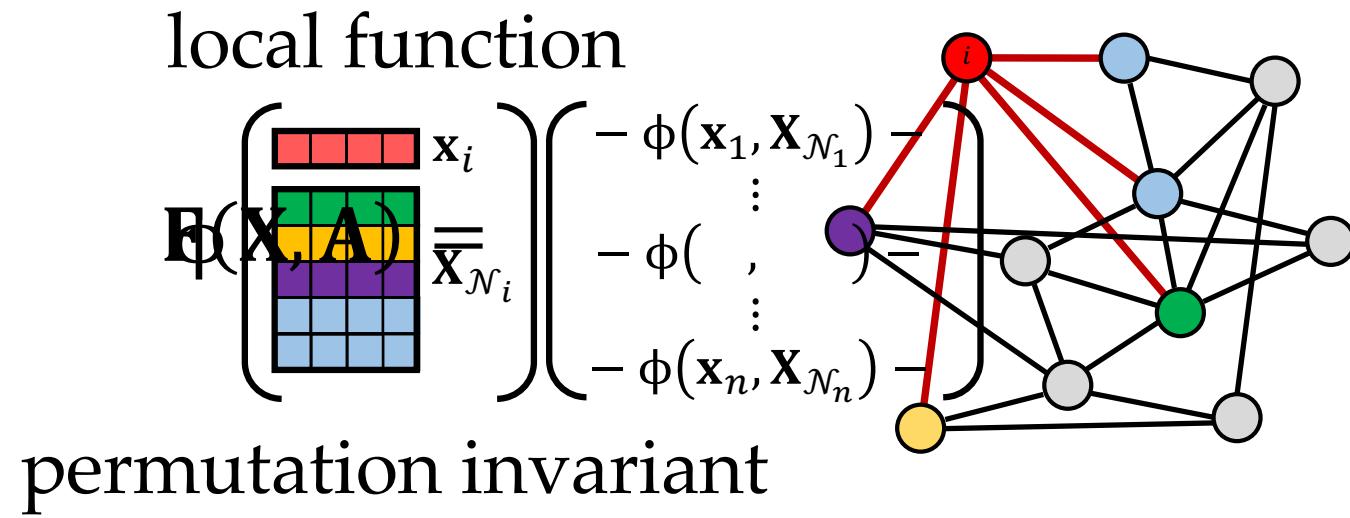
local function



A General Blueprint for Constructing Graph Functions



A General Blueprint for Constructing Graph Functions



A General Blueprint for Constructing Graph Functions

$$\mathbf{F}(\mathbf{X}, \mathbf{A}) = \begin{pmatrix} -\phi(\mathbf{x}_1, \mathbf{x}_{\mathcal{N}_1}) - \\ \vdots \\ -\phi(\mathbf{x}_i, \mathbf{x}_{\mathcal{N}_i}) - \\ \vdots \\ -\phi(\mathbf{x}_n, \mathbf{x}_{\mathcal{N}_n}) - \end{pmatrix}$$

permutation equivariant

"Flavours" of Graph Neural Networks

$$f(\mathbf{x}_i) = \phi \left(\mathbf{x}_i, \bigcup_{j \in \mathcal{N}_i} \psi(\mathbf{x}_j) \right)$$

permutation-invariant
aggregation operator, e.g. sum

new feature of
node i

learnable
functions

The diagram illustrates the update rule for a node i in a graph. It shows the new feature vector $f(\mathbf{x}_i)$ being computed as a function ϕ of the node's own feature \mathbf{x}_i and the features of its neighbors. The neighbors are represented by the set \mathcal{N}_i . The function ϕ is composed of a learnable function ψ applied to each neighbor's feature \mathbf{x}_j , and these results are aggregated using a permutation-invariant operator, such as a sum.

"Flavours" of Graph Neural Networks

$$f(\mathbf{x}_i) = \phi \left(\mathbf{x}_i, \bigcup_{j \in \mathcal{N}_i} c_{ij} \psi(\mathbf{x}_j) \right)$$

“convolutional”

importance of node j to
the representation of i

"Flavours" of Graph Neural Networks

$$f(\mathbf{x}_i) = \phi \left(\mathbf{x}_i, \bigcup_{j \in \mathcal{N}_i} a(\mathbf{x}_i, \mathbf{x}_j) \psi(\mathbf{x}_j) \right)$$

“attentional”

Monti et al. 2017; Veličković et al. 2018 (GAT)

Message Passing

$$f(\mathbf{x}_i) = \phi \left(\mathbf{x}_i, \bigcup_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j) \right)$$

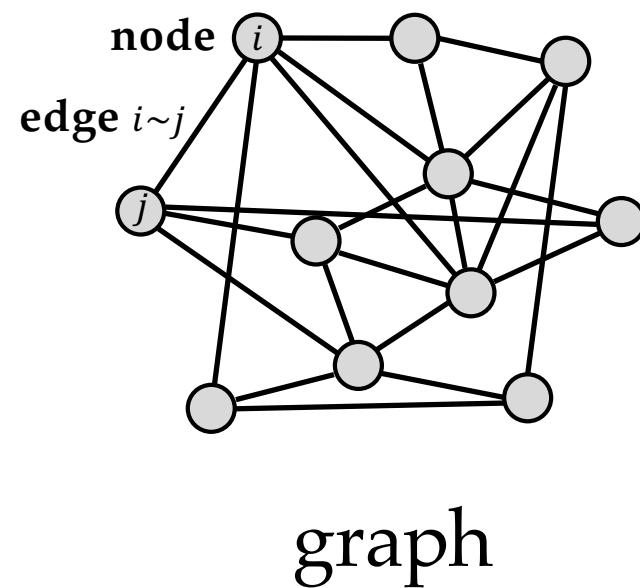
“message passing”

Message Passing

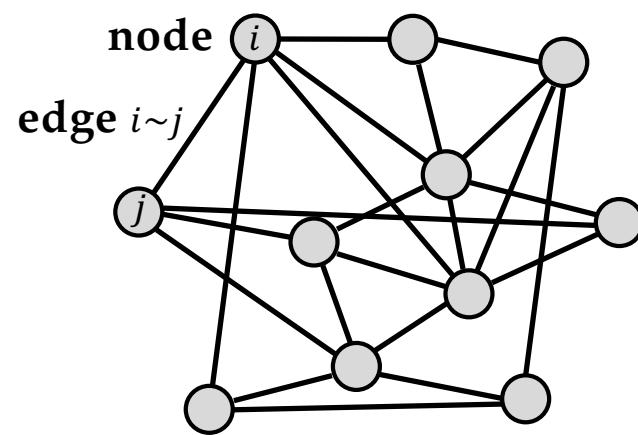
$$f(\mathbf{x}_i) = \phi \left(\mathbf{x}_i, \bigcup_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j) \right)$$

“message passing”

Special Cases of GNNs

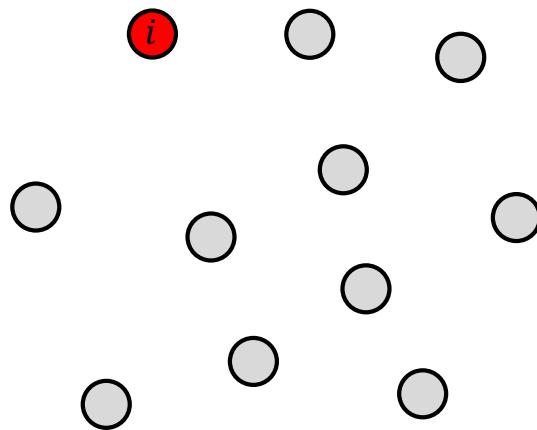


Special Cases of GNNs



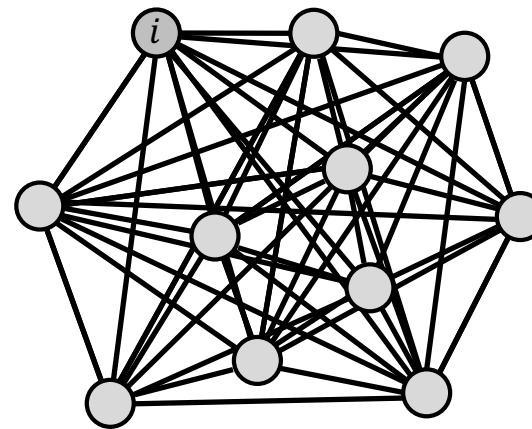
graph

DeepSets



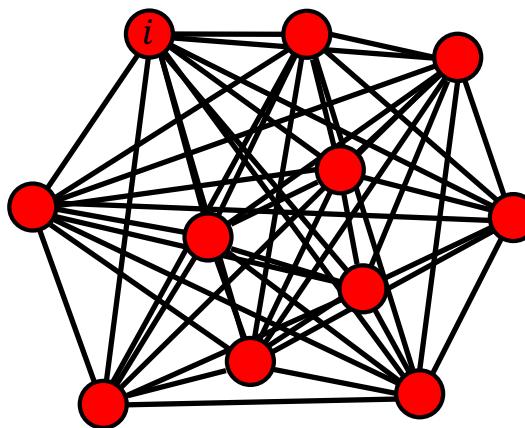
Zaheer et al 2017; Qi et al 2017

Transformers



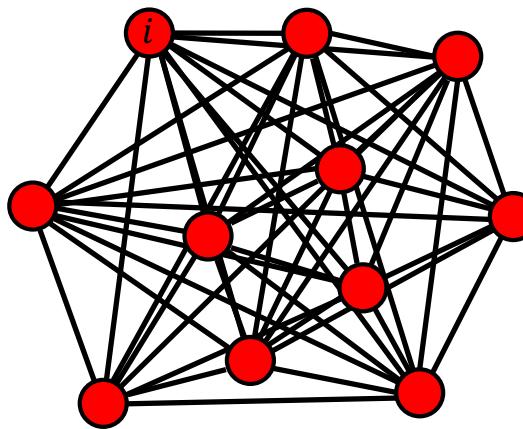
complete graph

Transformers



$$\phi \left(\mathbf{x}_i, \sum_{j=1}^n c_{ij} \psi(\mathbf{x}_j) \right)$$

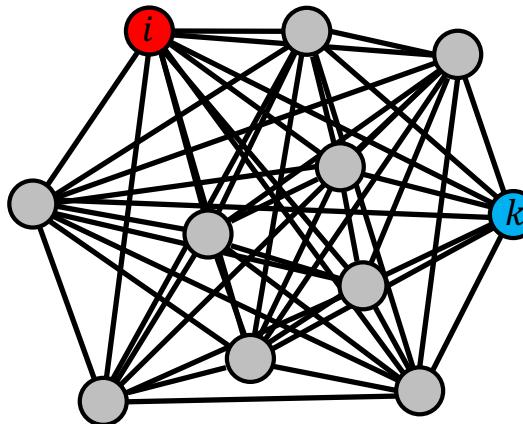
Transformers



$$\phi \left(\mathbf{x}_i, \bigcup_{j=1}^n a(\mathbf{x}_i, \mathbf{x}_j) \psi(\mathbf{x}_j) \right)$$

Vaswani et al. 2017

Transformers

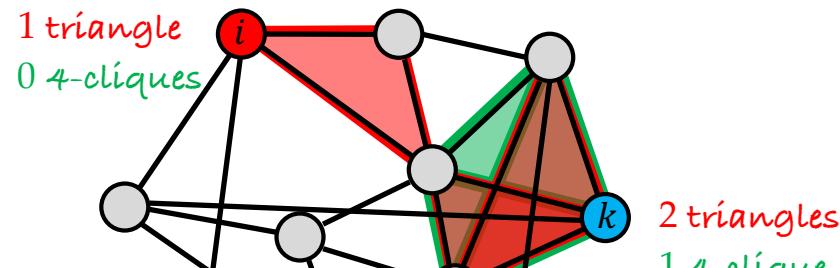


$$\phi \left(\mathbf{x}_i, \bigcup_{j=1}^n a(\mathbf{x}_i, \mathbf{x}_j, \mathbf{p}_i, \mathbf{p}_j) \psi(\mathbf{x}_j) \right)$$

↑
positional encoding

Vaswani et al. 2017

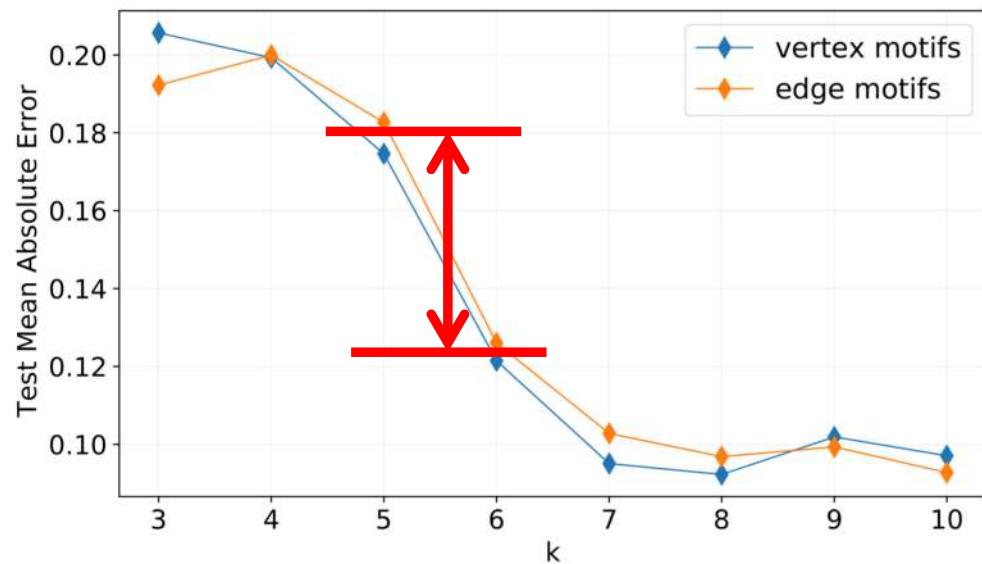
Graph Substructure Network



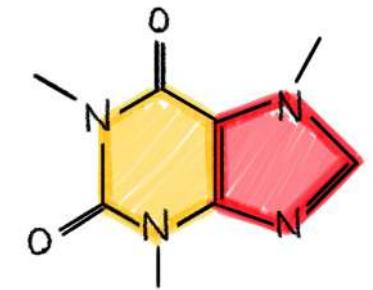
$$\phi \left(\mathbf{x}_i, \bigwedge_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j, \mathbf{p}_i) \right)$$

↑
structural encoding

Graph Substructure Network

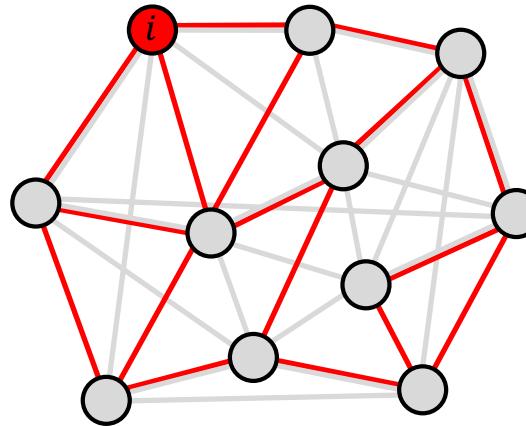


Molecule property prediction on ZINC
using GSN with k -cycles



Molecule of caffeine

decouple computational graph from the input graph



sampling

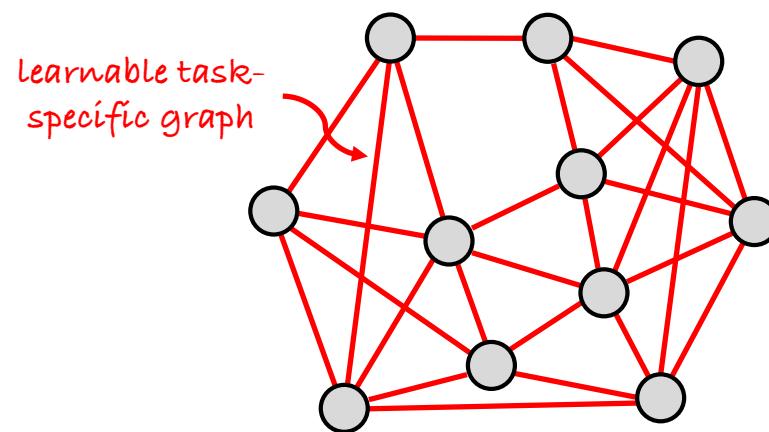
rewiring

multi-hop
filters

$$\phi \left(\mathbf{x}_i, \bigwedge_{j \in \mathcal{N}'_i} \psi(\mathbf{x}_i, \mathbf{x}_j) \right)$$

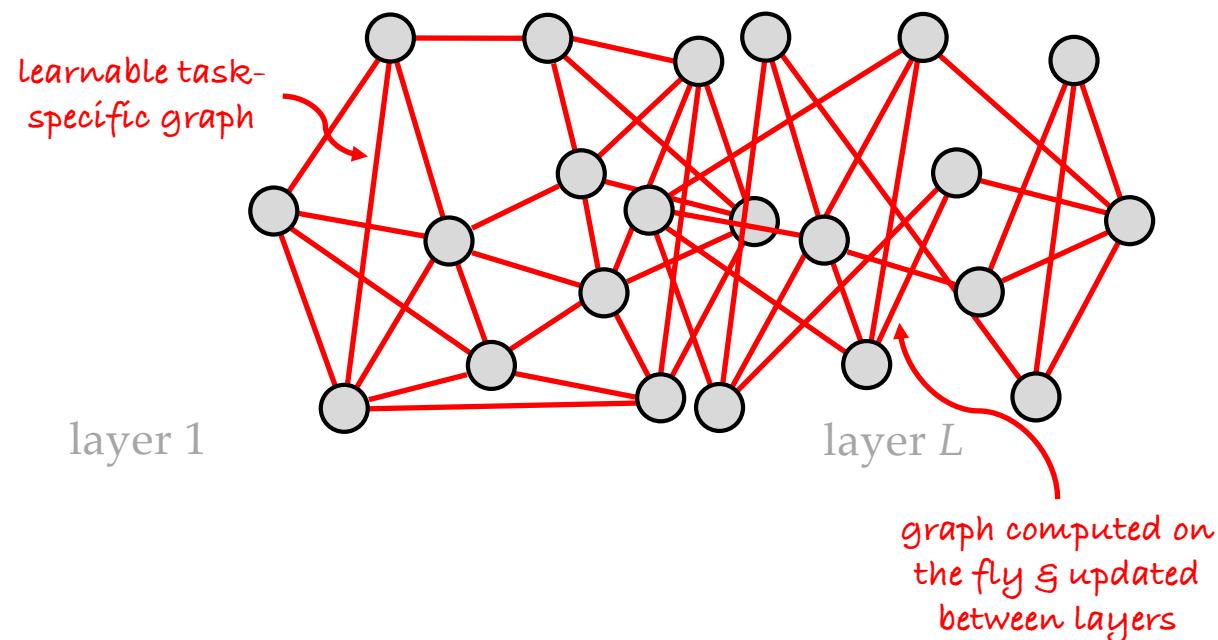
Hamilton et al 2017; Klicpera et al. 2019; Alon & Yahav 2020; Frasca, Rossi et B 2020

Latent Graph Learning

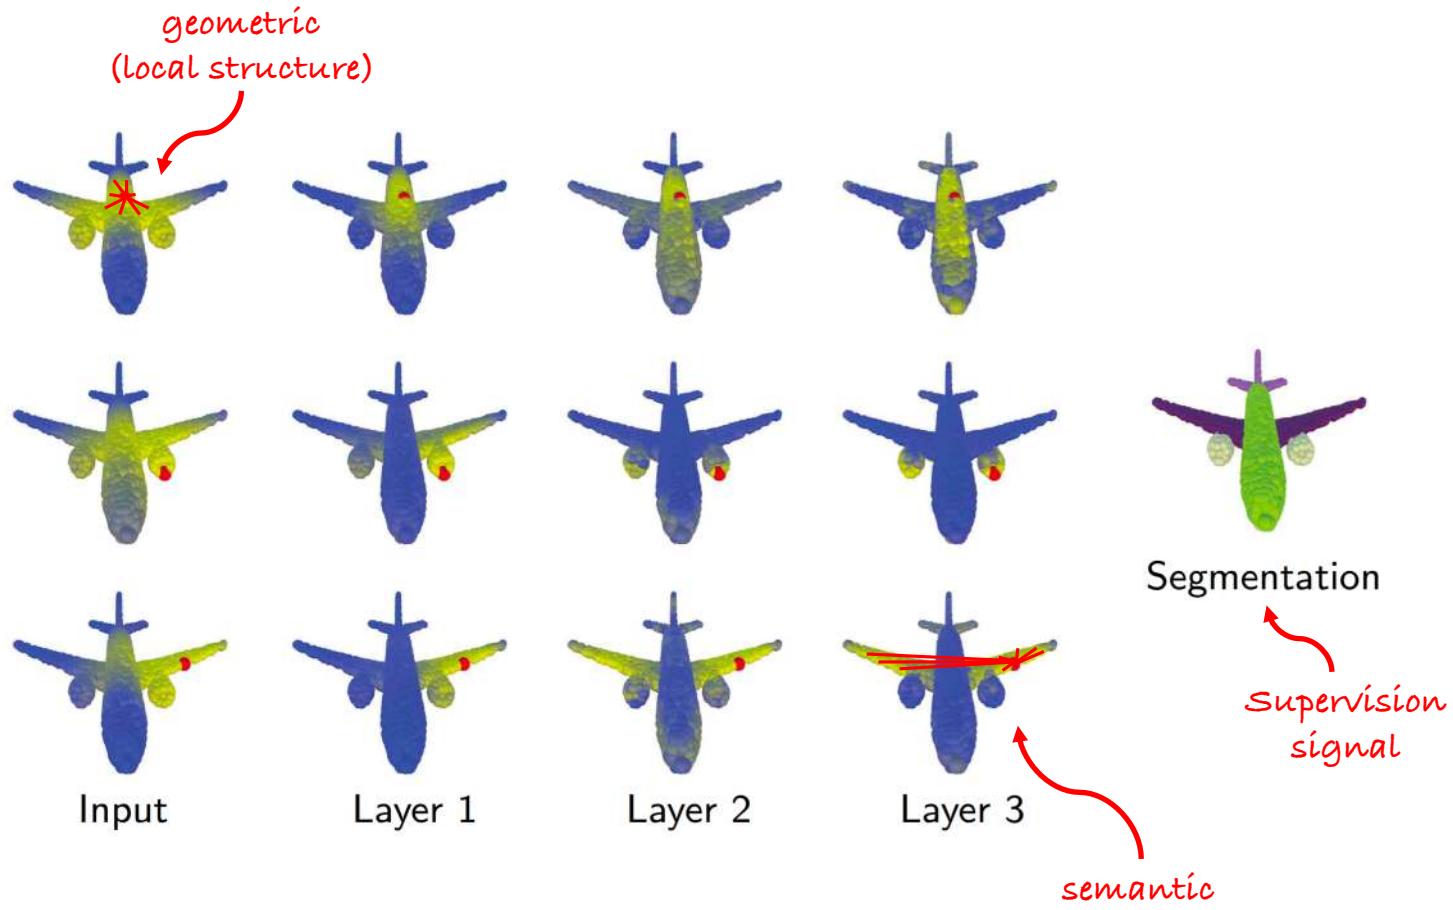


Wang et al 2018; Franceschi et al. 2019; Kipf et al. 2020; Kazi, Cosmo et al. 2020; Cranmer et al. 2020

Dynamic Graph CNN

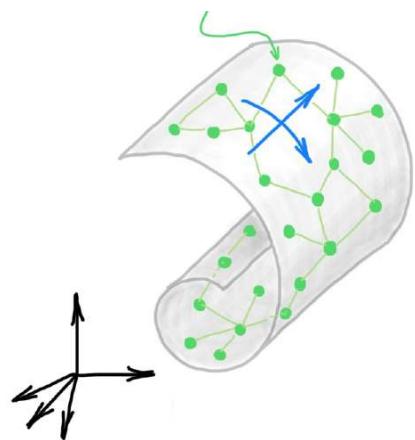


Wang et al 2018; Franceschi et al. 2019; Kipf et al. 2020; Kazi, Cosmo et al. 2020; Cranmer et al. 2020



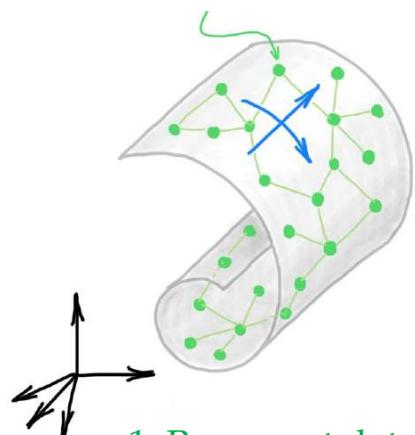
Manifold Learning

intrinsically low-dimensional
data in a high-dimensional space

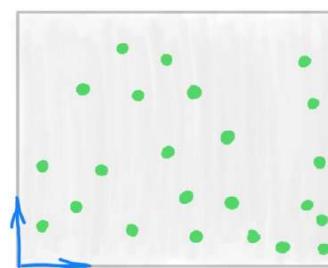


Manifold Learning

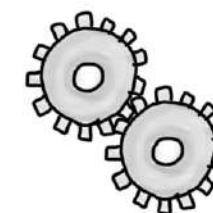
intrinsically low-dimensional
data in a high-dimensional space



1. Represent data
structure as a graph



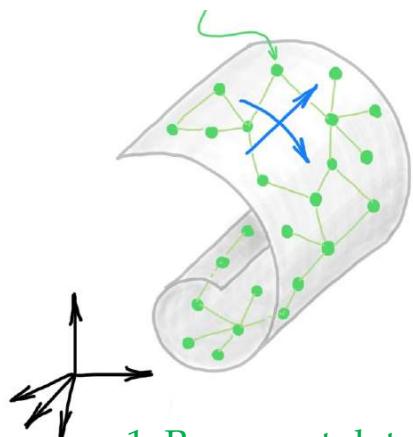
2. Compute low-
dimensional embedding



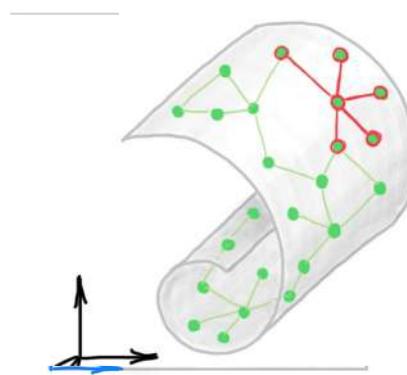
3. Apply ML

Manifold Learning 2.0

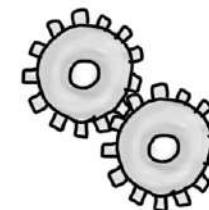
intrinsically low-dimensional
data in a high-dimensional space



1. Represent data
structure as a graph



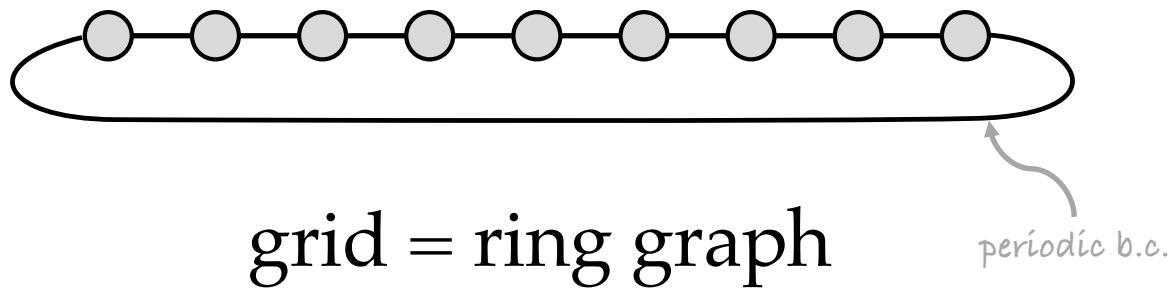
1. Represent data structure
as a graph **apply ML**
directly on the graph



