Gauge Equivariant Mesh CNN

Pim de Haan

PhD student — University of Amsterdam

Research Associate — Qualcomm Technologies Netherlands B.V.

CNNs on Meshes

- Representation of curved manifold
- E.g. human artery
- Predict wall stress vector due to blood flow
- No canonical orientation
- How to orient convolutional kernels?
- Gauge equivariance

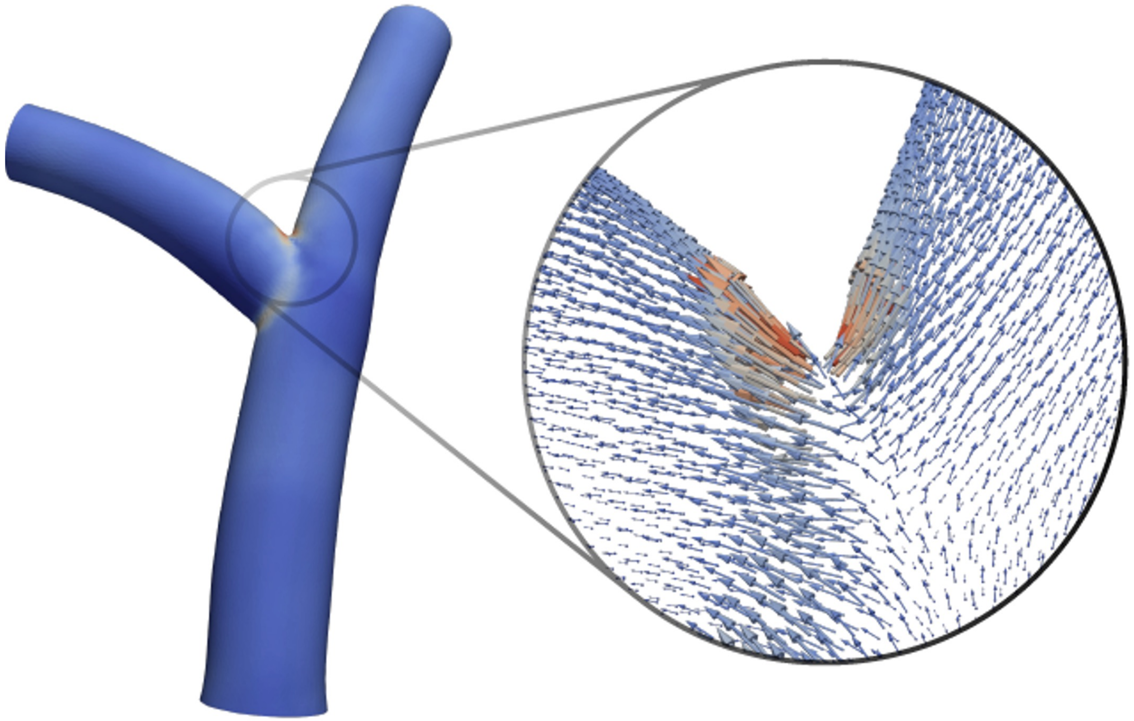

Collaboration

Maurice Weiler*
University of Amsterdam

*Joint first author

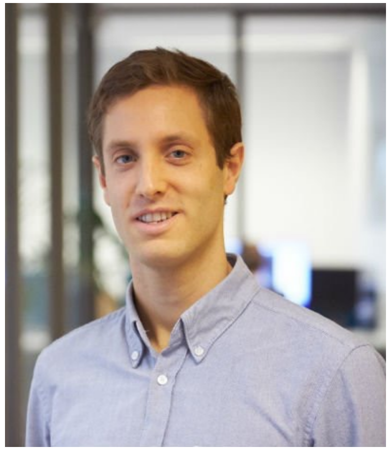

Taco Cohen
Qualcomm AI Research

Max Welling
University of Amsterdam

Outline

- CNN review
- Message passing on a mesh
- Scalar convolution
- Vector fields
- Gauge equivariant mesh convolutions
- Implementation
- Application to blood flow

Convolutional neural networks on images

- Image feature $f \in \mathbb{R}^{L \times L}$
- Kernel $k \in \mathbb{R}^{3 \times 3}$
- $f'_p = (k \star f)_p = \sum_q k(q) f(p - q)$
- Alternate convolutions with non-linearities
- Learn kernels k

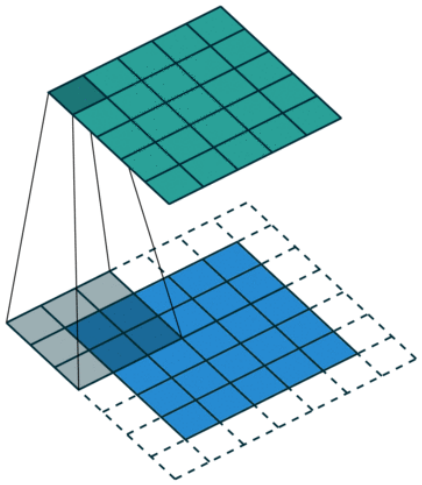

Anisotropy

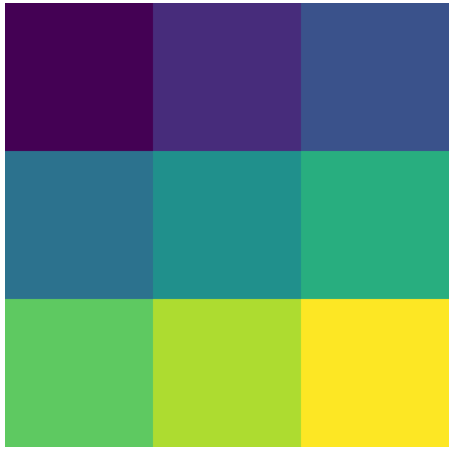

Anisotropic

- 9 parameters
- Detects any edge

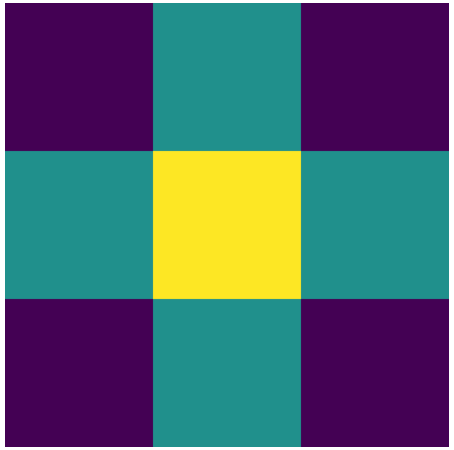

Isotropic

- 3 parameters
- Can not detect edges

Mesh

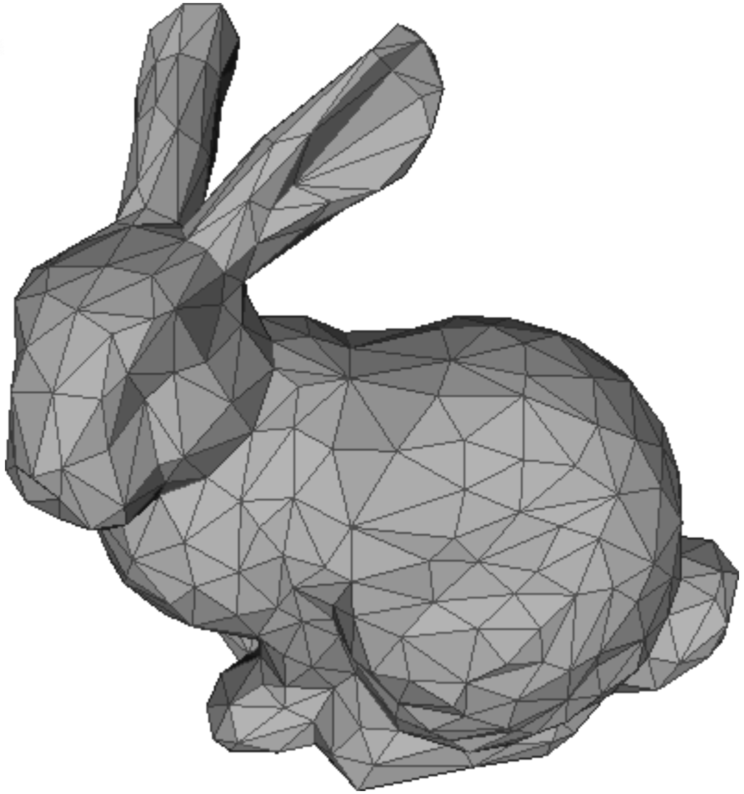

- Discretization of curved surface
- Triangular mesh = collection of triangular faces

Manifold Mesh

- Edge touches two faces
  - Or one if a boundary
- Vertex must be surrounded by plane of faces
  - Or half-plane if boundary
- Oriented manifold

