GRAPH NEURAL NETWORKS & ROTATIONAL EQUIVARIANCE

University of California, Berkeley Fall 2023, CS 189/289A: Introduction to Machine Learning

Rasmus Malik Hoeegh Lindrup

Postdoc, BAIR/ICSI

Building on slides originally made by Daniel Rothchild

LECTURE 2

OUTLINE

- Lecture 1
  - Graph data
  - Graph tasks
  - Message passing
  - Invariance and equivariance
- Lecture 2
  - Recap: Invariance and equivariance
  - Rotational equivariance
  - Equivariant neural networks

GEOMETRIC INFORMATION

GEOMETRIC INFORMATION

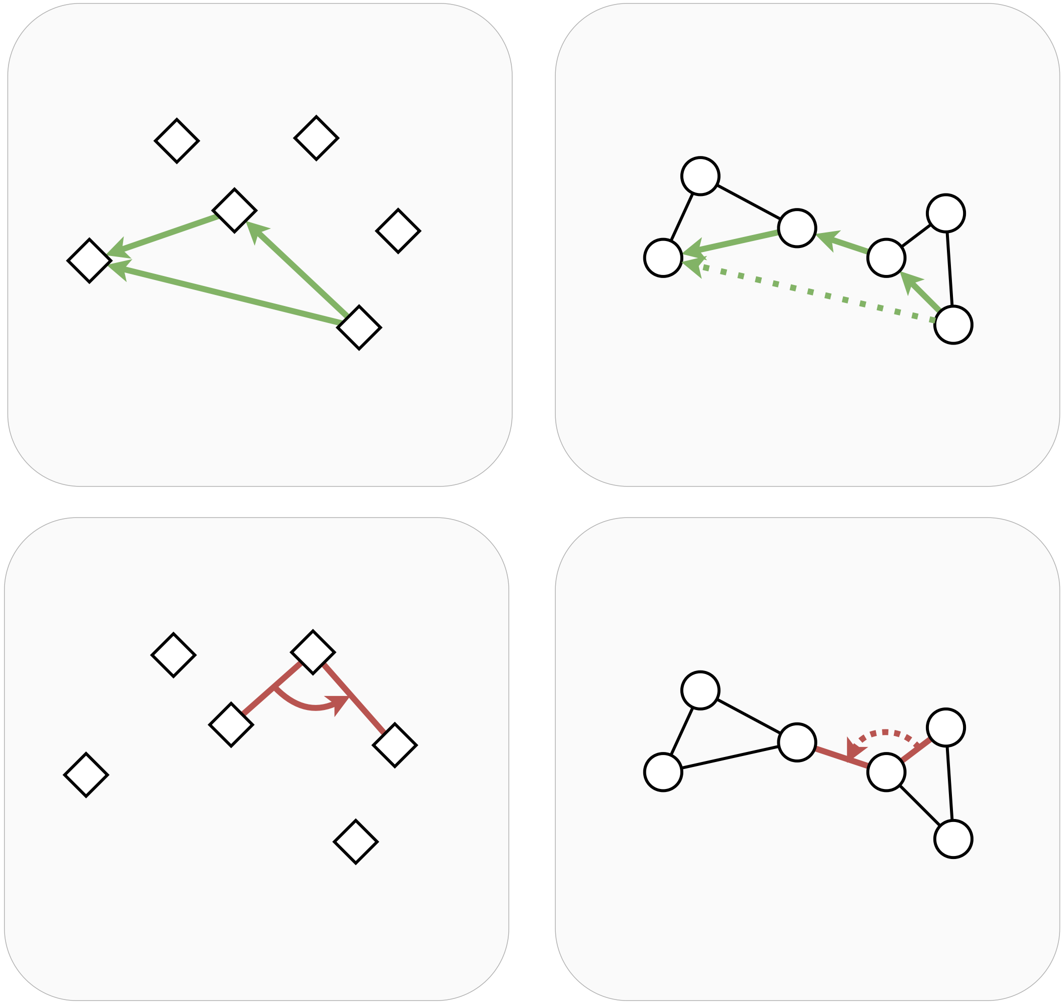

ROTATION INVARIANCE AND EQUIVARIANCE

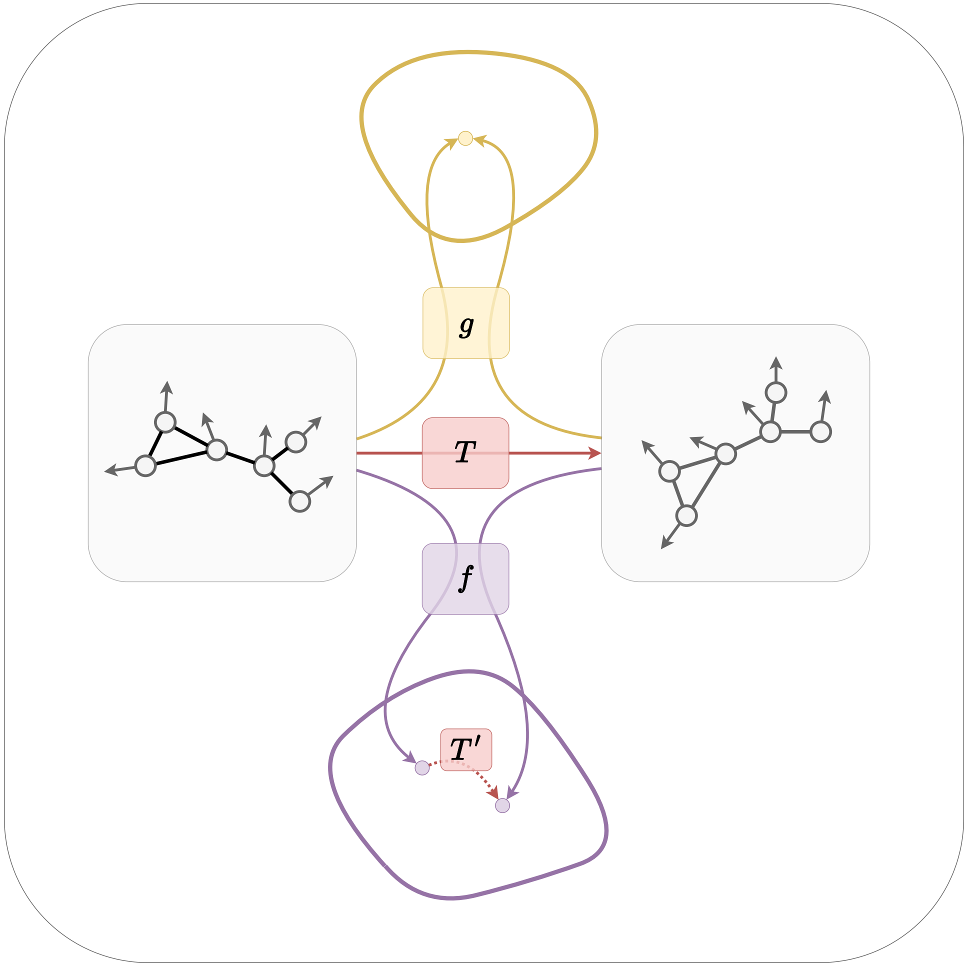