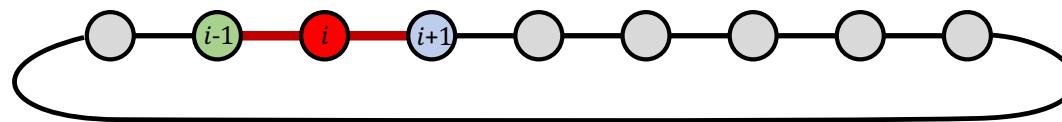
3. Apply ML

GRIDS

Grids vs Graphs

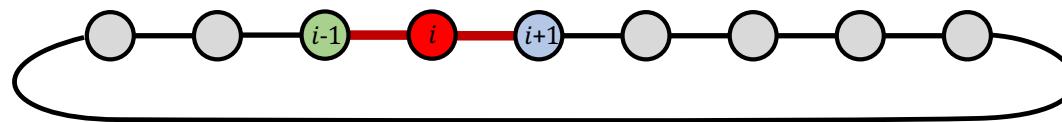


Grids vs Graphs



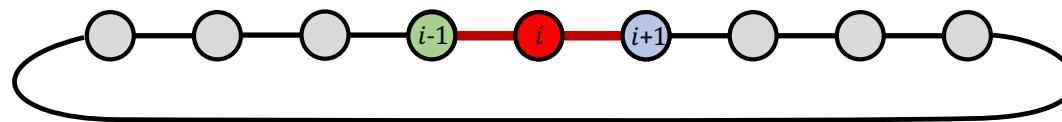
fixed neighbourhood structure

Grids vs Graphs



fixed neighbourhood structure

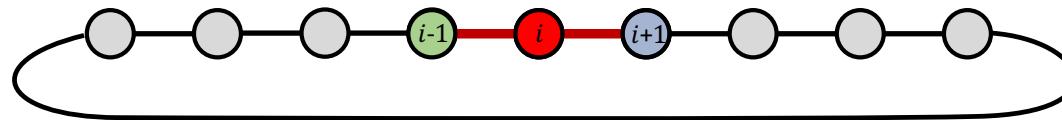
Grids vs Graphs



local aggregation function

$$f(\mathbf{x}_i) = \phi(\mathbf{x}_i, \{\mathbf{x}_{i-1}, \mathbf{x}_{i+1}\})$$

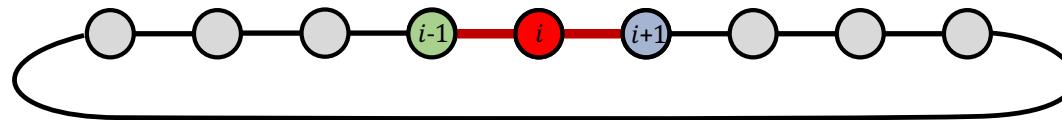
Grids vs Graphs



local aggregation function

$$f(\mathbf{x}_i) = \phi(\mathbf{x}_{i-1}, \mathbf{x}_i, \mathbf{x}_{i+1})$$

Grids vs Graphs



linear local aggregation function

$$f(\mathbf{x}_i) = a\mathbf{x}_{i-1} + b\mathbf{x}_i + c\mathbf{x}_{i+1}$$

Convolution

$$f(\mathbf{X}) = \begin{bmatrix} b & c \\ a & b & c \\ & a & b & c \\ c & a & b \end{bmatrix} \mathbf{X}$$

circulant matrix = convolution

Convolution

vector of parameters θ

$$f(\mathbf{X}) = \begin{pmatrix} b & c & & \\ a & b & c & \\ & a & b & c \\ c & & a & b \end{pmatrix} \mathbf{X}$$

circulant matrix $\mathbf{C}(\theta)$

Deriving Convolution from Symmetry

vector of parameters θ

$$\begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix} f(\mathbf{X}) = \mathbf{C} \begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix} = \mathbf{C} \begin{bmatrix} a & x & y & z \\ x & a & y & z \\ y & z & a & x \\ z & x & y & a \end{bmatrix} \mathbf{X}$$

circulant matrices a commute
circulant matrix $\mathbf{C}(\theta)$

Deriving Convolution from Symmetry

$$\begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix} \begin{bmatrix} 1 & & \\ & 1 & \\ & & 1 \end{bmatrix} = \begin{bmatrix} 1 & & \\ & 1 & \\ & & 1 \end{bmatrix} \begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix}$$

circulant matrix \Rightarrow commutes with shift

Deriving Convolution from Symmetry

$$\begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix} \left[\begin{array}{c} 1 \\ 1 \\ 1 \end{array} \right] = \left[\begin{array}{c} 1 \\ 1 \\ 1 \end{array} \right] \begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix}$$

shift \mathbf{S}

convolution \Rightarrow shift-equivariant

$$\mathbf{CS} = \mathbf{SC}$$

Deriving Convolution from Symmetry

$$\begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix} \left[\begin{array}{c} 1 \\ 1 \\ 1 \end{array} \right] \xrightarrow{\text{shift } S} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix} \left[\begin{array}{c} 1 \\ 1 \\ 1 \end{array} \right]$$

convolution \Leftrightarrow shift-equivariant

convolution emerges from translation symmetry

Deriving Fourier Transform from Symmetry

same eigenbasis for all convolutions

different eigenvalues for each convolution

$$\begin{bmatrix} b & c & & a \\ a & b & c & \\ & a & b & c \\ c & & a & b \end{bmatrix} = \begin{bmatrix} | & & | & & | & & | \\ u_1 & \dots & u_n & & & & u_n^* \\ | & & | & & | & & | \end{bmatrix} \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix} \begin{bmatrix} --- u_1^* --- \\ \vdots \\ --- u_n^* --- \end{bmatrix}$$

commuting matrices are jointly diagonalizable
vectors of \mathbf{S}

Deriving Fourier Transform from Symmetry

Fourier basis

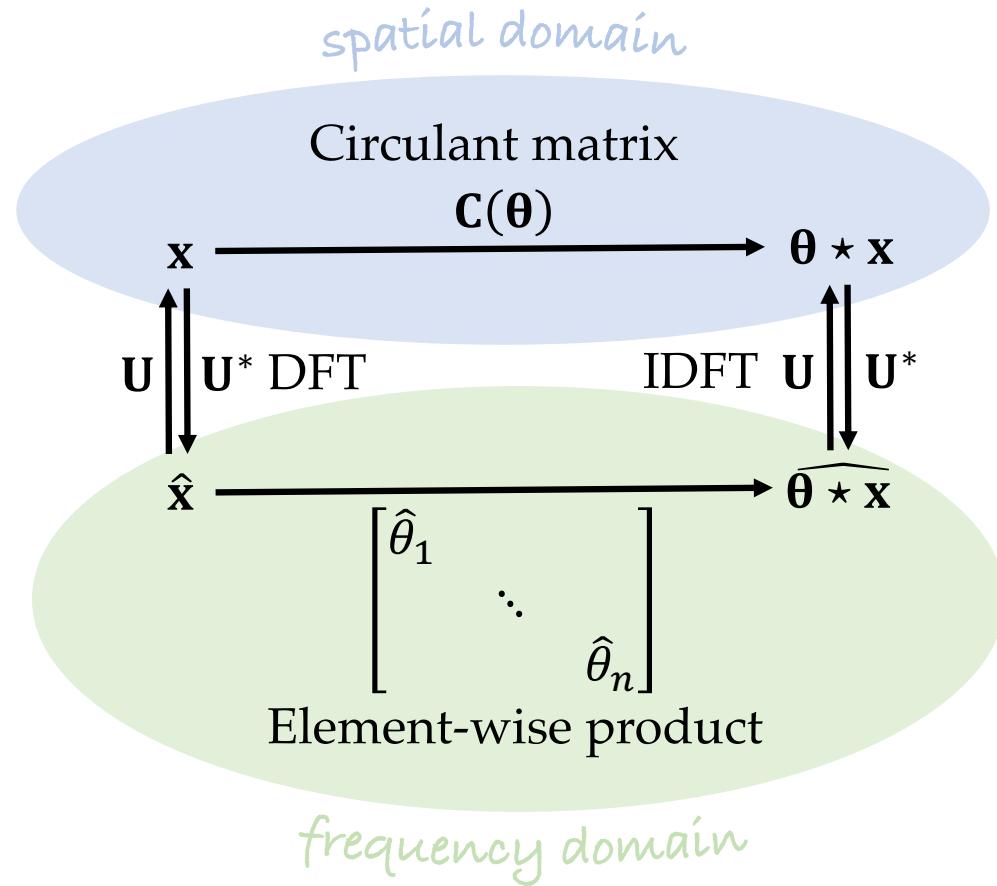
$$\mathbf{u}_k = \frac{1}{\sqrt{n}} \begin{bmatrix} 1 \\ e^{i\frac{2\pi}{n}k} \\ \vdots \\ e^{i\frac{2\pi}{n}(n-1)k} \end{bmatrix}$$

$$\begin{bmatrix} b & c & a & & & \\ a & b & c & & & \\ & a & b & c & & \\ & & a & b & c & \\ & & & a & b & c \\ & & & & a & b \end{bmatrix} = \begin{bmatrix} | & & | & & | & & | \\ \mathbf{u}_1 & \dots & \mathbf{u}_n & & & & \mathbf{u}_1^* \\ | & & | & & | & & | \end{bmatrix} \begin{bmatrix} \hat{\theta}_1 & & & & & \\ & \hat{\theta}_2 & & & & \\ & & \ddots & & & \\ & & & \hat{\theta}_n & & \end{bmatrix} \begin{bmatrix} | & & | & & | & & | \\ \mathbf{u}_1^* & \dots & \mathbf{u}_n^* & & & & \mathbf{u}_1 & \dots & \mathbf{u}_n \\ | & & | & & | & & | \end{bmatrix}$$

Fourier transform

$$\widehat{\theta} = \mathbf{U}^* \theta$$

commuting matrices are jointly
diagonalisable by Fourier Transform



Reminder: Euclidean case

$$\begin{bmatrix} b & c & & a \\ a & b & c & \\ & a & b & c \\ c & & a & b \end{bmatrix} \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{bmatrix} = \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{bmatrix} \begin{bmatrix} b & c & & a \\ a & b & c & \\ & a & b & c \\ c & & a & b \end{bmatrix}$$

- Convolutions commute
- Commuting matrices are jointly diagonalizable (“have same eigenvectors”)
- Pick up the shift operator to show eigenvectors = DFT \Rightarrow convolutions are diagonalised by the Fourier transform

Reminder: Euclidean case

$$\begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix} \begin{bmatrix} -2 & 1 & & & 1 \\ 1 & -2 & 1 & & \\ & 1 & -2 & 1 & \\ & & 1 & -2 & 1 \\ & & & 1 & -2 \end{bmatrix} = \begin{bmatrix} -2 & 1 & & & 1 \\ 1 & -2 & 1 & & \\ & 1 & -2 & 1 & \\ & & 1 & -2 & 1 \\ & & & 1 & -2 \end{bmatrix} \begin{bmatrix} b & c & a \\ a & b & c \\ c & a & b \end{bmatrix}$$

- Convolutions commute
- Commuting matrices are jointly diagonalizable (“have same eigenvectors”)
- Pick up the *Laplacian operator* whose eigenvectors are the DFT \Rightarrow convolutions are diagonalised by the Fourier transform

Laplacian Operator

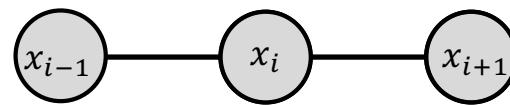
- Fourier basis functions = Laplacian eigenfunctions

$$\Delta e^{+i\omega u} = -\omega^2 e^{+i\omega u}$$

- Fourier basis = orthogonal basis minimising the *Dirichlet energy*

$$\varphi_{k+1} = \operatorname{argmin}_{\varphi} \int_{-\infty}^{+\infty} \|\nabla \varphi_k(u)\|^2 du \quad \text{s.t. } \|\varphi\| = 1 \quad \text{and} \quad \langle \varphi, \varphi_j \rangle = 0 \text{ where } j = 1, \dots, k$$

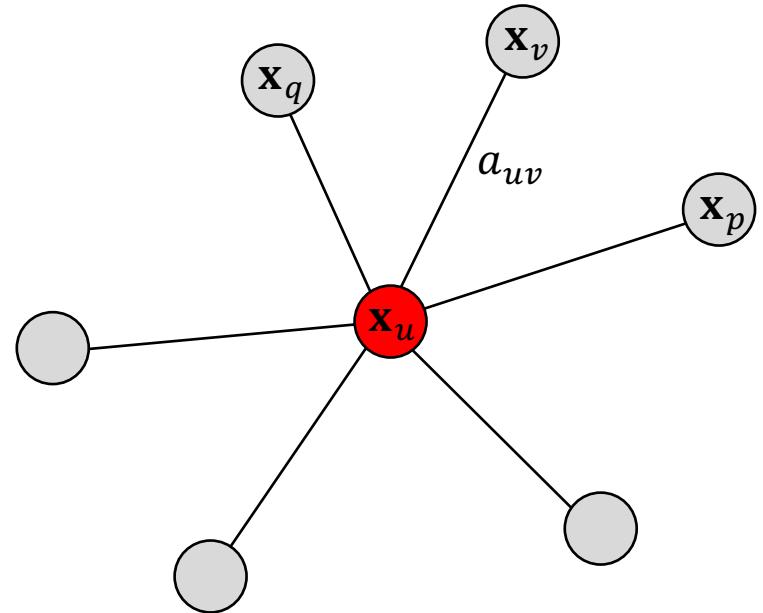
- “smoothest orthogonal basis”
- “local difference from neighbours”



$$(x_{i-1} + x_{i+1}) - 2x_i$$

Graph Laplacian

$$\begin{aligned}(\Delta \mathbf{x})_u &= \sum_{v \in \mathcal{N}_u} a_{uv} (\mathbf{x}_u - \mathbf{x}_v) \\&= d_u \mathbf{x}_u - \sum_{v \in \mathcal{N}_u} a_{uv} \mathbf{x}_v \\&= \phi(\mathbf{x}_u, \mathbf{X}_{\mathcal{N}_u})\end{aligned}$$



- Linear local *permutation-invariant* aggregator $\phi(\mathbf{x}_u, \mathbf{X}_{\mathcal{N}_u})$
- Node-wise *permutation-equivariant* linear function $\mathbf{F}(\mathbf{X}) = \Delta \mathbf{X}$

Graph Fourier Transform & Convolution

On undirected graph, the Graph Laplacian is a symmetric matrix admitting an orthogonal eigendecomposition $\Delta = \Phi \Lambda \Phi^T$ with $\Phi^T \Phi = I$

- *Graph Fourier Transform:* $\hat{\mathbf{x}} = \Phi^T \mathbf{x}$ $O(n^2)$ complexity
- *Inverse Graph Fourier Transform:* $\mathbf{x} = \Phi \hat{\mathbf{x}}$ $O(n^2)$ complexity
- *Spectral Convolution:*
$$\begin{aligned} \mathbf{x} * \Psi &= \Phi \left((\Phi^T \mathbf{x}) \cdot (\Phi^T \Psi) \right) \\ &= \Phi \operatorname{diag}(\widehat{\Psi}) \hat{\mathbf{x}} \end{aligned}$$
 $O(n^2)$ complexity

ChebNet

- *Polynomial spectral filter* $\hat{p}(\lambda) = \sum_{l=0}^r \alpha_l \lambda^l$ applied to the Graph Laplacian

$$\hat{p}(\Delta) = \Phi \hat{p}(\Lambda) \Phi^T = \Phi \sum_{l=0}^r \alpha_l \Lambda^l \Phi^T$$

ChebNet

- *Polynomial spectral filter* $\hat{p}(\lambda) = \sum_{l=0}^r \alpha_l \lambda^l$ applied to the Graph Laplacian

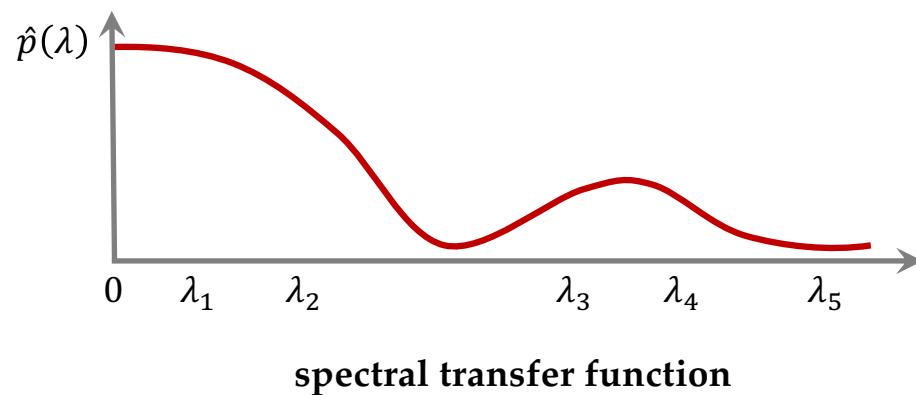
$$\hat{p}(\Delta) = \Phi \hat{p}(\Lambda) \Phi^T = \sum_{l=0}^r \alpha_l \Phi \Lambda^l \Phi^T = \sum_{l=0}^r \alpha_l \Delta^l$$

- No explicit eigendecomposition!
- Apply powers of the Graph Laplacian – complexity $O(r|\mathcal{E}|) \sim O(n)$
- Stable under graph deformations
- “Spectral” filter boils down to simple local averaging – “convolutional flavor” of GNN

Spectral vs Spatial Filters

$$(\hat{p}(\Delta)x)(u) = \sum_{k \geq 0} \hat{p}(\lambda_k) \langle x, \varphi_k \rangle \varphi_k(u)$$

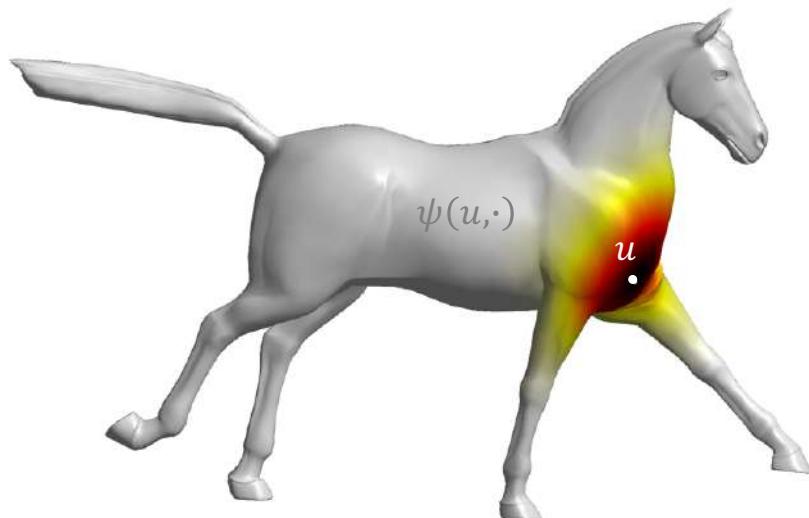
$\hat{\psi}_k$ \hat{x}_k



Stable Spectral Filters

$$\begin{aligned}(\hat{p}(\Delta)x)(u) &= \sum_{k \geq 0} \hat{p}(\lambda_k) \langle x, \varphi_k \rangle \varphi_k(u) \\&= \int_{\Omega} x(u) \sum_{k \geq 0} \hat{p}(\lambda_k) \varphi_k(v) \varphi_k(u) d\nu\end{aligned}$$

 $\psi(u, v)$

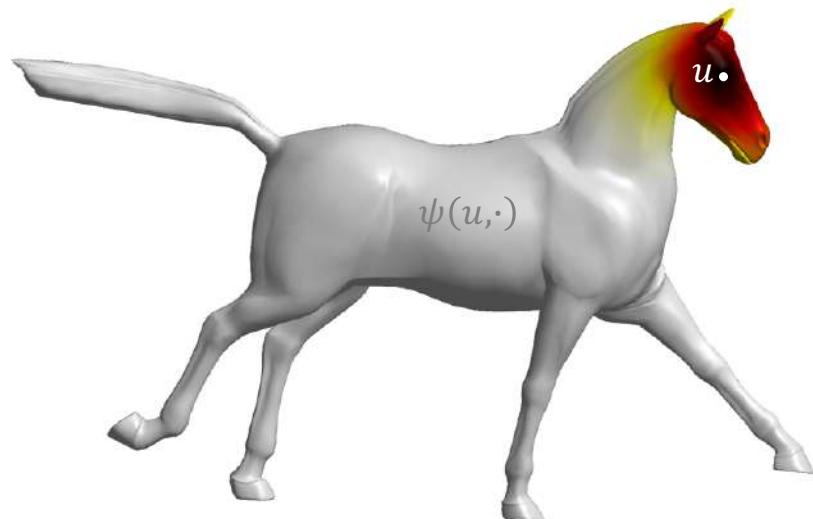


position-dependent spatial kernel

Stable Spectral Filters

$$(\hat{p}(\Delta)x)(u) = \sum_{k \geq 0} \hat{p}(\lambda_k) \langle x, \varphi_k \rangle \varphi_k(u)$$

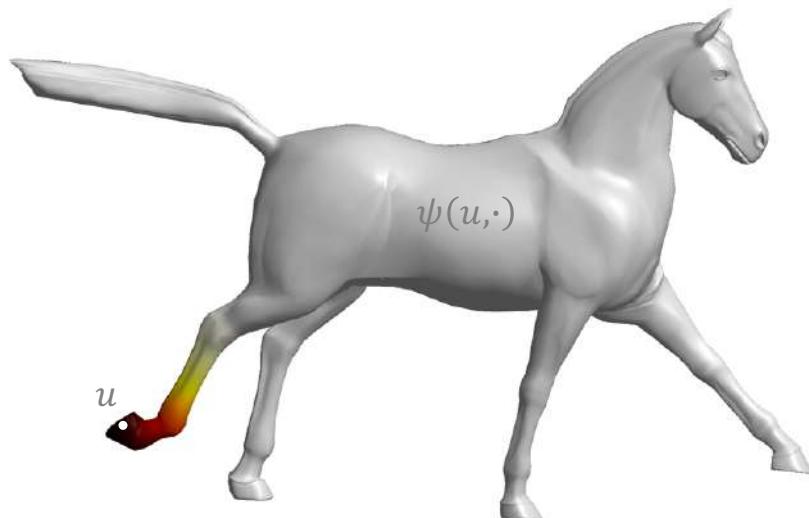
$$= \int_{\Omega} x(u) \sum_{k \geq 0} \hat{p}(\lambda_k) \varphi_k(v) \varphi_k(u) d\nu$$



position-dependent spatial kernel

Stable Spectral Filters

$$\begin{aligned}(\hat{p}(\Delta)x)(u) &= \sum_{k \geq 0} \hat{p}(\lambda_k) \langle x, \varphi_k \rangle \varphi_k(u) \\&= \int_{\Omega} x(u) \sum_{k \geq 0} \hat{p}(\lambda_k) \varphi_k(v) \varphi_k(u) d\nu\end{aligned}$$



position-dependent spatial kernel

Sanity Check

$$\begin{aligned} (\hat{p}(\Delta)x)(u) &= \sum_{k \geq 0} \hat{p}(\lambda_k) \langle x, \varphi_k \rangle \varphi_k(u) \\ &= \int_{-\pi}^{+\pi} x(u) \sum_{k \geq 0} \hat{p}(k) e^{-ikv} e^{+iku} dv \\ &\quad \underbrace{\phantom{\int_{-\pi}^{+\pi} x(u)}_{\psi(u-v)}} \end{aligned}$$

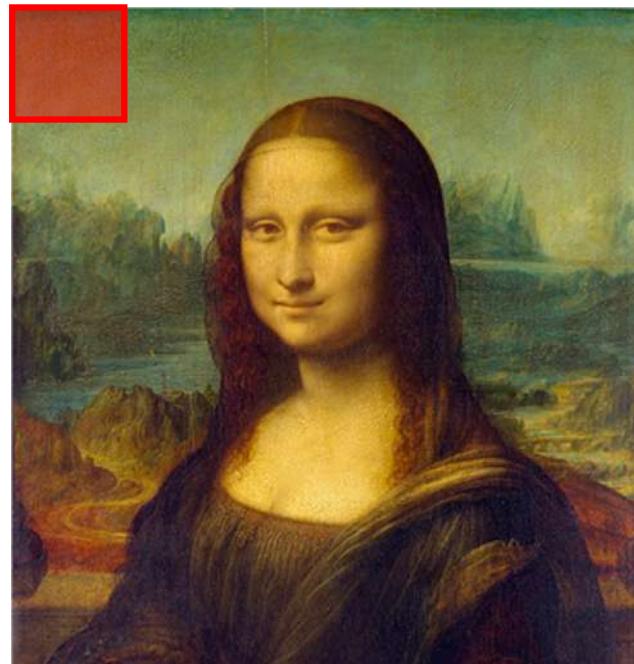
standard convolution in the Euclidean case

“Spectral GNNs”

- Spectral mesh filters [Taubin et al. 1996; Karni, Gotsman 2000]
- Graph Signal Processing [Shuman et al. 2013; Sandryhaila, Moura 2013]
- *Graph Fourier Transform* [Bruna et al. 2013]
 - Revival of the interest in Graph Neural Networks
 - High complexity – $O(n^2)$
 - Graph FT unstable under perturbations
- *Polynomial spectral filters a.k.a. ChebNet* [Defferrard et al. 2016]: $\hat{p}(\lambda) = \sum_{l=0}^r \alpha_l \lambda^l$
- *Rational spectral filters a.k.a. CayleyNet* [Levie et al. 2018]
- Spectral filter stability results [Gama et al. 2019; Levie et al. 2019; Kenlay et al. 2021]

GROUPS

Convolution, revisited



Convolution, revisited

convolution = matching shifted filter

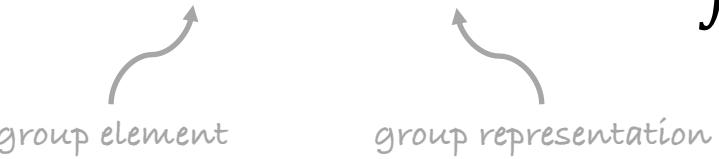
$$(x \star \psi)(u) = \langle x, T_u \psi \rangle = \int_{-\infty}^{+\infty} x(v) \psi(u - v) dv$$


A diagram illustrating the components of the convolution formula. Two grey arrows point from the labels "shift vector" and "shift operator" to the arguments of the inner product in the formula. The "shift vector" arrow points to the variable u in $T_u \psi$. The "shift operator" arrow points to the term $T_u \psi$ itself.

domain Ω = symmetry group \mathfrak{G}

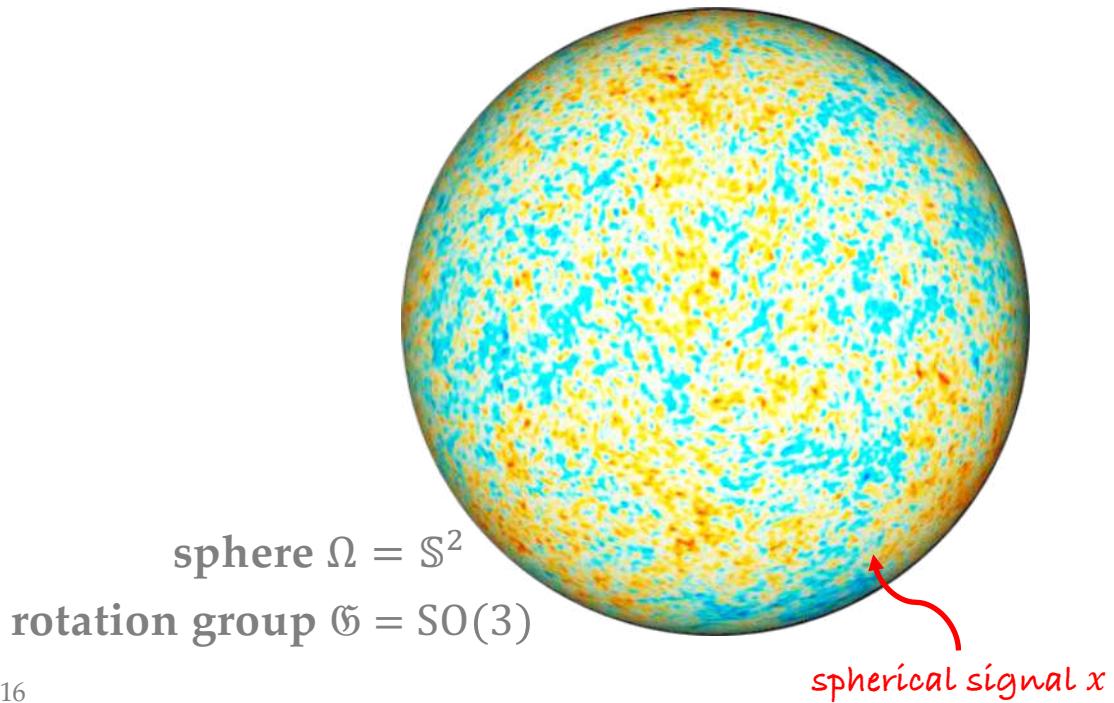
Group Convolution

convolution = matching transformed filter

$$(x * \psi)(g) = \langle x, \rho(g)\psi \rangle = \int_{\Omega} x(v)\psi(g^{-1}v)dv$$


group element group representation

Convolution on the Sphere



Cohen, Welling 2016

Convolution on the Sphere

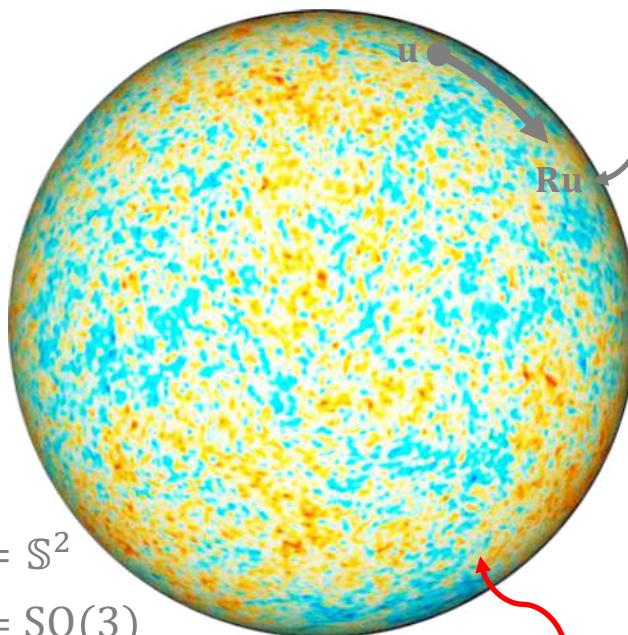
$$(x \star \psi)(R) = \int_{\mathbb{S}^2} x(u)\psi(R^{-1}u)du$$

signal on $SO(3)$

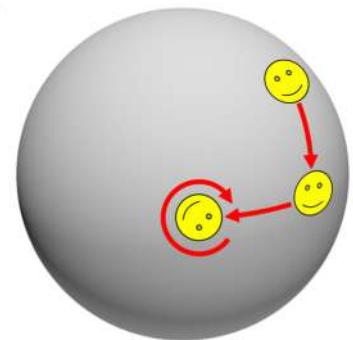
sphere $\Omega = \mathbb{S}^2$

rotation group $G = SO(3)$

Cohen, Welling 2016

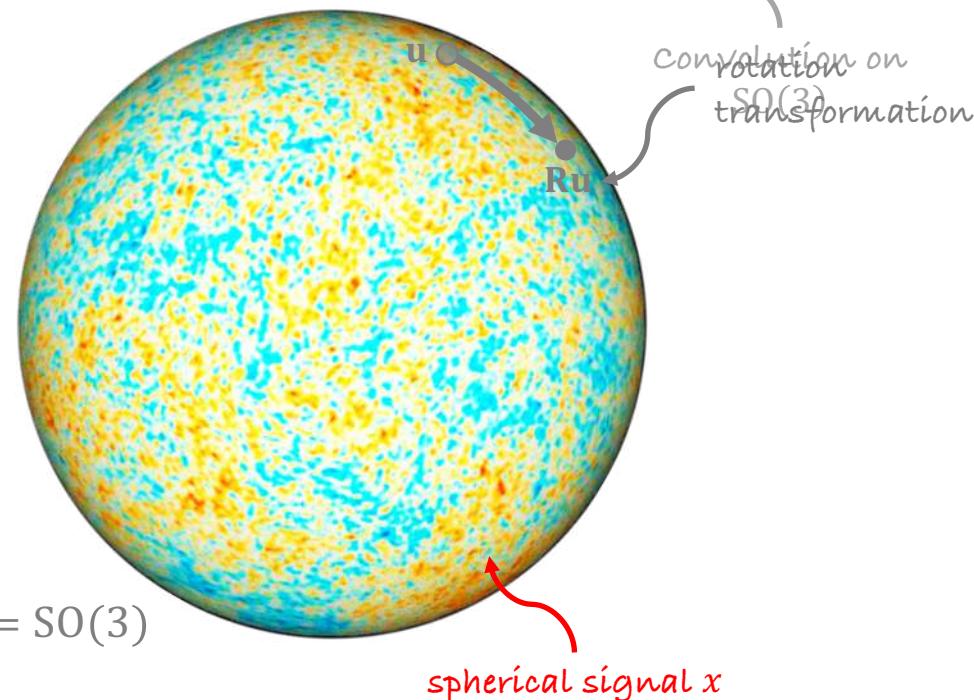


spherical signal x



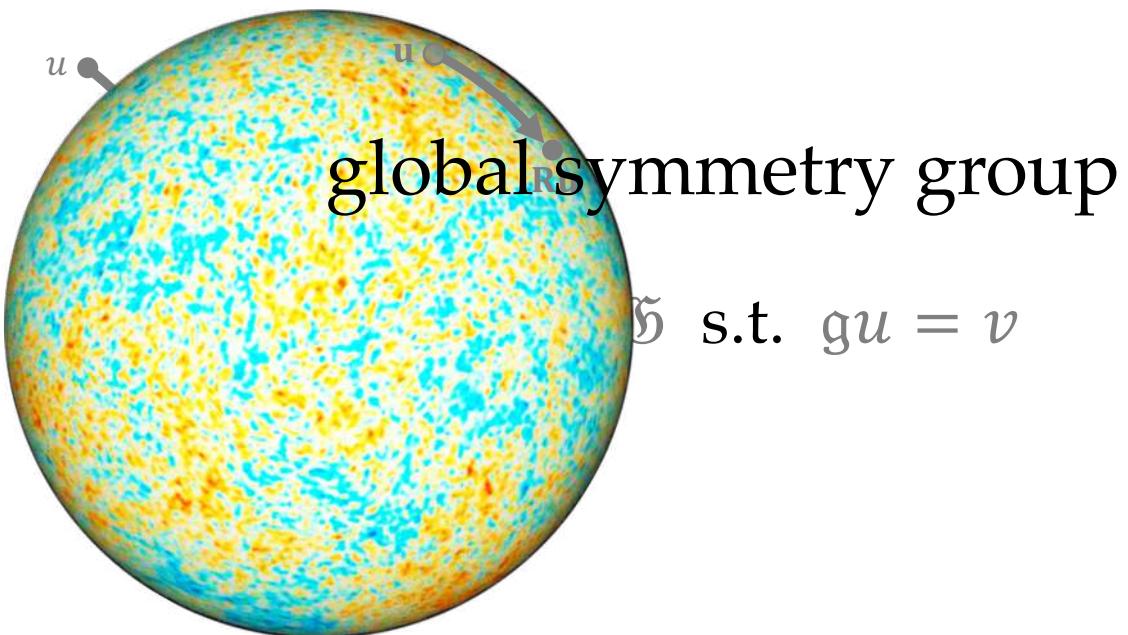
Convolution on the Sphere

$$((x \star \psi) \star \phi)(R) = \int_{SO(3)} (x \star \psi)(Q) \phi(R^{-1}Q)dQ$$



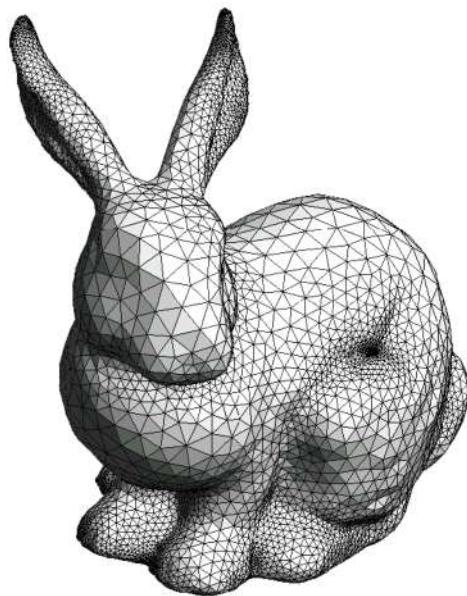
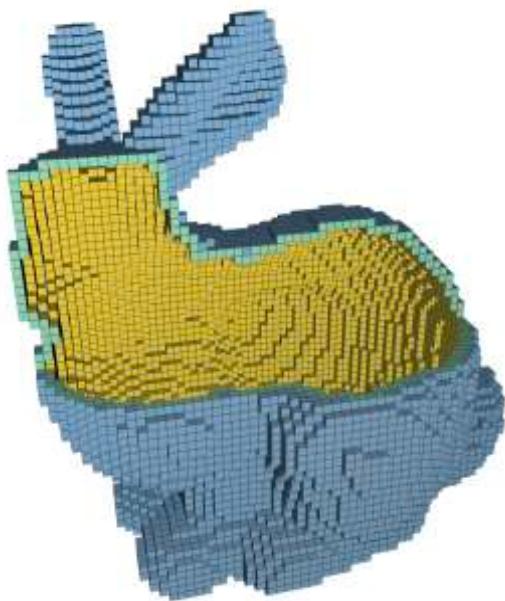
Cohen, Welling 2016

Homogeneous Spaces



GEODESICS & GAUGES

Why manifolds?