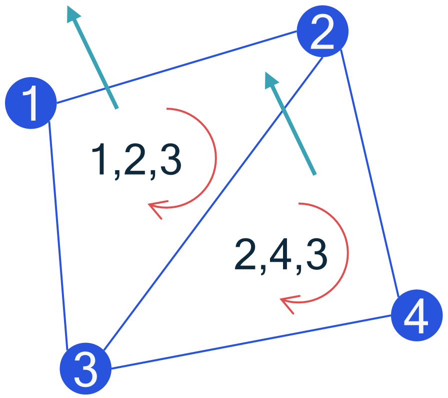

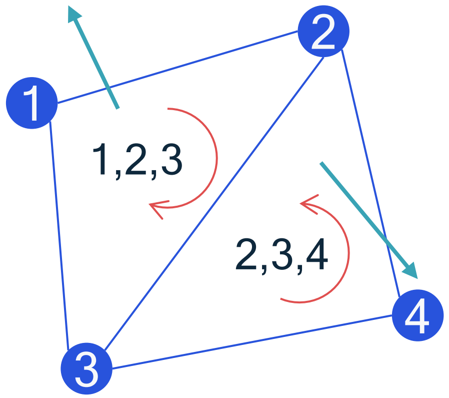

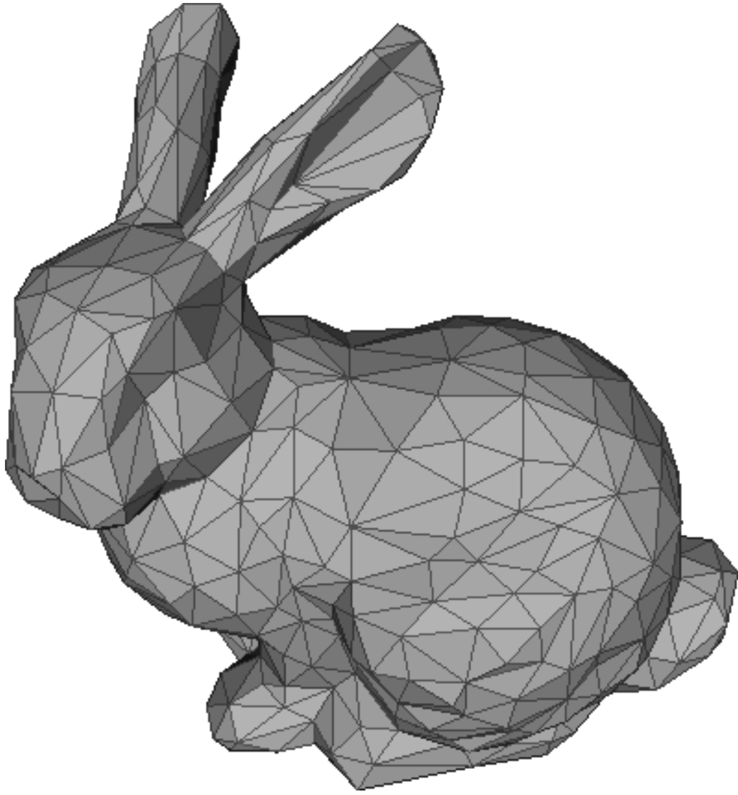

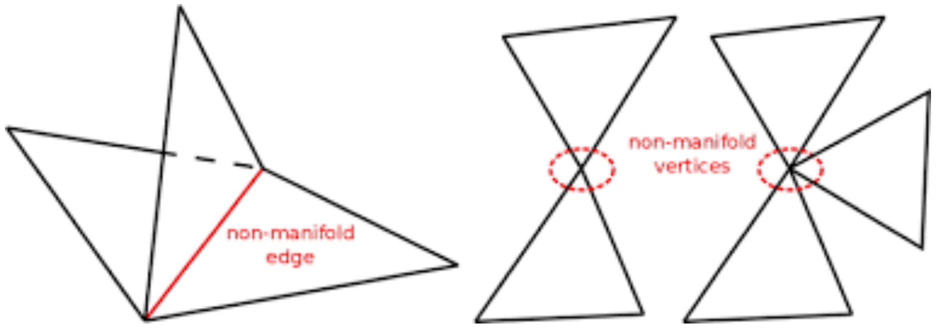

Message passing on a mesh

- Feature at vertices
- Message passing

$$f'_p = \sum_{q \in \mathcal{N}(p)} k(q \to p) f_q$$

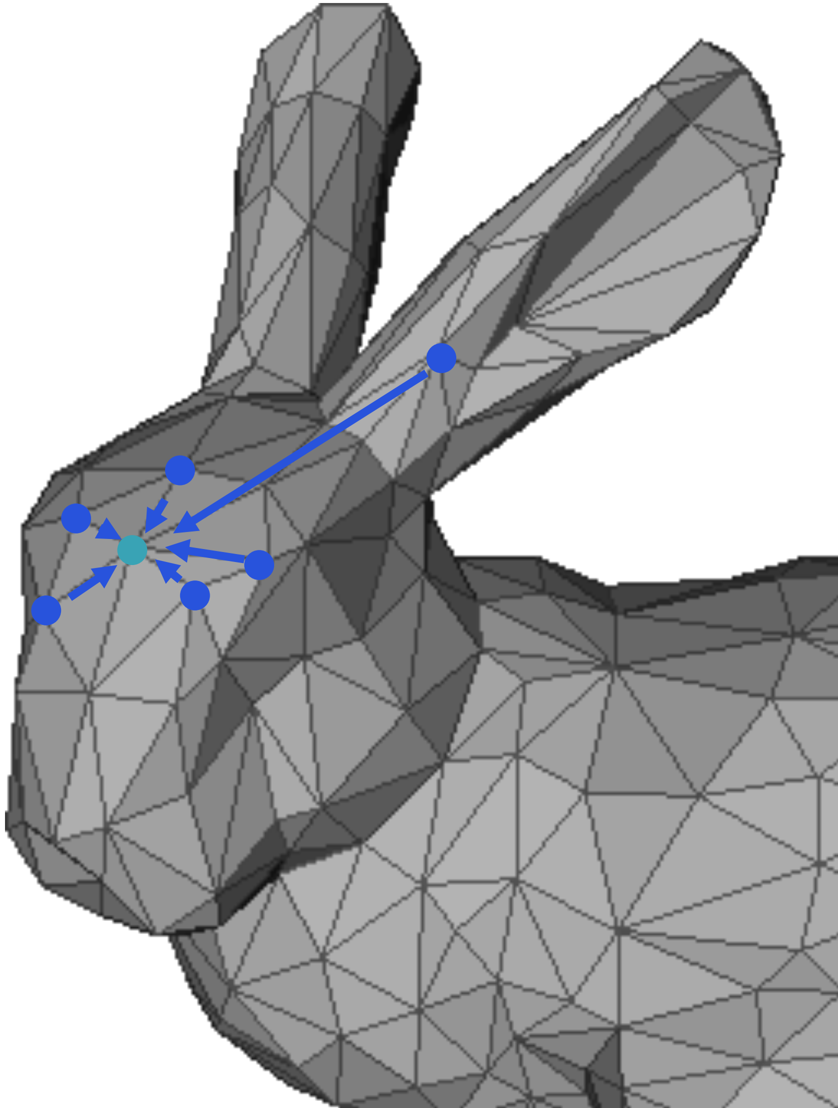

Convolutions on a mesh

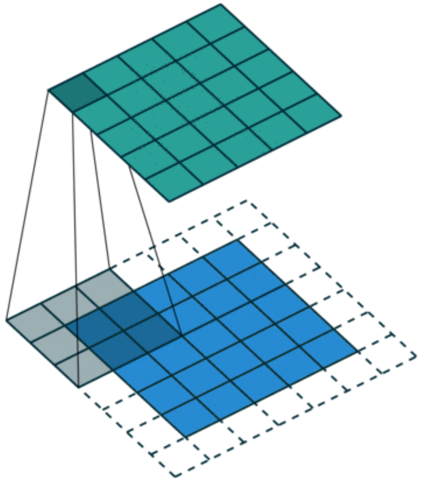

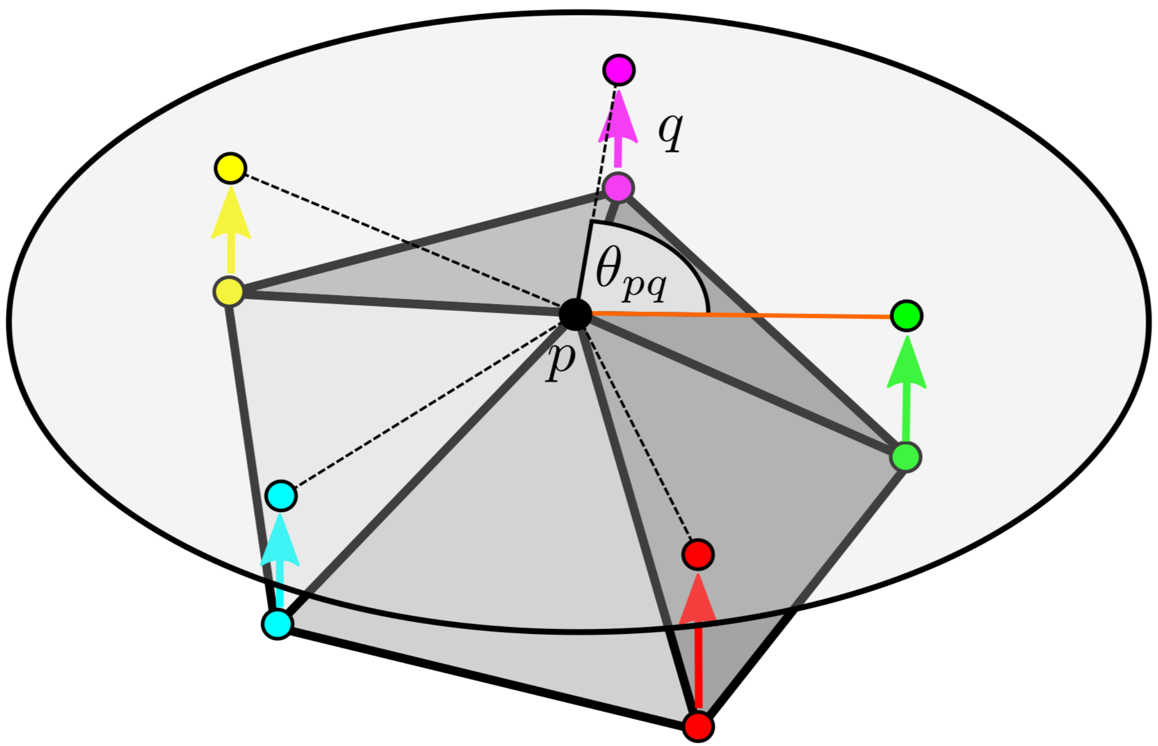

- Canonical relative (x, y) coordinates of neighbours

- Log map to tangent plane $\log_p q \in T_pM$
- Polar coordinates
- What is $\theta = 0$?
- Choice of coordinates: gauge

Mesh Tangent Planes & Gauges

- Oriented Manifold Mesh
- Faces have normal vector n_f
- Vertex normal n_p area-weighted average of adjacent face normals
- Tangent planes T_pM at vertex p normal to n_p
- Inherit metric tensor from $\mathbb{R}^3$ ambient space
- Gauge is basis $w_p: T_pM \to \mathbb{R}^2$
- Gauge defined up to $SO(2)$
- Polar coordinates

$w_p(v) = (r, \theta), w_p'(v) = (r, \theta + g_p),\ g_p \in SO(2)$

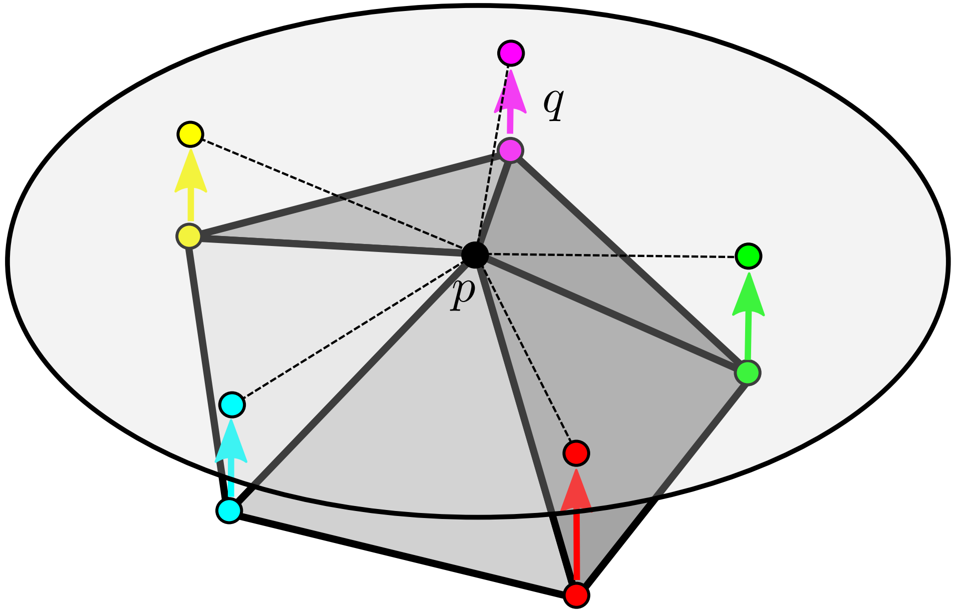

Gauge invariance, fixing & equivariance

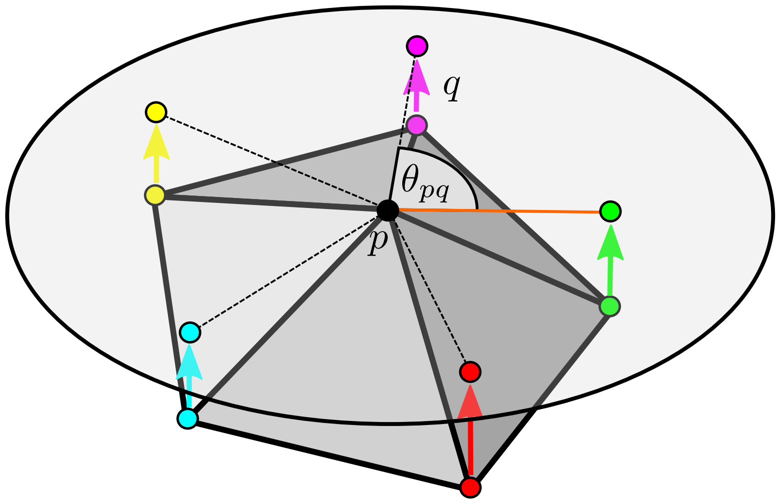

- Gauge: choice of basis for each tangent plane
  - Reference neighbour
- Option 1: gauge invariance
  - Message $q \to p$ independent of θ_{pq}
  - But: isotropic
- Option 2: gauge fixing
  - Principal curvature direction
  - But: ill-defined
- Option 3: Gauge equivariance [Cohen et al. 2019, Weiler et al. 2021]:
  - Features transform by known rule under gauge change
  - Network equivariant

$$
\begin{tikzcd}
V \arrow[r, "\text{Network}"] \arrow[d, "\text{Gauge transformation}"'] & V' \arrow[d, "\text{Gauge transformation}"] \\
V \arrow[r, "\text{Network}"'] & V'
\end{tikzcd}
$$

Invariance: Geodesic CNN

[Masci et al 2015]

- Fix any gauge
- Define kernel $K(r, \theta) \in \mathbb{R}$, unconstrained
- Compute convolution

$$(K \star f)_p = \max_{g \in SO(2)} \sum_{q \in \mathcal{N}(p)} K(r_q, \theta_q + g) f_q$$

- Invariant under gauge change

[Fig from paper]

Gauge equivariance with scalar features

- Feature $f: M \to \mathbb{R}$
- Transformation rule under gauge transformation: invariant
- Gauges w, w', polar coordinates of neighbour q of p:
$$w_p\left(\log_p q\right) = \left(r_q, \theta_q\right), \quad w'_p\left(\log_p q\right) = \left(r_q, \theta_q + g_p\right)$$
- Kernel $K(r, \theta) \in \mathbb{R}$
- In gauge w: $(K \star f)_p = \sum_{q \in \mathcal{N}(p)} K\left(r_q, \theta_q\right) f_q$
- In gauge w': $(K \star f)_p = \sum_{q \in \mathcal{N}(p)} K\left(r_q, \theta_q + g_p\right) f_q$
- Equality for any angle g_p implies $K\left(r_q, \theta_q + g_p\right) = K\left(r_q, \theta_q\right)$
- Kernel isotropic

Gauge equivariant convolutions on scalar fields

Scalar convolutions are isotropic

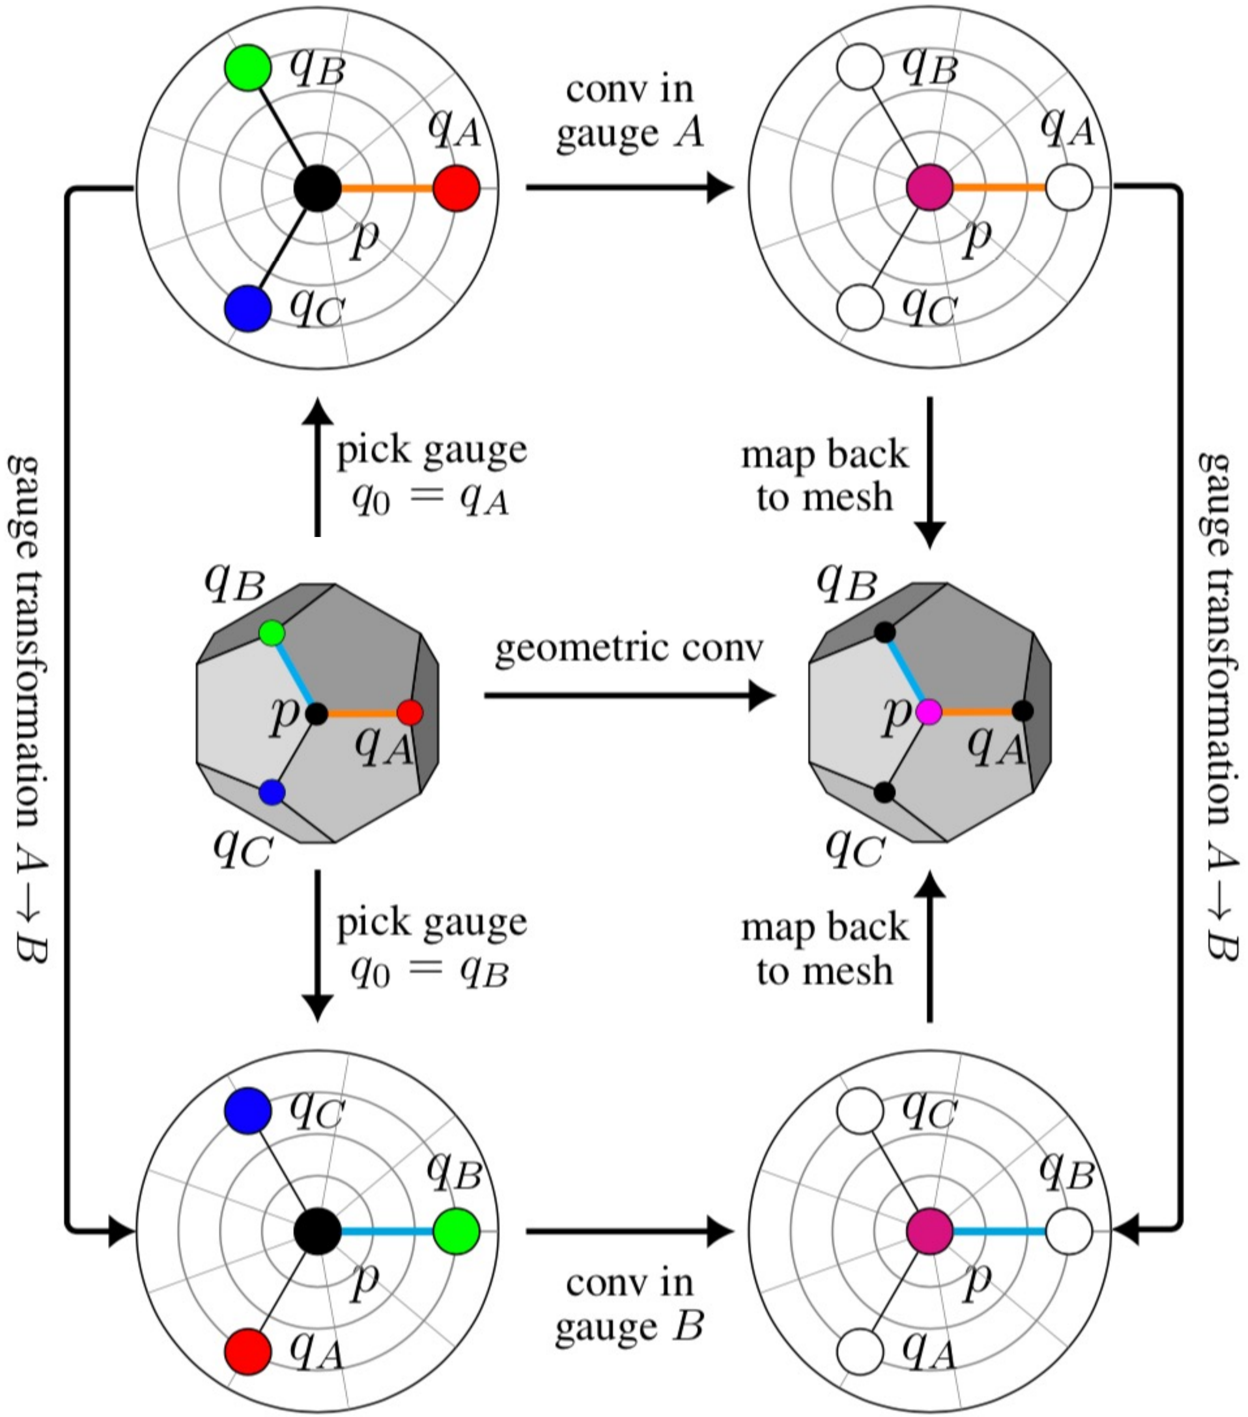

Vector features

- Vector space $V_p = \mathbb{R}^d$ that transforms under gauge transformation $g \in SO(2)$
- Transformation rule $\rho(g) \in \mathbb{R}^{d \times d}$
- Respects group structure: $\rho(g)\rho(g') = \rho(gg')$
- Group representation $\rho : G \to \mathrm{Aut}(\mathbb{R}^d)$
- Examples:
  - Scalar feature $\rho(g) = 1$
  - Tangent vector feature $\rho(g) = \begin{pmatrix} \cos(g) & -\sin(g) \\ \sin(g) & \cos(g) \end{pmatrix}$
- In general, concatenation of irreducible representations

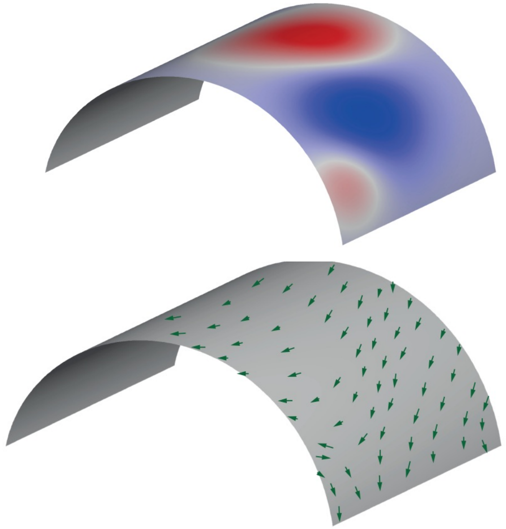

$$\rho_0(g) = 1, \quad \rho_n(g) = \begin{pmatrix} \cos ng & -\sin ng \\ \sin ng & \cos ng \end{pmatrix}$$