ROTATION INVARIANCE AND EQUIVARIANCE

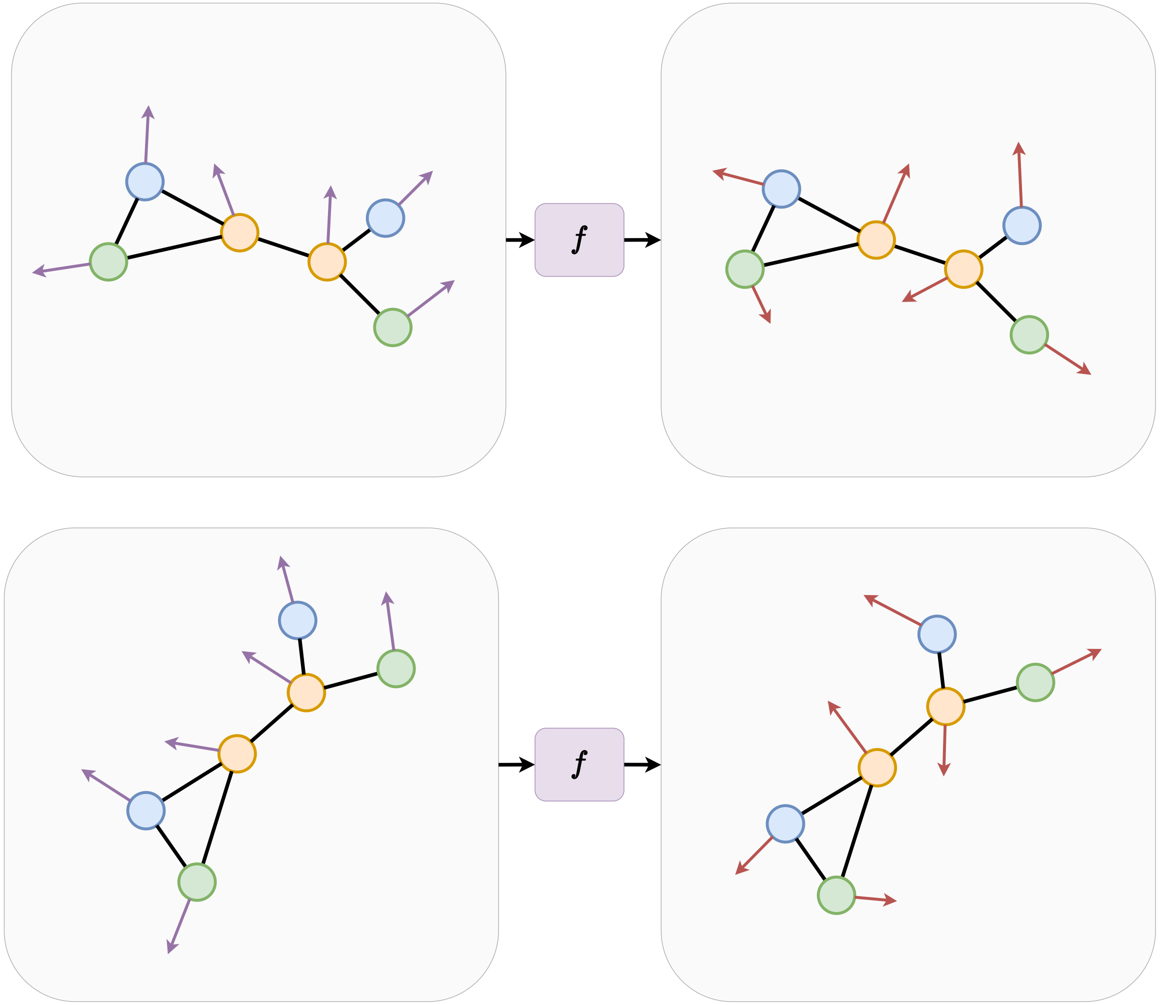

ROTATION INVARIANCE AND EQUIVARIANCE

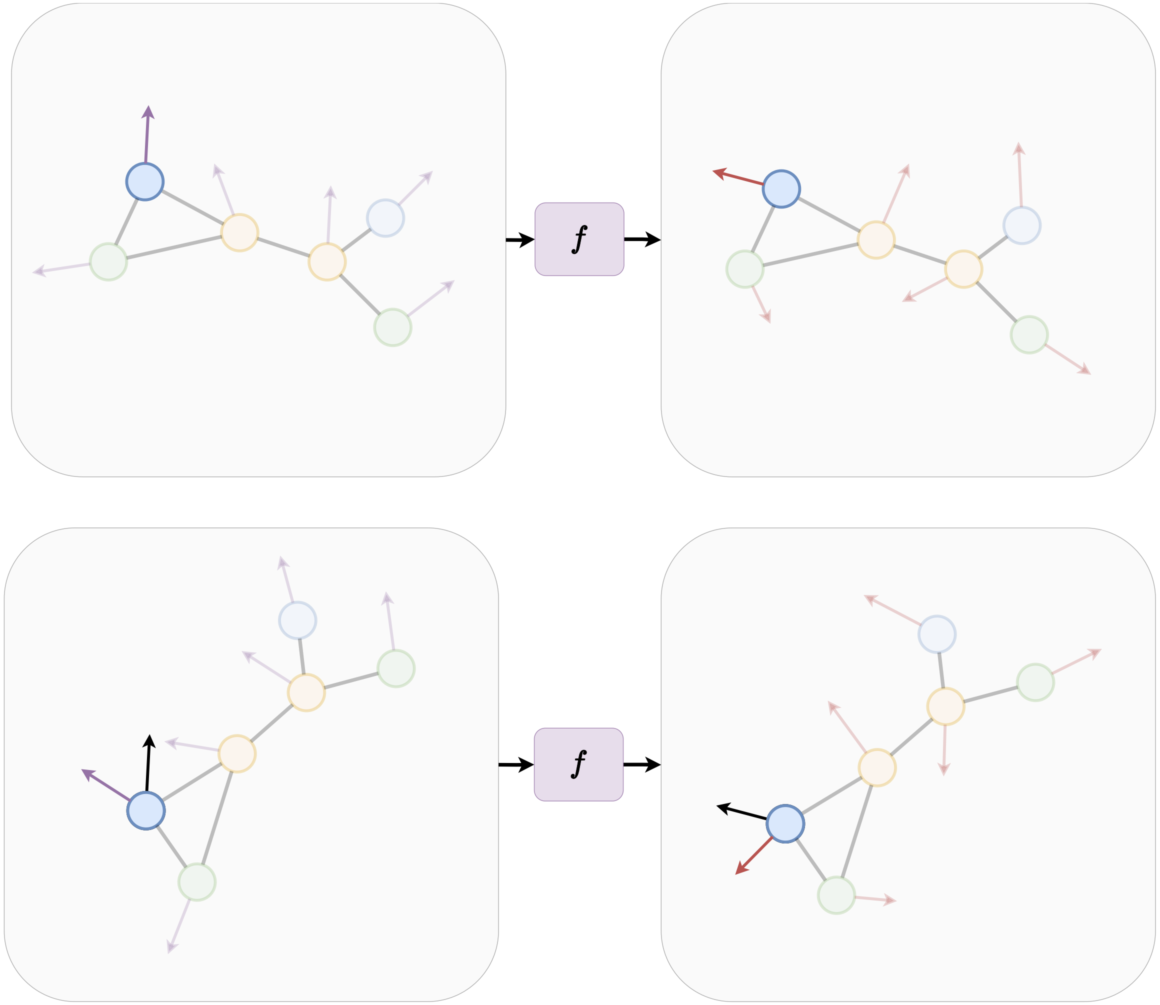

CLASSIFICATION EXAMPLE

CLASSIFICATION

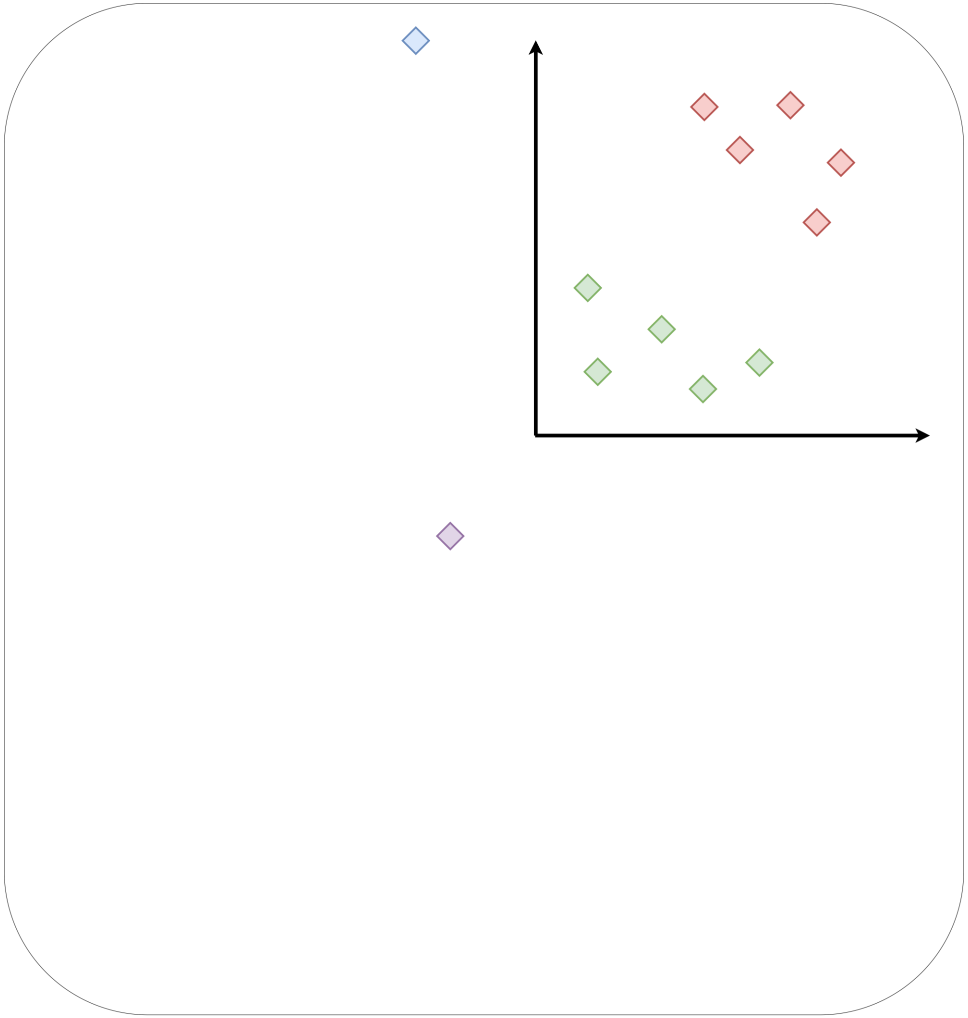

CLASSIFICATION

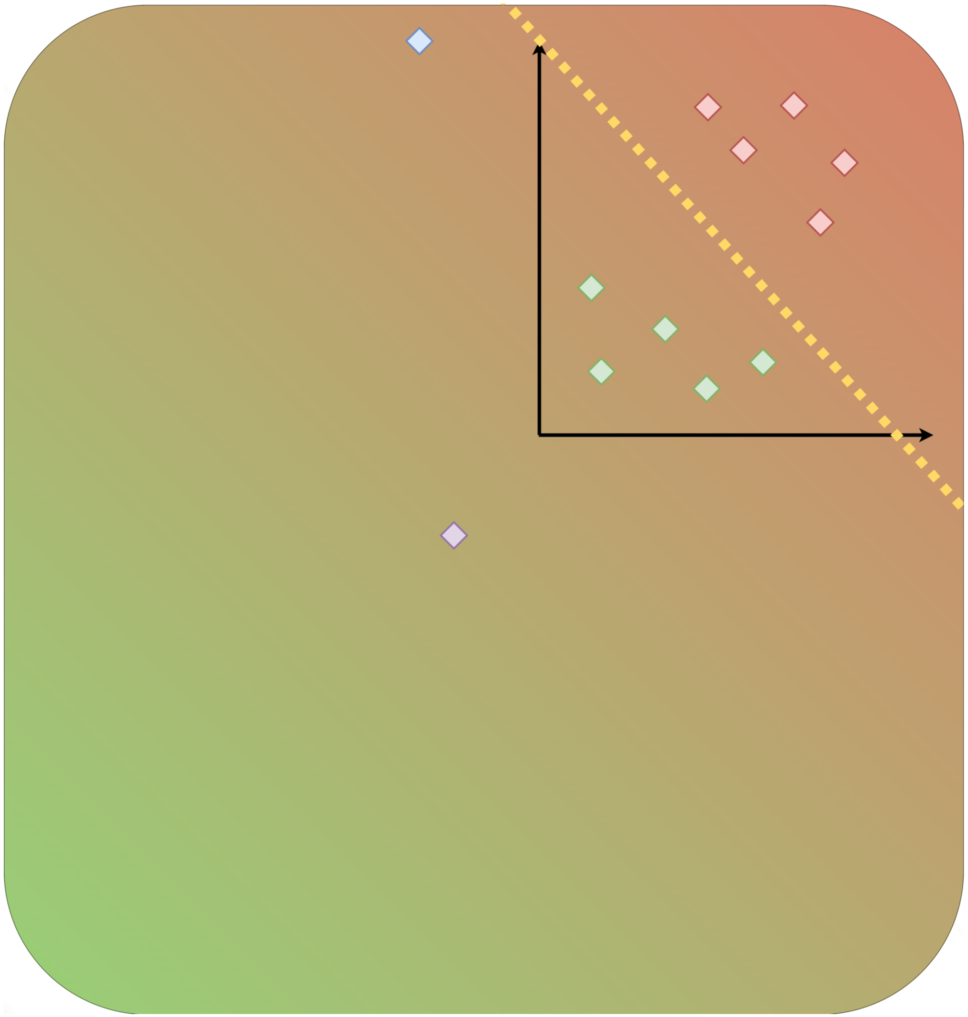