More efficient representation: no “waste”
for internal structures

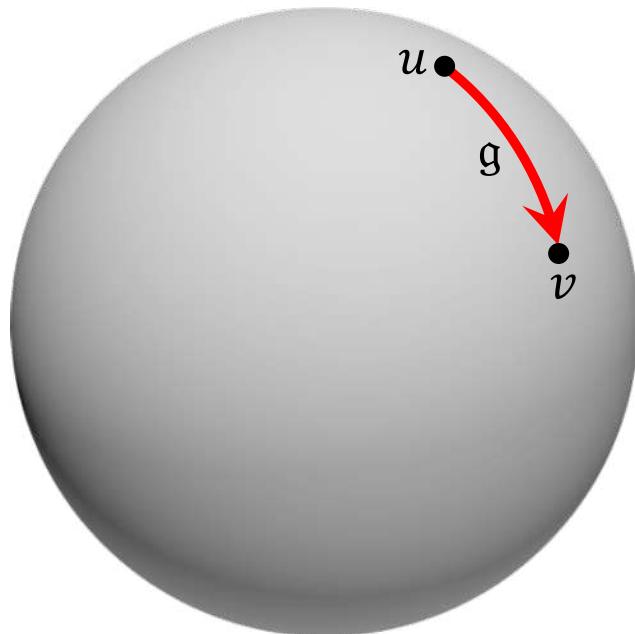
Natural model for
deformable shapes

Why manifolds?

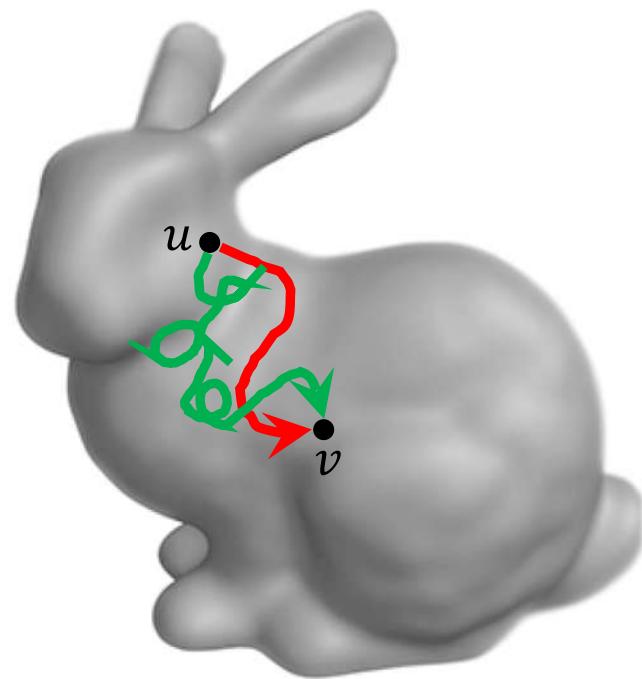
In protein modeling,
abstract out internal
structure that is irrelevant
for interactions + allow
some conformation
changes



Homogeneous Spaces

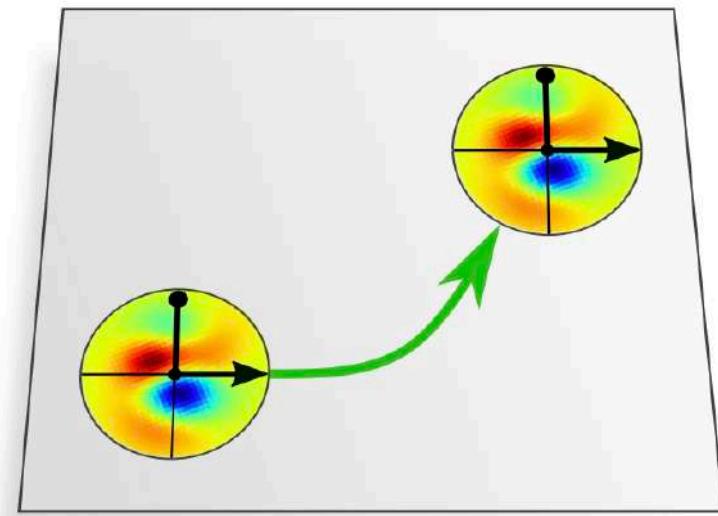
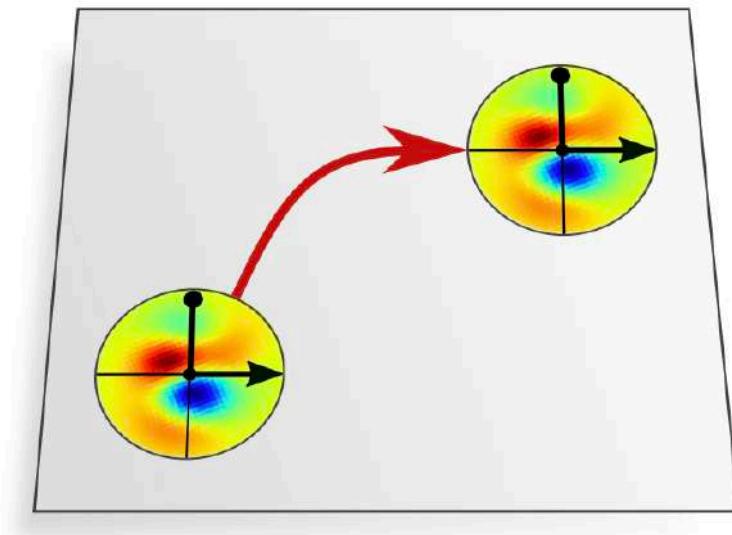


global symmetry group
 $\exists g \in \mathfrak{G} \text{ s.t. } gu = v$



no global symmetry group

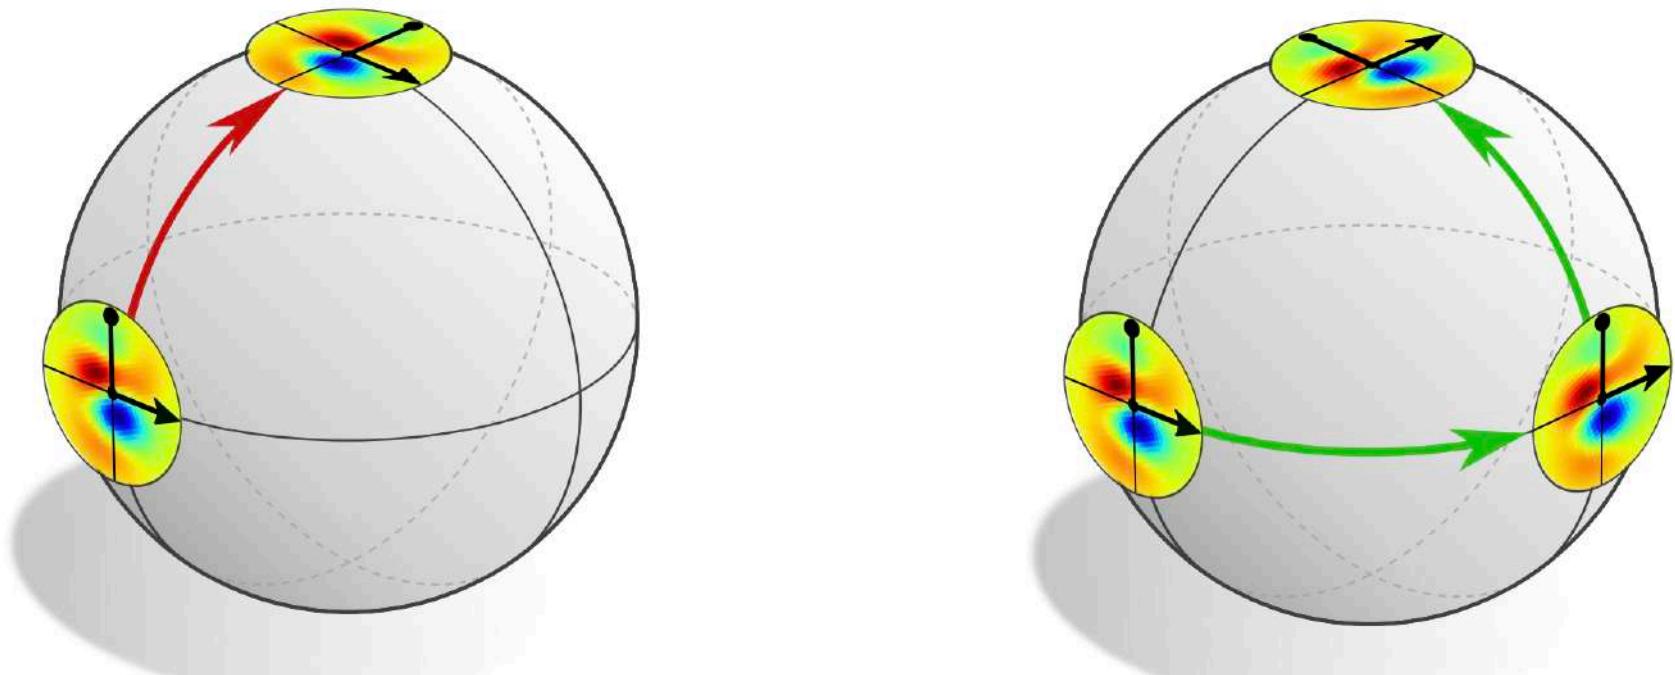
Euclidean Convolution



Euclidean space: Transport the filter around the domain

Figure: M. Weiler et al. 2021

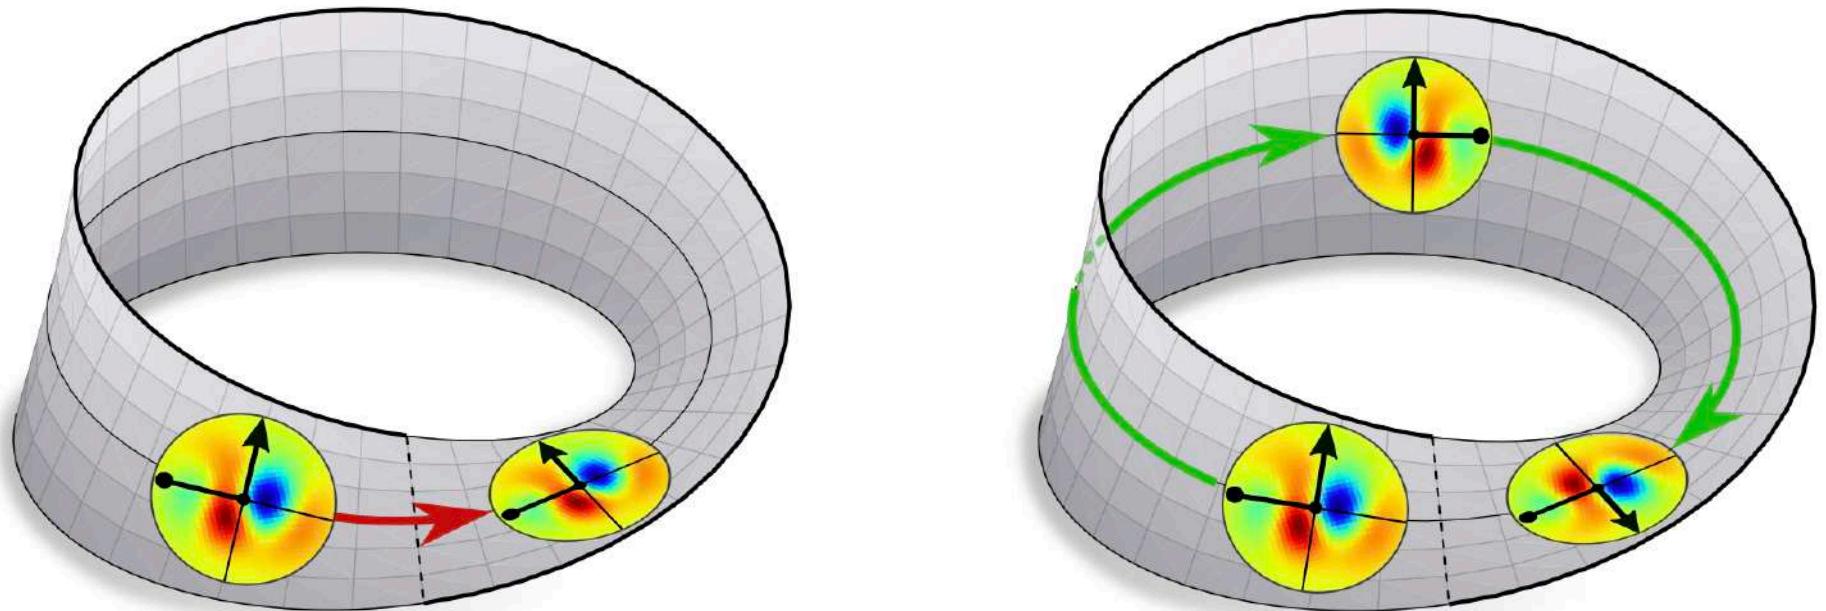
Non-Euclidean Convolution



Manifold: Result of transport is *path dependent*

Figure: M. Weiler et al. 2021

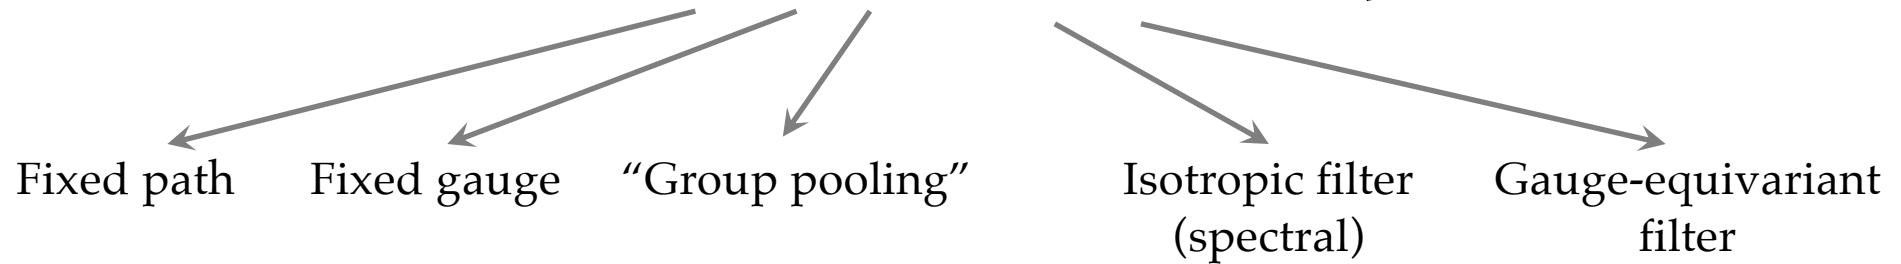
Non-Euclidean Convolution



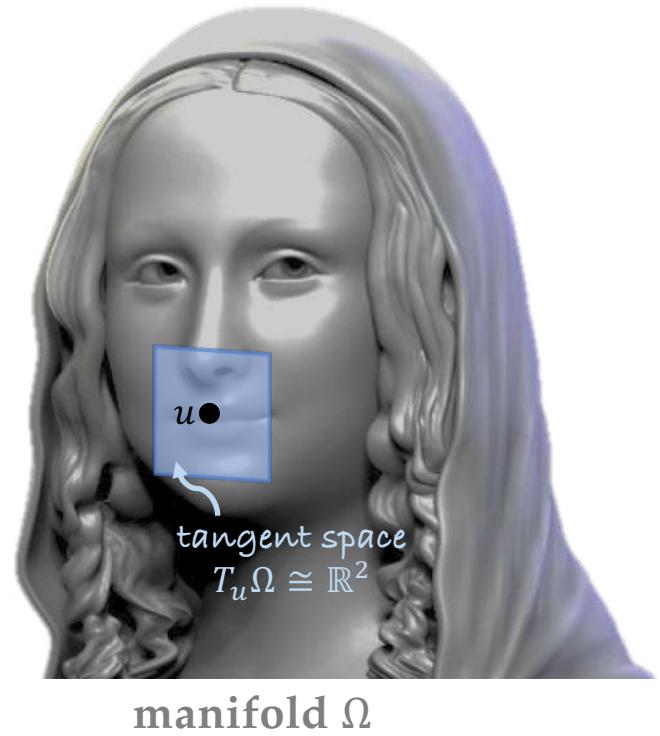
Manifold: Result of transport is *path dependent*

Figure: M. Weiler et al. 2021

Non-Euclidean Convolution Recipes



Manifolds

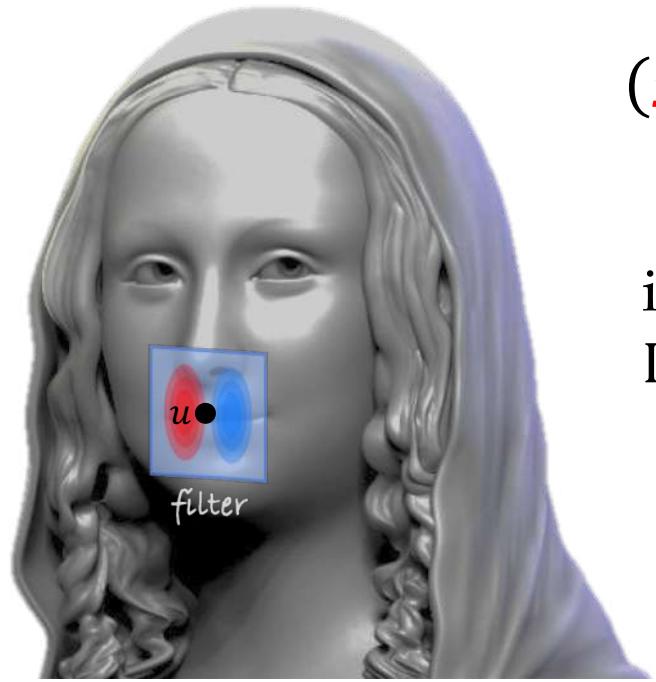


manifold = locally Euclidean space

Riemannian metric = local length / direction

Isometry = metric-preserving deformation

Geodesic CNNs



manifold Ω
isometry group $\text{Iso}(\Omega)$

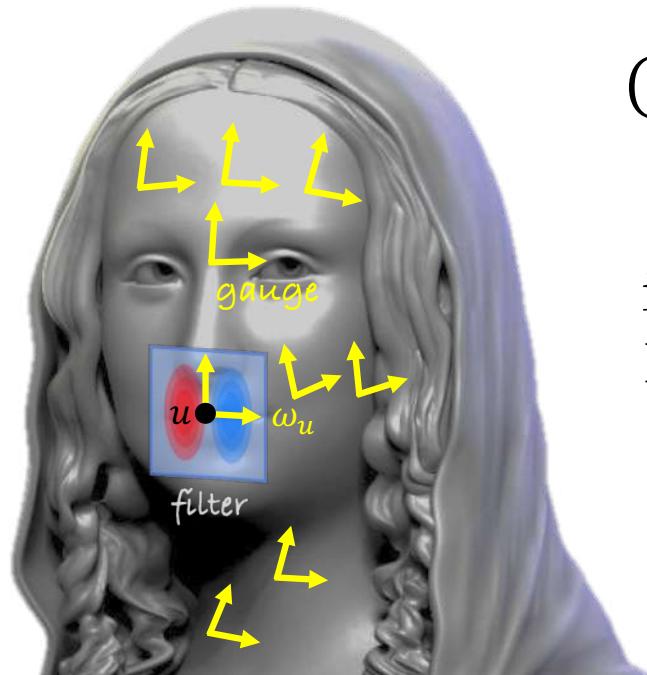
$$(x \star \psi)(u) = \int_{T_u \Omega} \psi(v) x(\exp_u v) dv$$

Exponential map
 $\exp_u: T_u \Omega \rightarrow \Omega$

intrinsic filter = invariant to isometry group
 $\text{Iso}(\Omega)$ of metric-preserving deformations

Masci et al. 2015; Boscaini et al. 2016; Monti et al. 2017

Geodesic CNNs



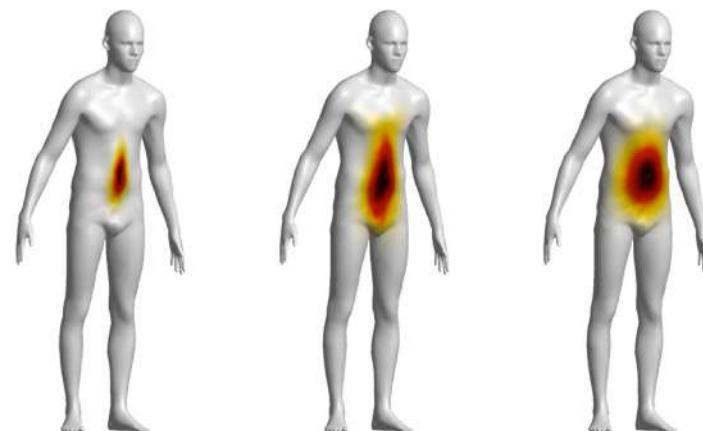
manifold Ω
isometry group $\text{Iso}(\Omega)$

Masci et al. 2015; Boscaini et al. 2016; Monti et al. 2017

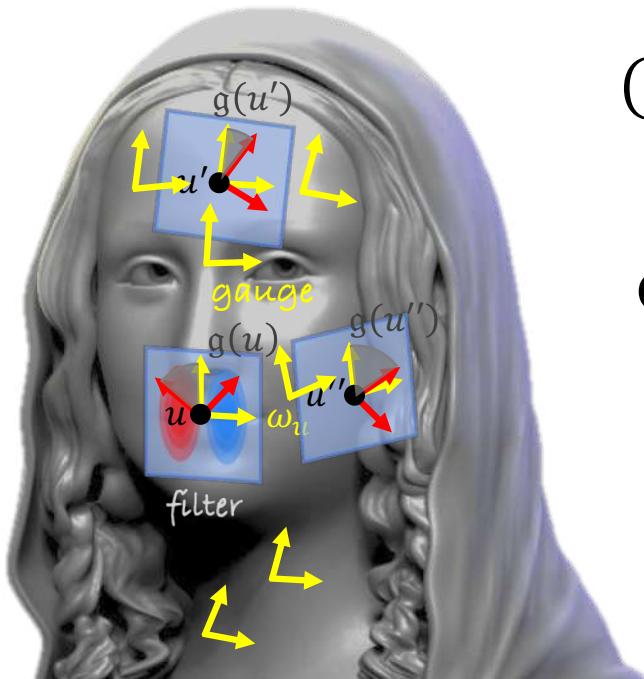
$$(x \star \psi)(u) = \int_{\mathbb{R}^2} \psi(v)x(\exp_u \omega_u v) dv$$

*local reference frame
 $\omega_u: \mathbb{R}^2 \rightarrow T_u \Omega$*

intrinsic filter = invariant to isometry group
 $\text{Iso}(\Omega)$ of metric-preserving deformations



Gauge-equivariant CNNs



manifold Ω
structure group \mathfrak{G}

$$(x \star \psi)(u) = \int_{\mathbb{R}^2} \psi(v) x(\exp_u \omega_u v) dv$$

Gauge defined up to gauge transformation

$$g: \Omega \rightarrow \mathfrak{G}$$

Structure Group

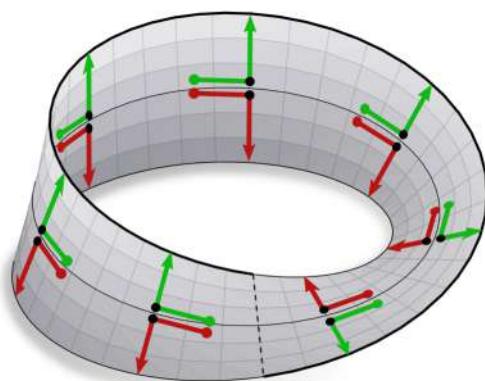
A gauge is defined up to a *gauge transformation* $g: \Omega \rightarrow \mathfrak{G}$

- | | | |
|-------------------------------|-----------------|-------------------------------------|
| • “Naked” manifold | $GL(s)$ | invertible matrices |
| • Manifold+orientation | $GL^+(s)$ | invertible matrices with $\det > 0$ |
| • Manifold+volume | $SL(s)$ | matrices with $\det = 1$ |
| • Manifold+metric | $O(s)$ | orthogonal matrices |
| • Manifold+metric+orientation | $SO(s)$ | orthogonal matrices with $\det = 1$ |
| • Manifold+frame field | $\{\text{id}\}$ | identity (no ambiguity) |

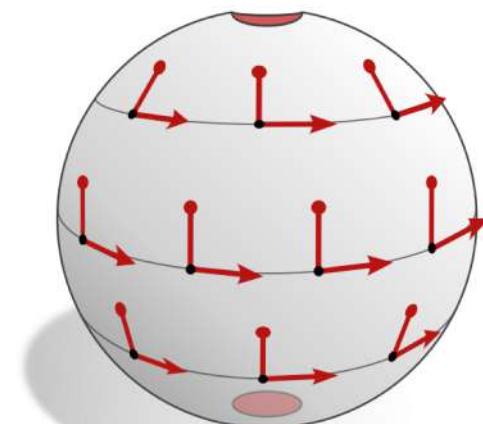
Structure Group



rotation
 $SO(2)$



reflection
 R



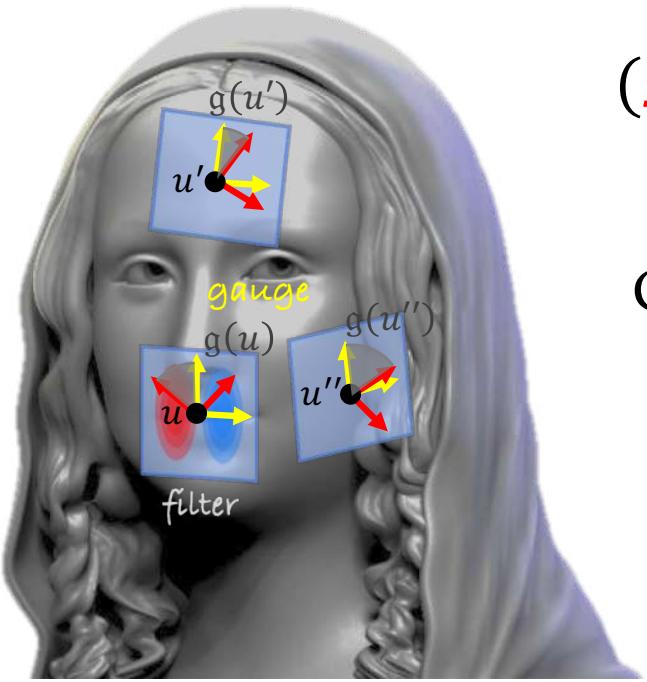
fixed gauge
 $\{id\}$

Structure Group

A gauge is defined up to a *gauge transformation* $g: \Omega \rightarrow \mathfrak{G}$

• “Naked” manifold	$GL(s)$	invertible matrices
• Manifold+orientation	$GL^+(s)$	invertible matrices with $\det > 0$
• Manifold+volume	$SL(s)$	matrices with $\det = 1$
• Manifold+metric	$O(s)$	orthogonal matrices
• Manifold+metric+orientation	$SO(s)$	orthogonal matrices with $\det = 1$
• Manifold+frame field	$\{\text{id}\}$	identity (no ambiguity)

Gauge-equivariant CNNs



manifold Ω
structure group $\mathfrak{G} = \text{SO}(2)$

Cohen et al. 2019

$$(x \star \psi)(u) = \int_{\mathbb{R}^2} \psi(v) \rho(\exp_u \omega_u v) dv$$

Gauge defined up to gauge transformation

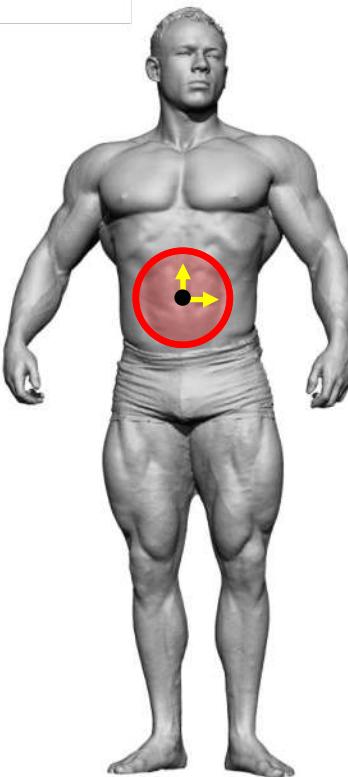
$$g: \Omega \rightarrow \text{SO}(2)$$

gauge-equivariant filter

$$\psi(g^{-1}v) = \rho(g^{-1})\psi(v)\rho(g)$$

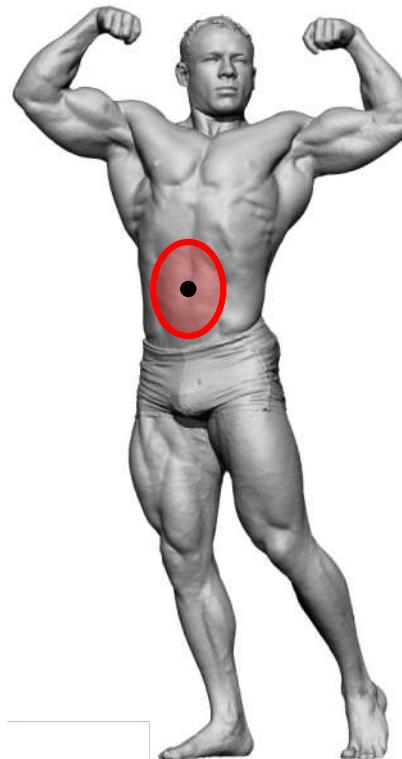
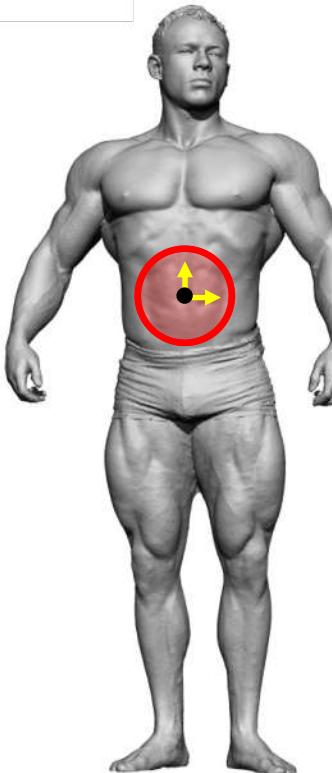
Two Types of Symmetry

**Local gauge
transformation**

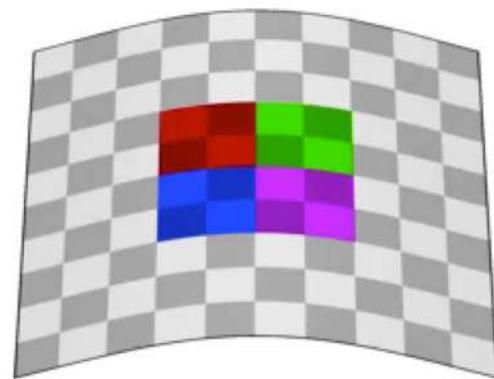
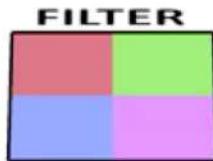