[Fig from Bronstein et al "Geometric Deep Learning"]

Parallel Transport

- Tangent planes not parallel
- Parallel transport of geodesic
- Transport gauge-defining X-axis
- Angle $g_{q \to p}$
- Any parallel transport by linearity

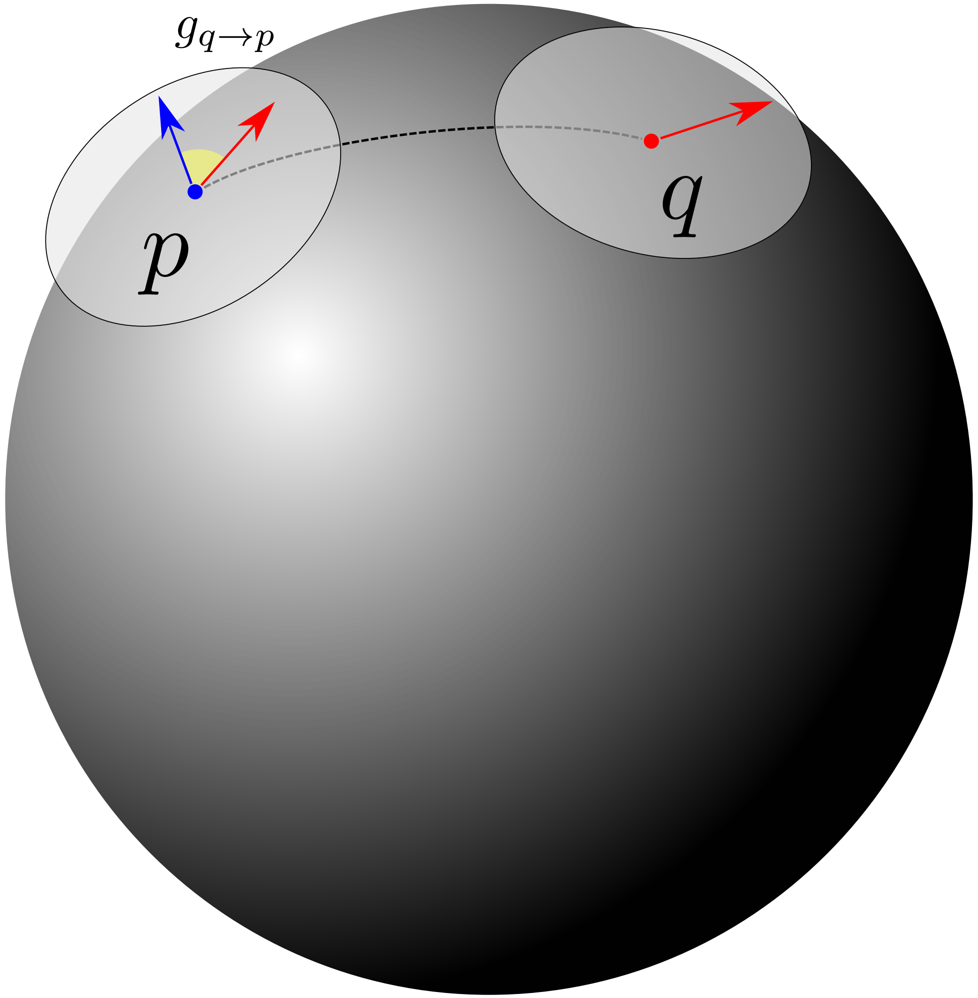

General Gauge Equivariant Convolution

- Input/output features $(\mathbb{R}^d, \rho), (\mathbb{R}^{d'}, \rho')$
- Kernel $K(r, \theta) \in \mathbb{R}^{d' \times d}$
- Convolution: $(K \star f)_p = \sum_{q \in \mathcal{N}(p)} K(r_q, \theta_q) \rho(g_{q \to p}) f_q$
- Equivariance if: $\rho'(g) K(r, \theta) = K(r, \theta + g) \rho(g)$

Gauge equivariant convolutions on vector fields

Vector convolutions are anisotropic

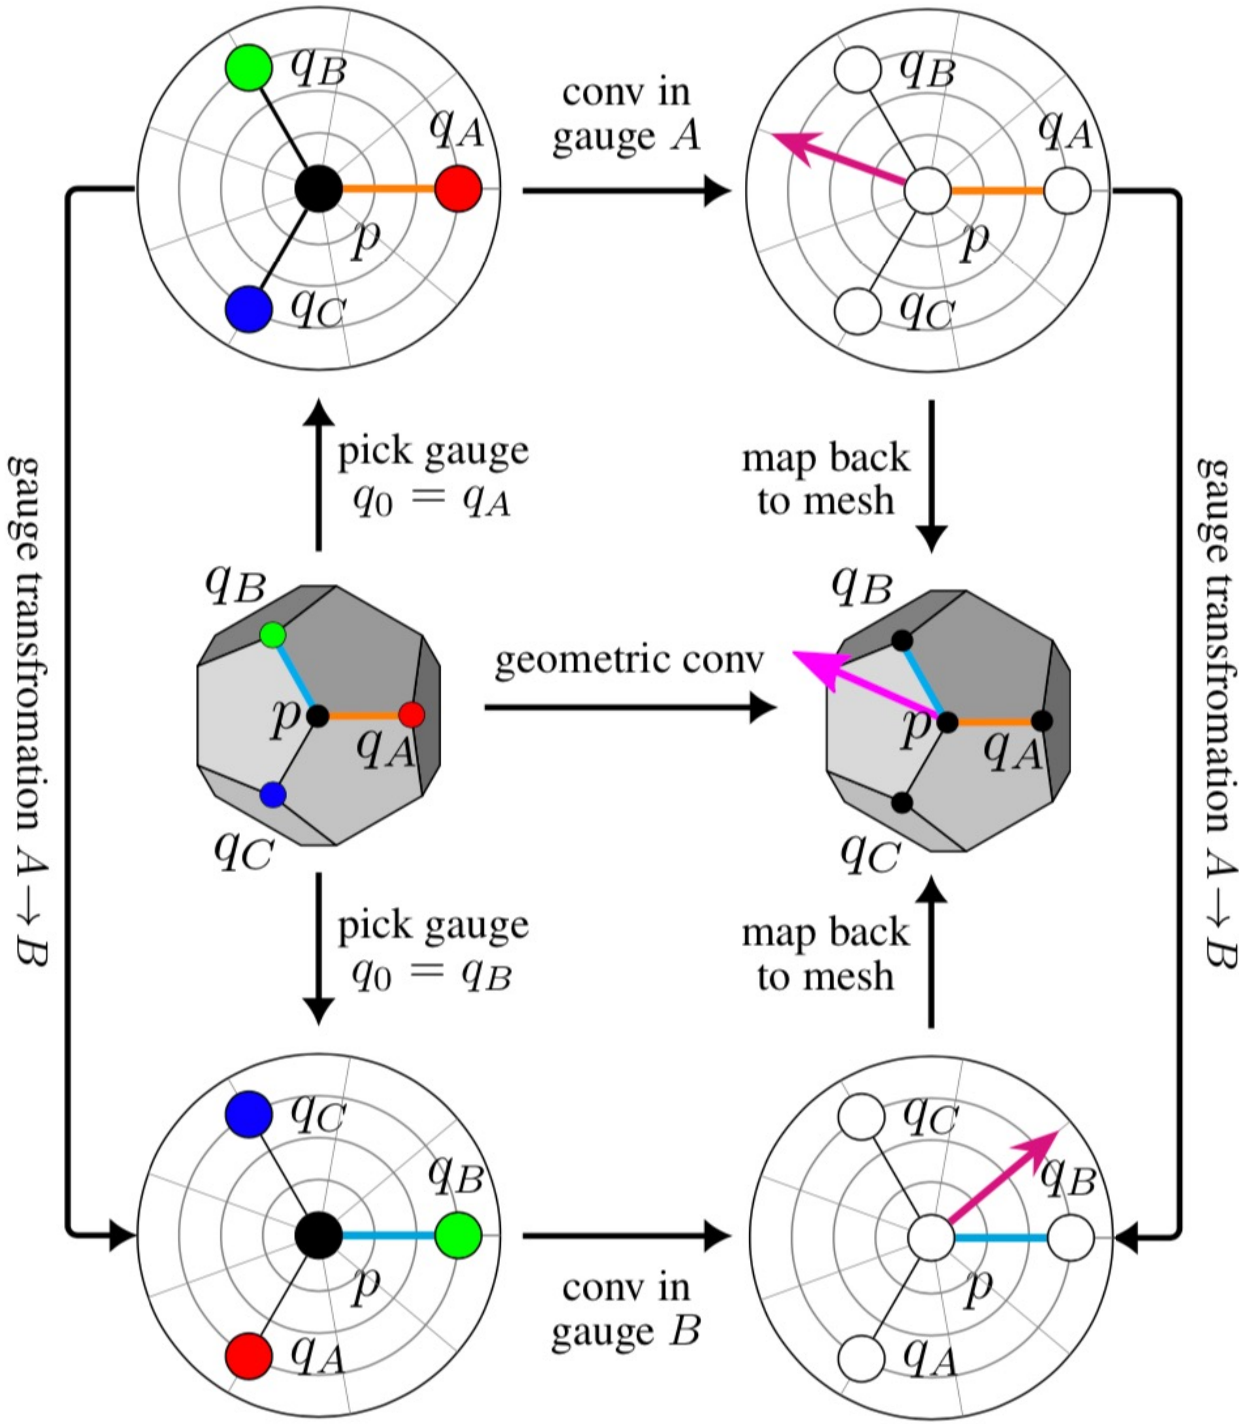

Solutions to kernel constraint

- $\rho'(g)K(r,\theta) = K(r,\theta+g)\rho(g)$
- $K(r,\theta) = K(r)K(\theta)$
- $K(r)$ unconstrained
- Example: between tangent vectors
- Angular component $K(\theta) \in \mathbb{R}^{2\times 2}$
- Four solutions
- Precomputed
- Linearly combined with learnable parameters

Example for $\rho_1 \to \rho_1$

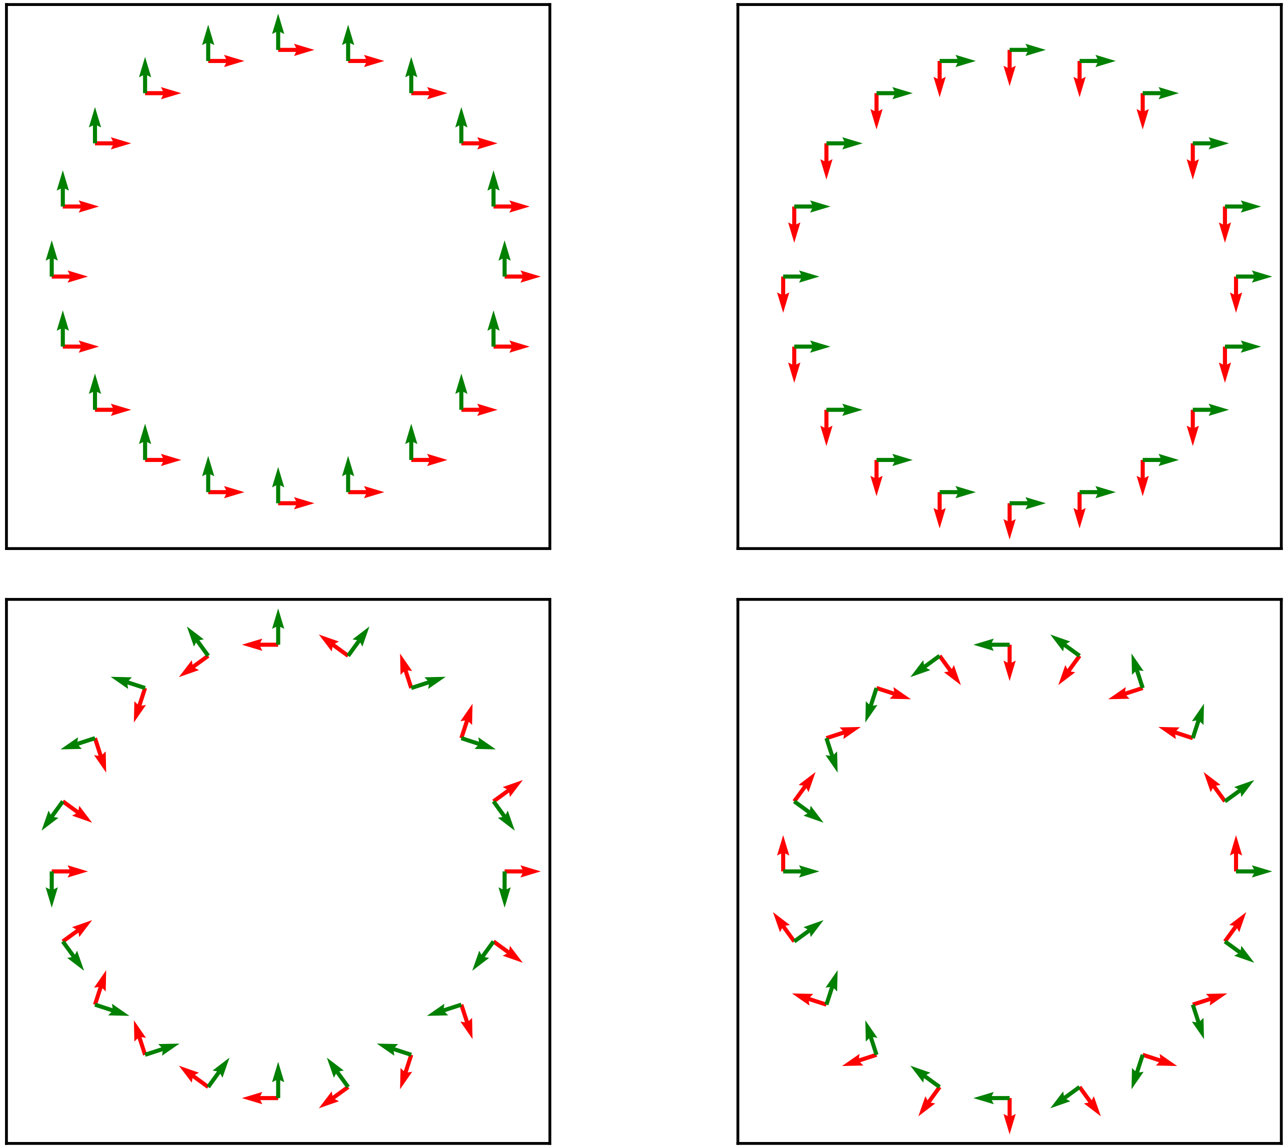

Solving the kernel constraint

- All representations can be characterised as copies of ($n \in \mathbb{N}$)

$$\rho_0(g) = 1, \quad \rho_n(g) = \begin{pmatrix} \cos ng & -\sin ng \\ \sin ng & \cos ng \end{pmatrix}$$

- Solutions for constraint:

$\rho_{in} \to \rho_{out}$	linearly independent solutions for $K_{\text{neigh}}(\theta)$
$\rho_0 \to \rho_0$	(1)
$\rho_n \to \rho_0$	$(\cos n\theta \;\; \sin n\theta)$, $(\sin n\theta \;\; -\cos n\theta)$
$\rho_0 \to \rho_m$	$\begin{pmatrix} \cos m\theta \\ \sin m\theta \end{pmatrix}$, $\begin{pmatrix} \sin m\theta \\ -\cos m\theta \end{pmatrix}$
$\rho_n \to \rho_m$	$\begin{pmatrix} c_- & -s_- \\ s_- & c_- \end{pmatrix}$, $\begin{pmatrix} s_- & c_- \\ -c_- & s_- \end{pmatrix}$, $\begin{pmatrix} c_+ & s_+ \\ s_+ & -c_+ \end{pmatrix}$, $\begin{pmatrix} -s_+ & c_+ \\ c_+ & s_+ \end{pmatrix}$