CLASSIFICATION

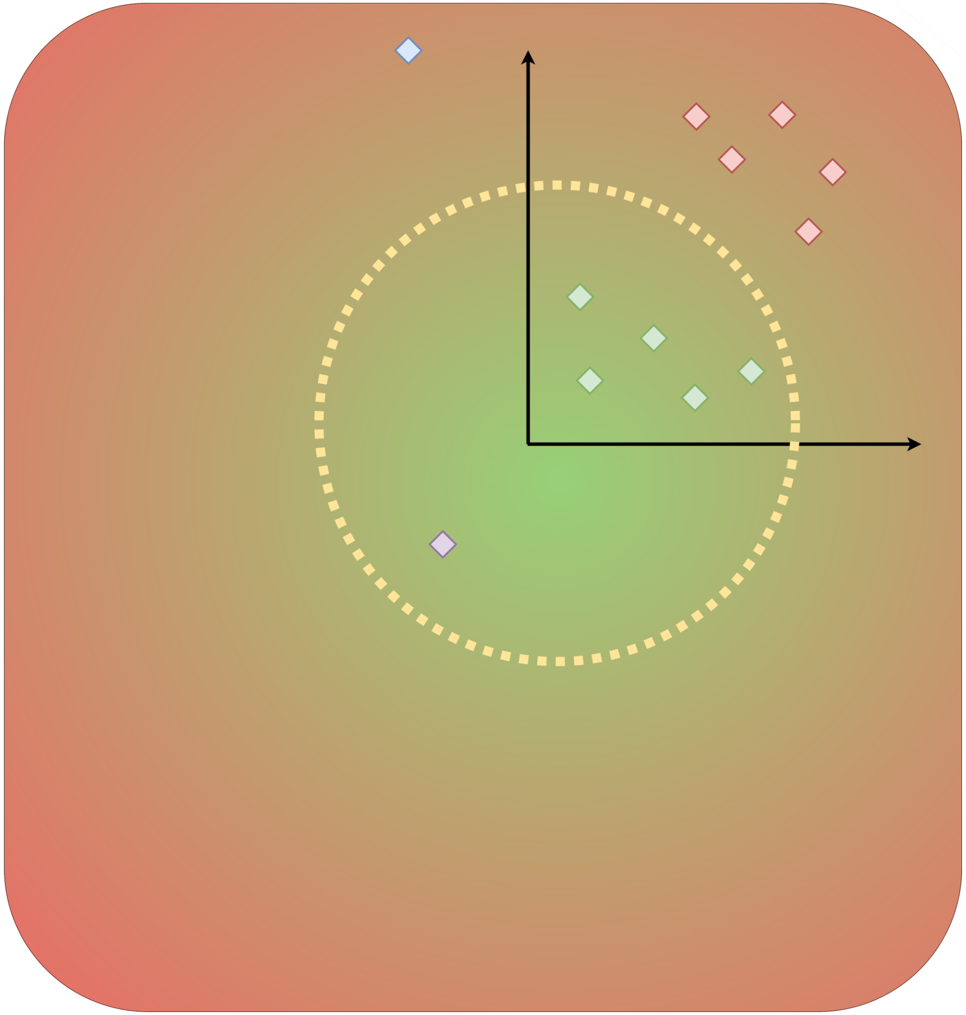

RADIUS GRAPHS

RADIUS GRAPHS

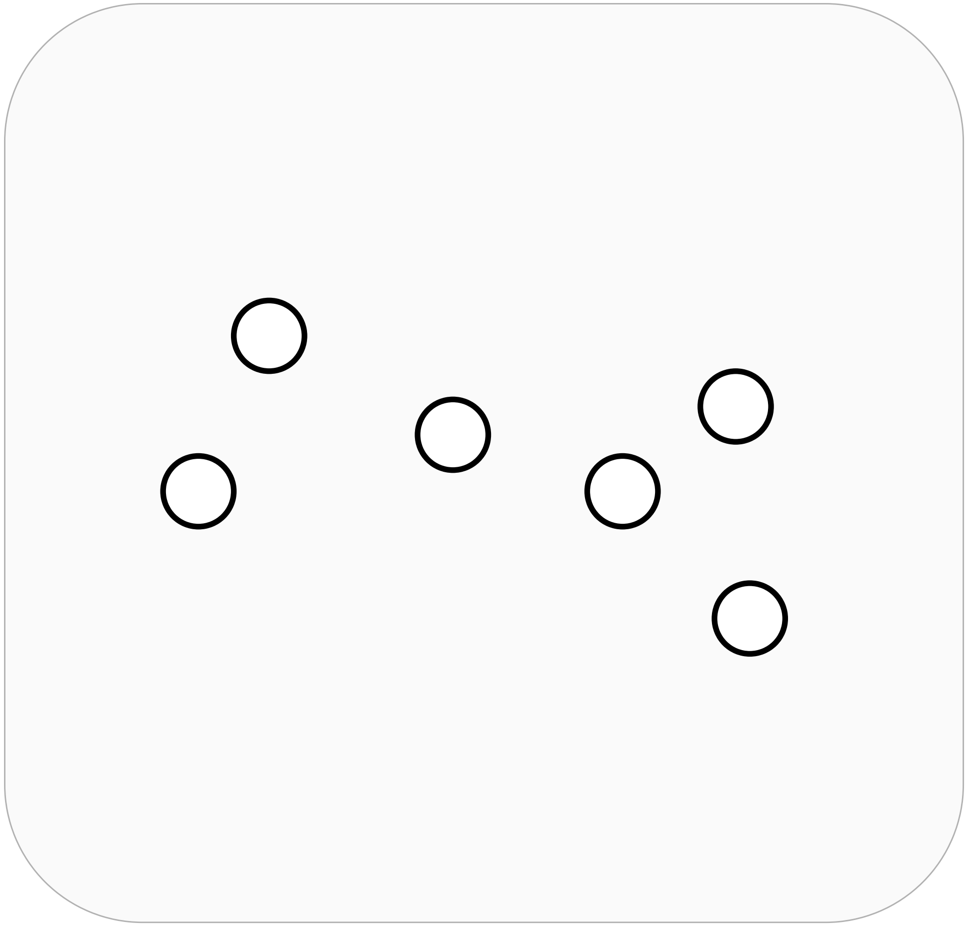

RADIUS GRAPHS

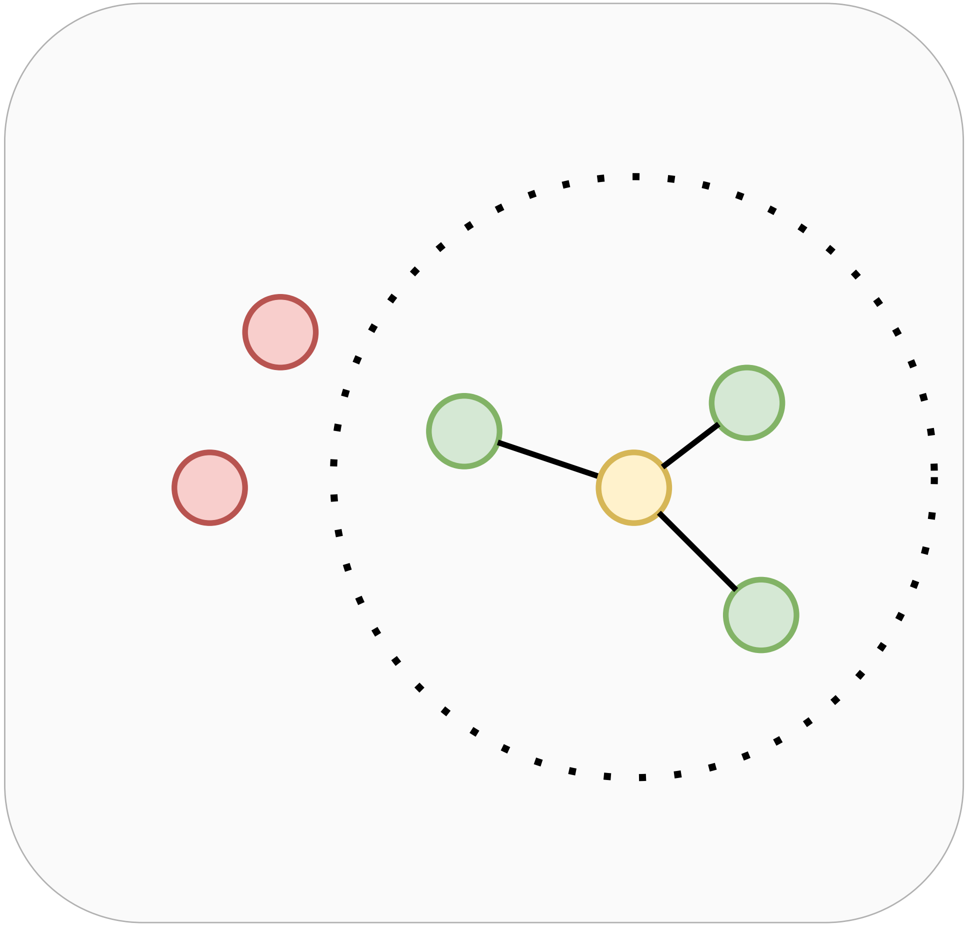

RADIUS GRAPHS

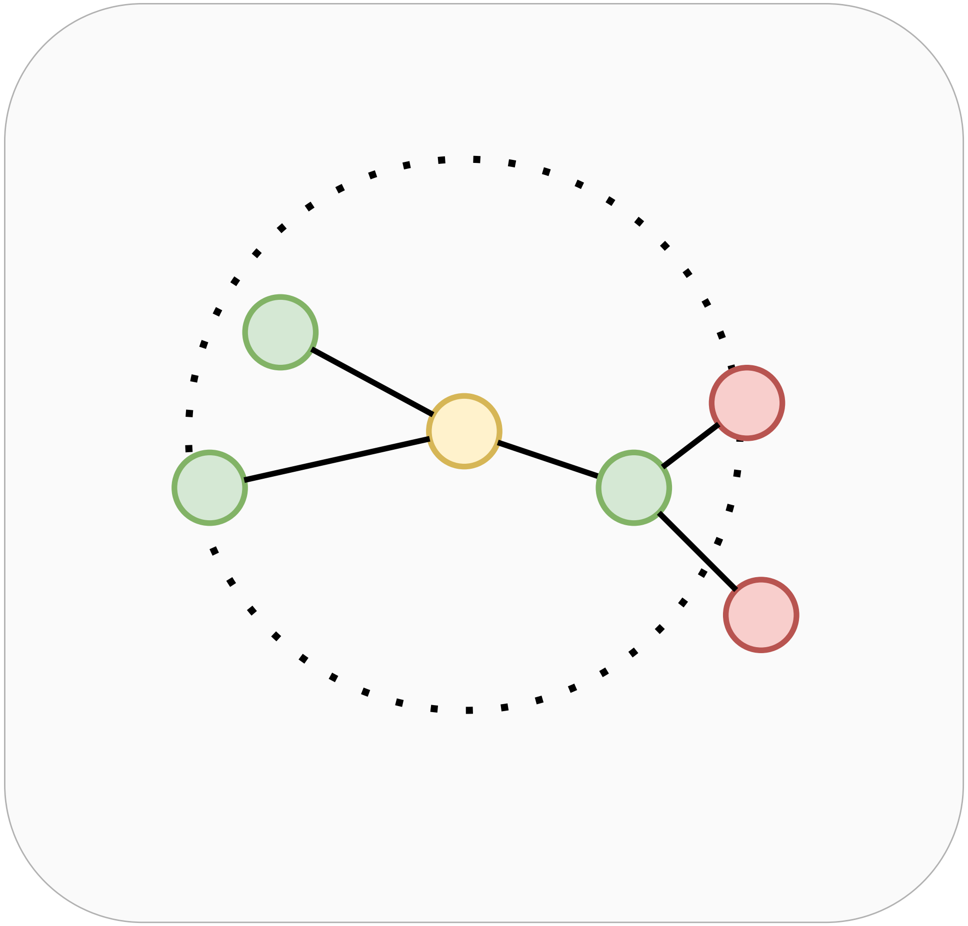

RADIUS GRAPHS

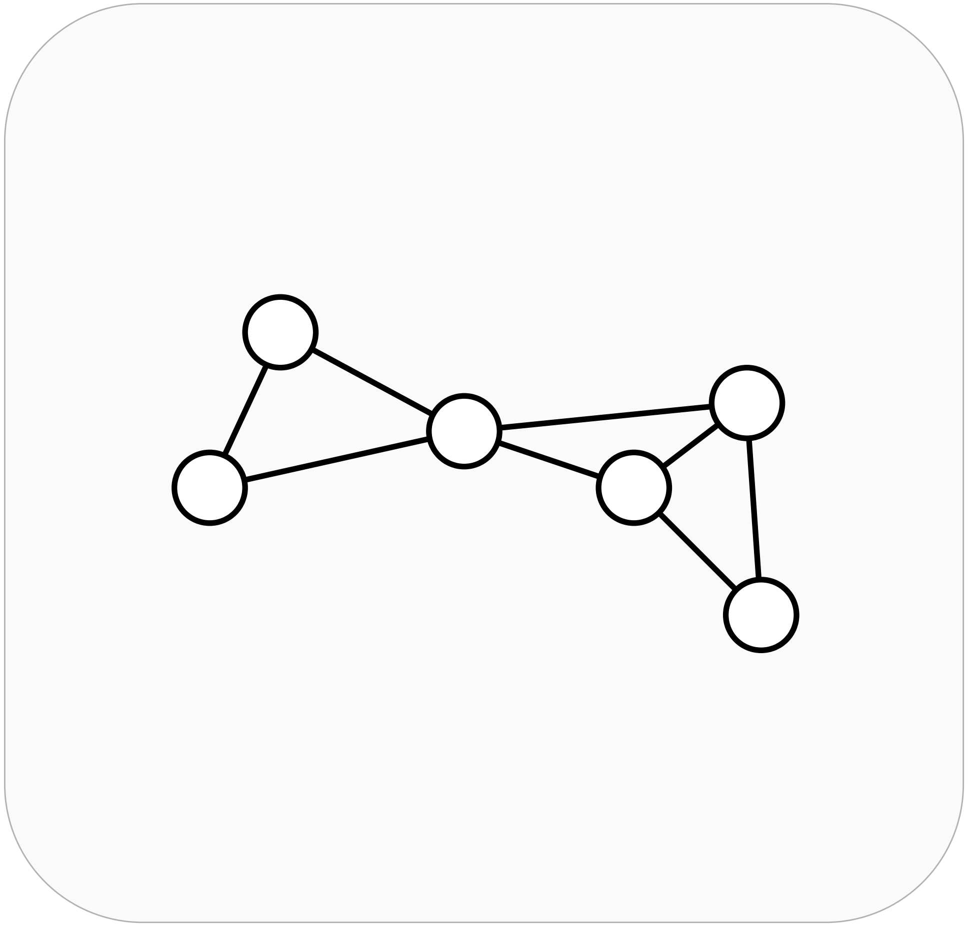