Two Types of Symmetry

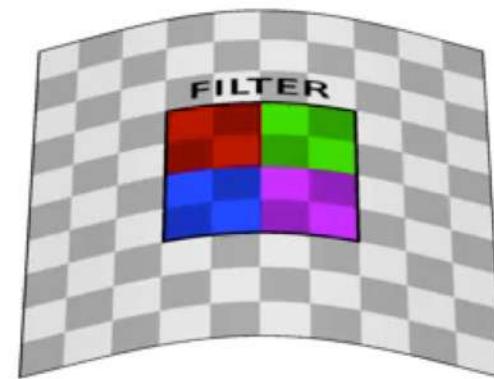
**Local gauge
transformation**



**Global
deformation**

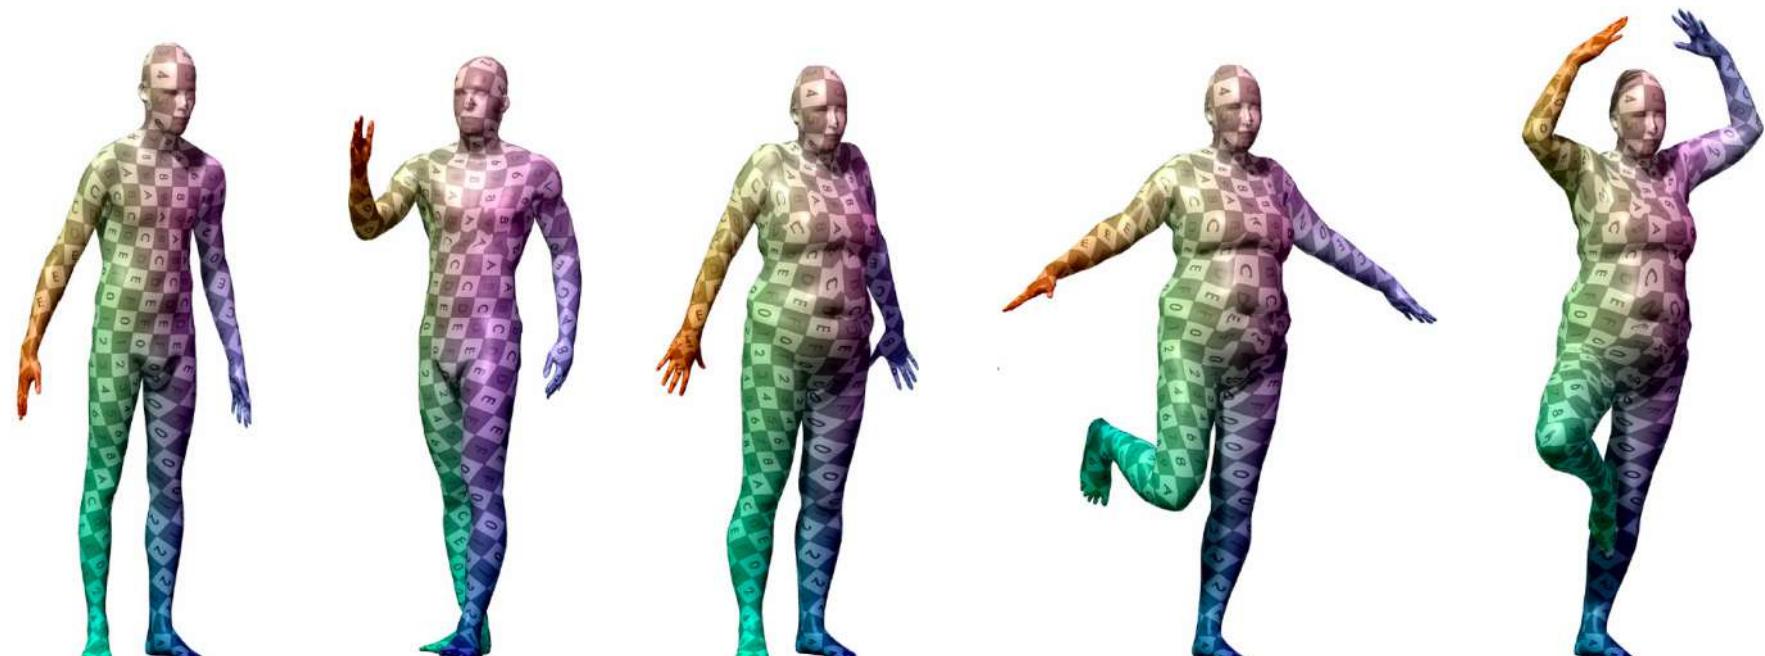


**Euclidean (extrinsic)
convolution**



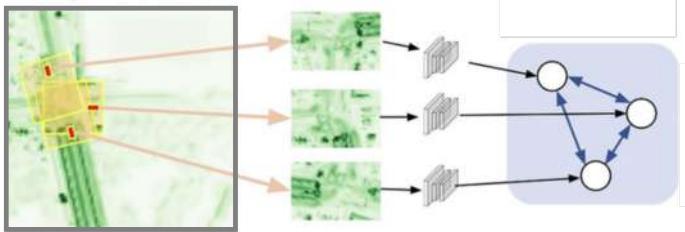
**Geometric (intrinsic)
convolution**

Applications in Computer Graphics



Correspondence of deformable 3D shapes

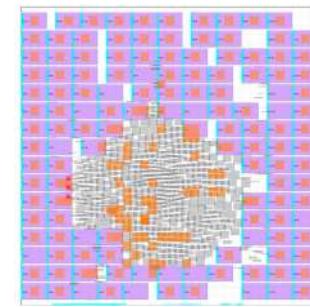
Wild zoo of Applications!



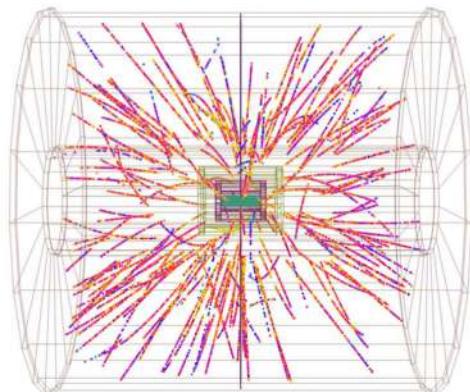
Self-driving cars



Navigation



Chip design



Particle physics



Protein science



Drug discovery

3D VISION & GRAPHICS

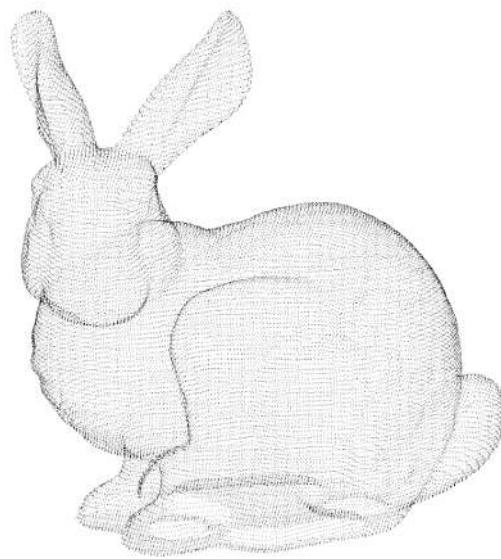
Geometric Data



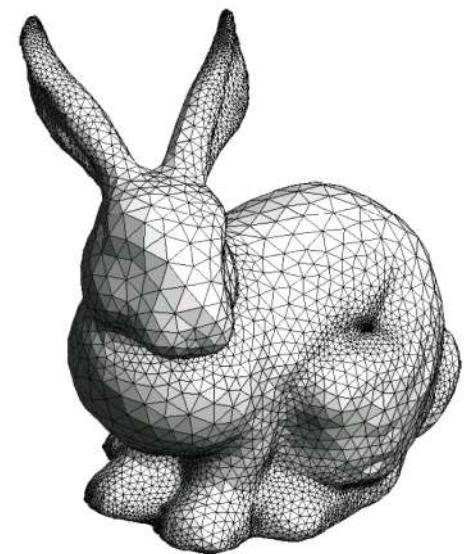
Image-based



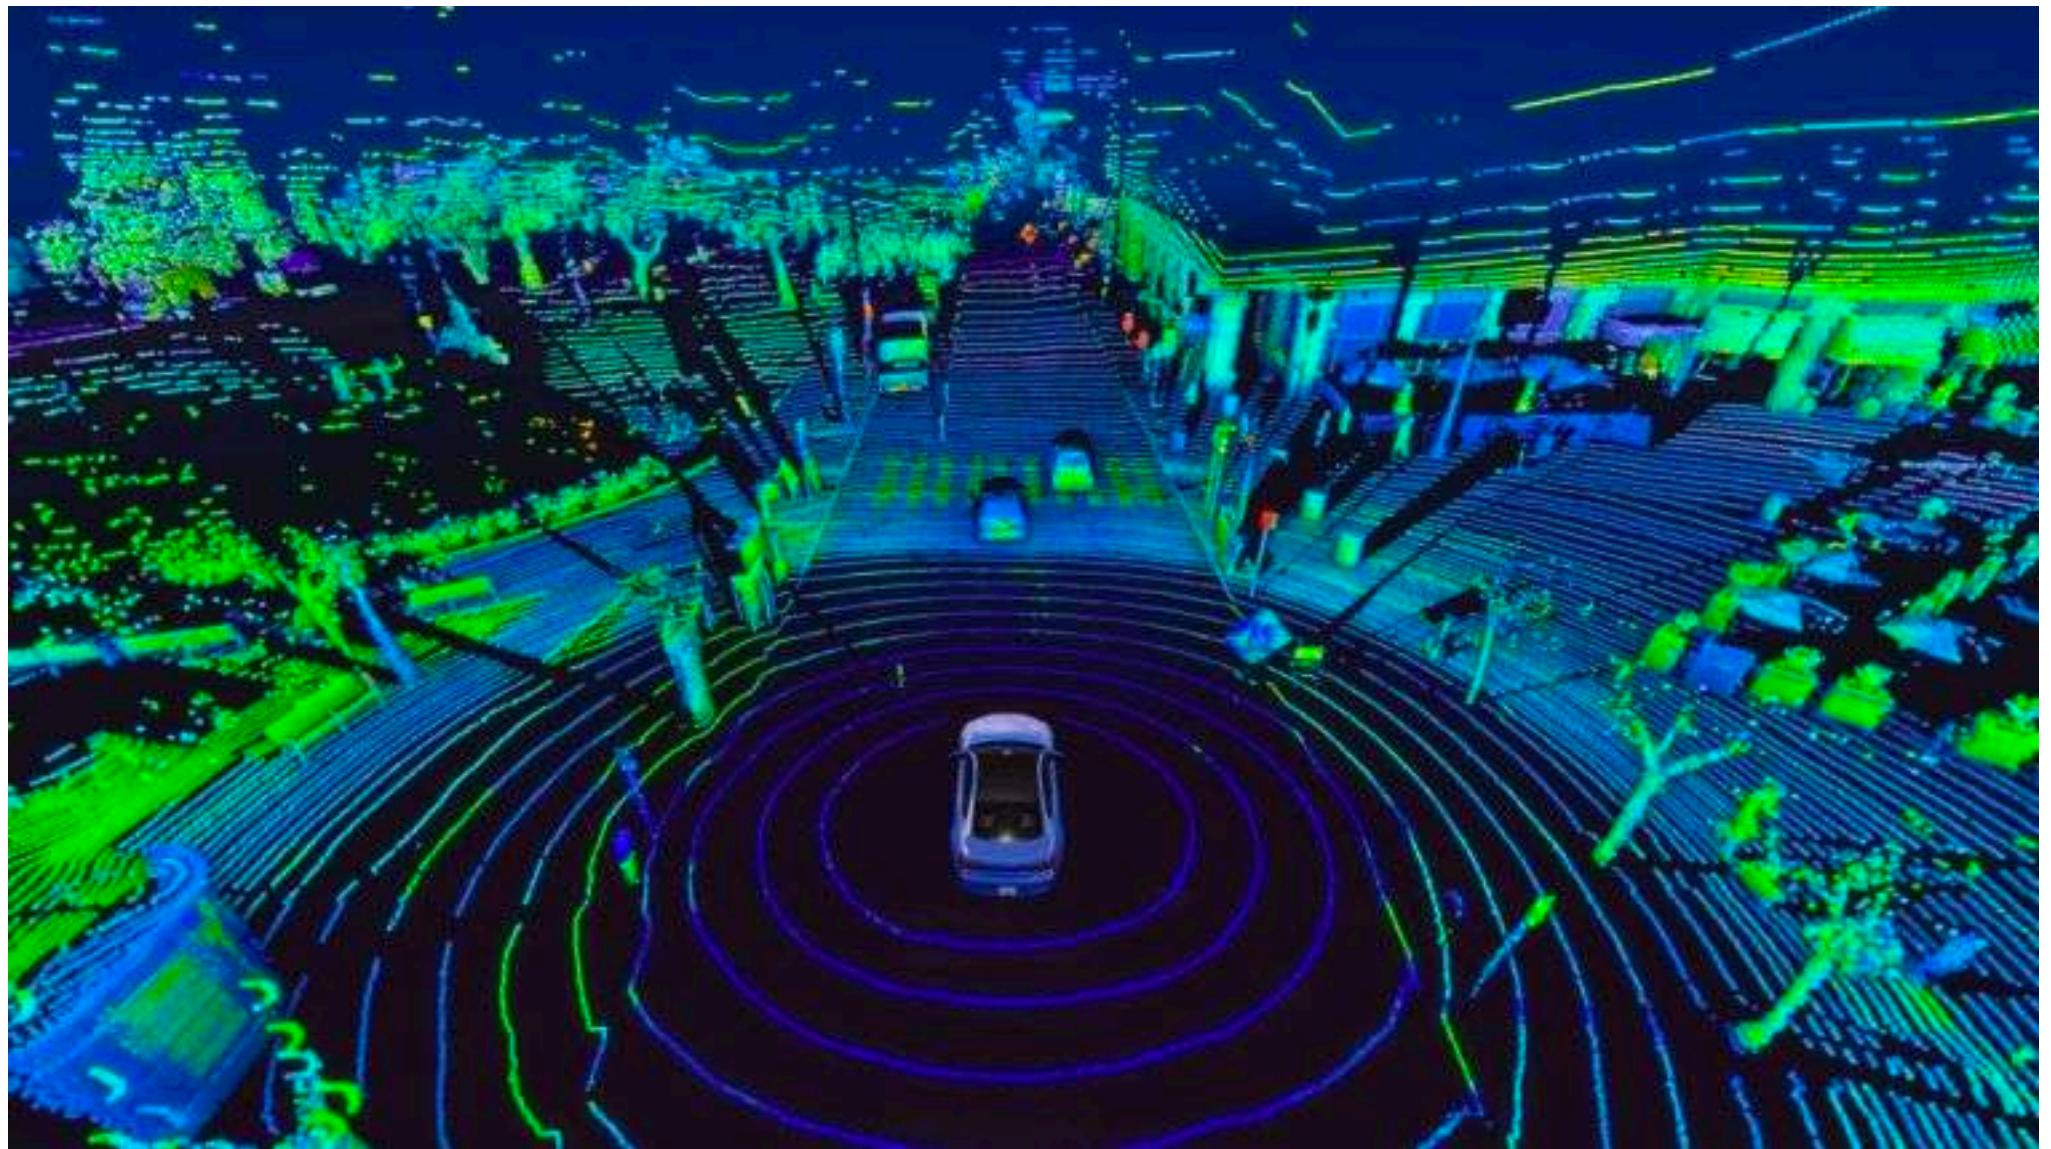
Volumetric



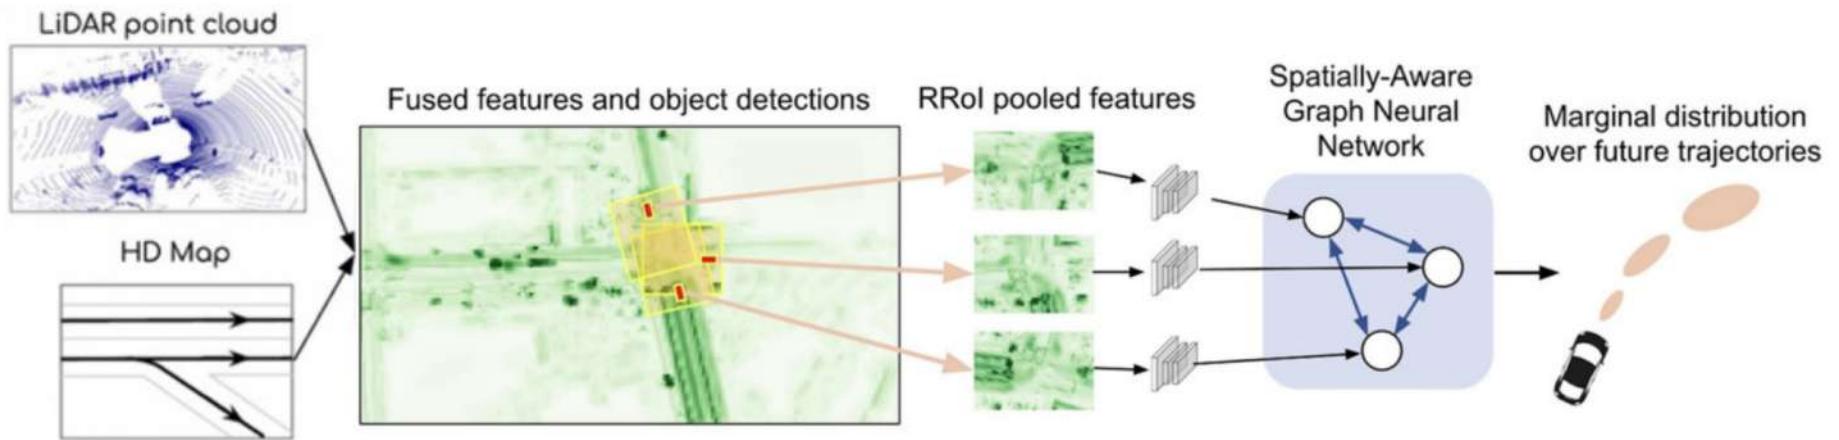
Point cloud



Mesh-based



Self-Driving Cars



Casas et al. 2020 (from Raquel Urtasun's talk)

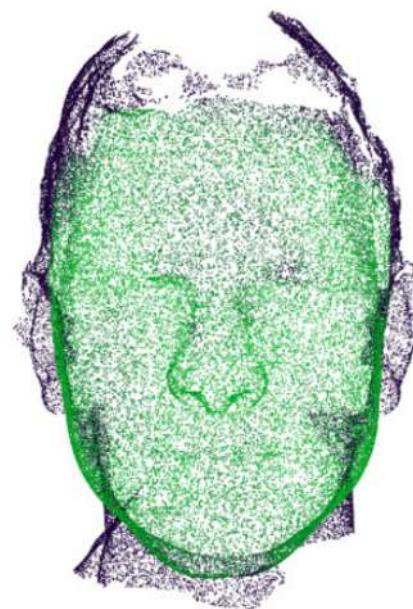
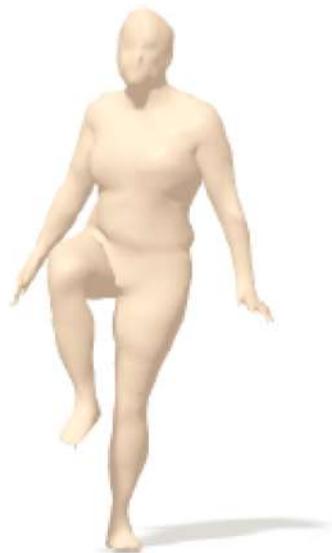
FaceShift 2015



GDC



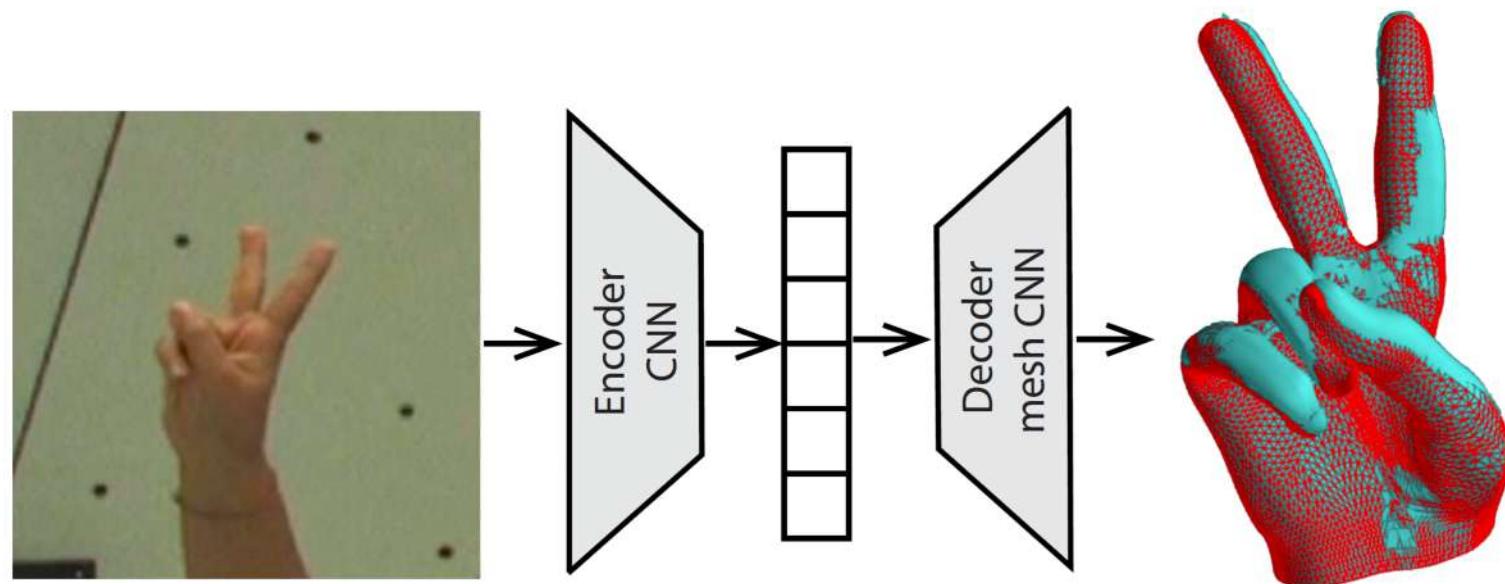
3D Avatars



Litany et B 2018

Bahri et B 2020

3D Hand Reconstruction





Snap Acquires Ariel AI To Enhance AR Features





Face from DNA



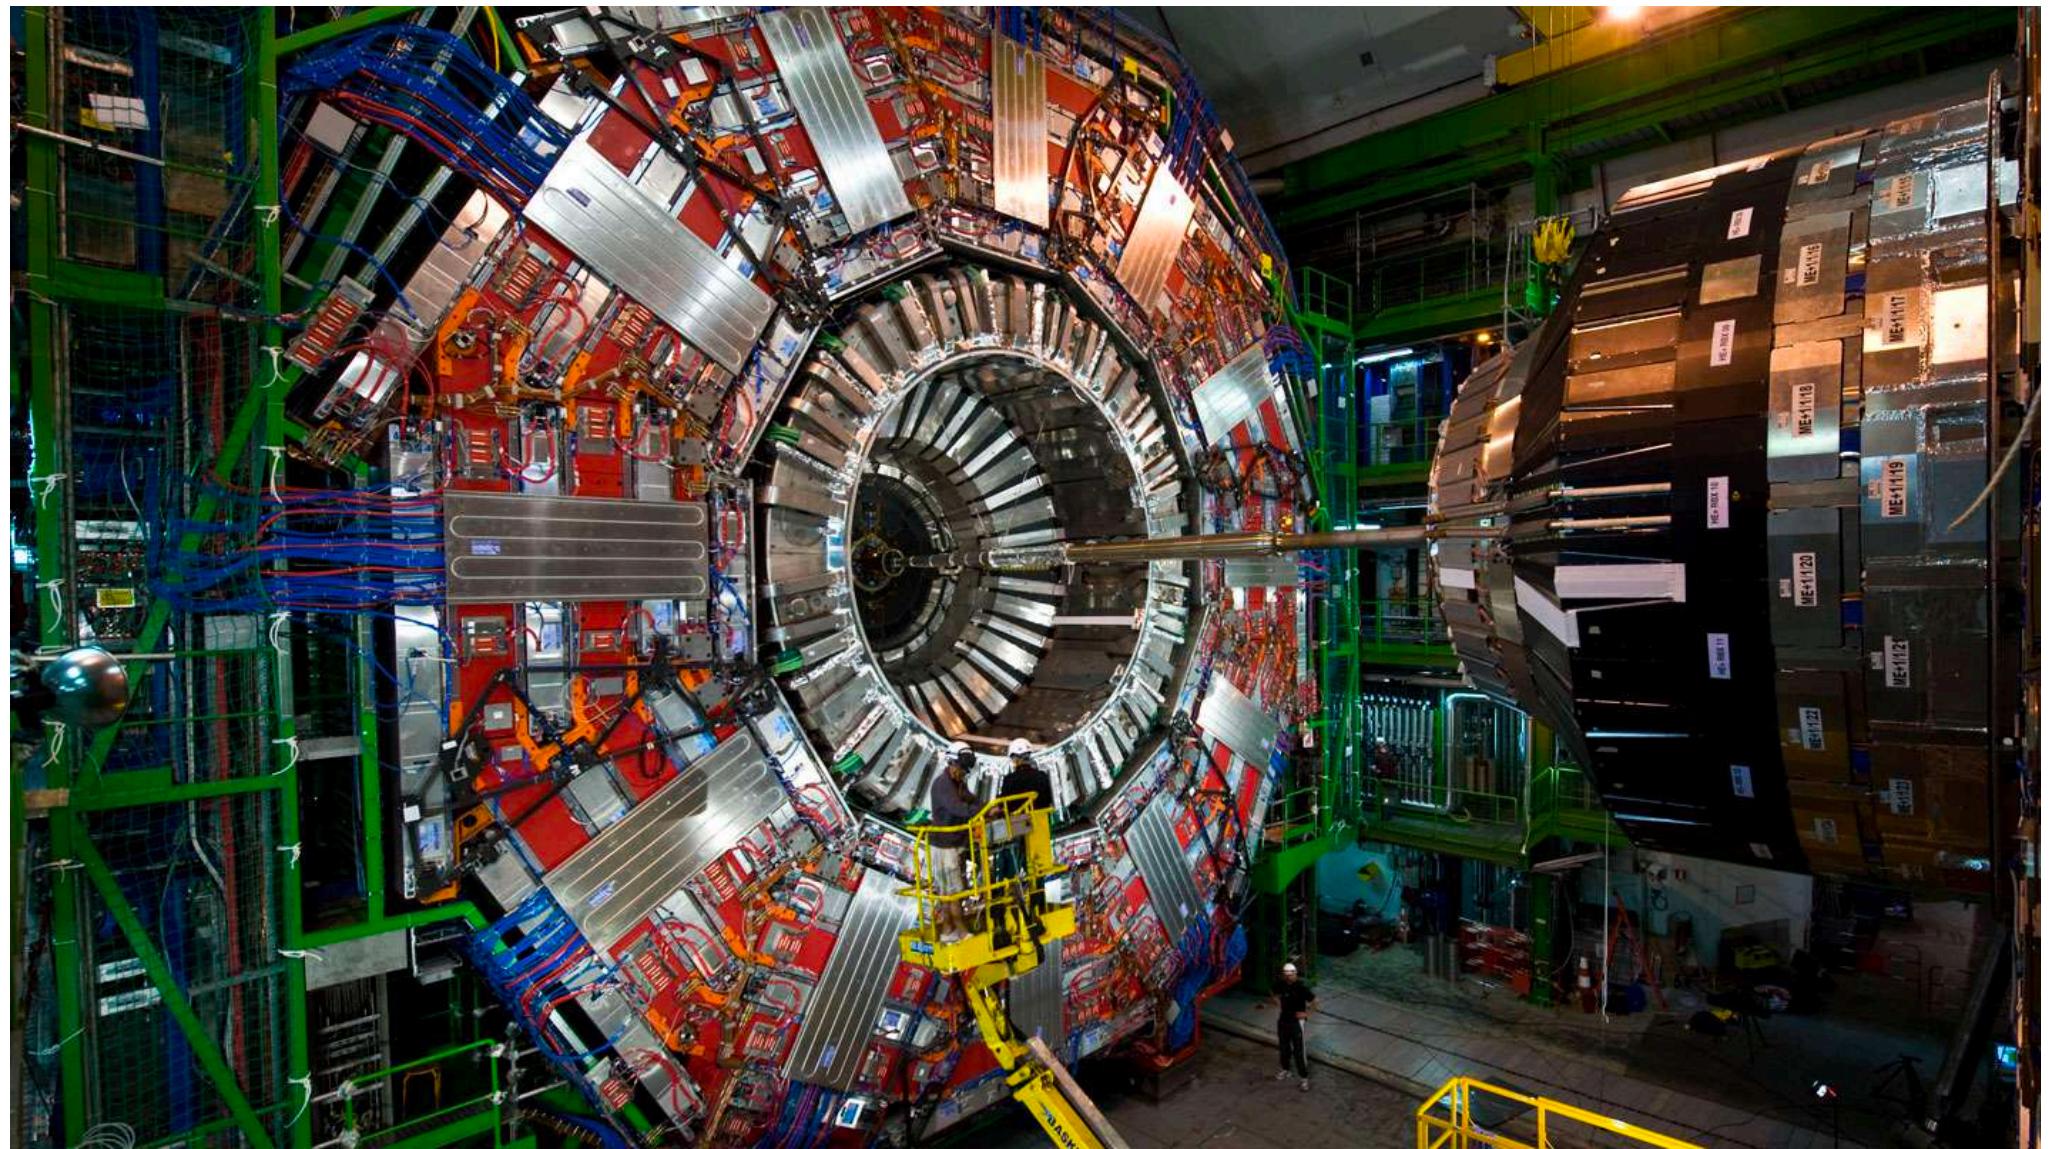
collaboration with P. Claes

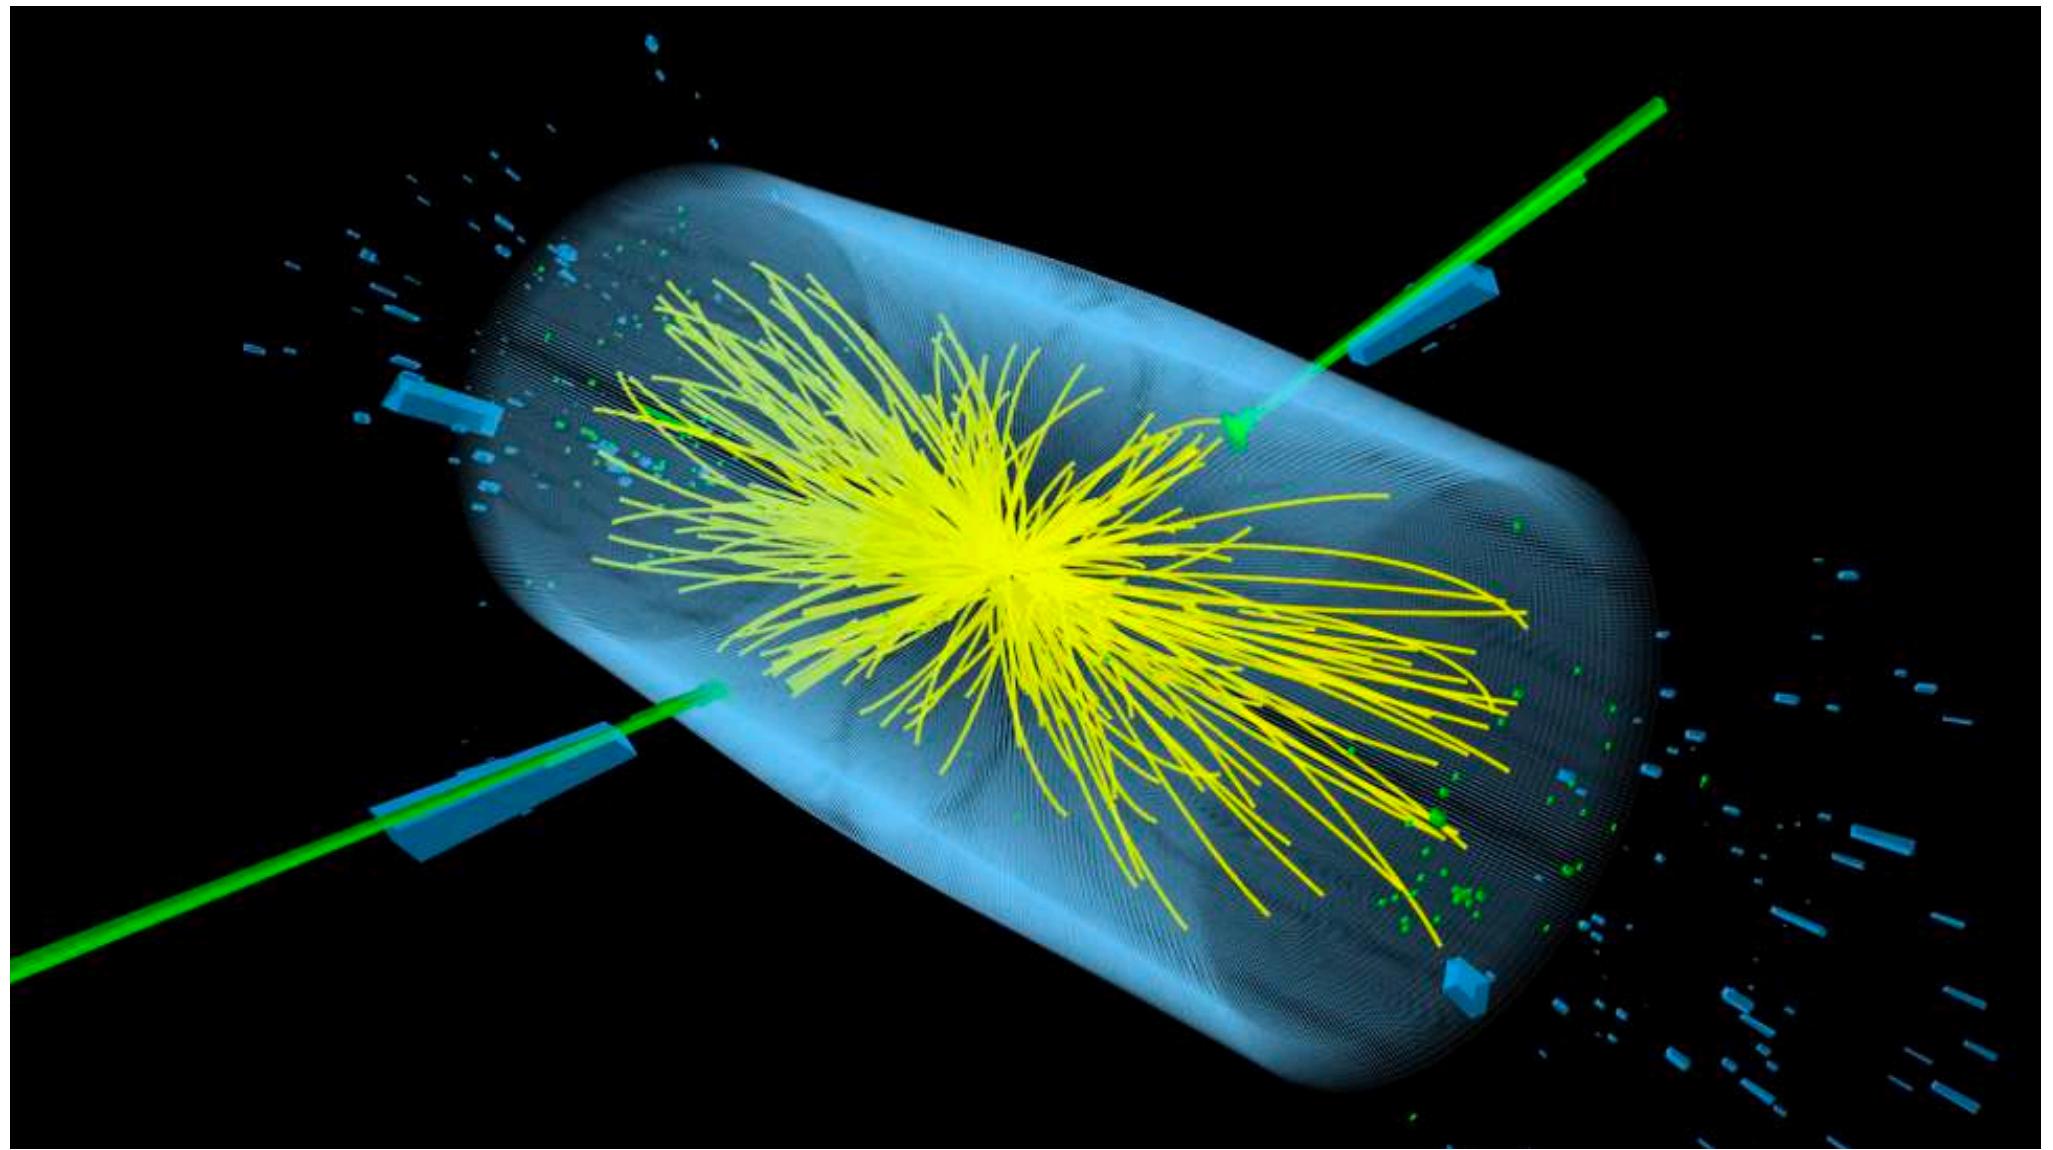
PHYSICS

“It has been amazing to see how in the last two years Graph ML has become very popular in the field of physics.”