$\rho_{in} \to \rho_{out}$	linearly independent solutions for K_{self}
$\rho_0 \to \rho_0$	(1)
$\rho_n \to \rho_n$	$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$

$$c_{\pm} = \cos(m \pm n)\theta$$

- Learned parameters linearly combine basis kernels:

$$K(\theta) = \sum_i K_i(\theta)\,\alpha_i$$

Implementation

- For each layer, pick an input and output representation
- Precomputation:
  - For each vertex, we define tangent plane
  - On each tangent plane, we pick gauge
  - Between neighbouring vertices, we compute parallel transport matrices $\rho(g_{p\to q})$
  - For each pair of neighbours, construct basis kernels $K_i(r_q, \theta_q)$
- During forward pass, combine basis kernels:

$$(K \star f)_p = \sum_i \sum_{q \in \mathcal{N}(p)} \alpha_i K_i(r_q, \theta_q) \rho(g_{q\to p}) f_q$$

Symmetry properties

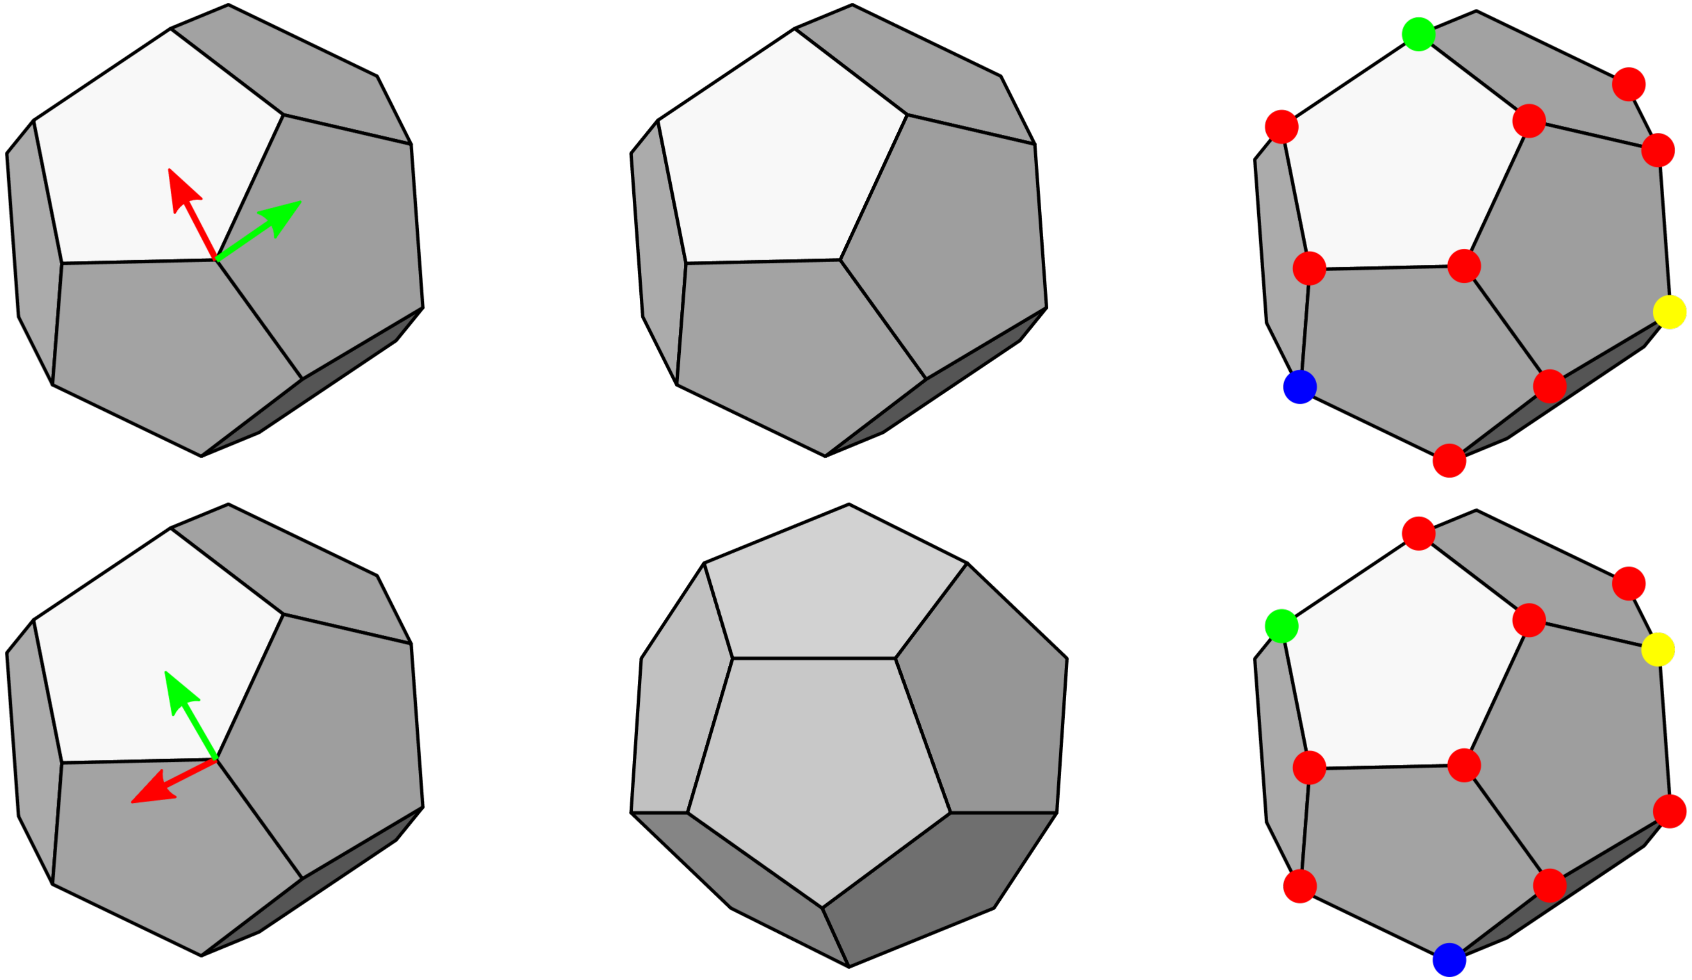

Gauge   Vertex coordinates   Mesh isometries

Reduction to Equivariant CNNs on images

CNN

Equivariant CNN

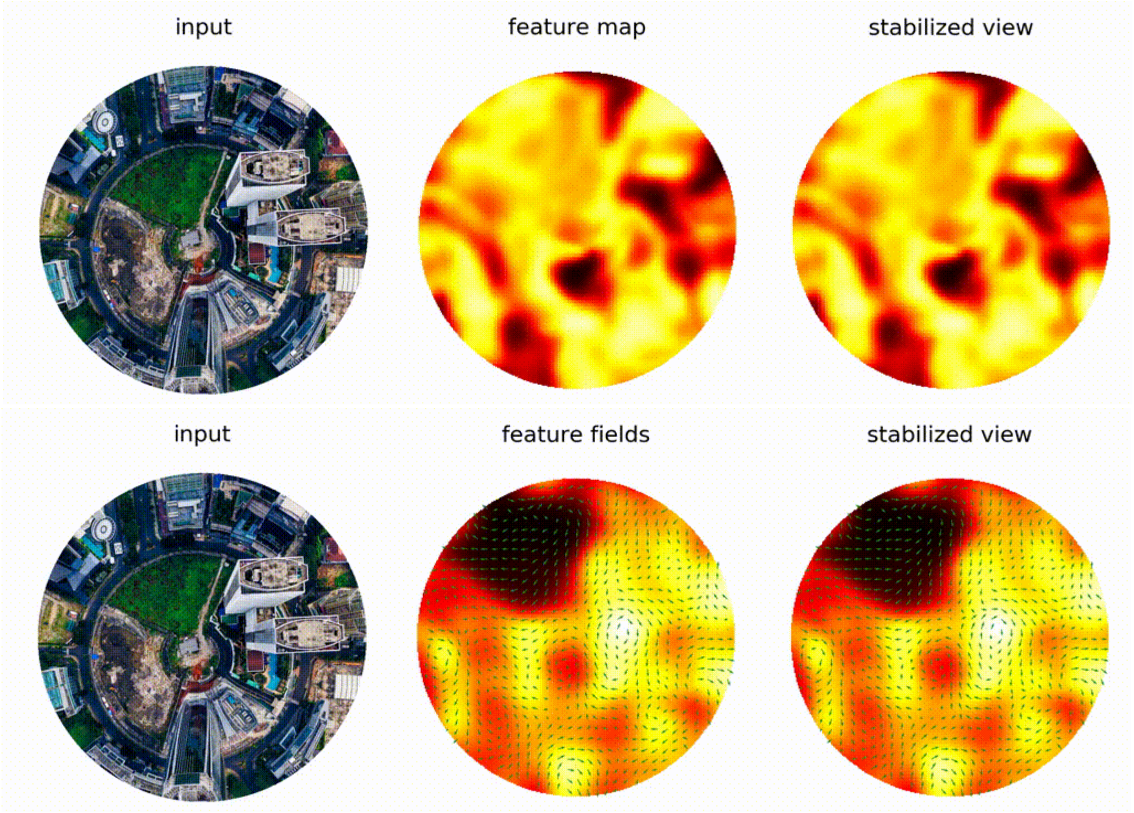

If mesh is planar grid, method reduces to equivariant planar CNN

[Cohen & Welling (2016), Weiler & Cesa (2019)]

Computing geometry geometry

- Geometric quantities
  - Logarithmic map
  - Parallel transport
- Edge walking
- Diffusion
- Spherical approximation

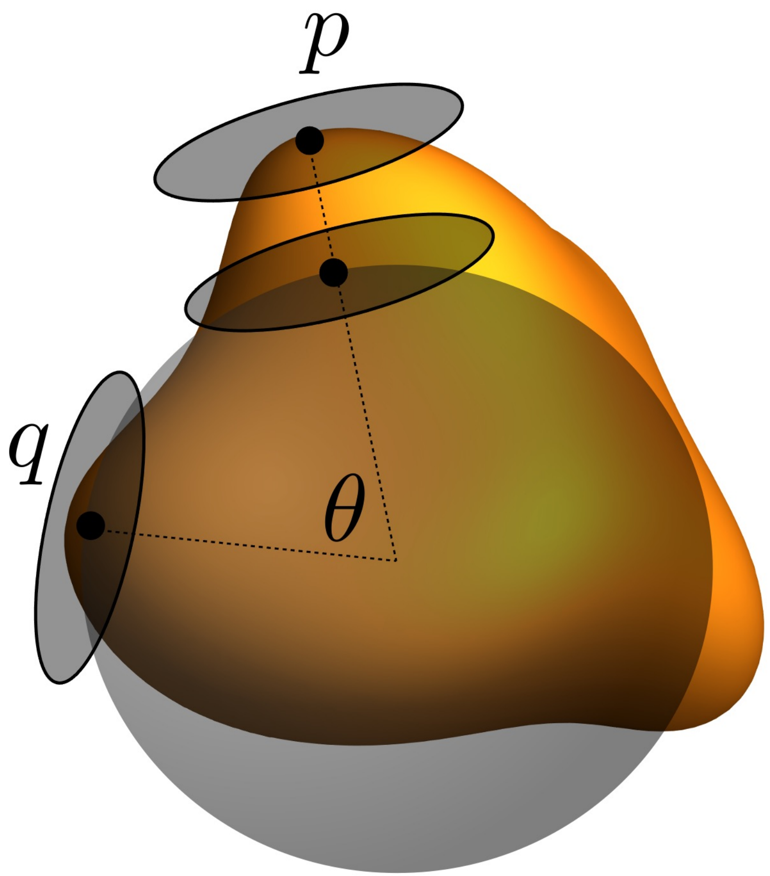

Edge walking

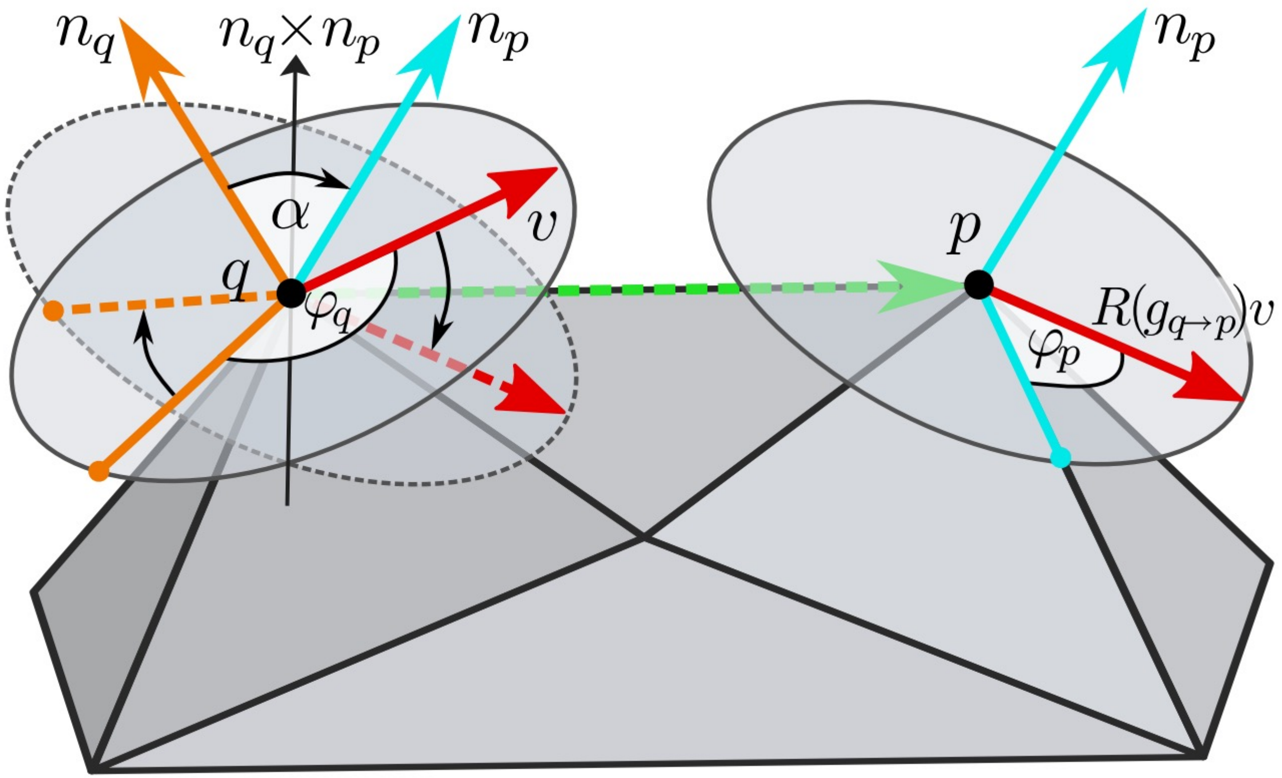

- Instable & expensive

Vector Heat Method

Sharp, Soliman, Krane (2019)

- Diffusion of vector fields
- Parallel transport
  - Vector field: x-axis at p, 0 elsewhere
  - Diffuse with infinitesimal time
  - Resulting vector at q is parallel transport $p \to q$
- Logarithmic map (r, θ)
  - r: length of transported R
  - θ: angle between H, R
- Discrete implementation
  - Complexity approaching $O(n)$ per node