AUGMENTATION AND CANONICALIZATION

ACHIEVING EQUIVARIANCE

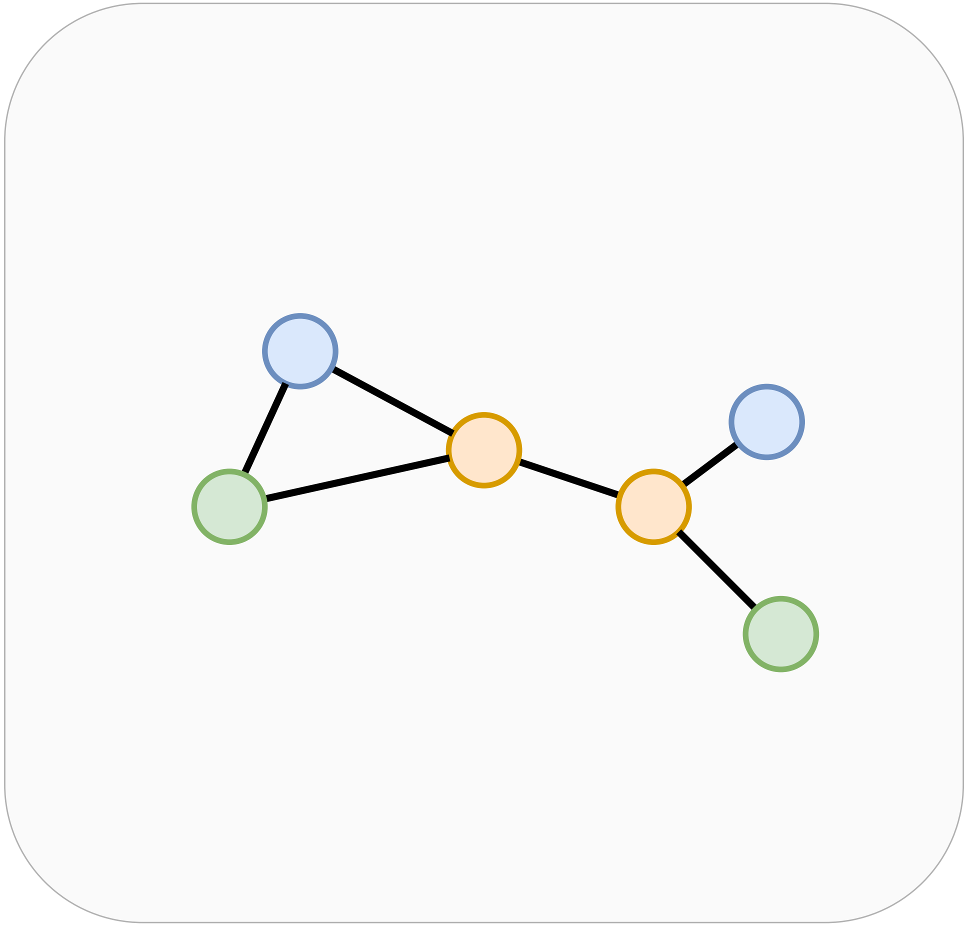

AUGMENTATION

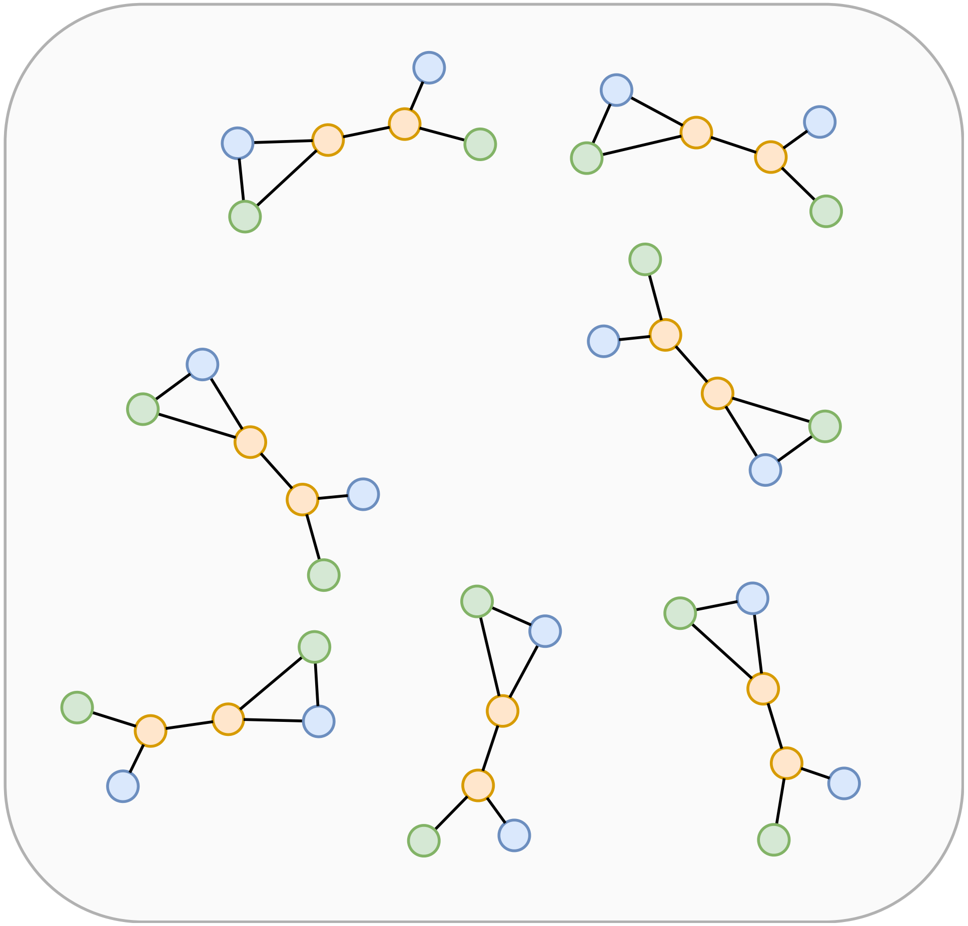

CANONICALIZATION

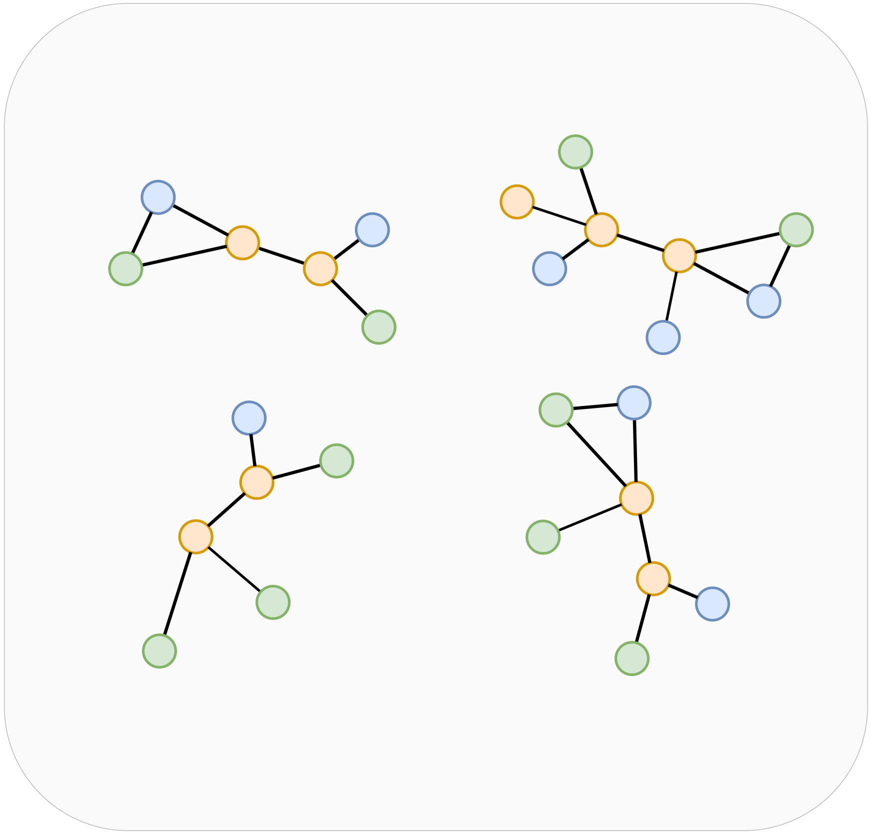

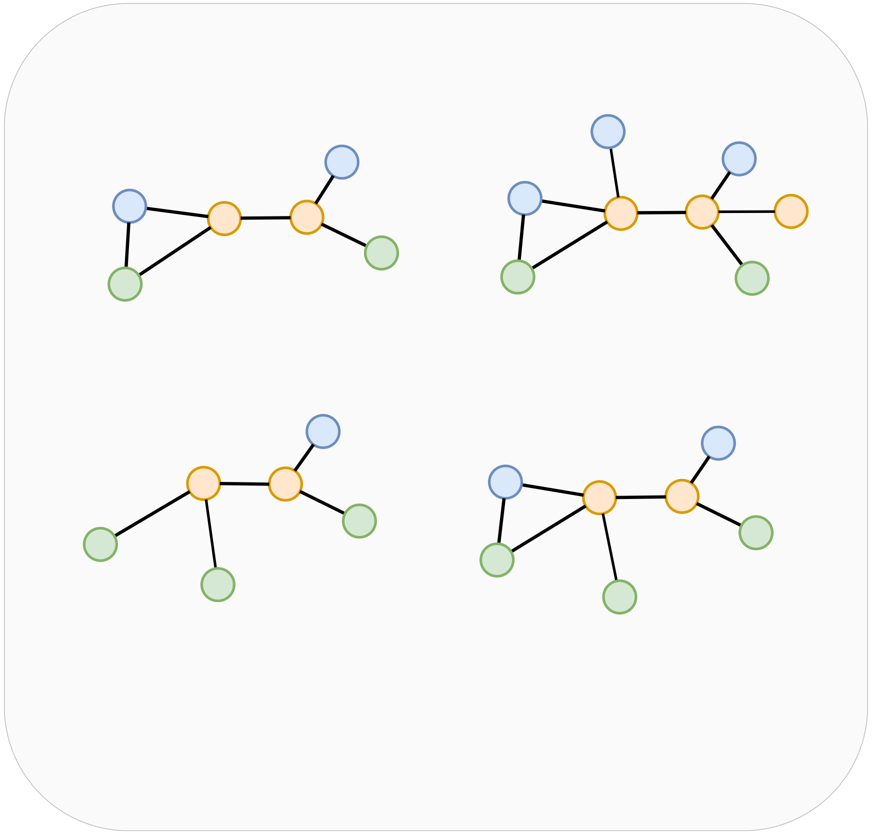

DISCUSSION: HOW MIGHT AUGMENTATION OR CANONICALIZATION FAIL?