—Kyle Cranmer (NYU)

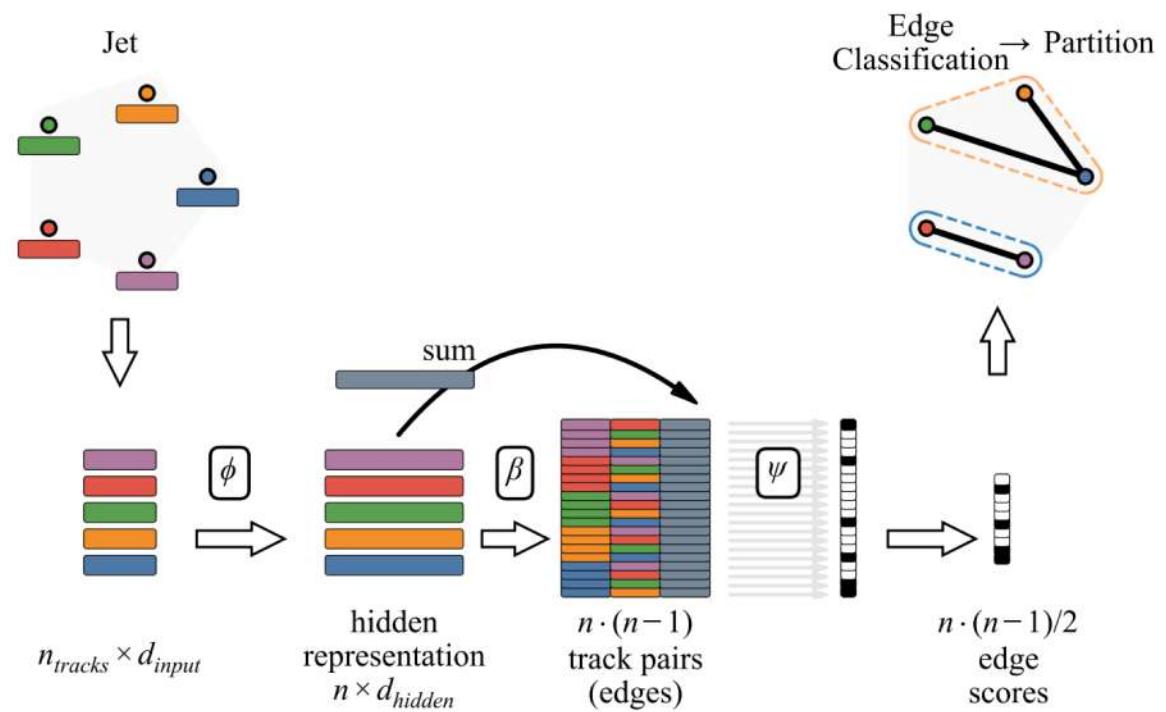
one of the discoverers of the Higgs boson





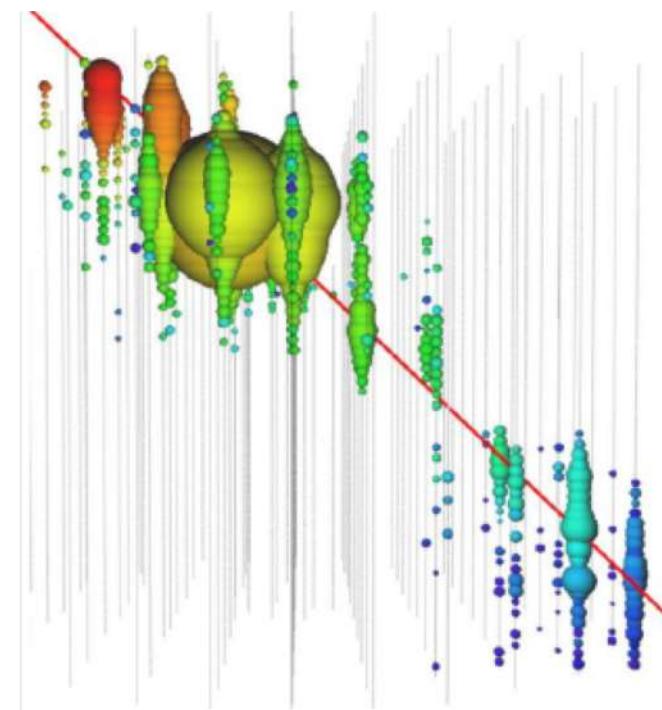
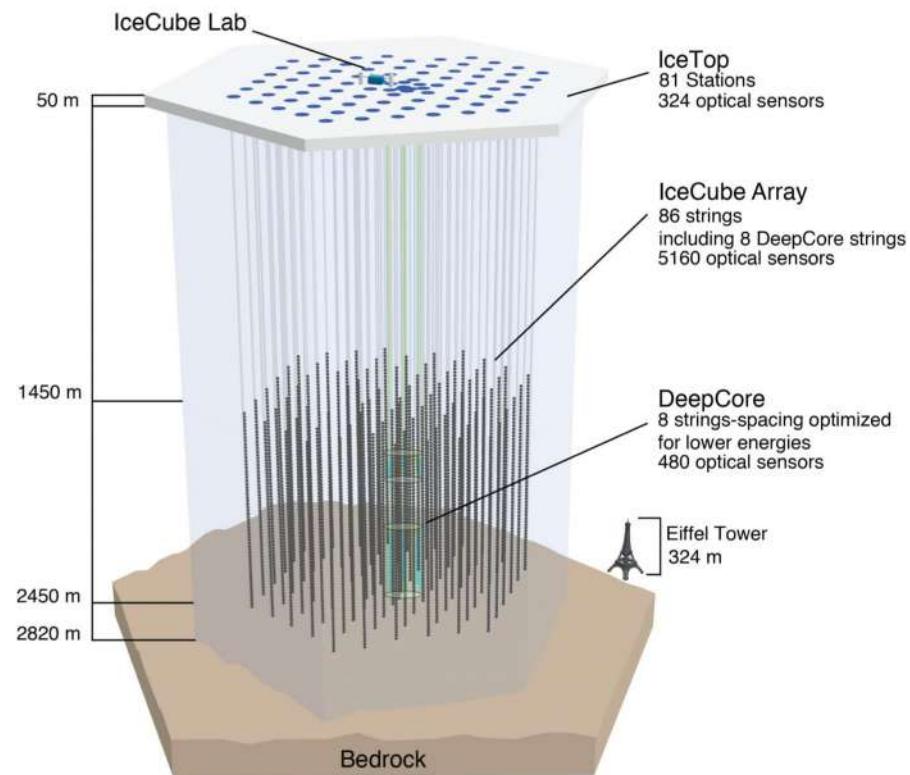
Particle Jet Reconstruction & Segmentation

Set2graph architecture for partitioning observed particles based on their point of origin





Detector Modeling



Light deposition for a high-energy single muon in IceCube detector

SOCIAL NETWORKS

Recommender Systems

PinSage: A new graph convolutional neural network for web-scale recommender systems



Pinterest Engineering Follow

Aug 15, 2018 · 8 min read



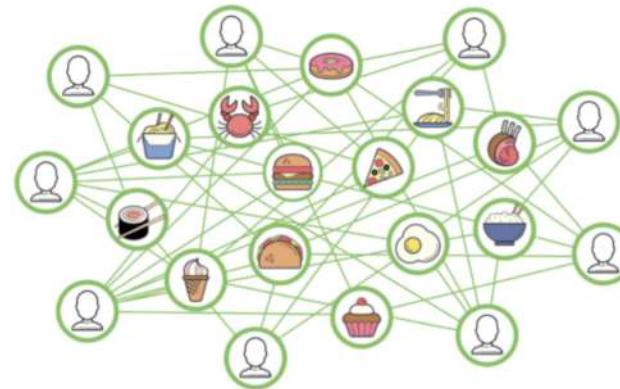
Ruining He | Pinterest engineer, Pinterest Labs

Food Discovery with Uber Eats: Using Graph Learning to Power Recommendations

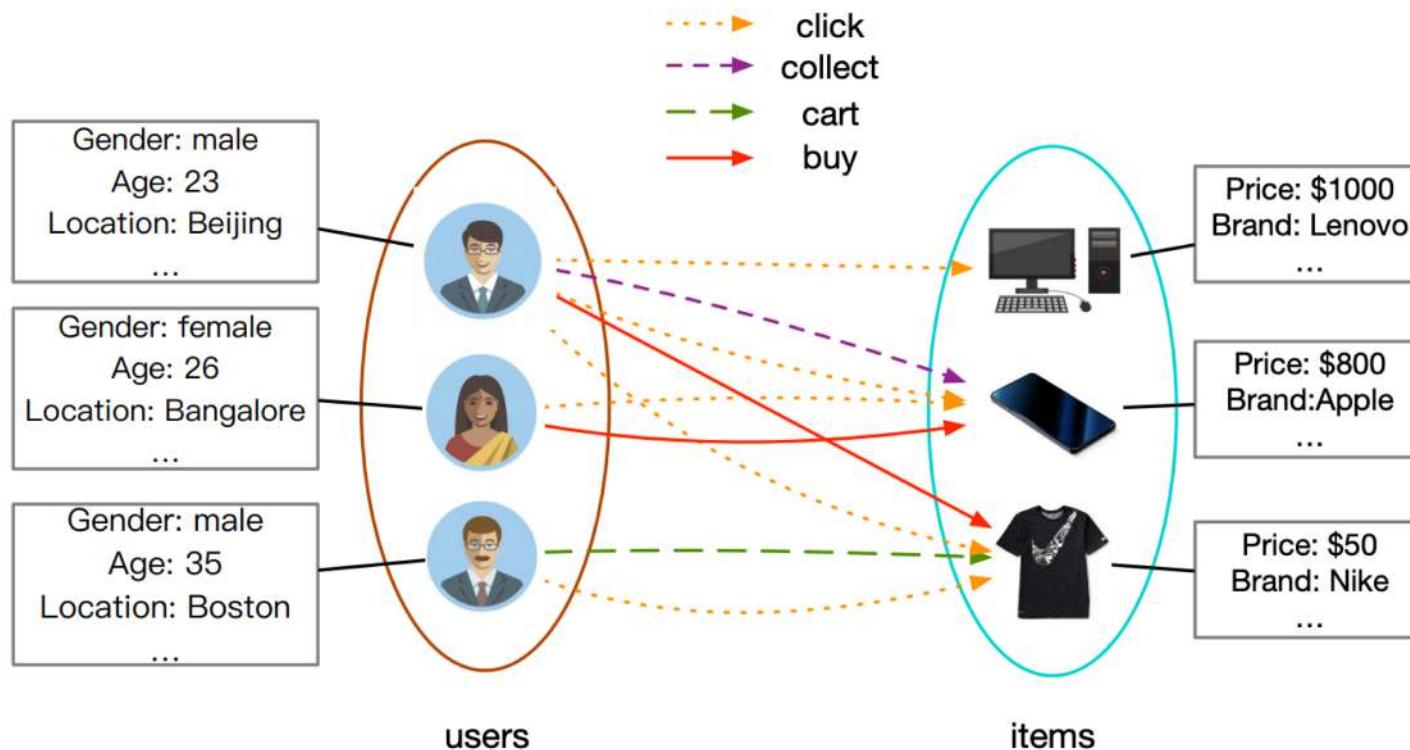
Ankit Jain, Isaac Liu, Ankur Sarda, and Piero Molino

0

December 4, 2019



Recommender Systems

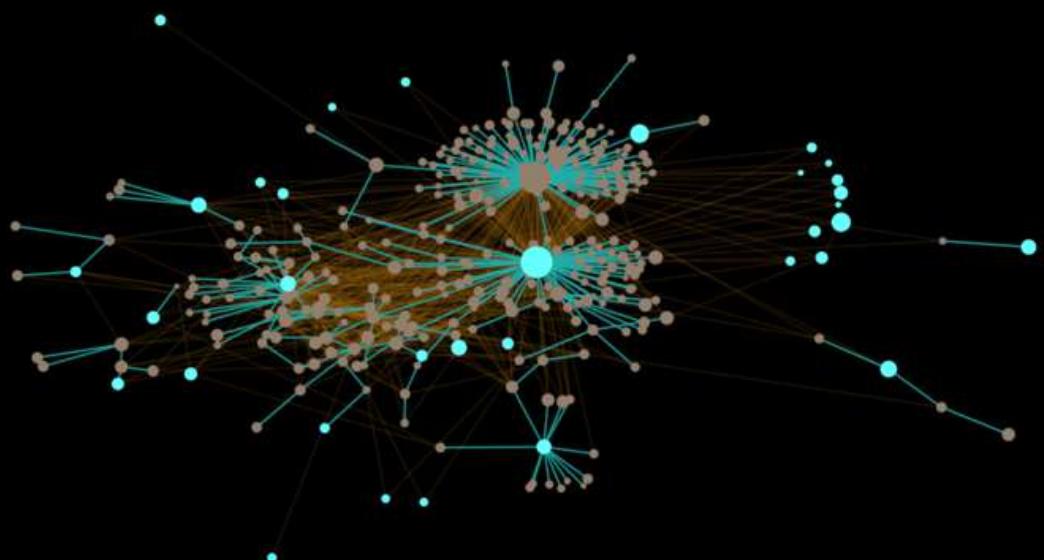


AliGraph (Zhu et al. 2019)

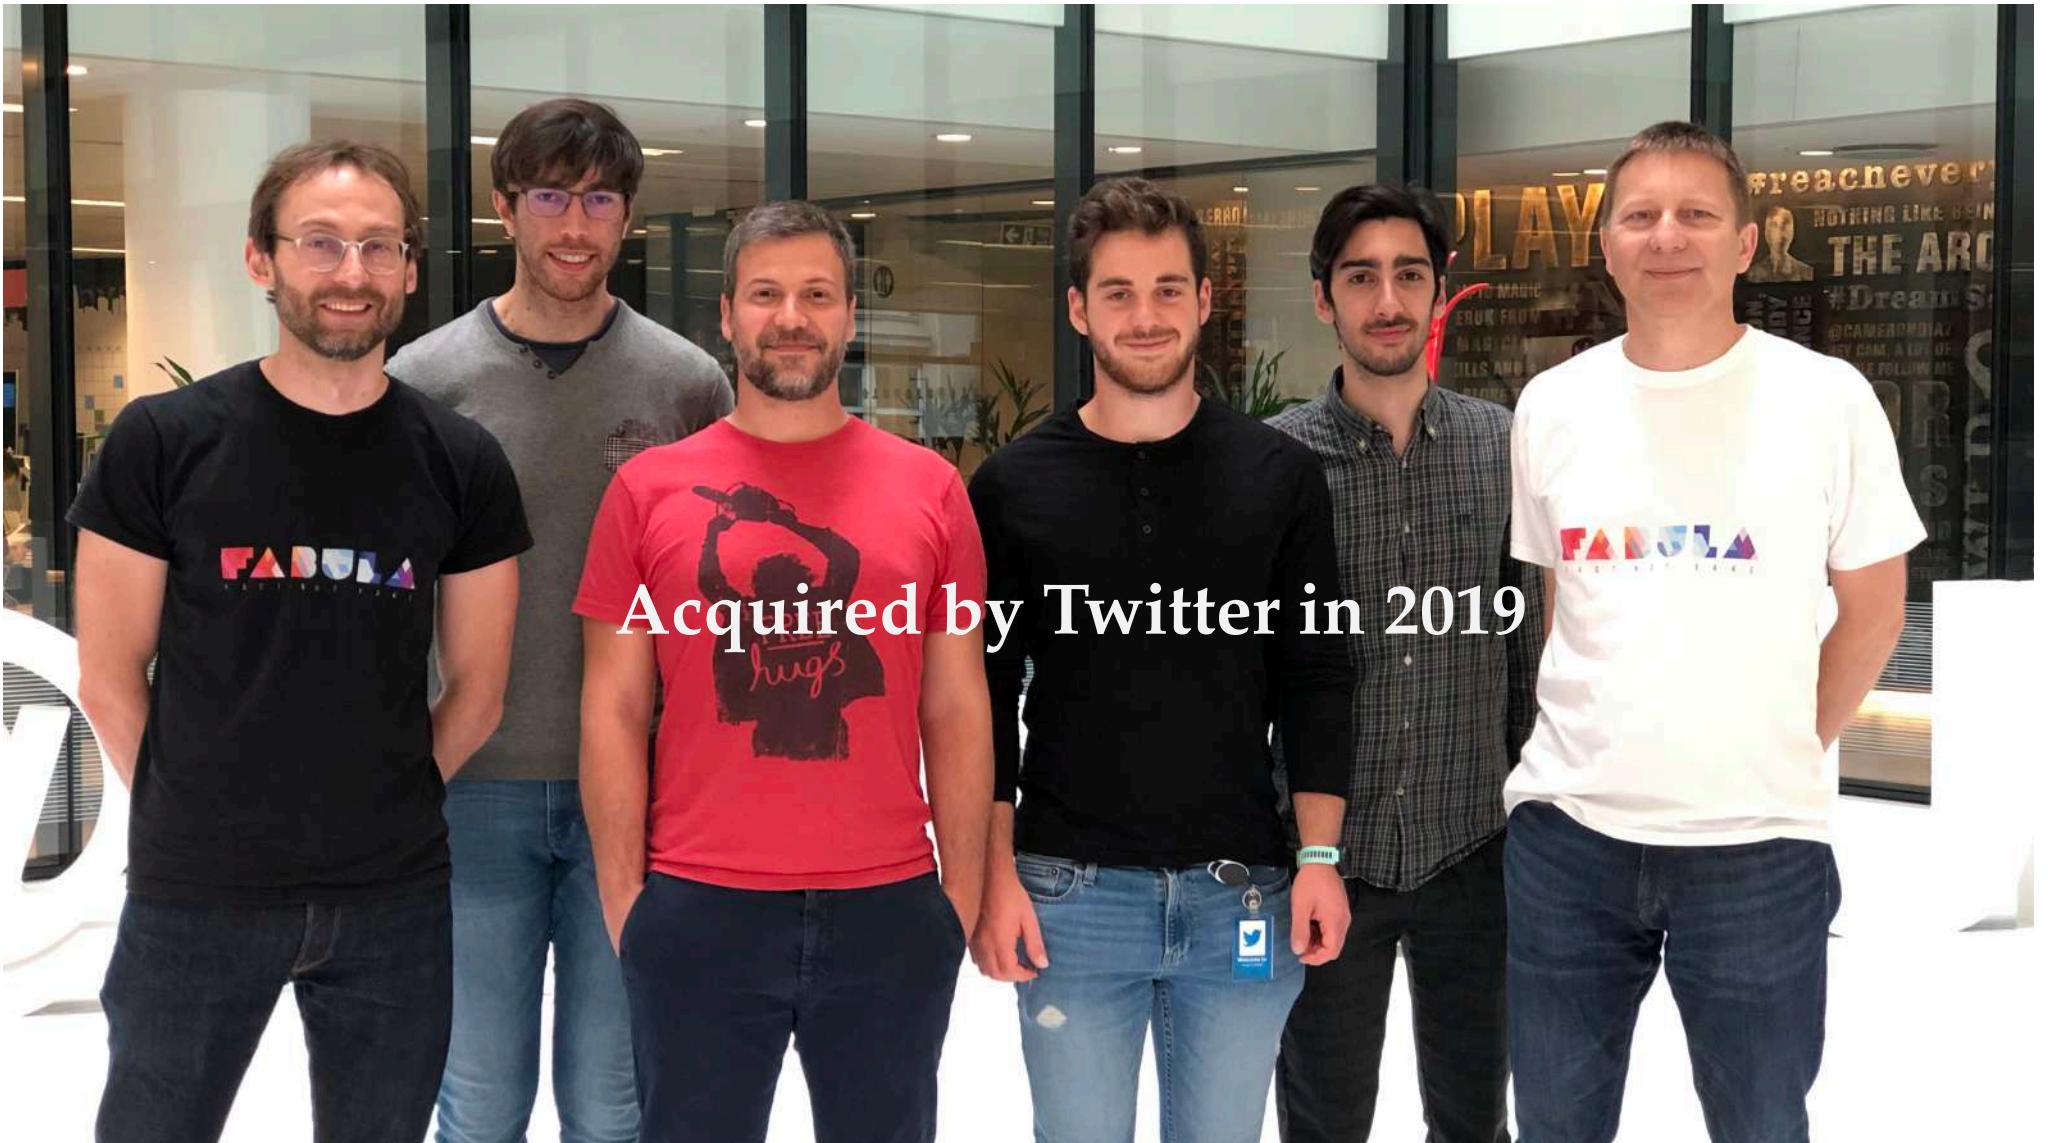
Safe spaces for South
Asia's vultures p. 1086

Synthetic Notch receptors
enhance T cell therapies p. 1112

Dietary fiber fights
diabetes p. 1151



Monti et B 2019



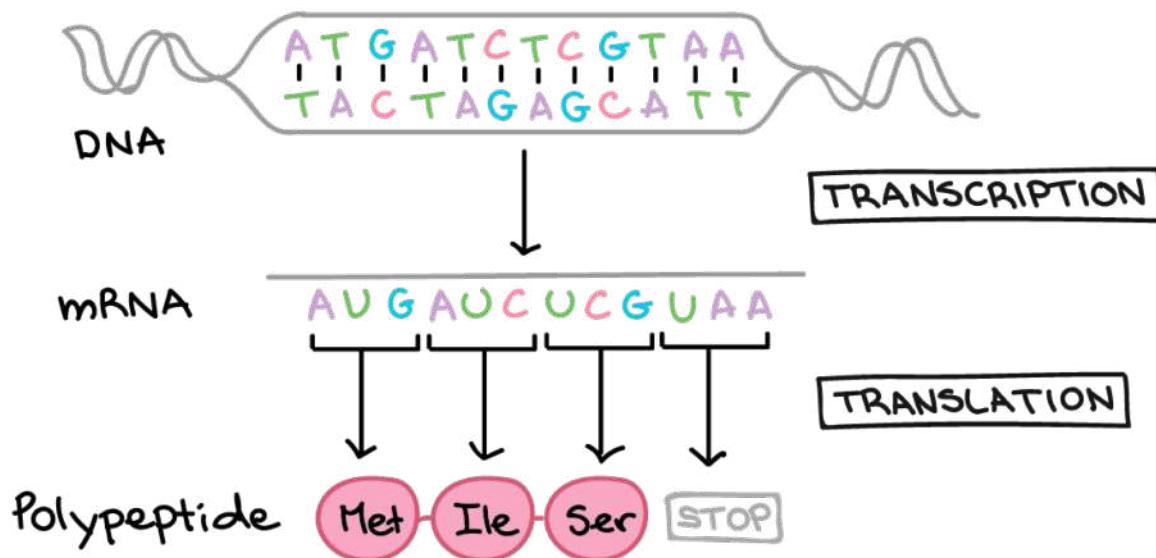
Acquired by Twitter in 2019

STRUCTURAL BIOLOGY

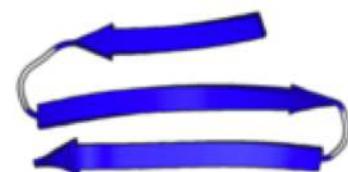
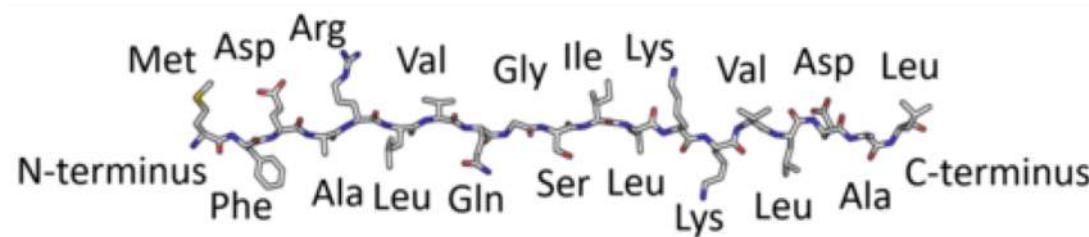
“In 2020, exciting progress has been made in protein structure prediction, a key problem in bioinformatics. Yet, ultimately the chemical and geometric pattern displayed at the surface of these molecules are critical for protein function.”