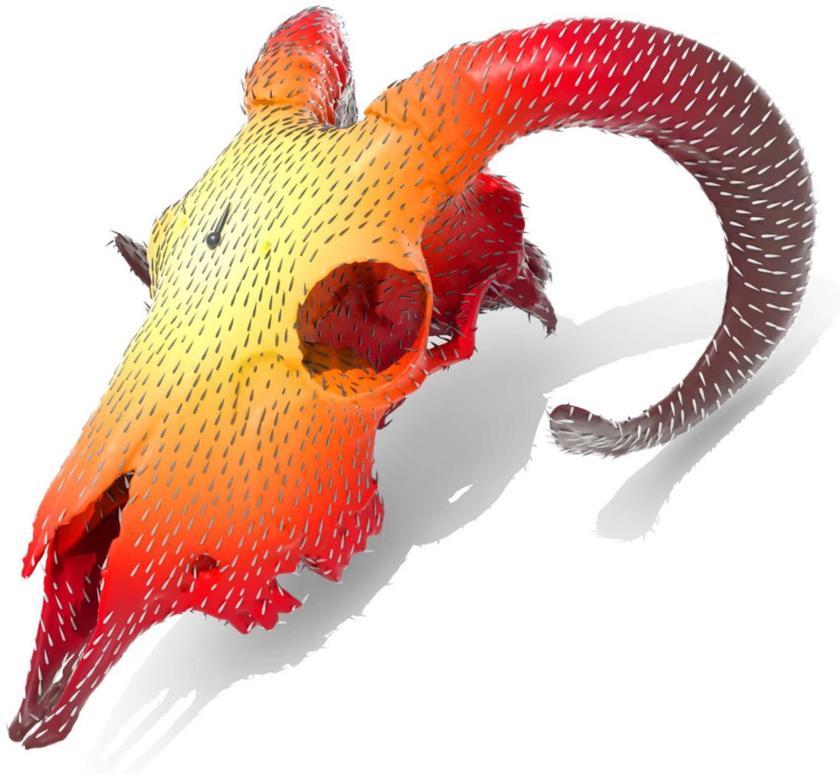

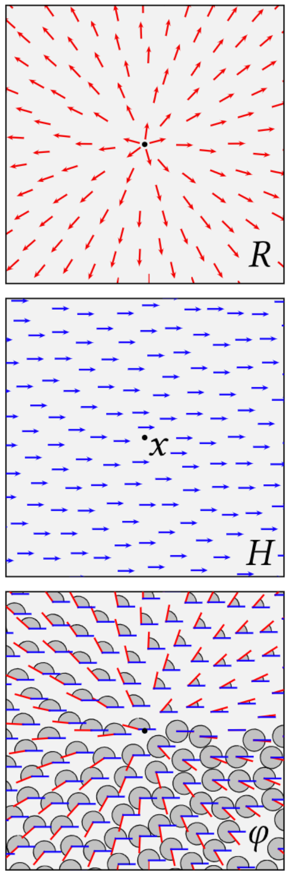

[Figs from paper]

Spherical Approximation

- Only depends on normals n_p, n_q (oriented point cloud)
- $R \in SO(3)$ rotate around $A = n_p \times n_q$ by $\theta = \arccos n_p^T n_q$
- Logarithmic map v s.t. $q = \exp_p v$:
  - $q - p \in \mathbb{R}^3$
  - Project to tangent plane $T_p M$
  - Rescale to length $||q - p||$
- Parallel transport:
  - $v \in T_p M$
  - Embed in $\mathbb{R}^3$
  - Rotate by $R \in SO(3)$
  - Normal to n_q
  - Reinterpret in $T_q M$

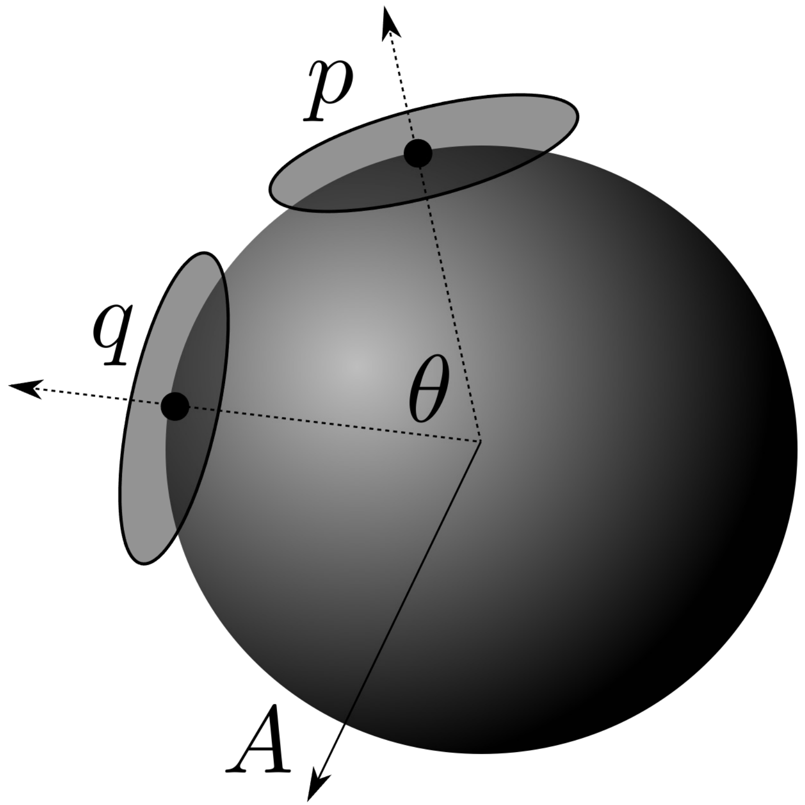

Gauge Equivariant Spherical CNN

Kicanaoglu, de Haan, Cohen [2019]

Non-linearity design

- Scalar multiplication $\phi: v \mapsto \alpha v$
- Nonlinearity $f: \mathbb{R} \to \mathbb{R}$
- Norm $\phi: v \mapsto \frac{f(|v|)}{|v|} v$
- Squash $\phi: v \mapsto \frac{|v|}{|v|+1} v$
- Gated nonlinearity $\phi: r, v \mapsto f(r)v$
- If $\rho(g)$ a permutation: $v \mapsto (gv)_i = v_{\sigma i}$
  - Pointwise equivariant $\phi(v)_i = f(v_i)$
- Infinite groups?

$$
\begin{tikzcd}
\mathbb{R}^d \arrow[r, "\phi"] \arrow[d, "\rho(g)"'] & \mathbb{R}^d \arrow[d, "\rho(g)"] \\
\mathbb{R}^d \arrow[r, "\phi"'] & \mathbb{R}^d
\end{tikzcd}
$$

Sampling non-linearity

- Continous signals on circle $C(S^1) = \{v: S^1 \to \mathbb{R}\}$
- Regular representation of $SO(2)$: $v \mapsto gv(x) = v(g^{-1}x)$
- Pointwise equivariant: $v \mapsto \phi(v)(x) = f\big(v(x)\big)$
- Infinite memory
- Band-limited fourier transform $\rho_0 + \rho_1 + \cdots + \rho_B \to C(S^1)$
- Finite sample: $\rho_0 + \rho_1 + \cdots + \rho_B \to C(S^1) \to \mathbb{R}^n$
- Point-wise on samples
- Equivariance error $\mathcal{O}(B^2/N)$

Input features

- XYZ coordinates are gauge-invariant, but not coordinate-free
- Ambient feature as surface feature
  - Vector: $T_p\mathbb{R}^3 \xrightarrow{\sim} T_pM + \mathbb{R}$
  - Matrix $T_p^2\mathbb{R}^3 \xrightarrow{\sim} T_p^2M + 2T_pM + \mathbb{R}$
- Ambient matrices
  - $M_p^1 = \frac{1}{\mathcal{N}(p)} \sum_{q\in\mathcal{N}(p)} n_p n_q^T$
  - $M_p^2 = \frac{1}{\mathcal{N}(p)} \sum_{q\in\mathcal{N}(p)} (q-p)(q-p)^T$
  - $M_p^3 = \frac{1}{\mathcal{N}(p)} \sum_{q\in\mathcal{N}(p)} (q-p)n_p^T$
- Initial features $3(3\,\rho_0 + 2\rho_1 + \rho_2)$

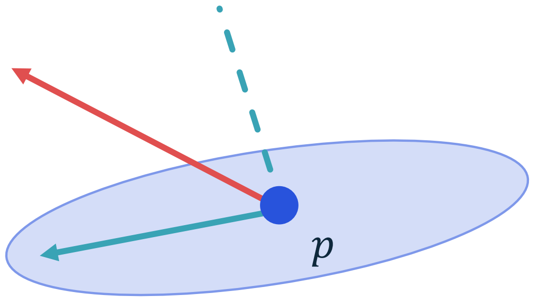

Pooling

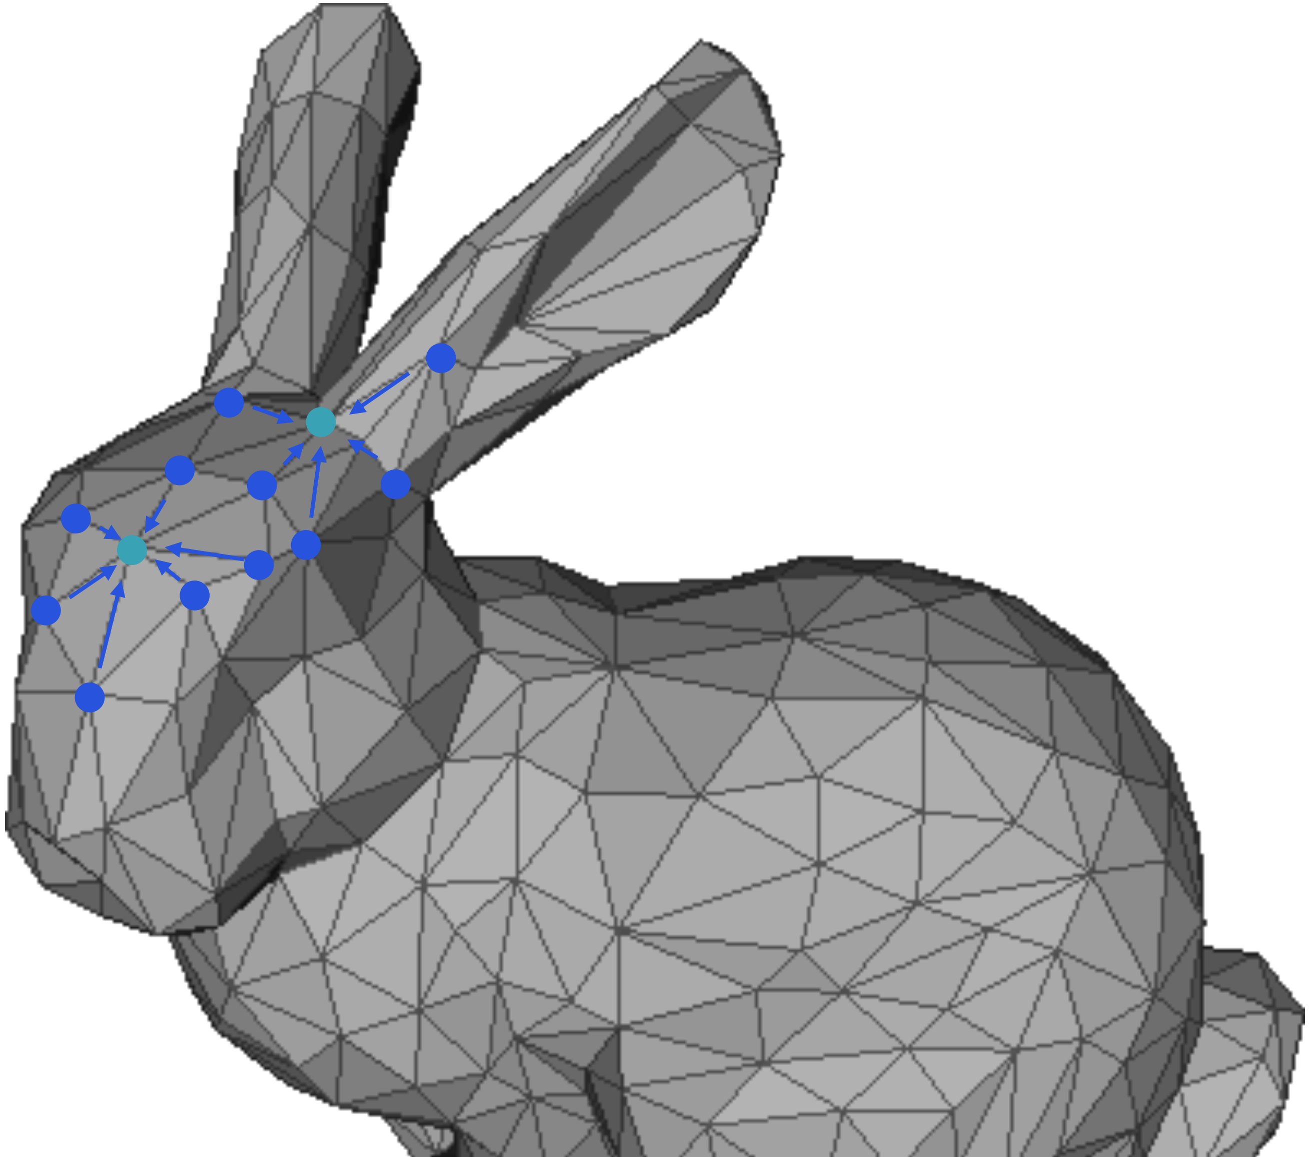

Application to blood flow

[Suk, de Haan, Lippe, Brune, Wolterink, 2021]

Julian Suk
University of Twente

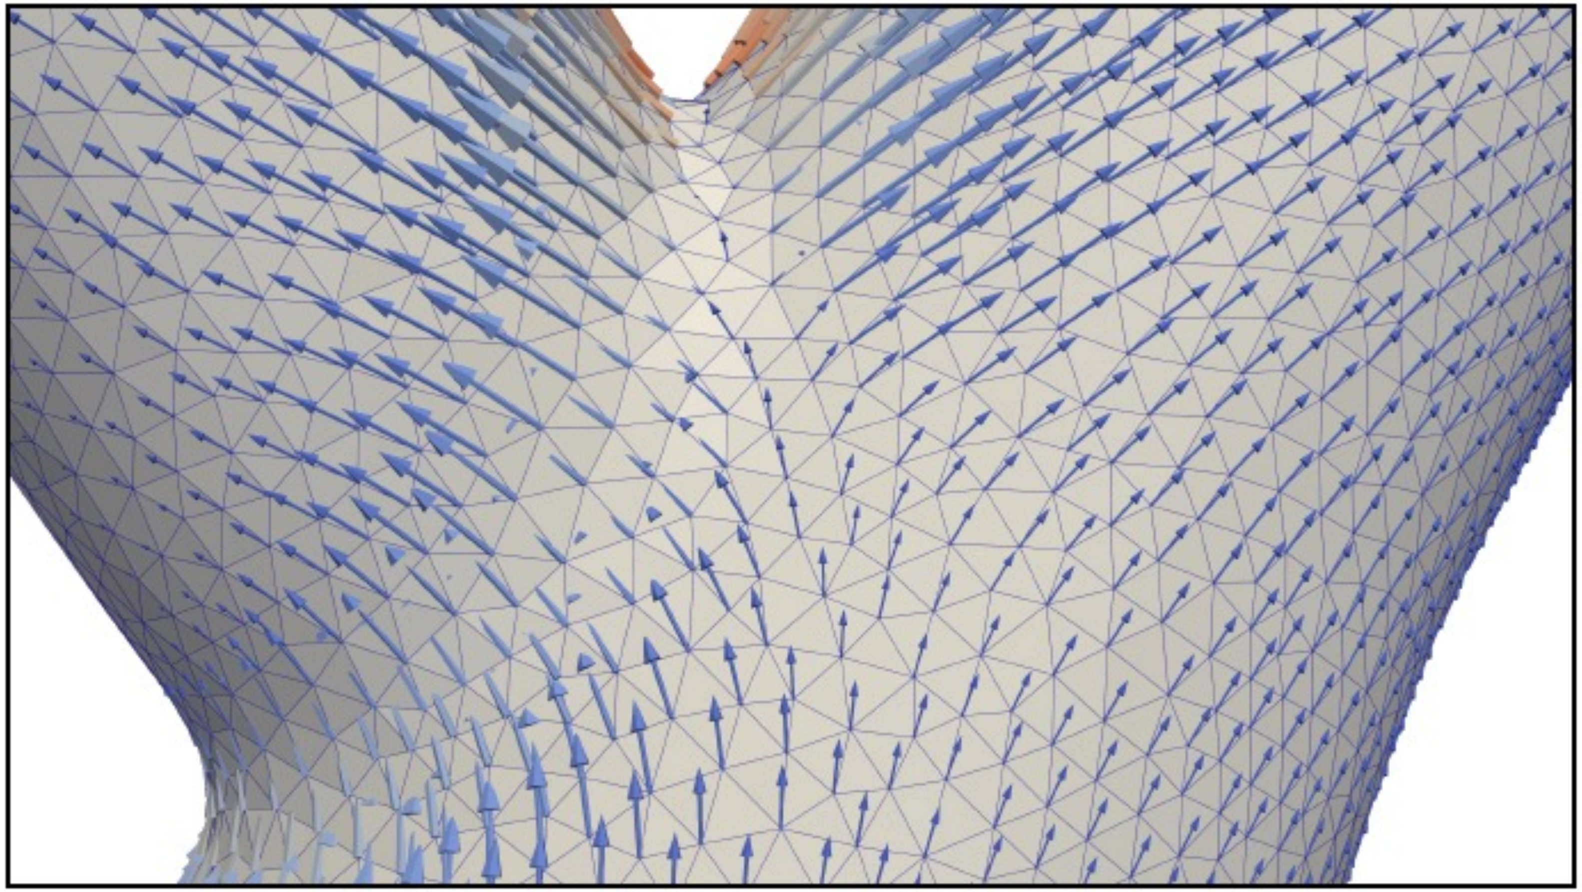

Problem formulation

- Shape of arteries in human body related to e.g. aneurysm
- Quantitative analysis of blood flow useful indicator – wall shear stress
- Non-invasively: model artery with MRI scanner
- Simulate blood flow with computational fluid dynamics (> 20h)
- Learn neural network surrogate to predict WSS on CFD ground truth
- Dataset of realistic random meshes
- Neural network inference in milliseconds