EQUIVARIANT NEURAL NETWORKS

MESSAGE PASSING AGAIN

$$\mathbf{h}_u = \phi\left(\mathbf{x}_u, \bigoplus_{v \in \mathcal{N}(u)} c_{uv}\psi\left(\mathbf{x}_v\right)\right)$$

$$\mathbf{h}_u = \phi\left(\mathbf{x}_u, \bigoplus_{v \in \mathcal{N}(u)} a\left(\mathbf{x}_u, \mathbf{x}_v\right)\psi\left(\mathbf{x}_v\right)\right)$$

$$\mathbf{h}_u = \phi\left(\mathbf{x}_u, \bigoplus_{v \in \mathcal{N}(u)} \psi\left(\mathbf{x}_u, \mathbf{x}_v\right)\right)$$

Based on Chap 5. of: Bronstein, Michael M., et al. "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges." arXiv preprint arXiv:2104.13478 (2021).

EQUIVARIANT MESSAGE PASSING

Features:

$$\mathbf{f}_u \in \mathbb{R}^d$$

Spatial coordinates:

$$\mathbf{x}_u \in \mathbb{R}^3$$

Equivariance to rotations, reflections, and translations:

$$\mathbf{Rx} + \mathbf{b}$$

Based on Chap 5. of: Bronstein, Michael M., et al. "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges." arXiv preprint arXiv:2104.13478 (2021).

EXAMPLE VECTOR DATA

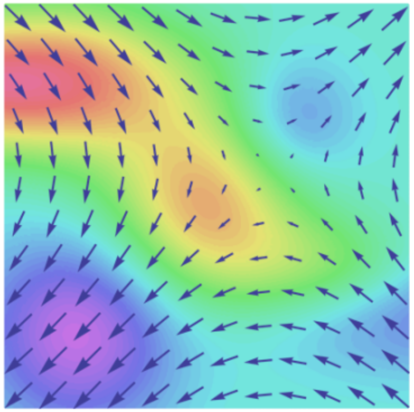

From Chap 5. of: Bronstein, Michael M., et al. "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges." arXiv preprint arXiv:2104.13478 (2021).

EXAMPLE VECTOR DATA

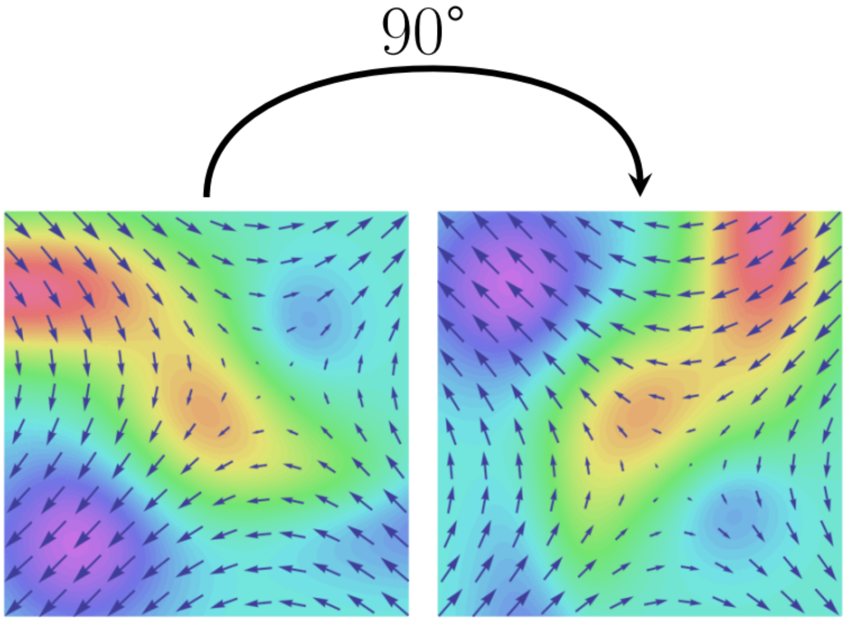

From Chap 5. of: Bronstein, Michael M., et al. "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges." arXiv preprint arXiv:2104.13478 (2021).

DISCUSSION: EDGES BETWEEN SPATIAL POINTS

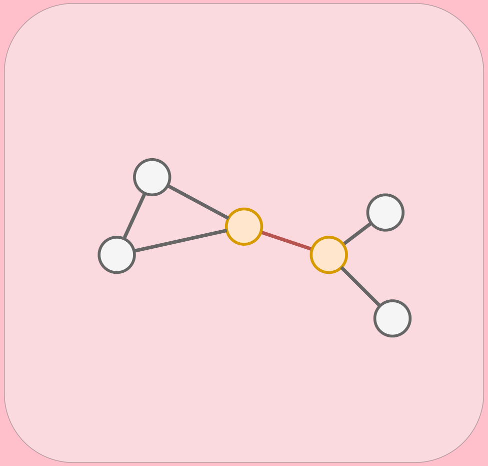

DISCUSSION: EDGES BETWEEN SPATIAL POINTS

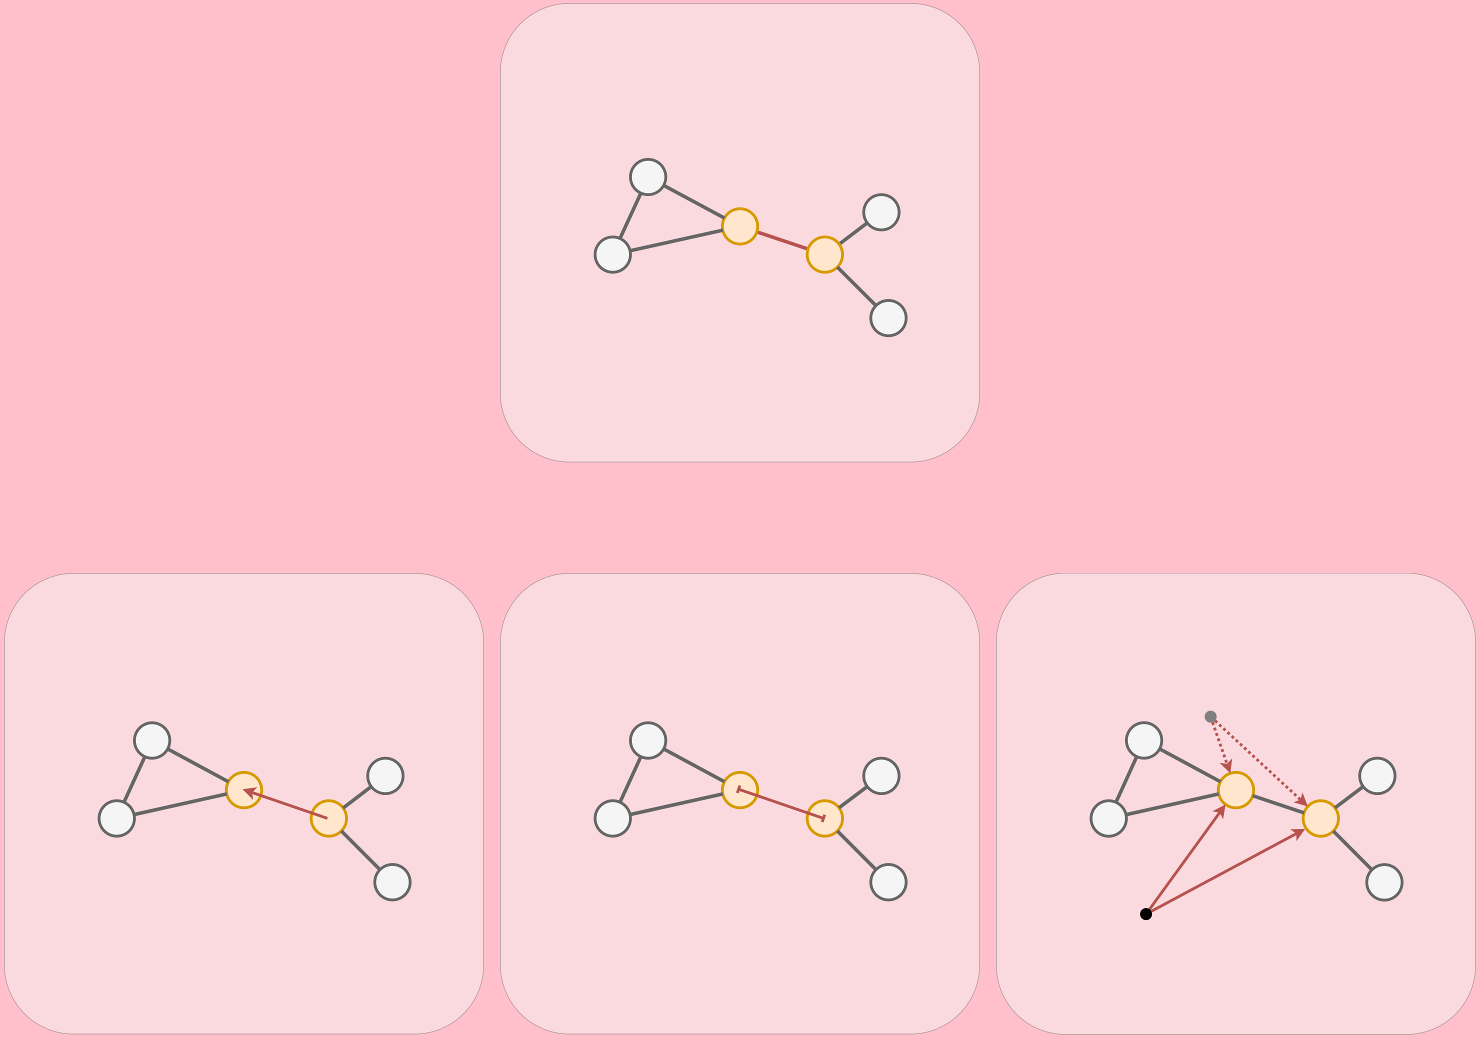

EQUIVARIANT MESSAGE PASSING

$$\mathbf{f}_u' = \phi\left(\mathbf{f}_u, \bigoplus_{v \in \mathcal{N}(u)} \psi_f\left(\mathbf{f}_u, \mathbf{f}_v, ||\mathbf{x}_u - \mathbf{x}_v||^2\right)\right)$$

$$\mathbf{x}_u' = \mathbf{x}_u + \sum_{v \neq u} (\mathbf{x}_u - \mathbf{x}_v)\, \psi_c\left(\mathbf{f}_u, \mathbf{f}_v, ||\mathbf{x}_u - \mathbf{x}_v||^2\right)$$

Based on Chap 5. of: Bronstein, Michael M., et al. "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges." arXiv preprint arXiv:2104.13478 (2021).