—Bruno Correia (EPFL Lab of Protein Design)

“Central Dogma”



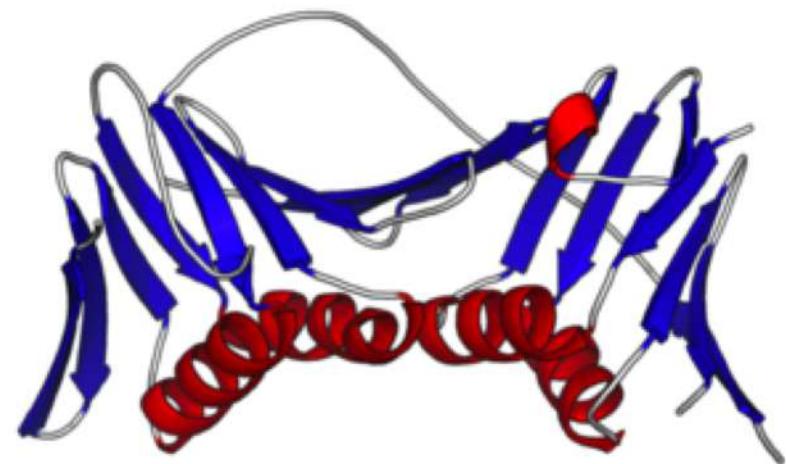
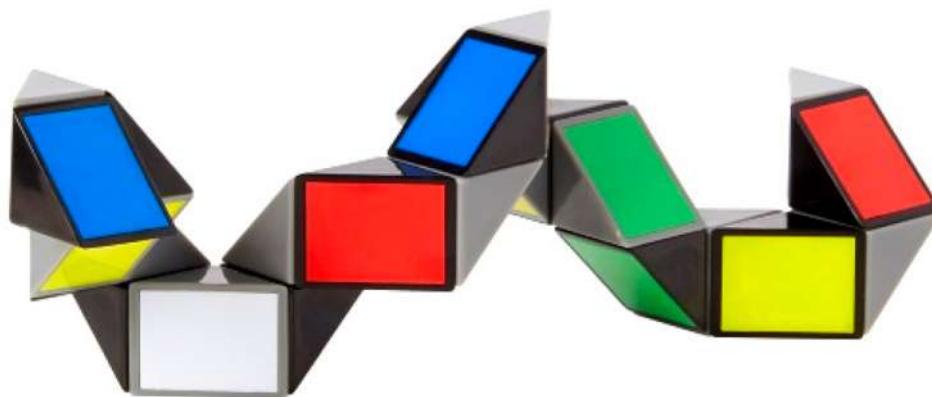
TTT	Phe	TCT	Ser	TAT	Tyr	TGT	Cys
TTC	Phe	TCC	Ser	TAC	Tyr	TGC	Cys
TTA	Leu	TCA	Ser	TAA	stop	TGA	stop
TTG	Leu	TCG	Ser	TAG	stop	TGG	Trp
CTT	Leu	CCT	Pro	CAT	His	CGT	Arg
CTC	Leu	CCC	Pro	CAC	His	CGC	Arg
CTA	Leu	CCA	Pro	CAA	Gln	CGA	Arg
CTG	Leu	CCG	Pro	CAG	Gln	CGG	Arg
ATT	Ile	ACT	Thr	AAT	Asn	AGT	Ser
ATC	Ile	ACC	Thr	AAC	Asn	AGC	Ser
ATA	Ile	ACA	Thr	AAA	Lys	AGA	Arg
ATG	Met	ACG	Thr	AAG	Lys	AGG	Arg
GTT	Val	GCT	Ala	GAT	Asp	GGT	Gly
GTC	Val	GCC	Ala	GAC	Asp	GGC	Gly
GTA	Val	GCA	Ala	GAA	Glu	GGA	Gly
GTG	Val	GCG	Ala	GAG	Glu	GGG	Gly



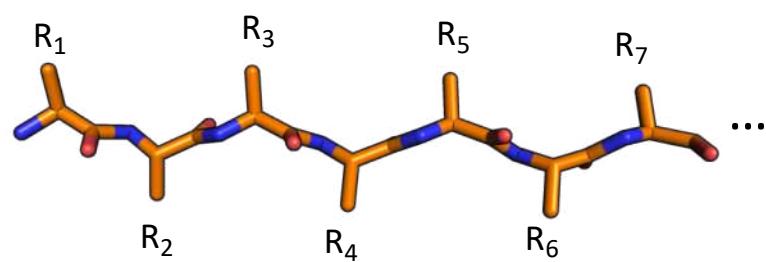
β -Sheet (3 strands)



α -helix

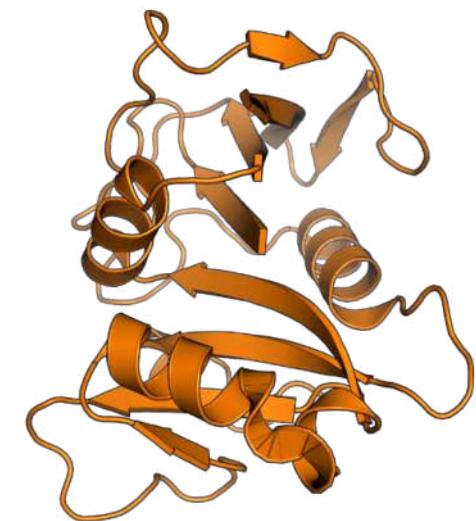


Protein Folding



Primary protein
structure

Protein folding

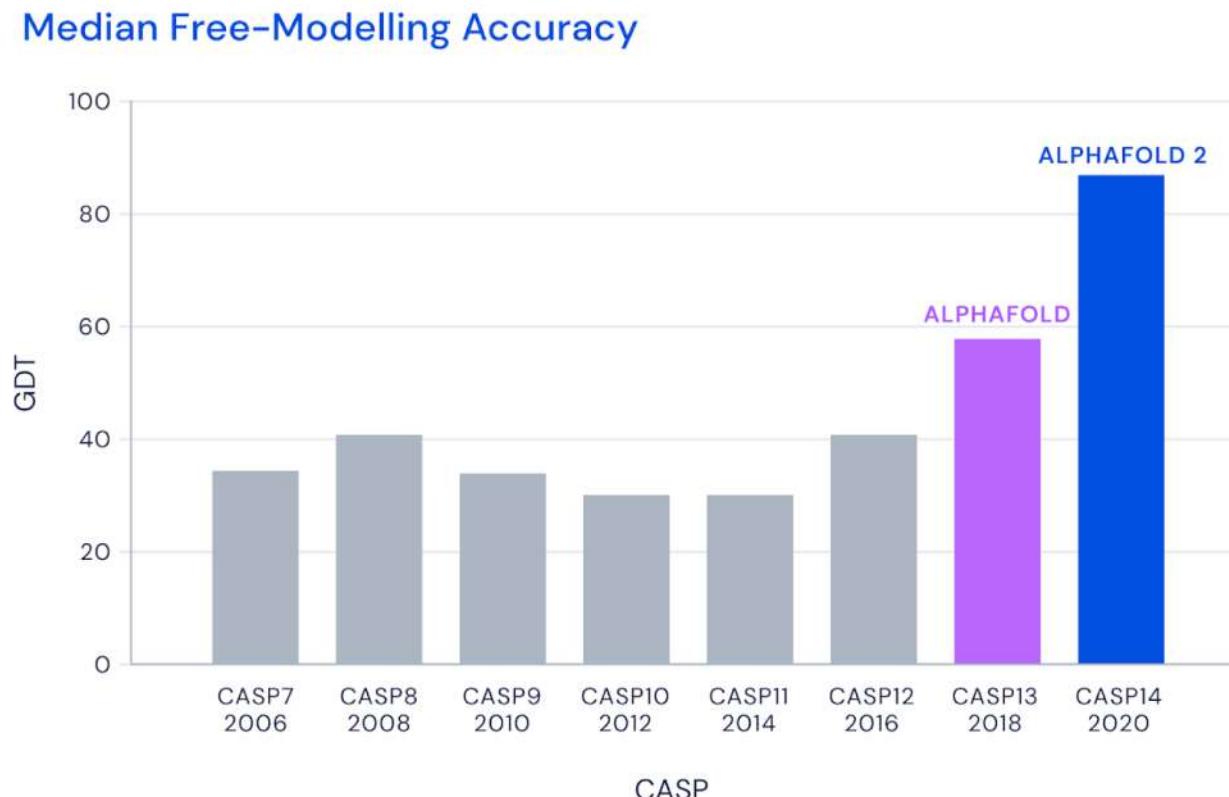


Tertiary protein
structure

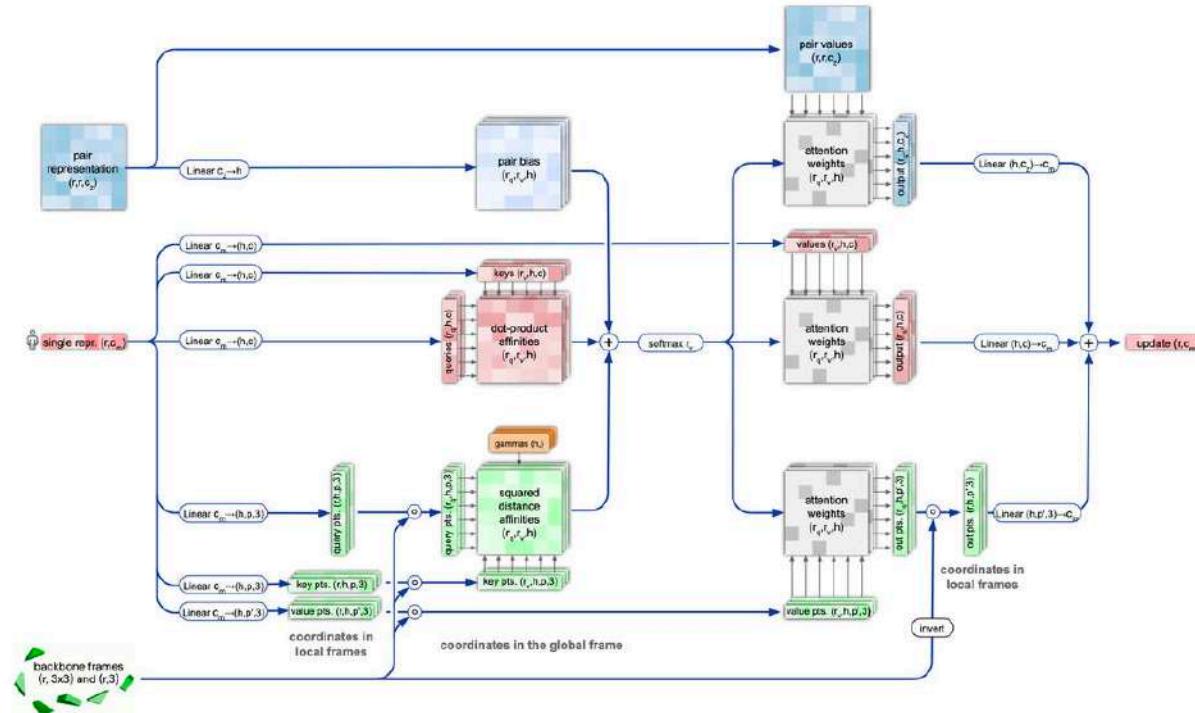
“[protein folding] is determined by [...] the aminoacid sequence in a given environment”

—Christian Anfinsen (1972 Nobel Laureate in Chemistry)

“ImageNet Moment” of Structural Biology



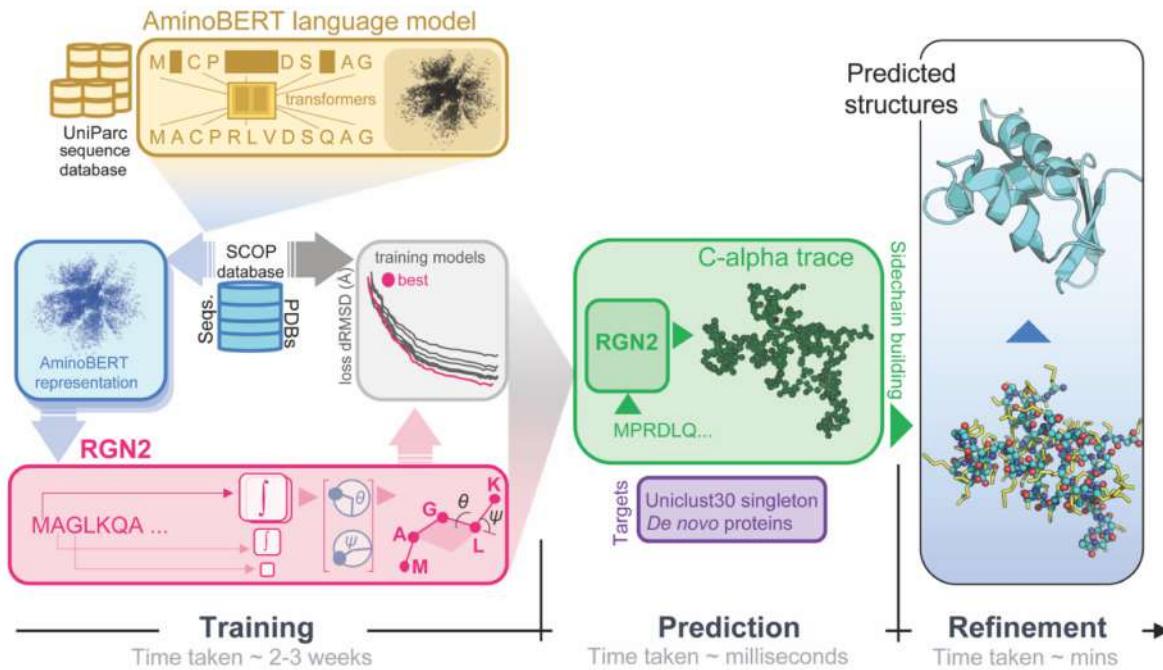
AlphaFold 2



Invariant Point Attention

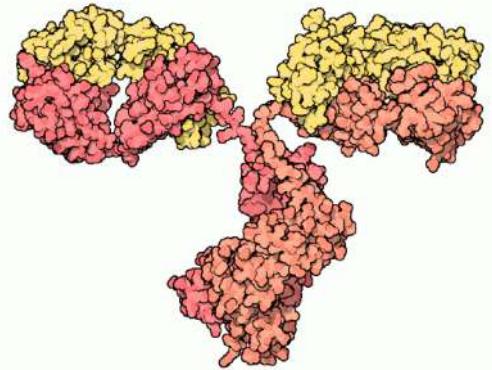
Jumper et al. 2021

Recurrent Geometric Network

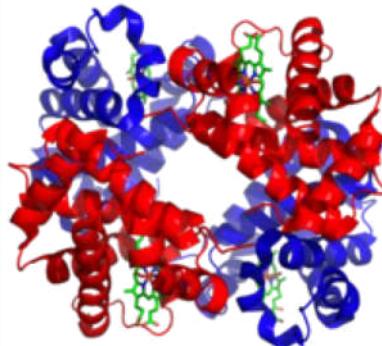


Single sequence protein structure prediction

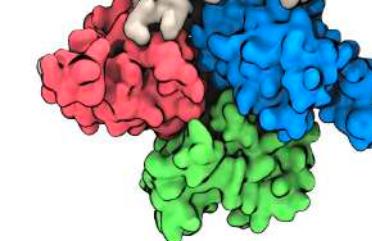
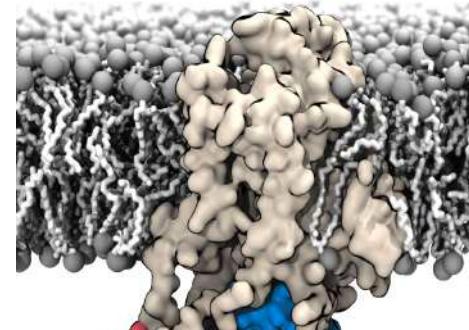
Chowdhury et al. 2021



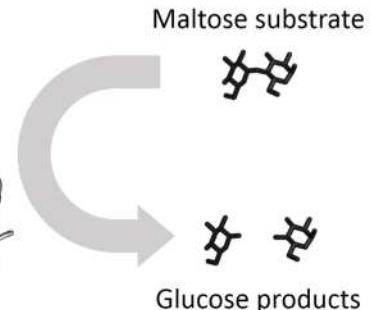
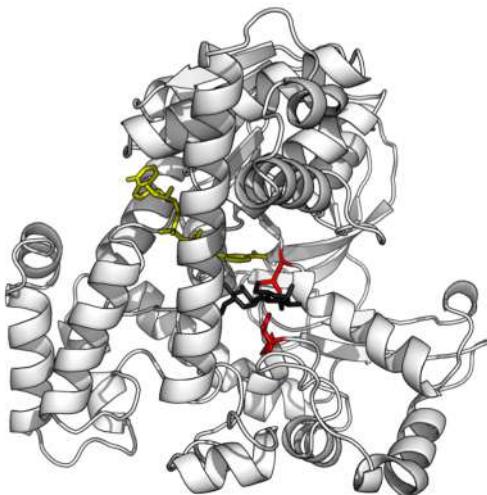
Defense (antibody)



Storage (haemoglobin)



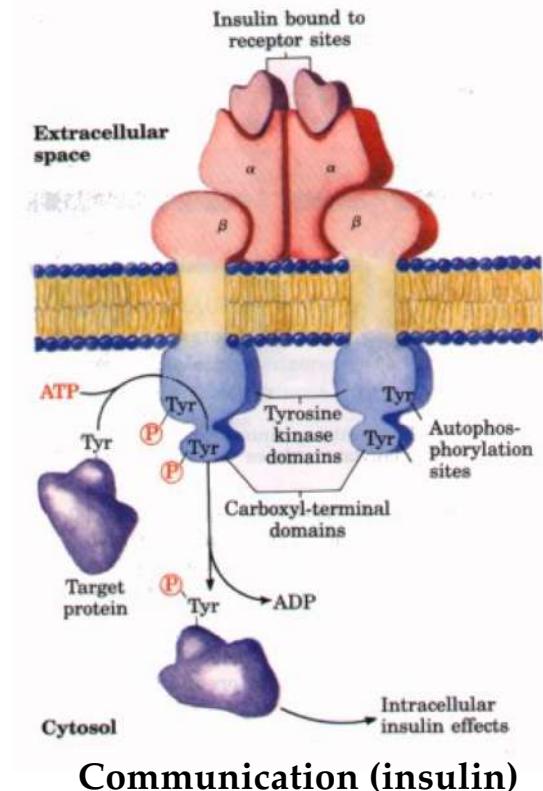
Transport (calcium pump)



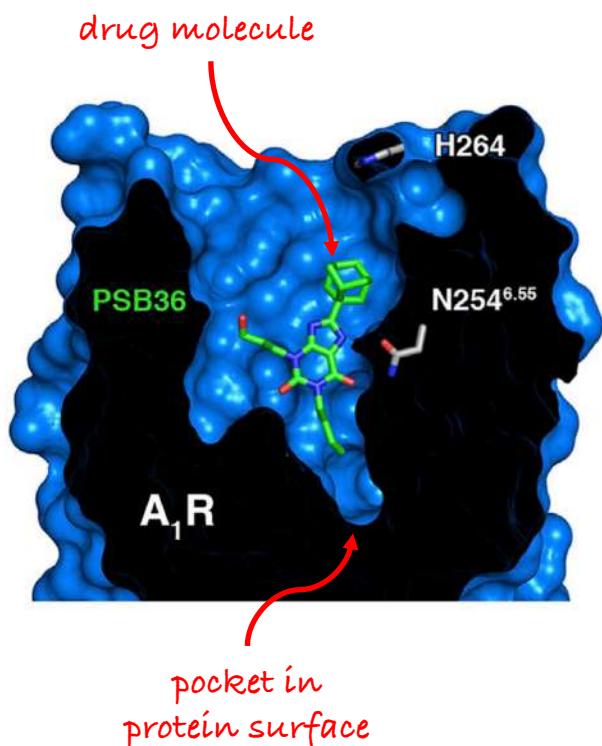
Catalysis (enzyme)



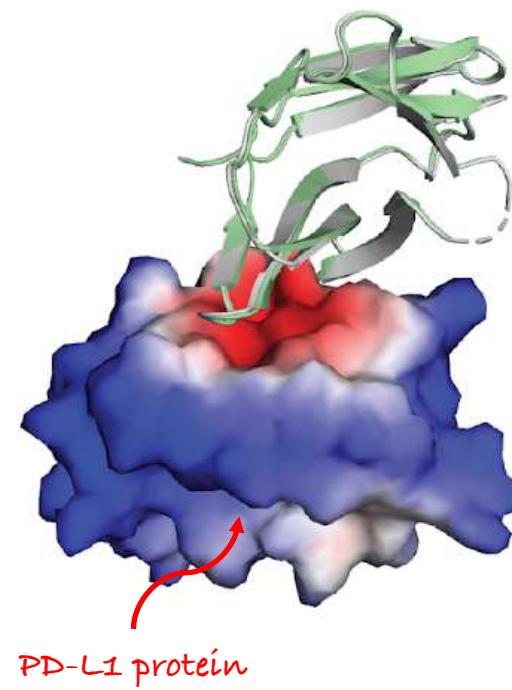
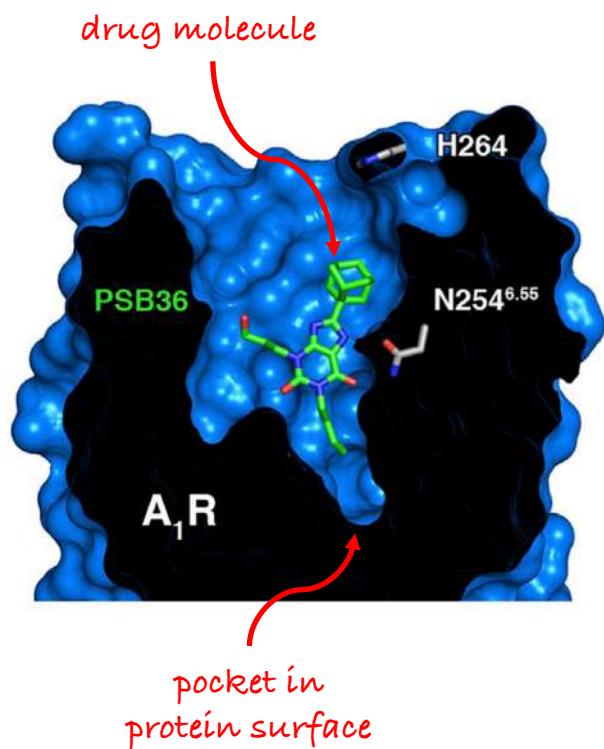
Structure (collagen)



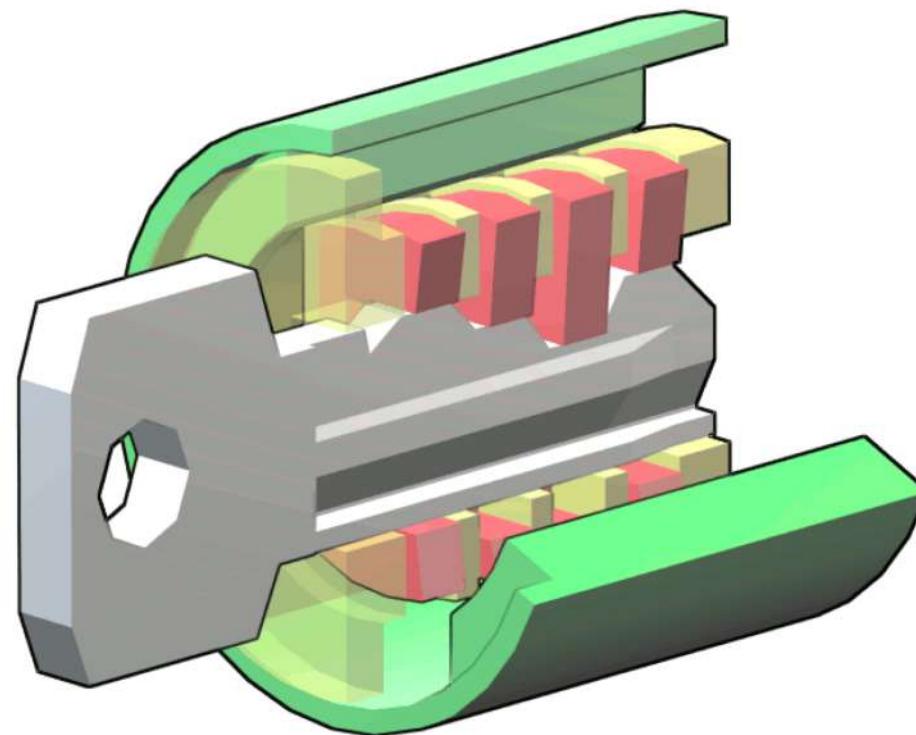
Small Molecule Drugs



Protein-Protein Interactions



Lock-Key Metaphor



Emil Fischer "Schlüssel-Schloss-Prinzip" 1894

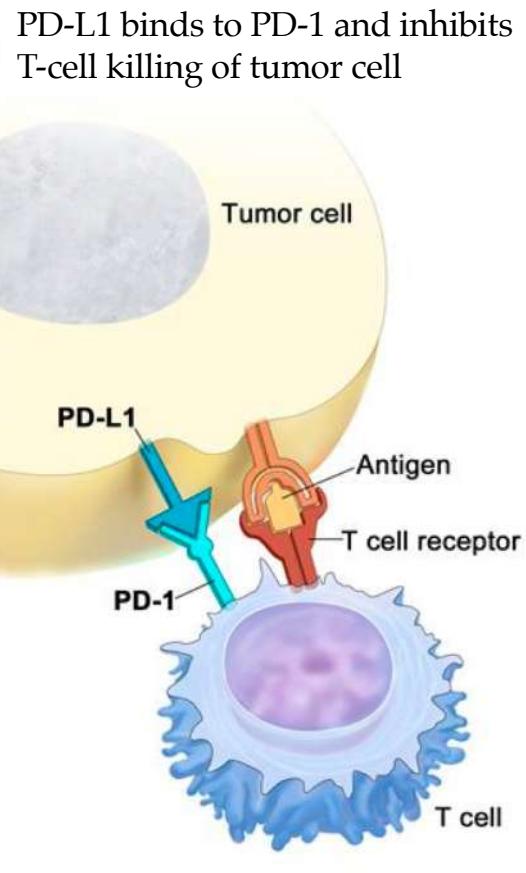


PD-1

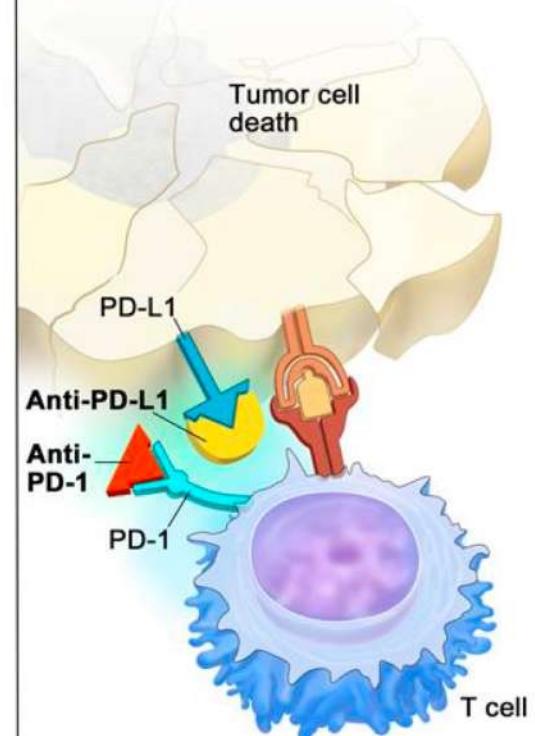
PD-L1



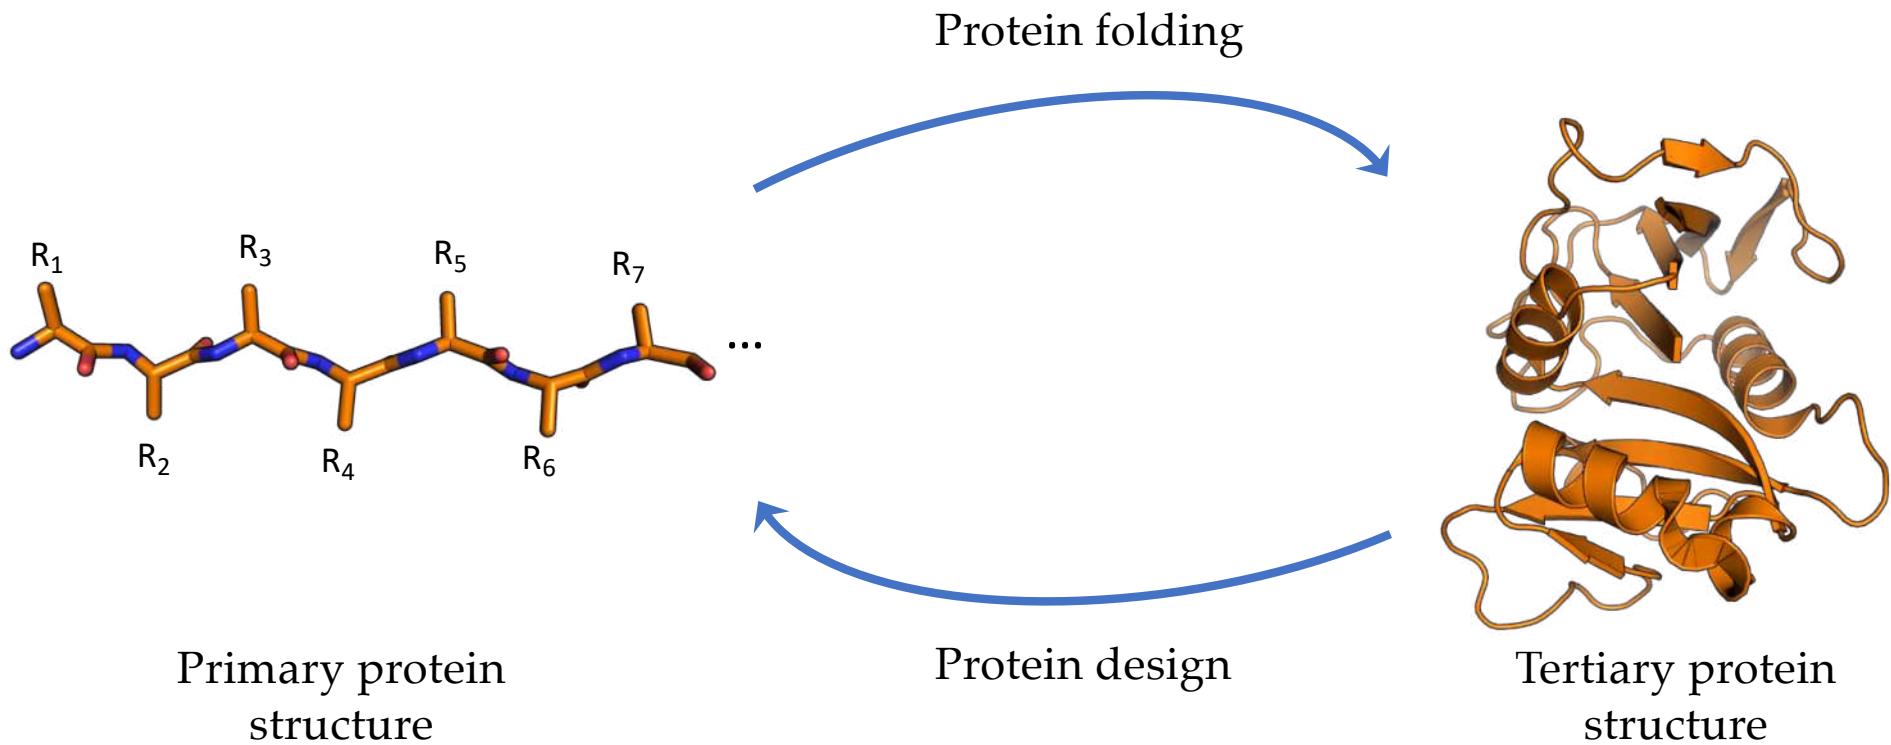
2018 Nobel Prize
PD-proteins role in
immunotherapy



Blocking PD-L1 or PD-1 allows T-cell killing of tumor cell

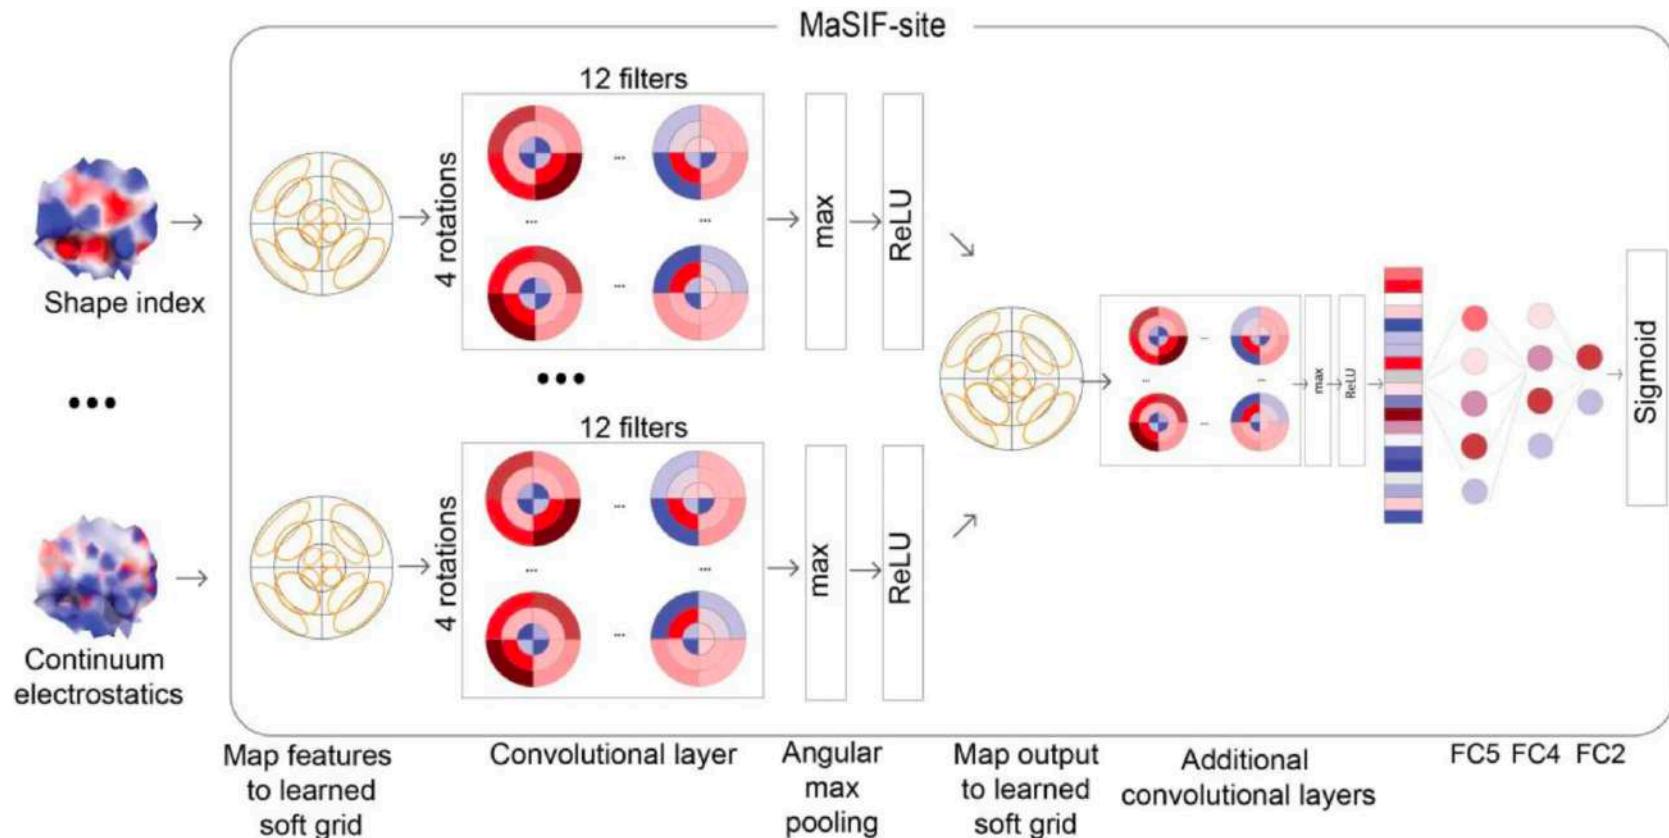


Protein Design = “Inverse Folding”