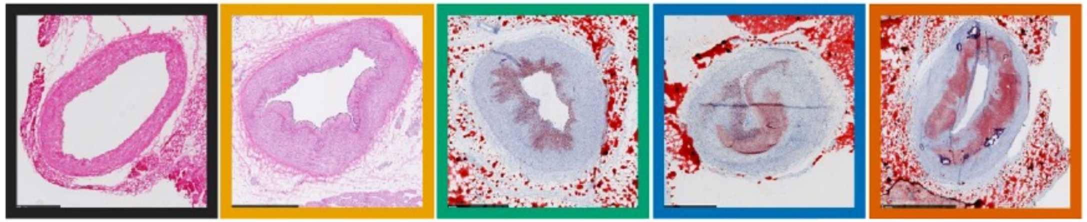

Equivariance

- Arteries not in canonical orientation
- Equivariance to global transformations
- Gauge Equivariant Mesh CNN

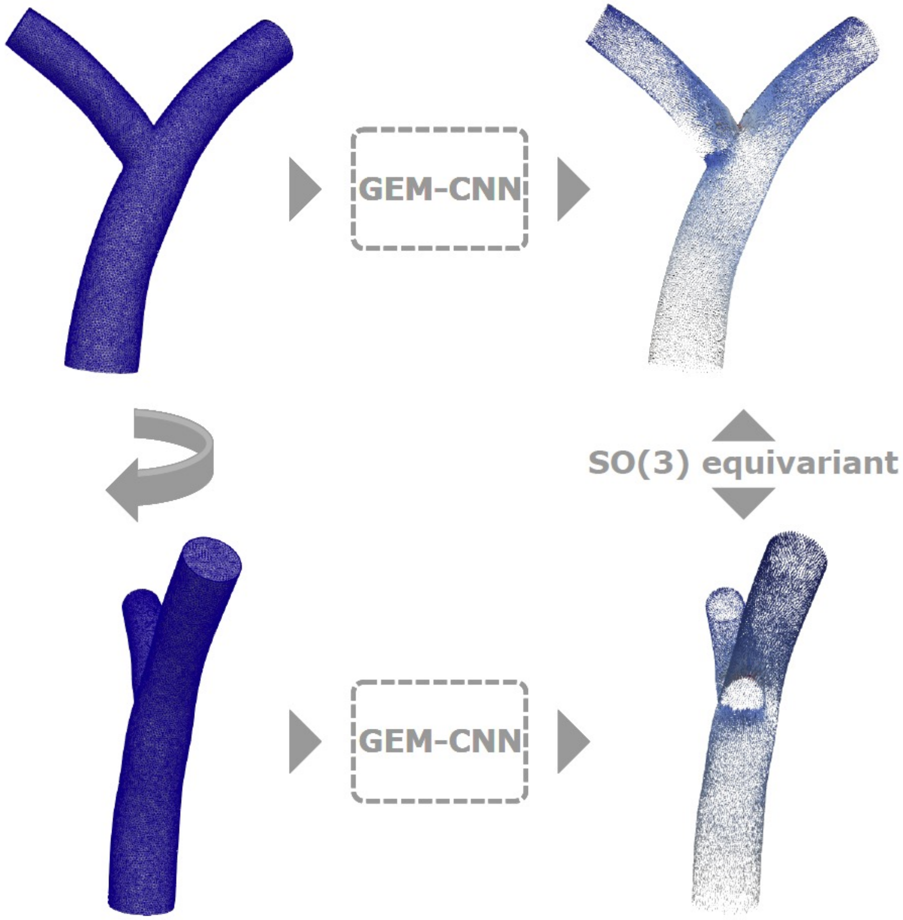

Network Architecture

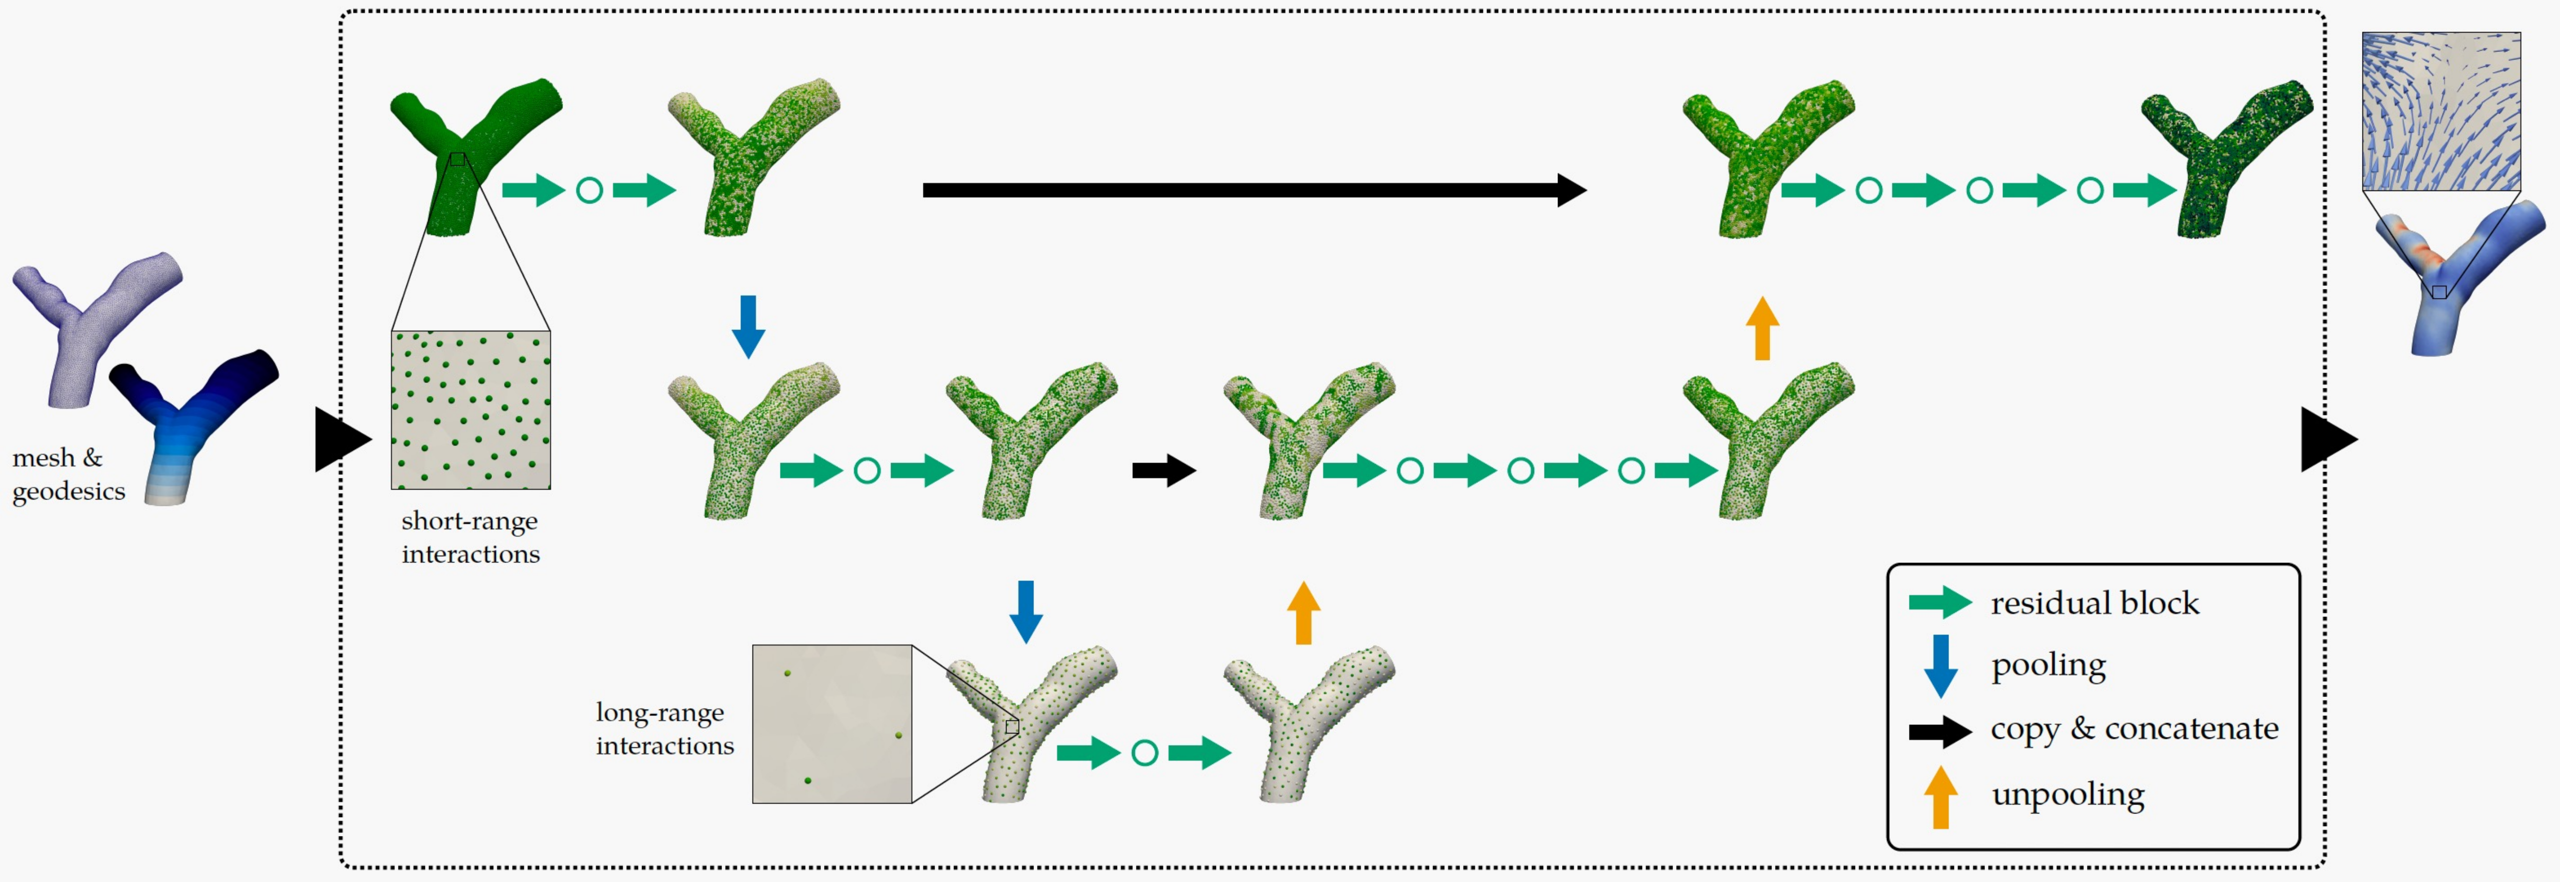

Predicting WSS across the Cardiac Cycle

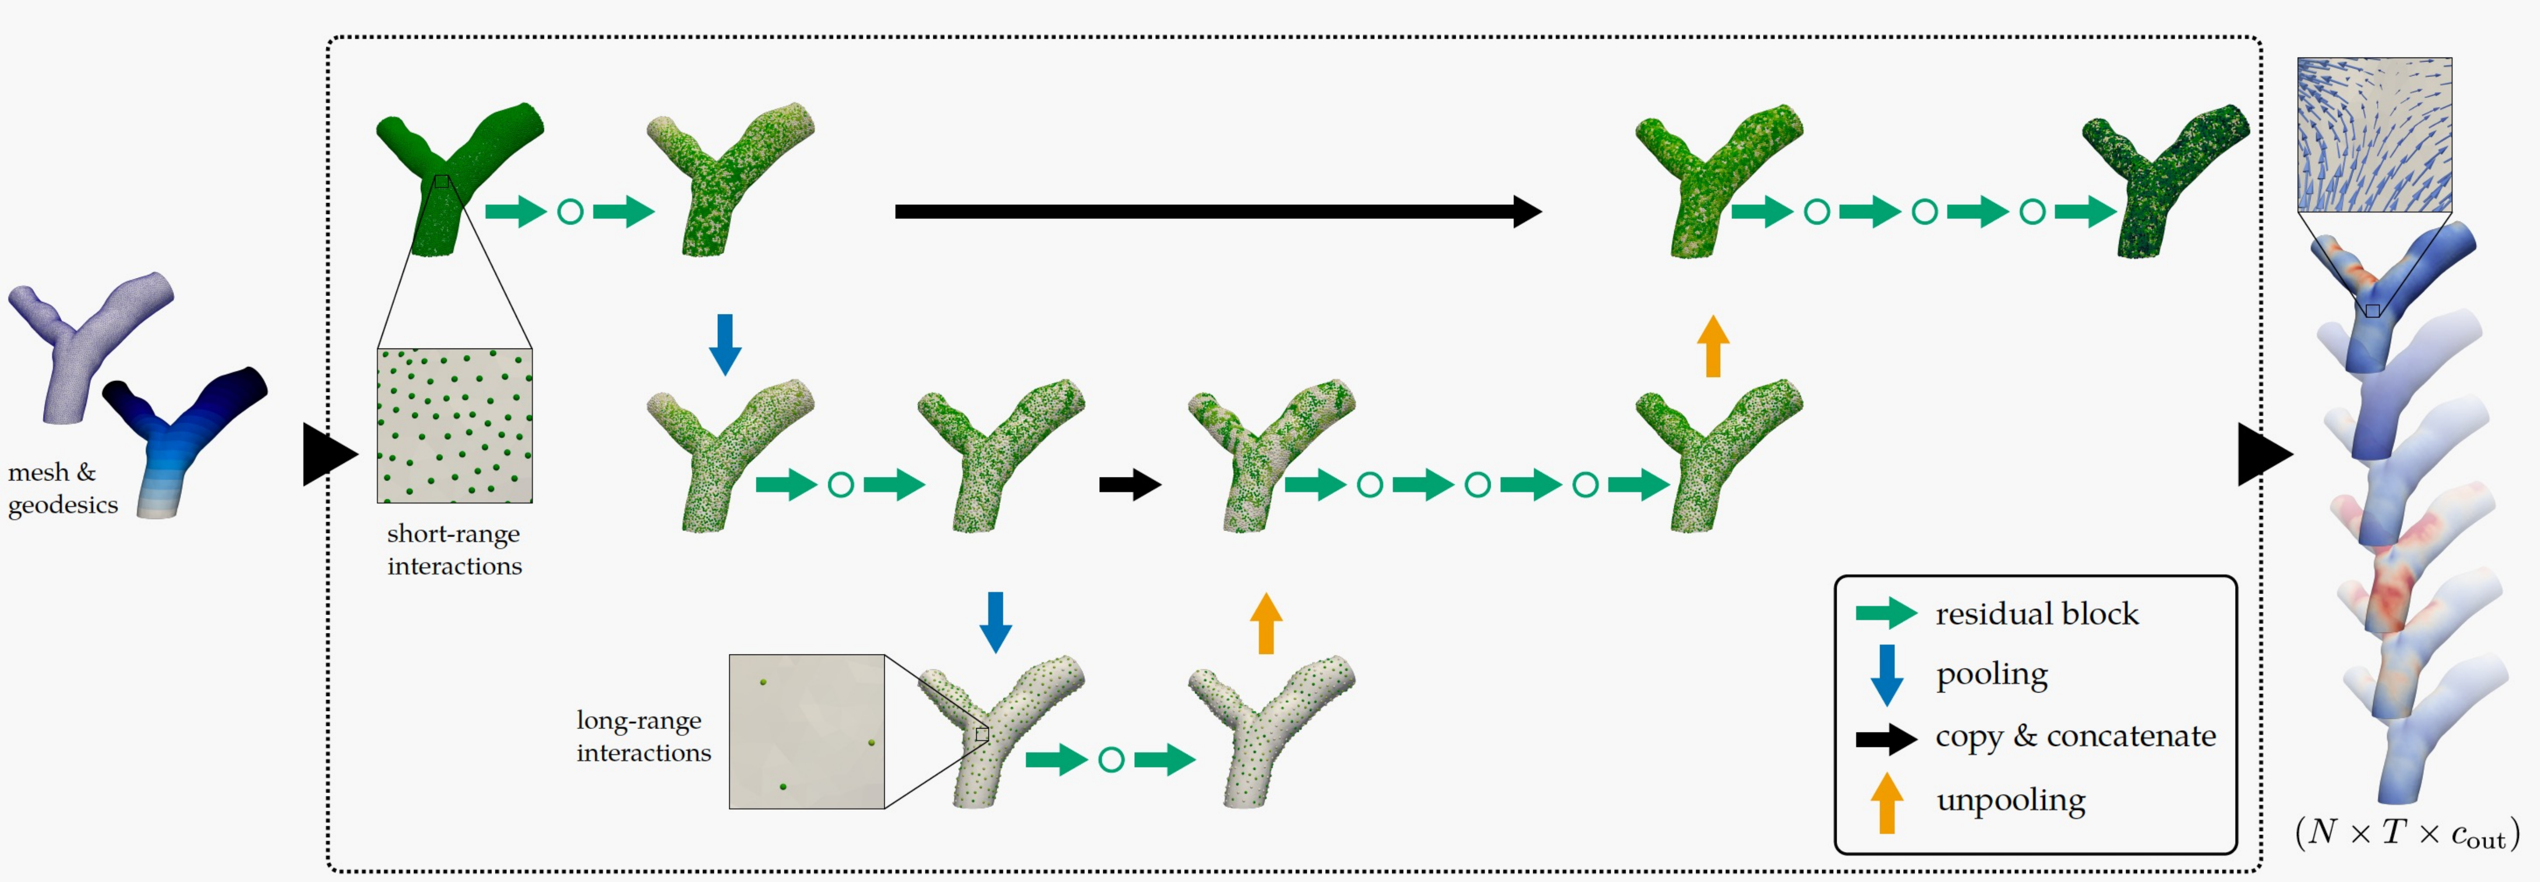

Qualitative Results on Steady Flow

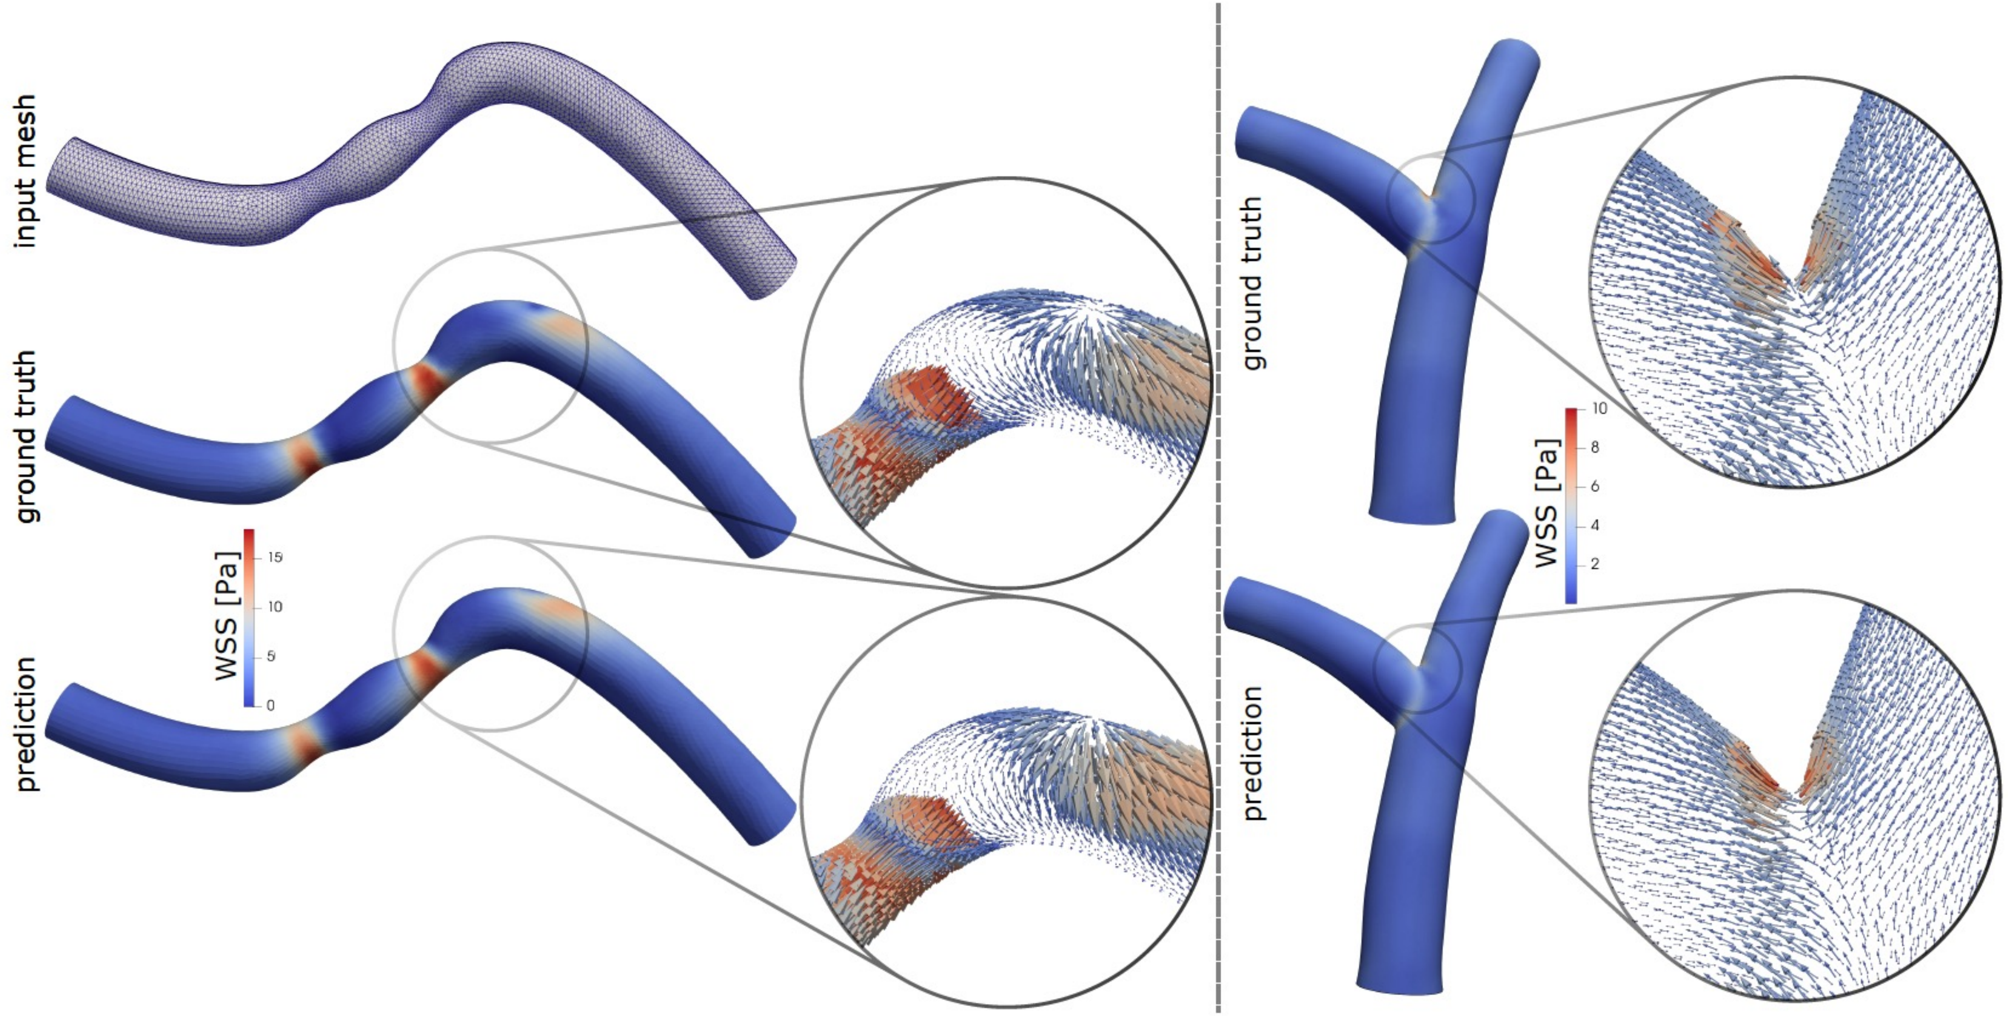

Quantitative Results on Steady Flow

		NMAE [%]		
		mean	median	75th
Single arteries (steady flow)	SAGE-CNN	0.9	0.9	1.2
	FeaSt-CNN	0.6	0.6	0.8
	PointNet++	**0.5**	**0.4**	0.7
	GEM-CNN	**0.5**	**0.4**	**0.6**
(randomly rotated in 3D)	PointNet++†	10.1	10.0	11.9
	PointNet++‡	0.7	0.6	1.0
	GEM-CNN†	**0.5**	**0.4**	**0.6**
Bifurcating arteries (steady flow)	SAGE-CNN	1.0	0.9	1.0
	FeaSt-CNN	0.7	0.6	0.7
	PointNet++	**0.6**	**0.5**	**0.6**
	GEM-CNN	**0.6**	0.6	0.7
(randomly rotated in 3D)	PointNet++†	7.8	7.6	11.0
	PointNet++‡	**0.6**	**0.6**	**0.7**
	GEM-CNN†	**0.6**	**0.6**	**0.7**

† trained on canonically oriented samples

‡ trained under data augmentation (random rotation in 3D)