REPRESENTING VECTORS

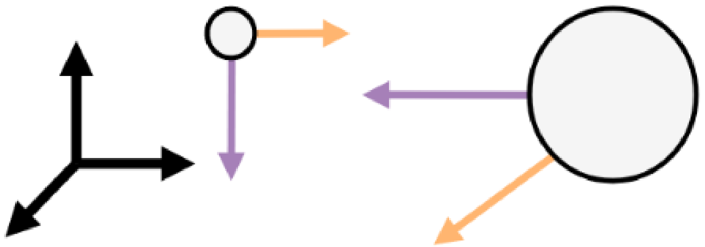

```
geometry = [[x0, y0, z0],[x1, y1, z1]]
features = [
    [m0, v0y, v0z, v0x, a0y, a0z, a0x]
    [m1, v1y, v1z, v1x, a1y, a1z, a1x]
]
...
```

(from "Neural networks with Euclidean Symmetry for the Physical Sciences", lecture by Tess Smidt: https://tinyurl.com/e3nn-physics-meets-ml)

REPRESENTING VECTORS

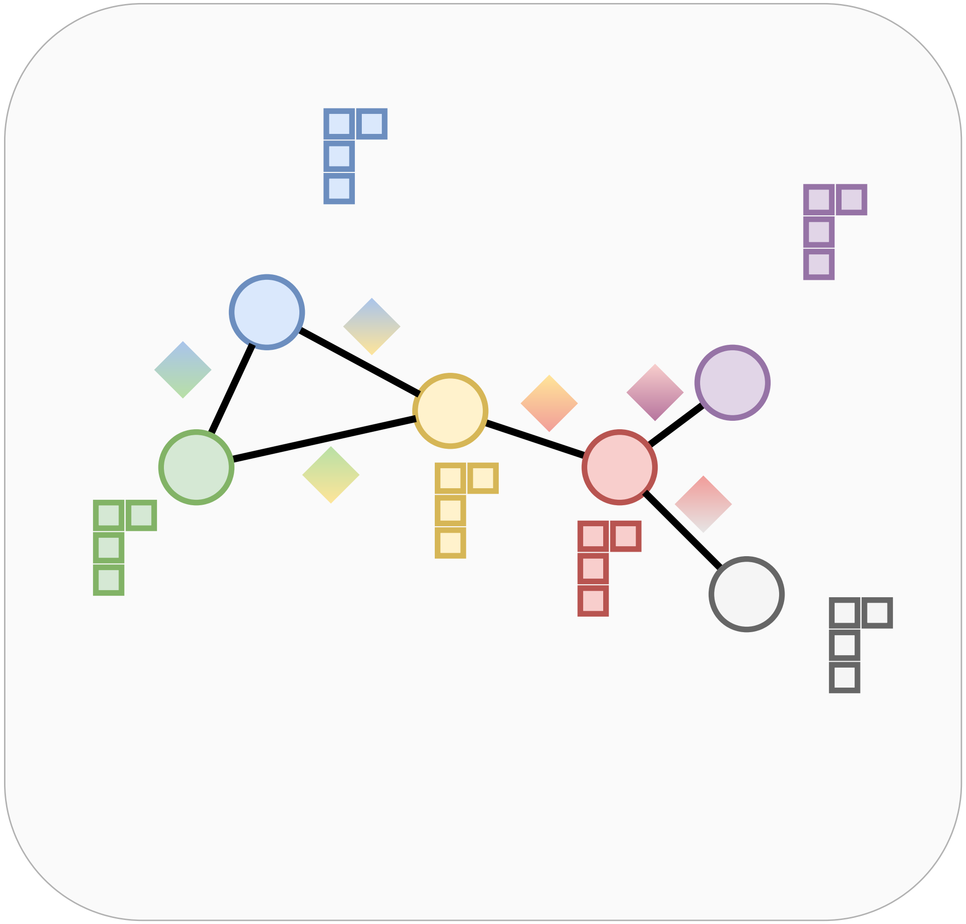

GEOMETRIC VECTOR PERCEPTRON

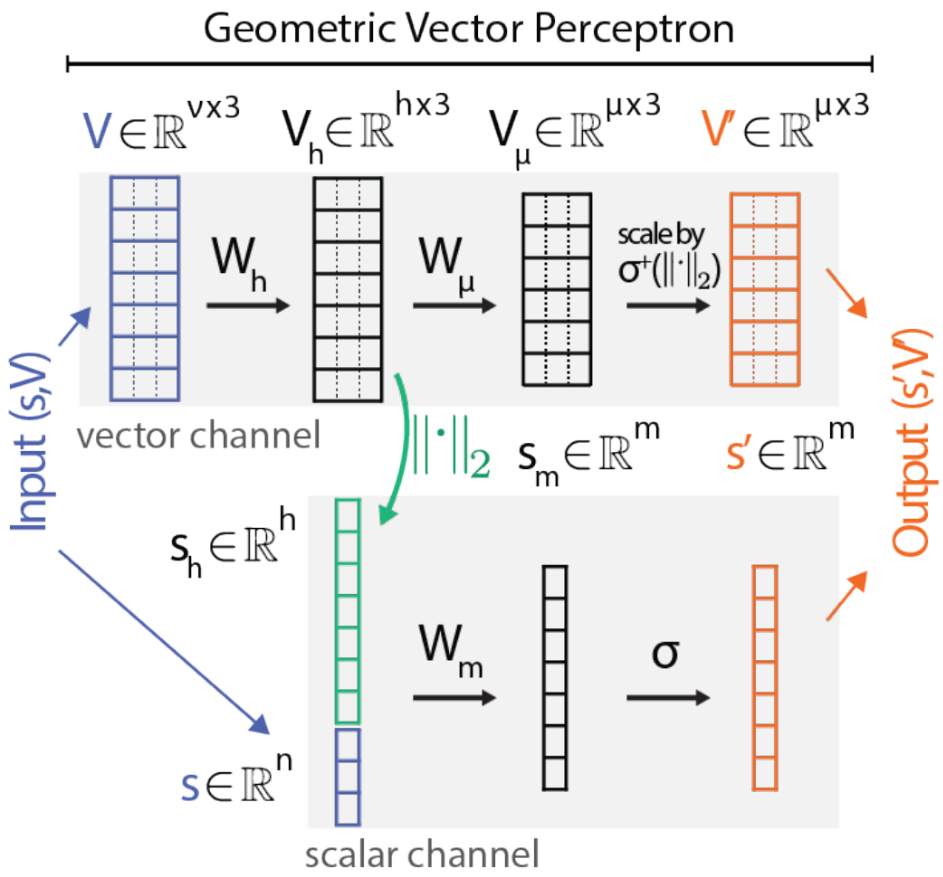

(Jing, Bowen, et al. "Learning from protein structure with geometric vector perceptrons." International Conference on Learning Representations (2021).))

GEOMETRIC VECTOR PERCEPTRON

Algorithm 1 Geometric vector perceptron

Input: Scalar and vector features $(\mathbf{s}, \mathbf{V}) \in \mathbb{R}^n \times \mathbb{R}^{\nu \times 3}$.
Output: Scalar and vector features $(\mathbf{s}', \mathbf{V}') \in \mathbb{R}^m \times \mathbb{R}^{\mu \times 3}$.
$h \leftarrow \max(\nu, \mu)$
GVP:
$\quad \mathrm{V}_h \leftarrow \mathbf{W}_h \mathbf{V} \quad \in \mathbb{R}^{h \times 3}$
$\quad \mathrm{V}_\mu \leftarrow \mathbf{W}_\mu \mathrm{V}_h \quad \in \mathbb{R}^{\mu \times 3}$
$\quad \mathrm{s}_h \leftarrow \|\mathrm{V}_h\|_2$ (row-wise) $\quad \in \mathbb{R}^h$
$\quad \mathrm{v}_\mu \leftarrow \|\mathrm{V}_\mu\|_2$ (row-wise) $\quad \in \mathbb{R}^\mu$
$\quad \mathrm{s}_{h+n} \leftarrow \mathrm{concat}(\mathrm{s}_h, \mathbf{s}) \quad \in \mathbb{R}^{h+n}$
$\quad \mathrm{s}_m \leftarrow \mathbf{W}_m \mathrm{s}_{h+n} + \mathbf{b} \quad \in \mathbb{R}^m$
$\quad \mathbf{s}' \leftarrow \sigma(\mathrm{s}_m) \quad \in \mathbb{R}^m$
$\quad \mathbf{V}' \leftarrow \sigma^+(\mathrm{v}_\mu) \odot \mathrm{V}_\mu$ (row-wise multiplication) $\quad \in \mathbb{R}^{\mu \times 3}$
return $(\mathbf{s}', \mathbf{V}')$

(Jing, Bowen, et al. "Learning from protein structure with geometric vector perceptrons." International Conference on Learning Representations (2021).))