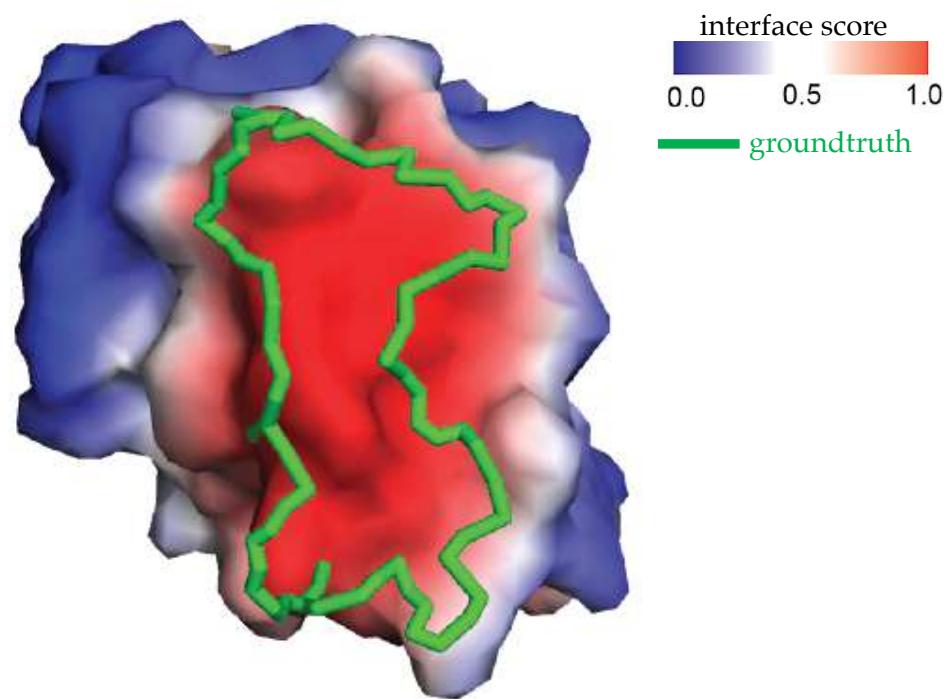


Sequence → Structure → Function

MaSIF: Protein Function Prediction

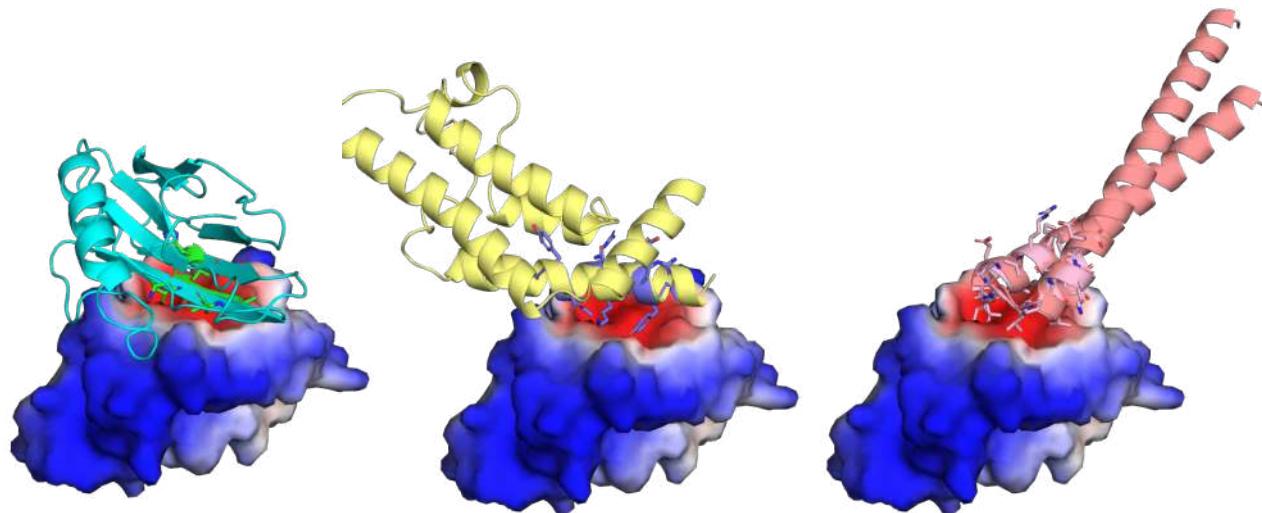


MaSIF: Protein Function Prediction

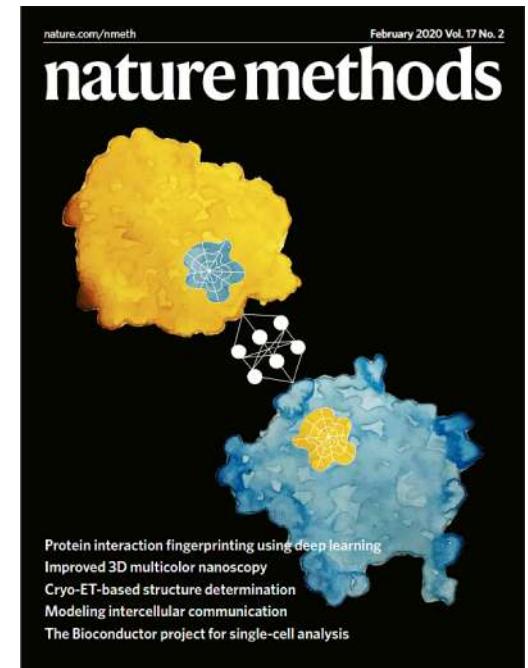


Gainza, Sverrisson et B 2020

De novo Protein Design



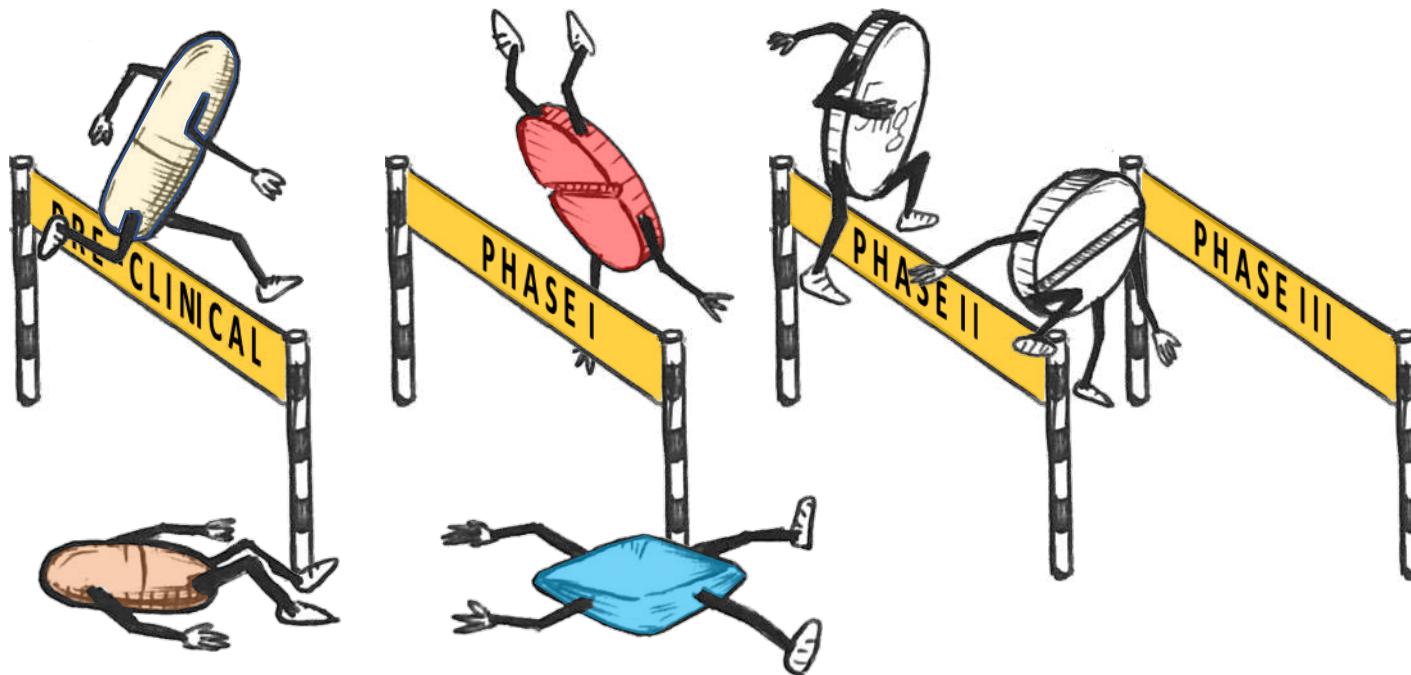
Binding site identification + designed binders
for the oncological target PD-L1



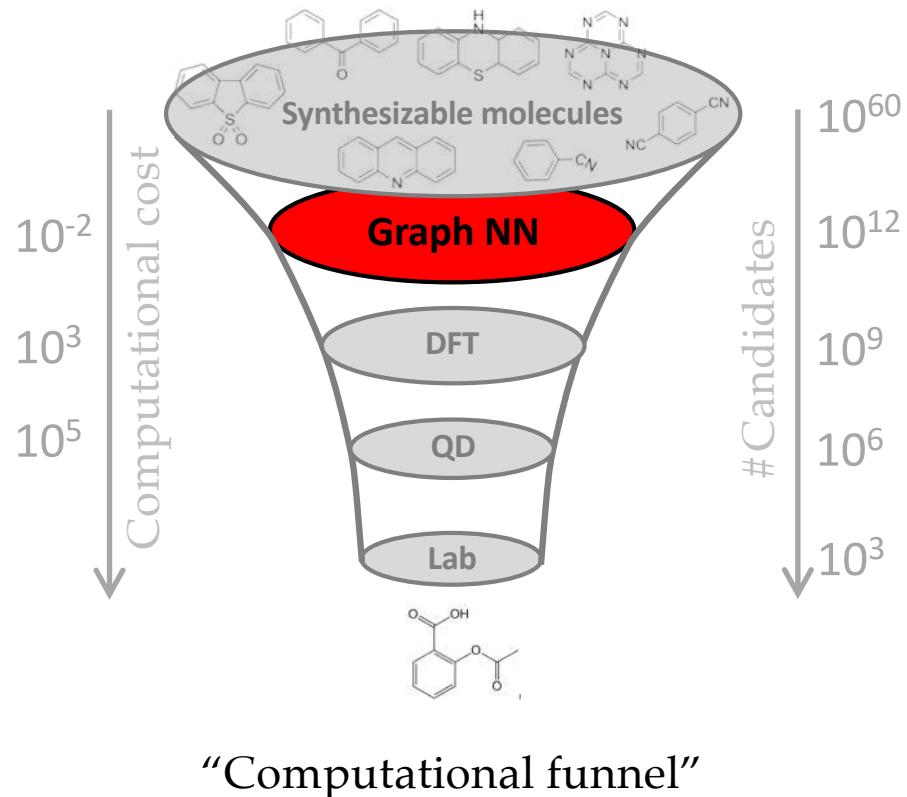
Gainza, Sverrisson et B 2020; Sverrisson, Feydy et B 2021

CHEMISTRY & DRUG DESIGN

Drug Discovery & Design



Virtual Drug Screening



Duvenaud et al. 2015; Gilmer et al. 2017; Jin et al. 2020

History of Graph Neural Networks according to Chemists



D. Kireev

ChemNet

1995



I. Baskin

Neural descriptors

1997



C. Merkwirth

Molecular graph net

2005



D. Duvenaud

Molecular fingerprints

2015

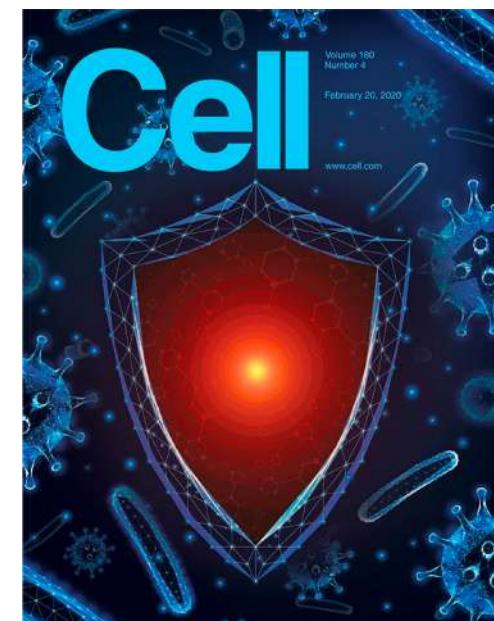
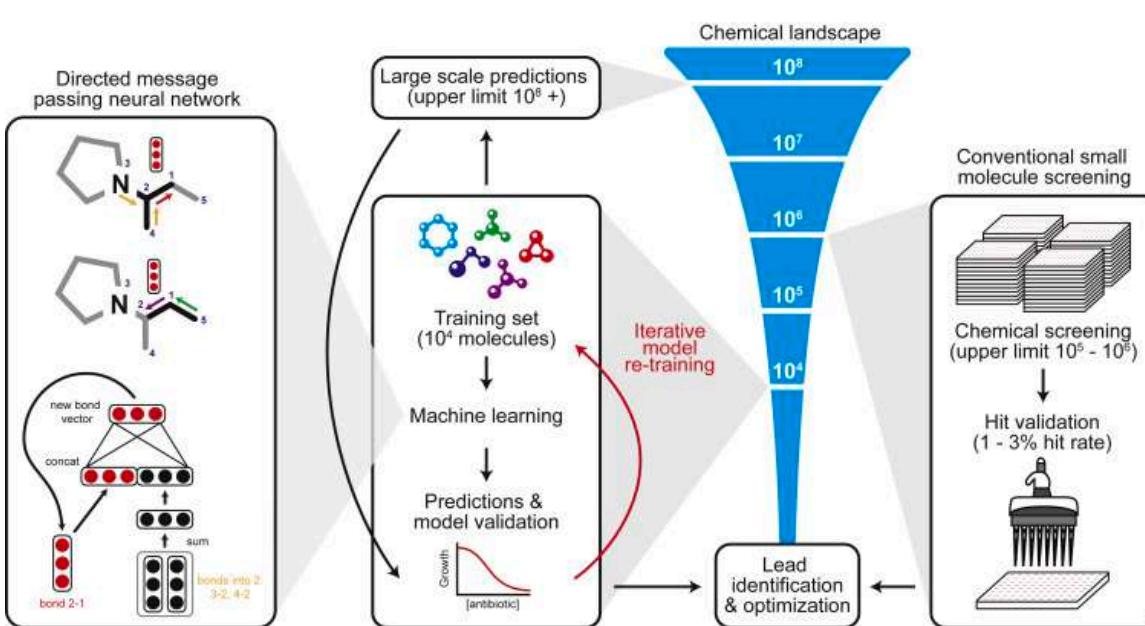


J. Gilmer

MPNNs

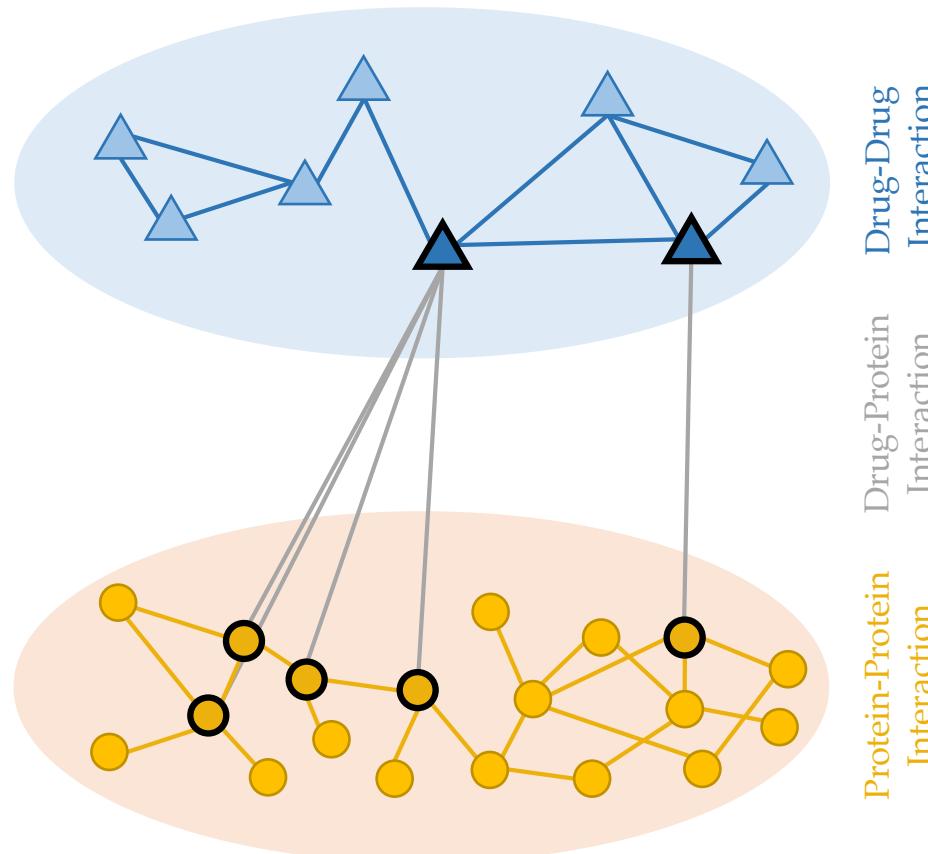
2017

New Antibiotic Discovery



Stokes et al. 2020

Drug repositioning & Combinatorial therapy

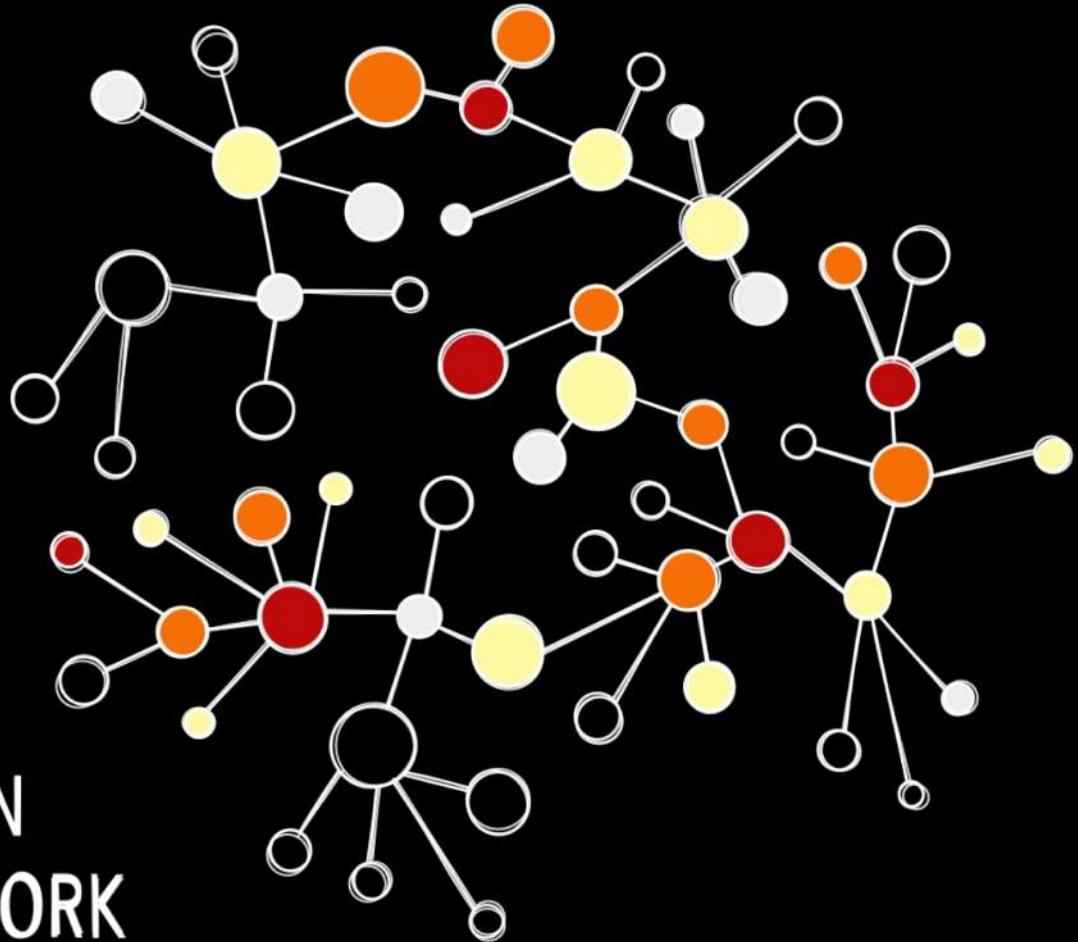


Zitnik et al. 2018



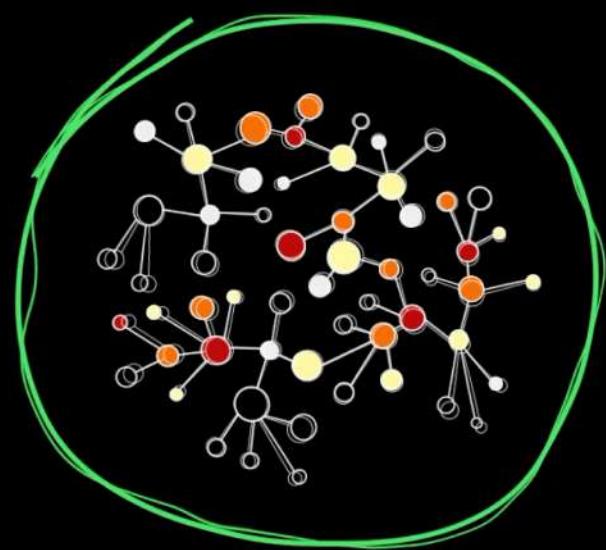
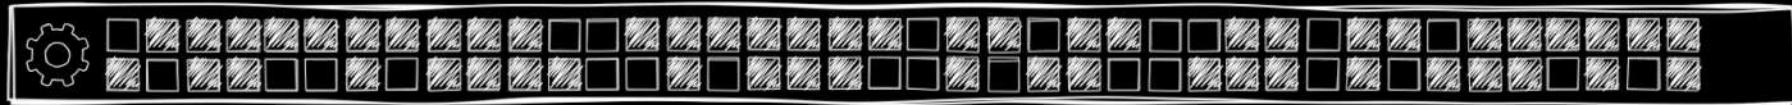
Veselkov et B 2019

Hyperfoods



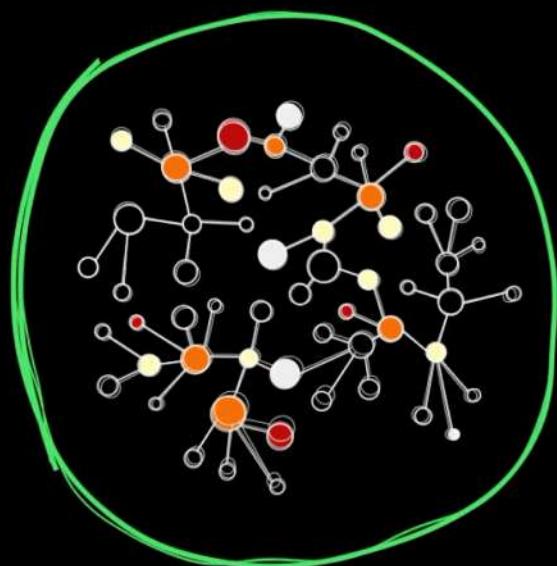
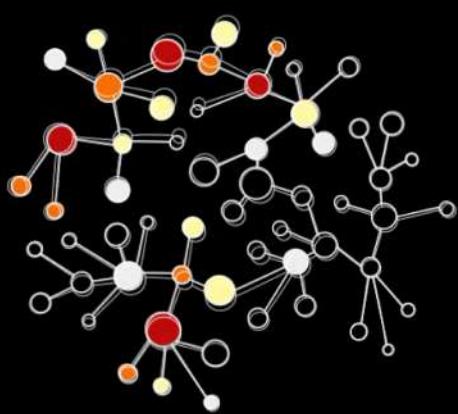
PROTEIN-PROTEIN INTERACTION NETWORK

Veselkov et B 2019



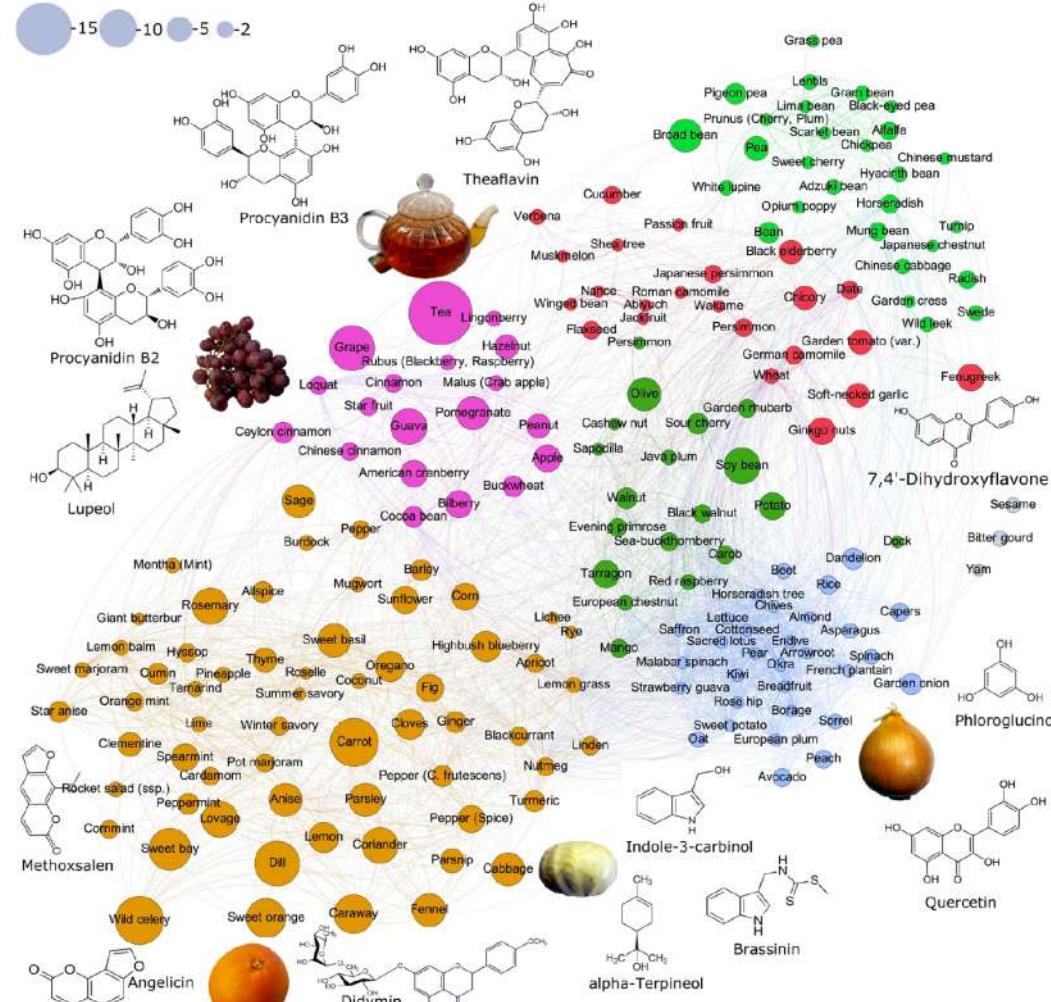
ANTICANCER

non ANTICANCER



ANTICANCER

Hyperfoods



Veselkov et B 2019

Video: Vodafone Foundation / Recipes: Bruno Barbieri
Based on Laponogov et B 2020

WHAT'S NEXT?

YEAR 2020 IN REVIEW & PREDICTIONS FOR 2021

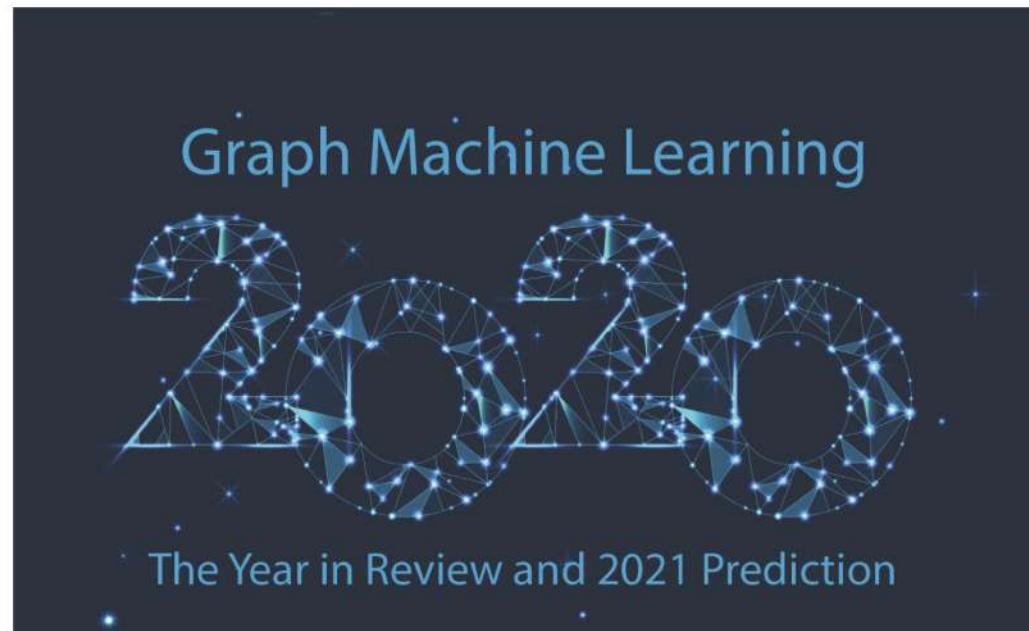
What 2021 holds for Graph ML?

The end of the year is a good time to recap and make predictions. 2020 has turned Graph ML into a celebrity of machine learning. For this post, I sought the opinion of prominent researchers in the field of graph ML and its applications trying to summarise the highlights of the past year and predict what is in store for 2021.



Michael Bronstein Jan 5 · 17 min read ★

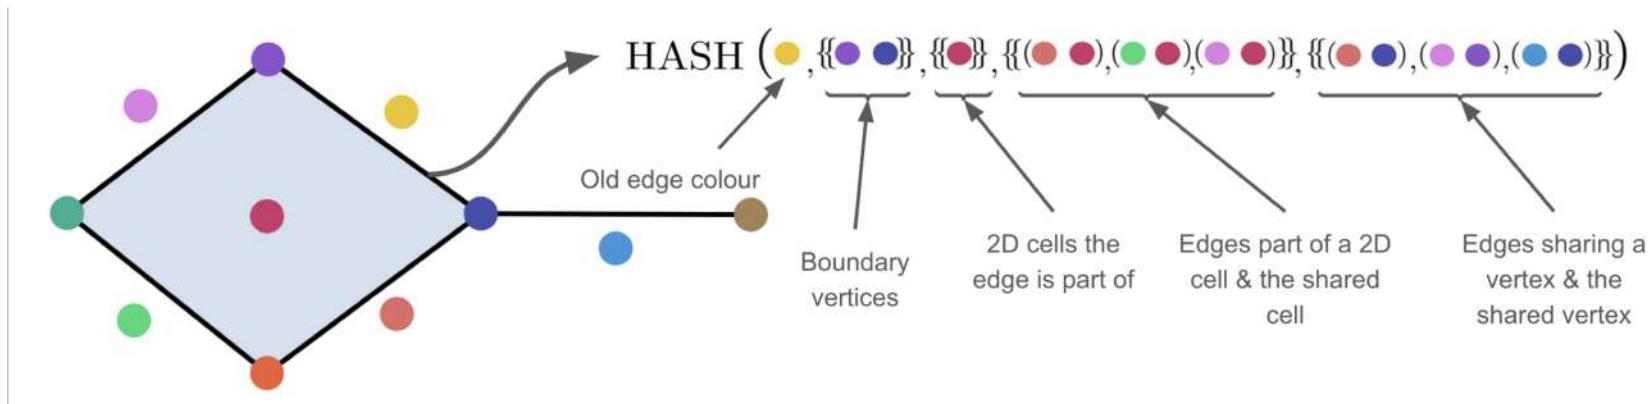
↑ ⌂ ...