Qualitative Results on Pulsatile Flow

Ground truth

Prediction

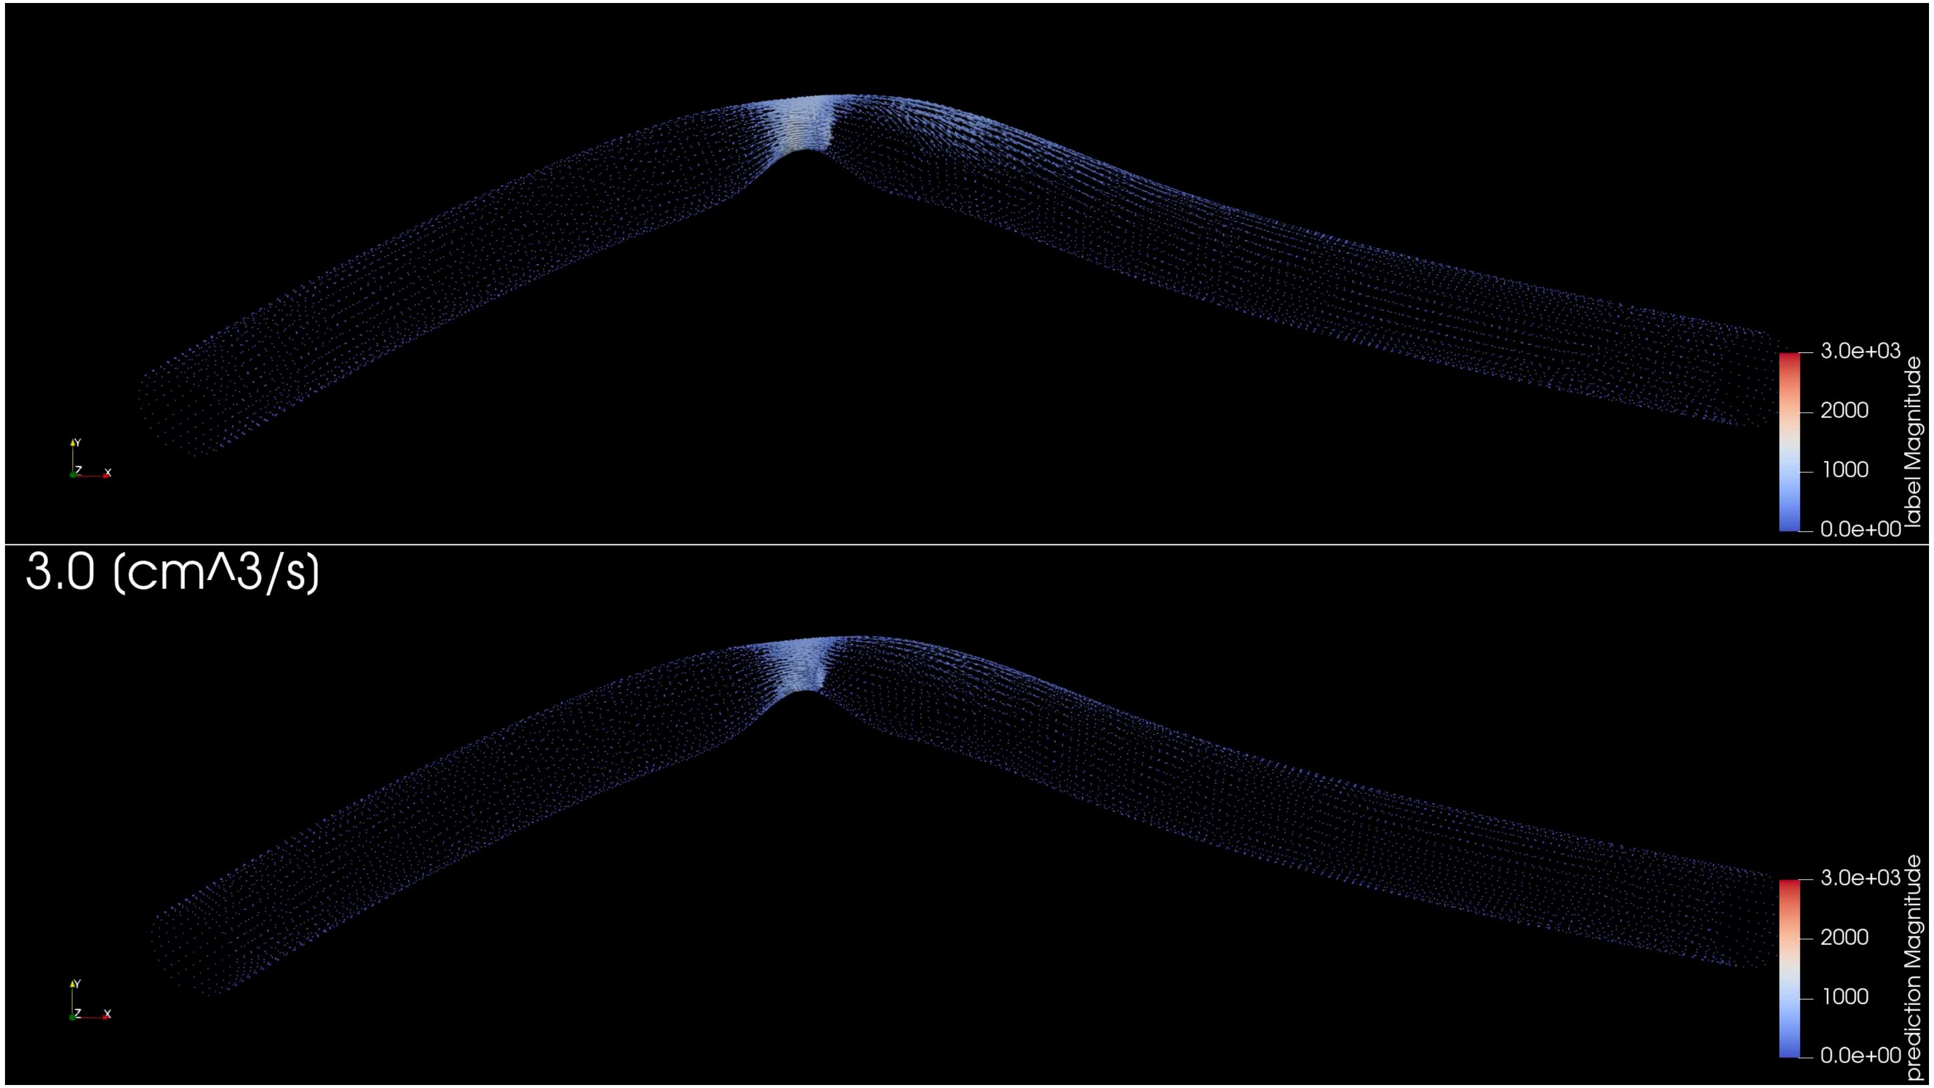

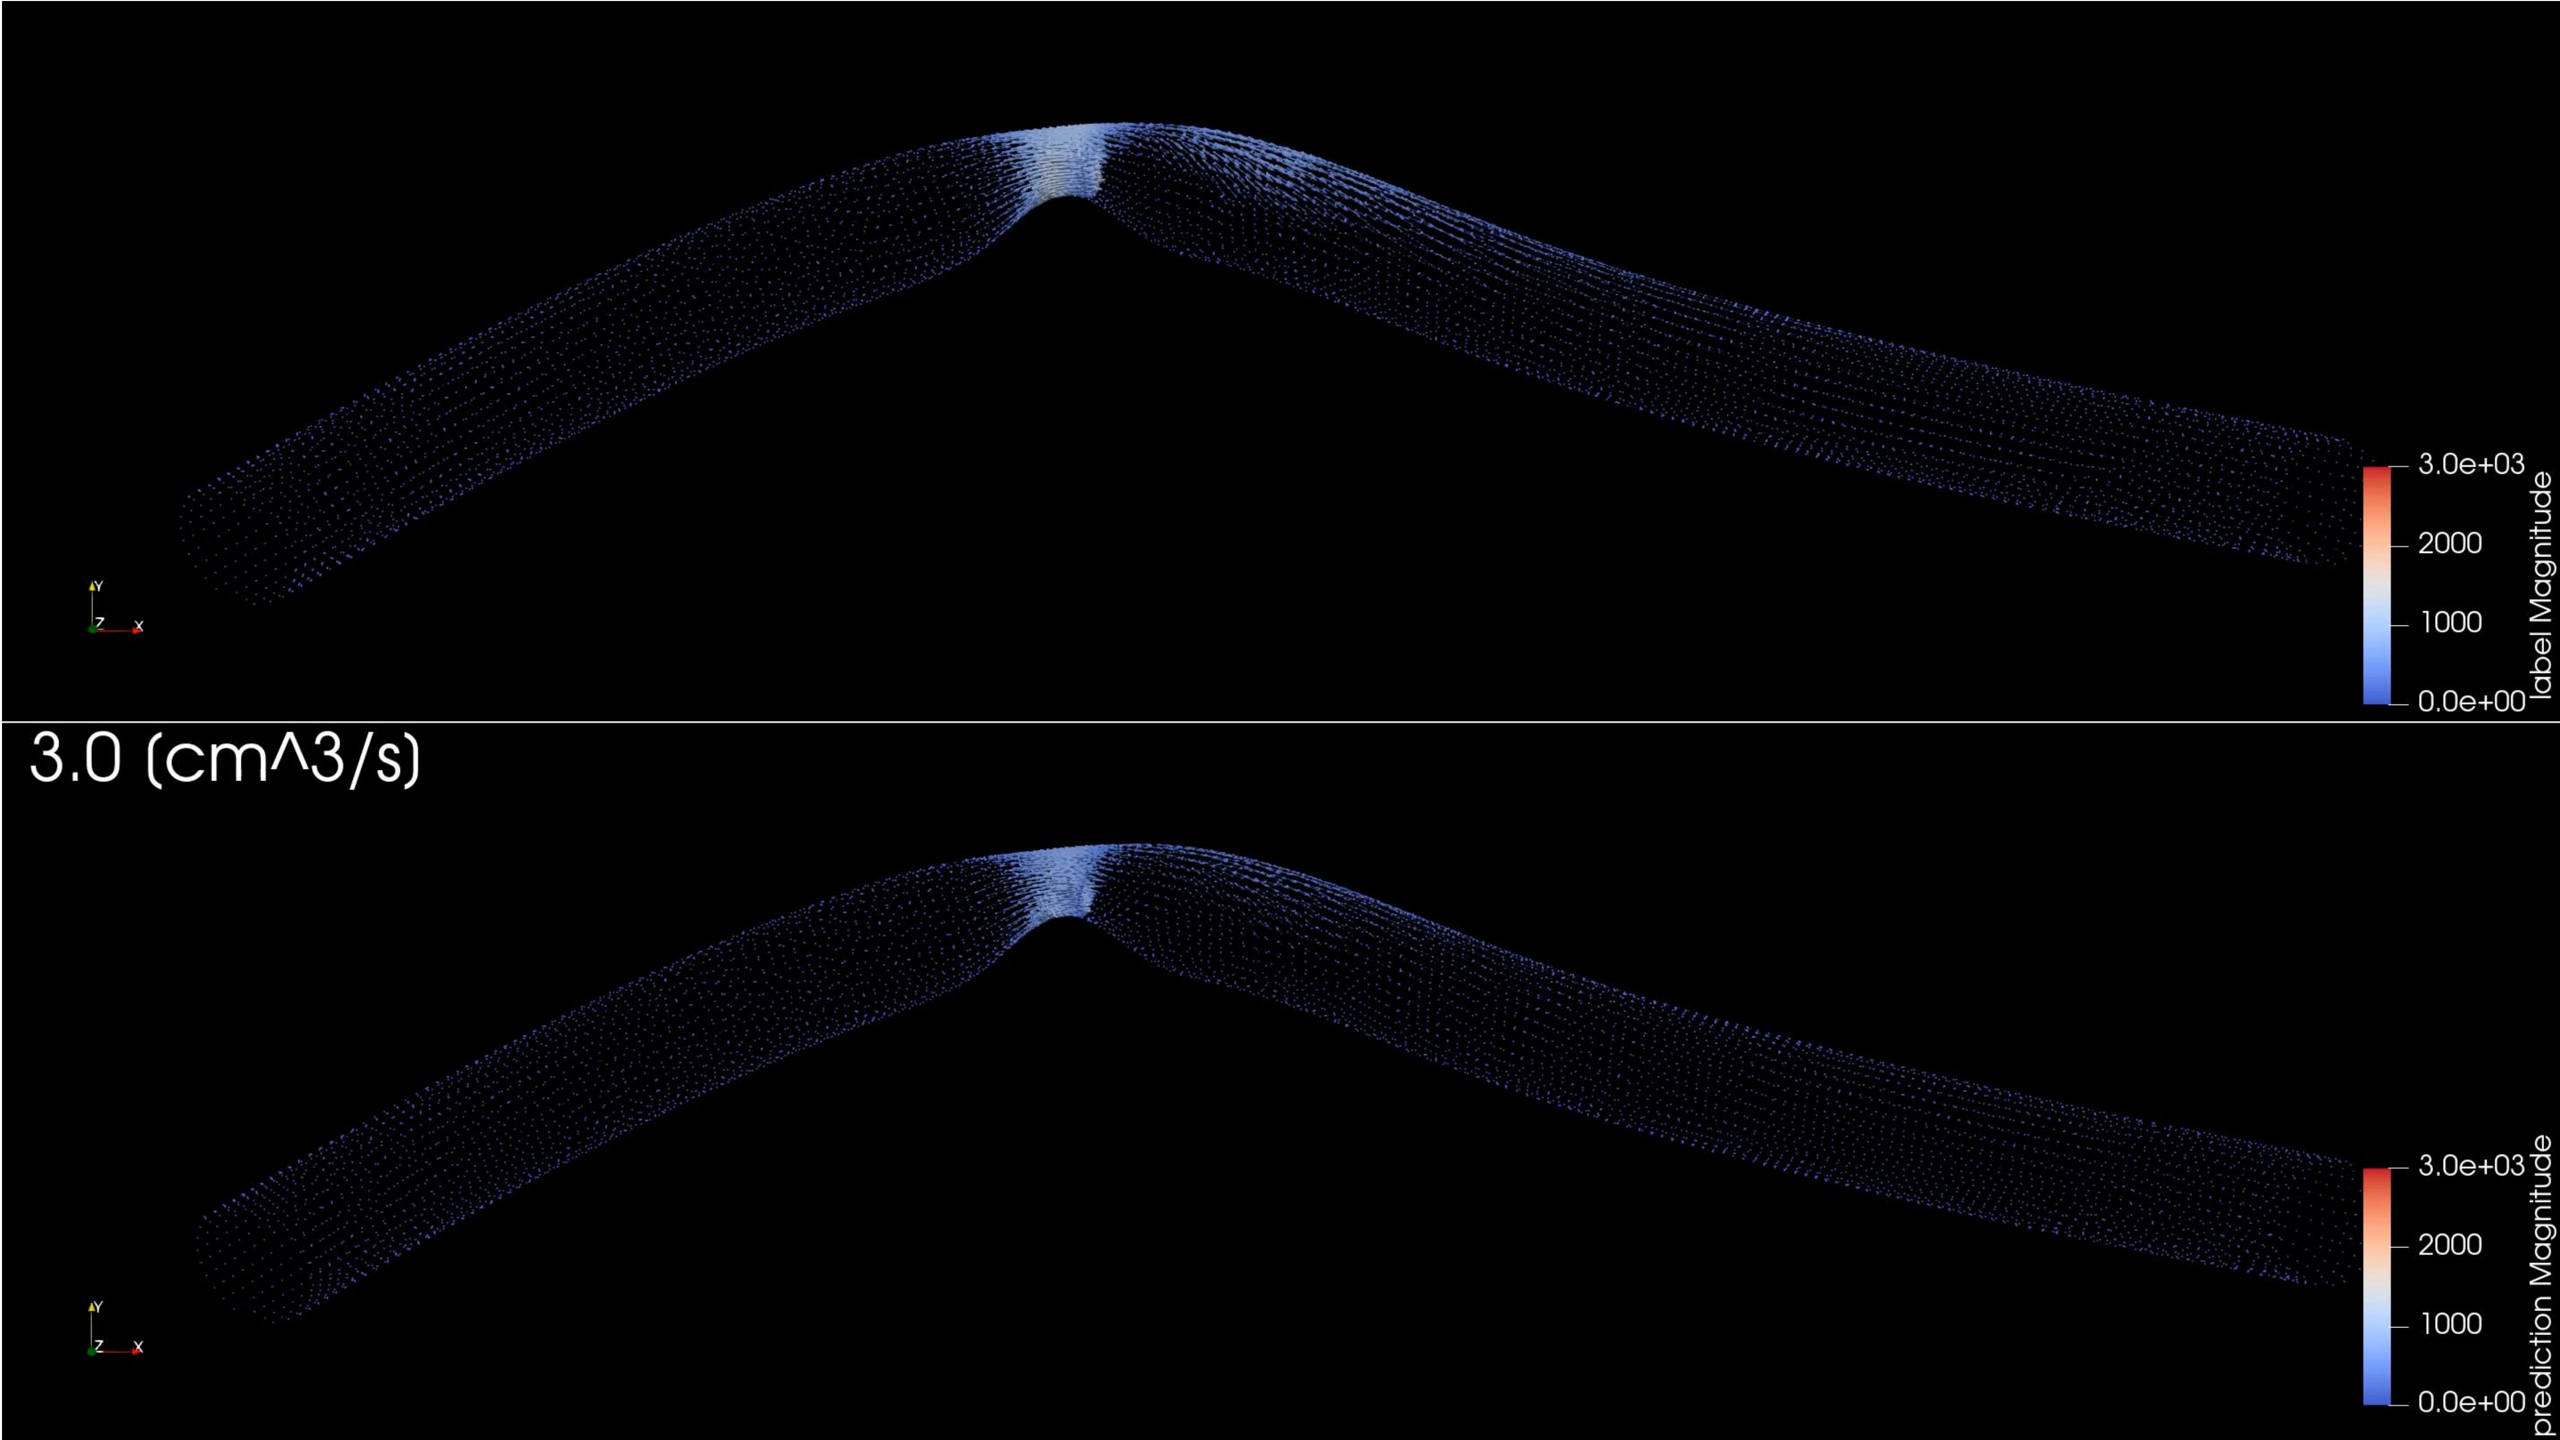
3.0e+03
2000
1000
0.0e+00
label Magnitude
3.0 (cm^3/s)
3.0e+03
2000
1000
0.0e+00
prediction Magnitude

Takeaway

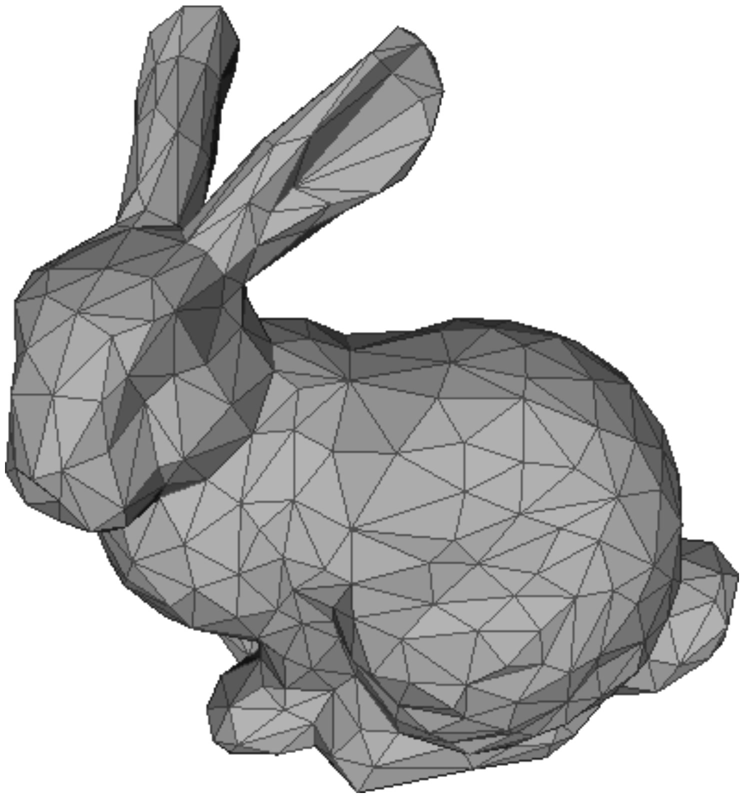

Gauge Equivariant Mesh CNN is:

- Simple
- Scalable
- Anisotropic ⇒ expressive
- Symmetry properties
- Try it out:
  github.com/Qualcomm-AI-research/gauge-equivariant-mesh-cnn

Qualcomm

Thank you

Follow us on:

For more information, visit us at:

www.qualcomm.com & www.qualcomm.com/blog