“2020 saw the field of Graph ML come to terms with the fundamental limitations of the message-passing paradigm.”

—Will Hamilton (Mila)

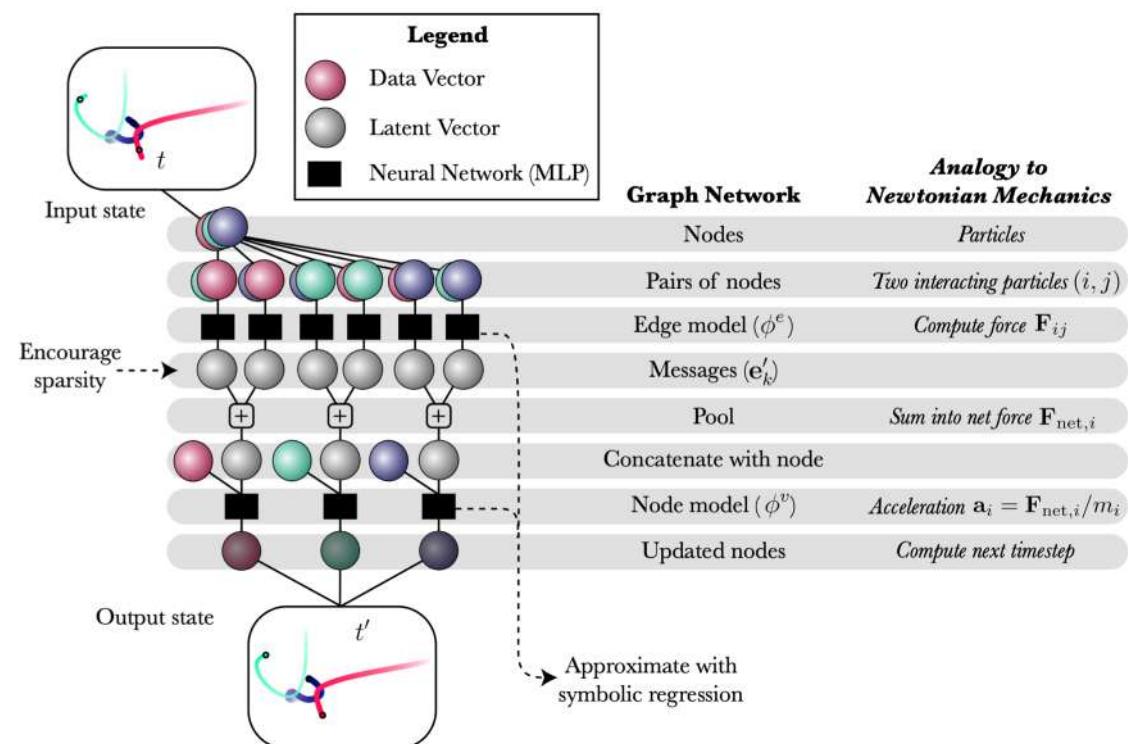
Beyond Traditional Message Passing



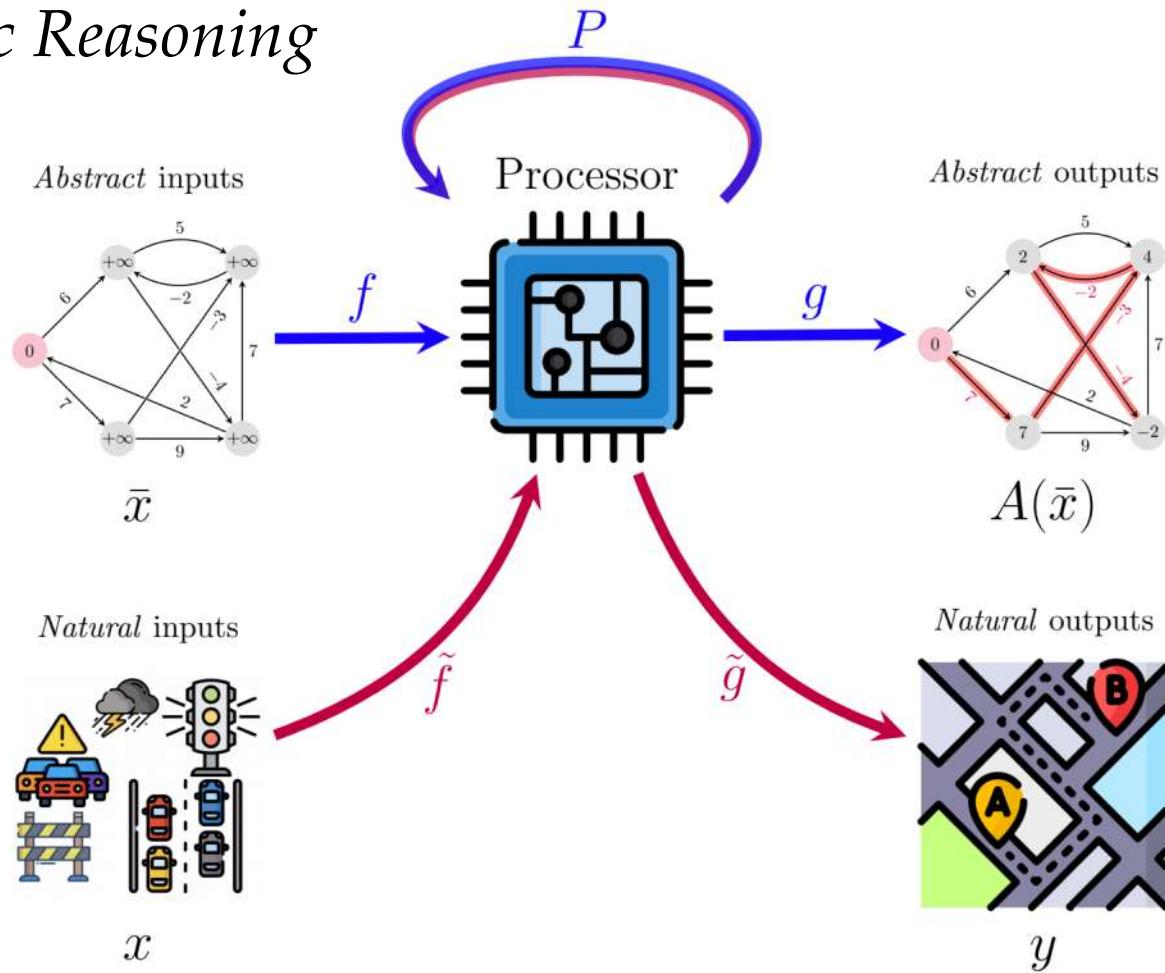
- Weisfeiler-Lehman-like schemes for higher order structures (simplicial/cell complexes)
- Strictly more powerful than MPNNs
- Links to computational topology
- Promising results in chemistry

"Explainable" GNNs

Learning equations of motion using GNN to learn a predictor for a dynamical systems + symbolic regression on the messages sent along the edges. Extracted symbolic equations re-introduced into the GNN, replacing the original learned components

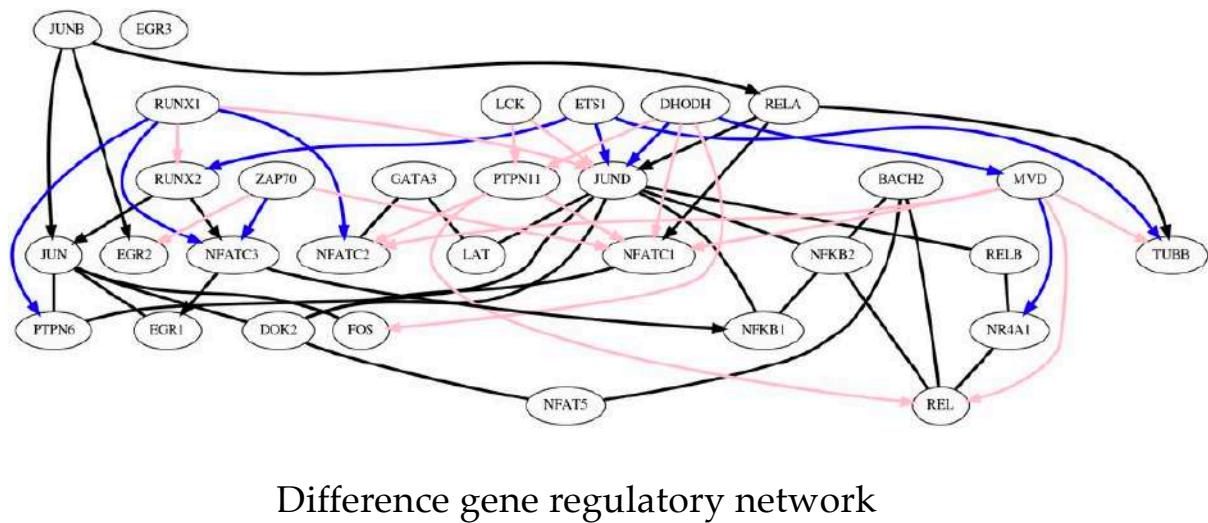
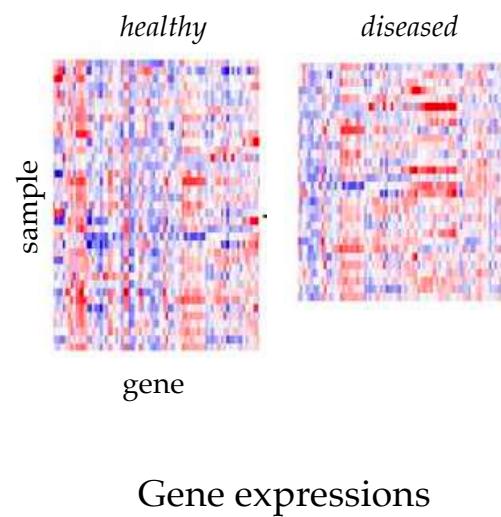


Algorithmic Reasoning



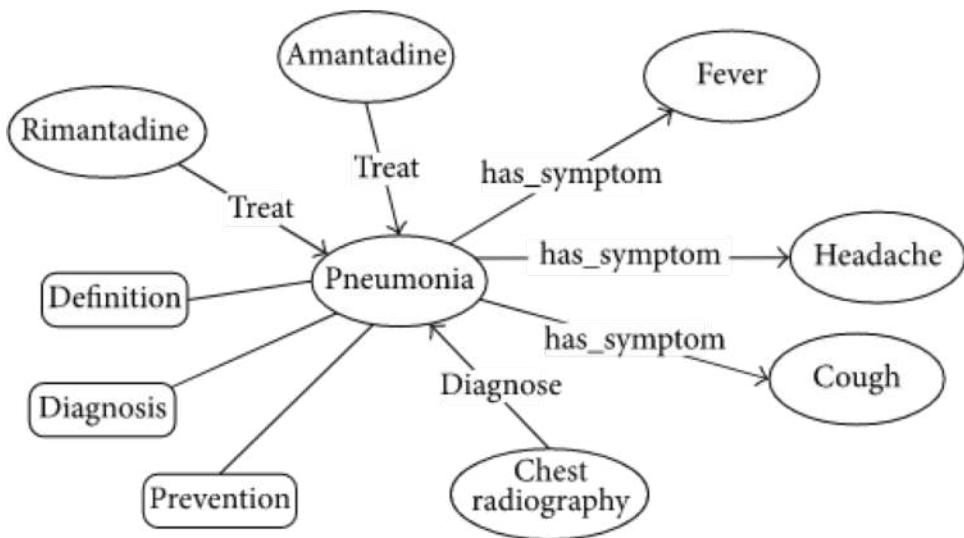
Velickovic et al. 2021

Causal Inference

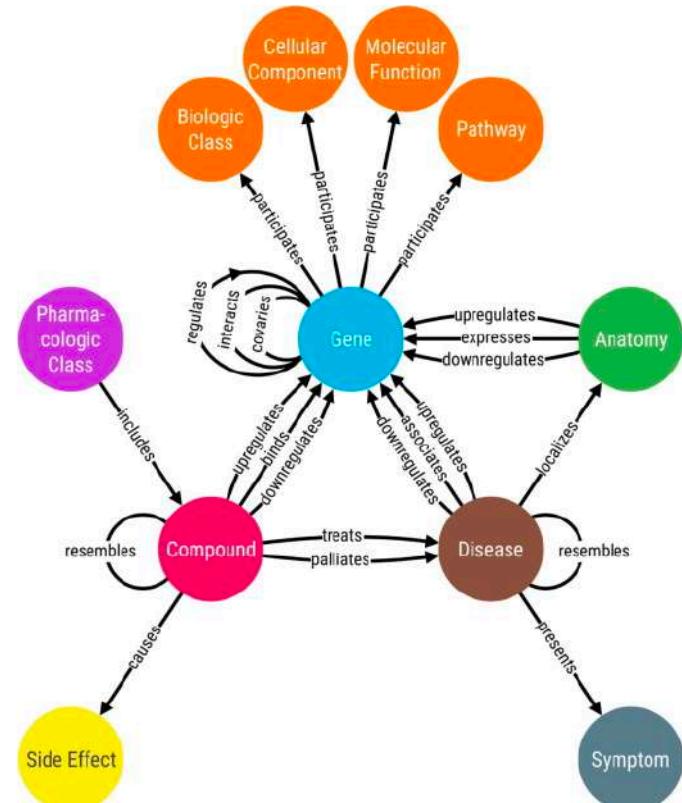


Belyaeva et al. 2020

Knowledge Graphs



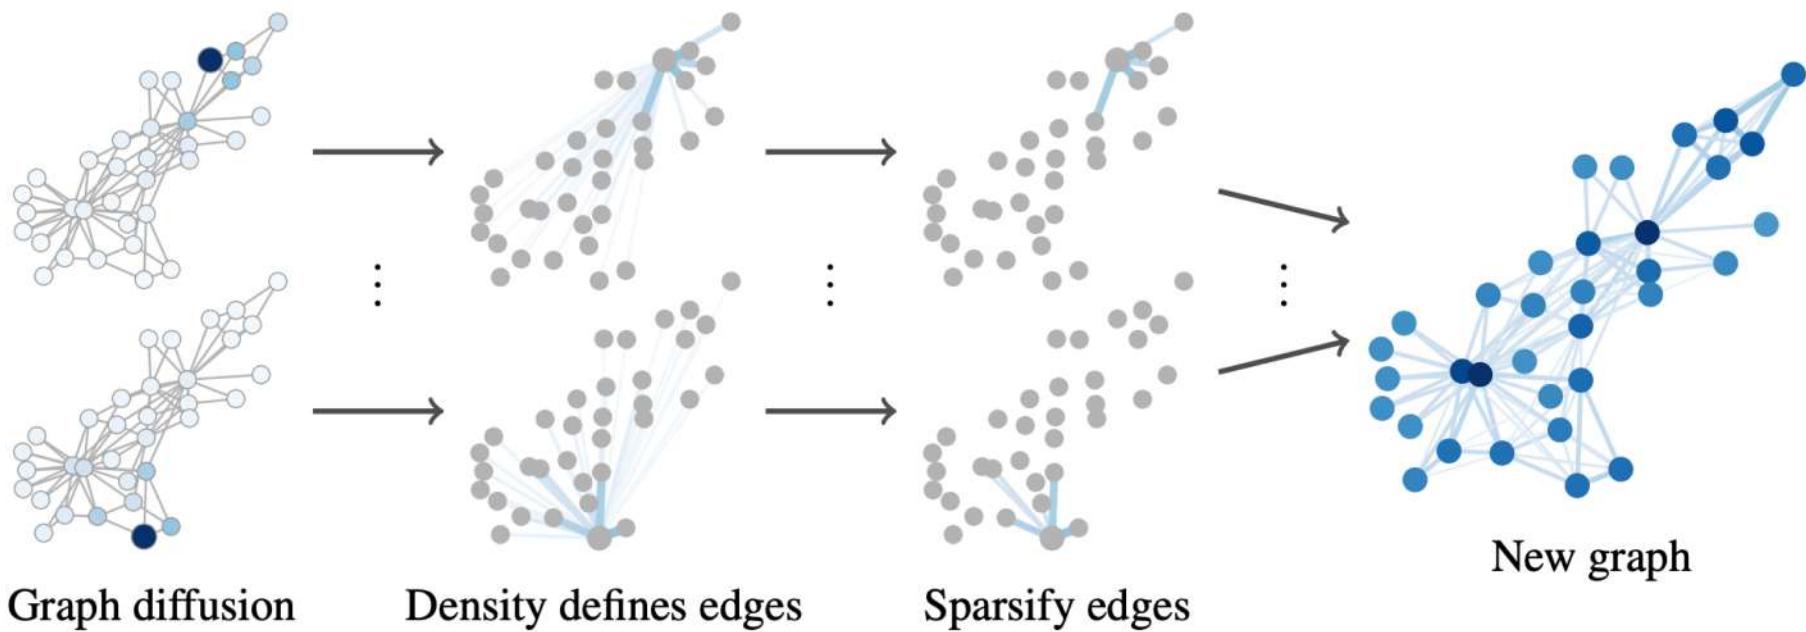
Entity
 Attribute
→ Relationship



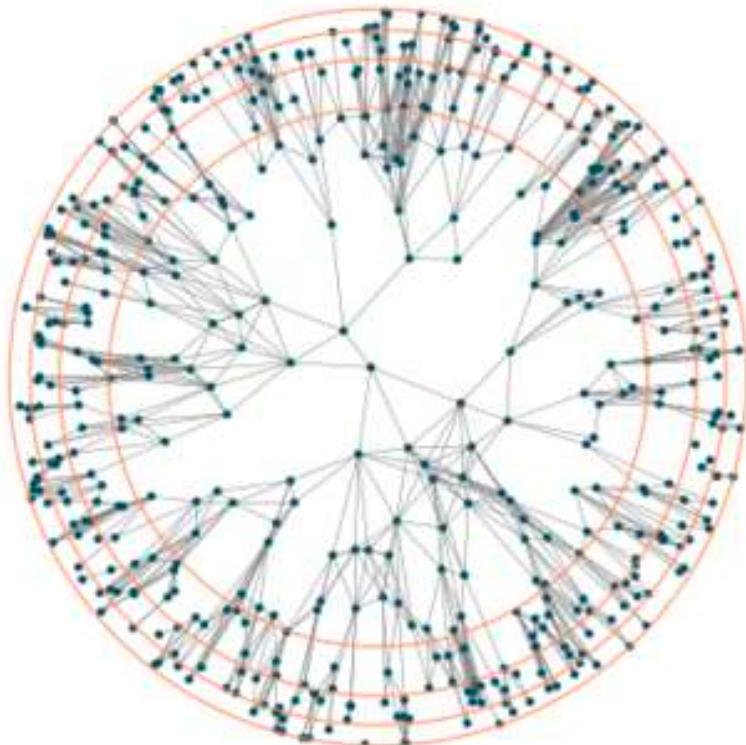
“One particularly noteworthy trend in the Graph ML community since the recent widespread adoption of GNN-based models is the separation of computation structure from the data structure.”

—Thomas Kipf (Google)

Data vs Computational Graphs: “Graph Rewiring”



“Network Geometry”



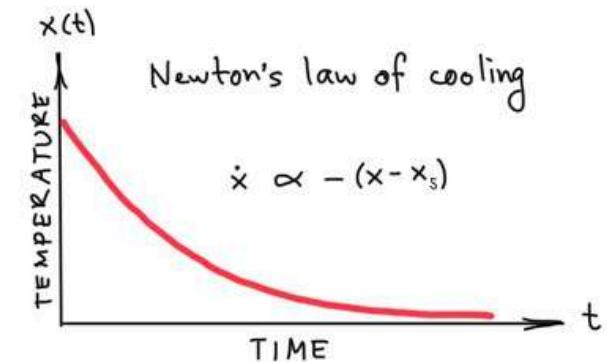
Boguna et al. 2020

Neural PDEs

- GNNs are discretised diffusion PDEs
- Use efficient numerical schemes
- Deep links to differential geometry

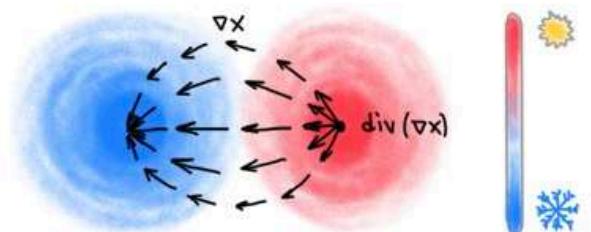


I. Newton

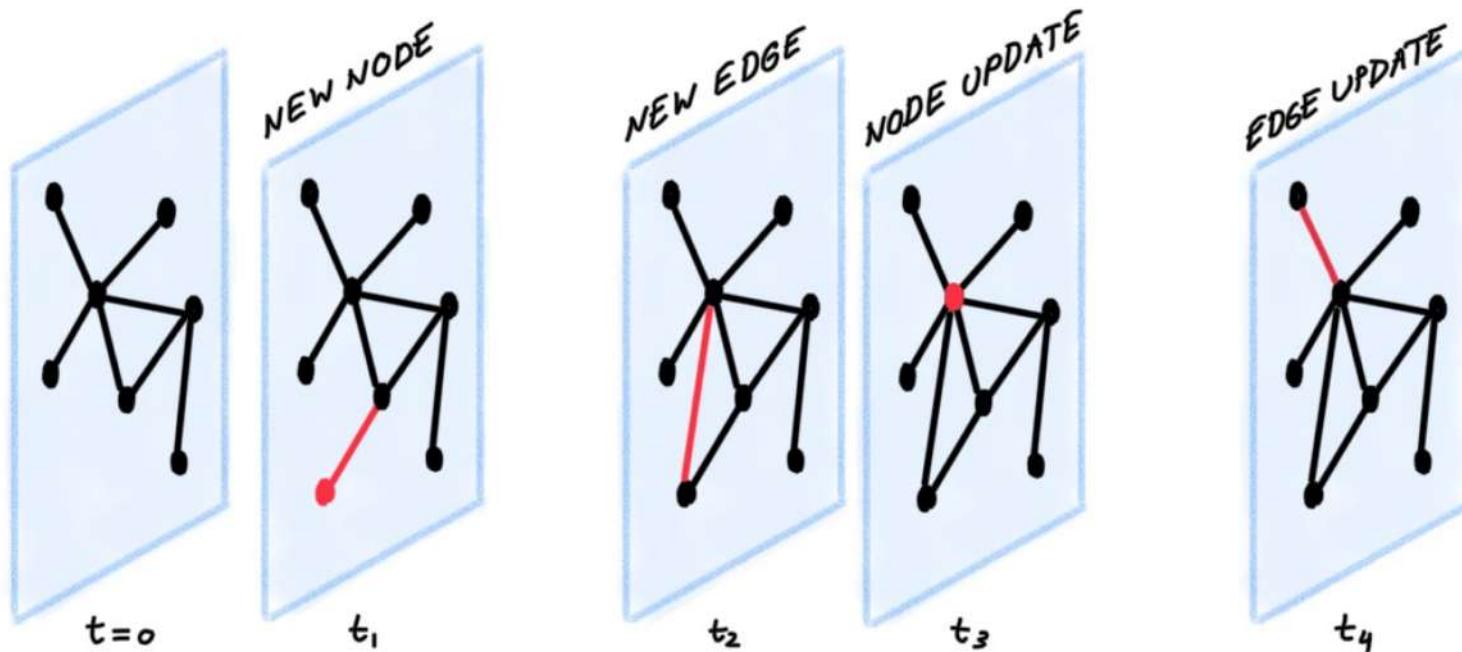


J. Fourier

Fourier's heat transfer law
$$h = -\alpha \nabla x$$



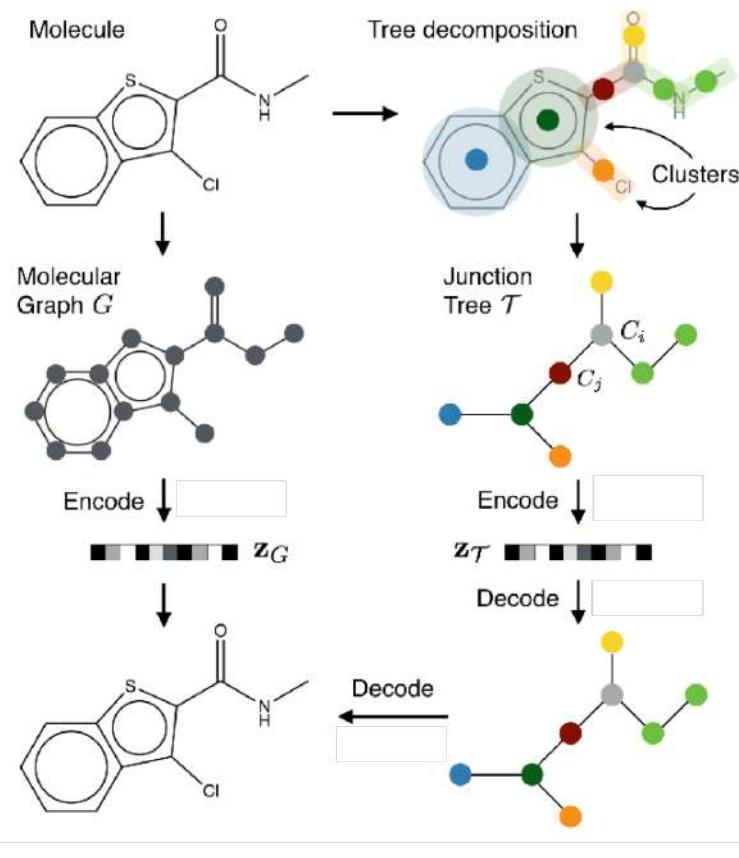
Dynamic Graphs



Generative Models

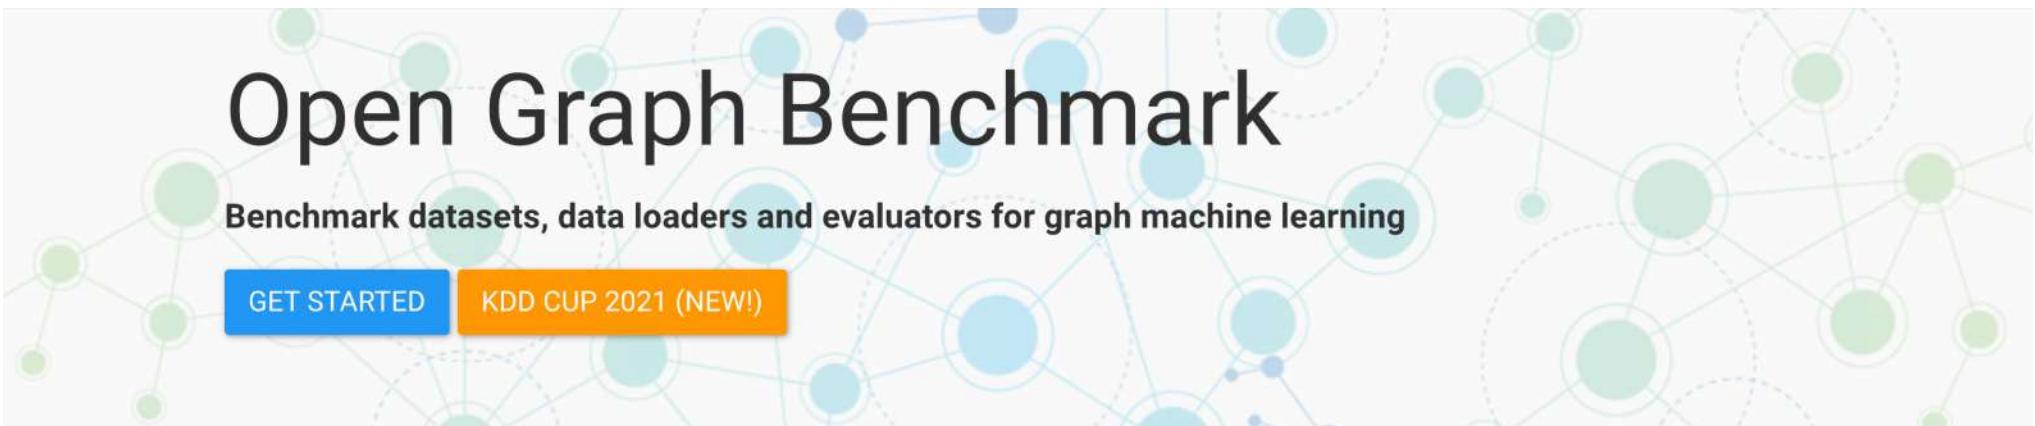


Generative Models for Graphs



Jin et al. 2018

Standardised Benchmarks



The Open Graph Benchmark (OGB) is a collection of realistic, large-scale, and diverse benchmark datasets for machine learning on graphs. OGB datasets are automatically downloaded, processed, and split using the [OGB Data Loader](#). The model performance can be evaluated using the [OGB Evaluator](#) in a unified manner.

OGB is a community-driven initiative in active development. We expect the benchmark datasets to



Industrial Applications



Google Invents AI That Learns a Key Part of Chip Design

AI helps designs AI chip that might help an AI design future AI chips

By Samuel K. Moore



Google is using AI to design chips that will accelerate AI



Google is using AI to design AI processors much faster than humans can

By Paul Lilly, 10 days ago
Chips making chips.

Comments



Google Proposes AI as Solution for Speedier AI Chip Design

Google uses artificial intelligence to optimize AI chip production

By Michael Kan, April 1, 2018

Facebook LinkedIn Twitter

Print

Email

Print

Share

Comment

Report

Feedback

Print

Google Hoping The Next AI Chips Will Be Designed By AI

Company researchers have come up with an AI system that can design other AI chips. The goal is to help improve AI with the help of AI.

By Michael Kan, April 1, 2018

Facebook LinkedIn Twitter

Print

Share

Comment

Report

Feedback

Print

Google Researchers Create AI-ception with an AI Chip That Speeds Up AI

Using a reinforcement learning algorithm, the AI has learned to optimize the placement of components on a computer chip.

By Michael Kan, April 1, 2018

Facebook LinkedIn Twitter

Print

Share

Comment

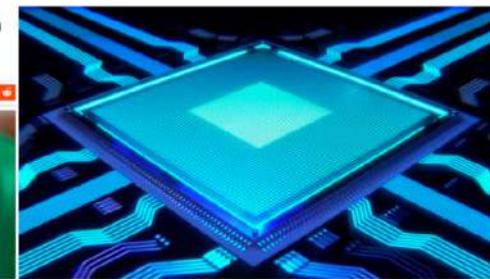
Report

Feedback

Print

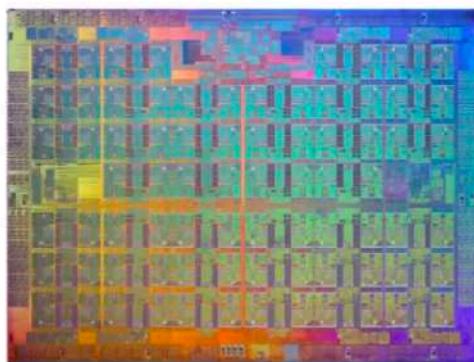
Google trains chips to design themselves

by Peter Grad, Tech Xplore



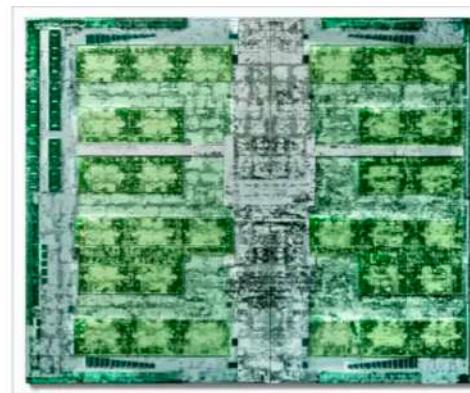
New Hardware: Beyond GPUs

Hardware Lottery: the phenomenon when certain algorithms win not because they are ideally suited to solve a certain problem, but because they run well on existing hardware



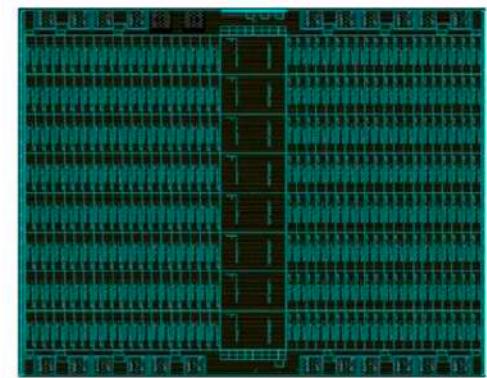
CPU

Scalar



GPU

Vector



IPU

Graph

Thank you!