EXAMPLE VECTOR DATA

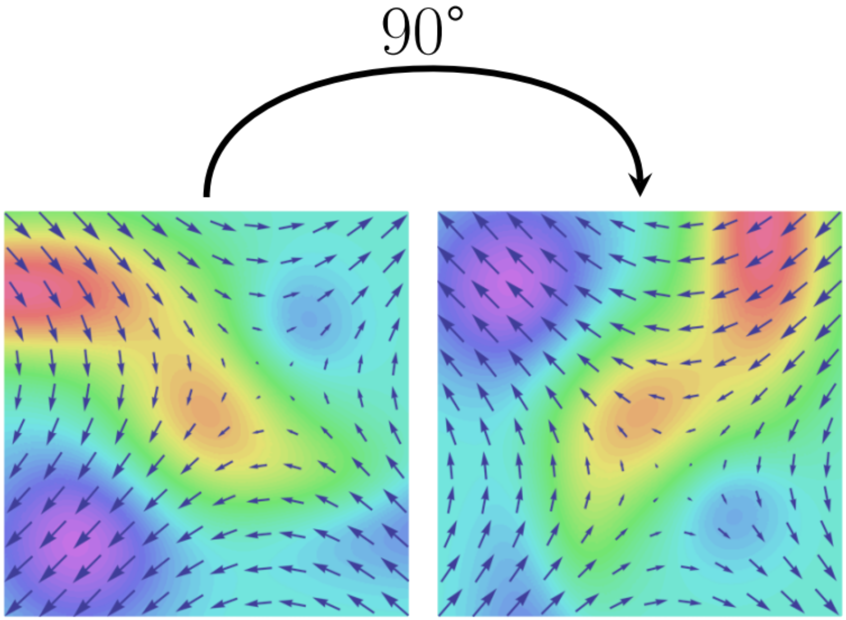

From Chap 5. of: Bronstein, Michael M., et al. "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges." arXiv preprint arXiv:2104.13478 (2021).

REPRESENTING VECTORS

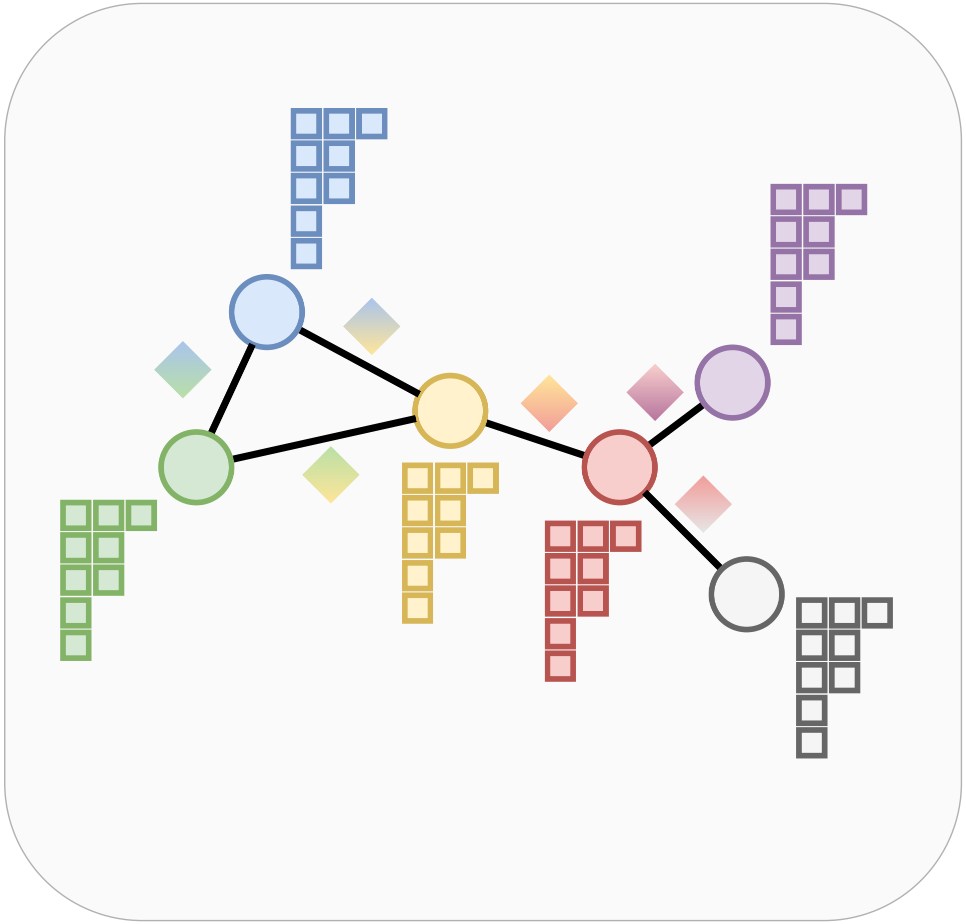

OBSERVED AND EXPECTED SHIFTS

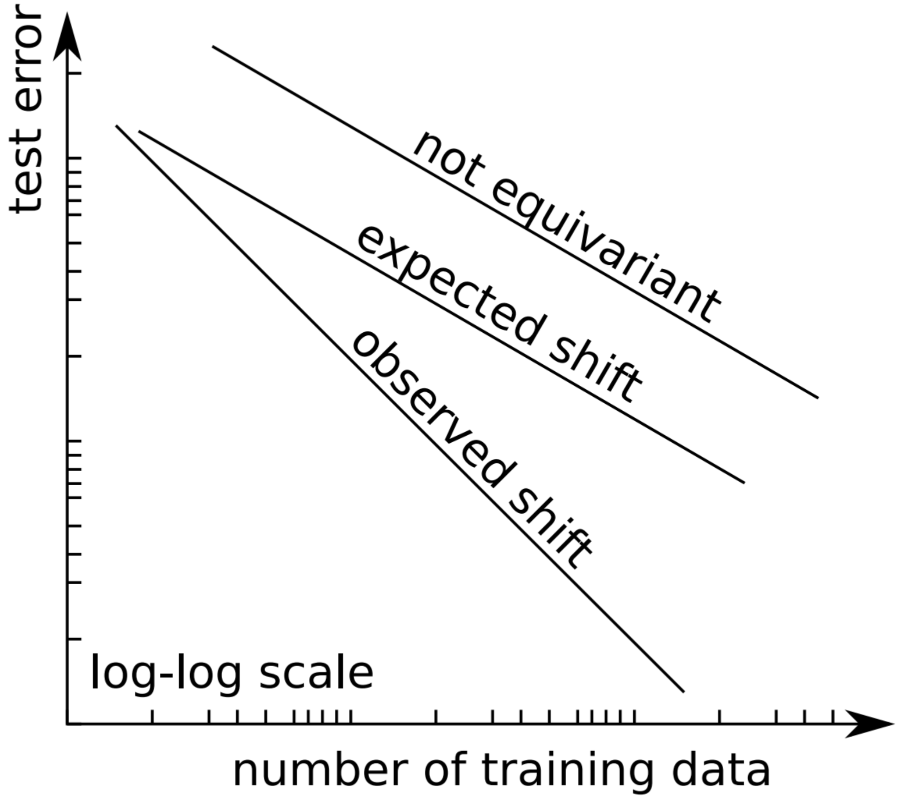

(Geiger, Mario, and Tess Smidt. "e3nn: Euclidean neural networks." arXiv preprint arXiv:2207.09453 (2022))

EXAMPLE: NEURAL NETWORK INTERATOMIC POTENTIALS

$$E_{pot} = \sum_{i \in N_{atoms}} E_{i,atomic}$$

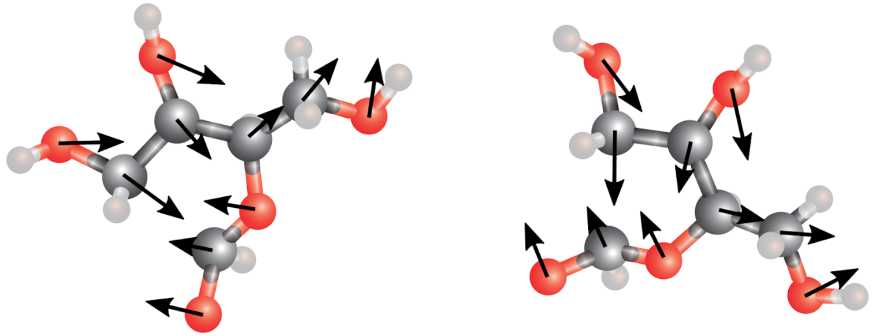

$$\vec{F}_i = -\nabla_i E_{pot}$$

(left: Batzner, Simon, et al. "E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials." Nature communications 13.1 (2022): 2453.)
(right: Neural networks with Euclidean Symmetry for the Physical Sciences, lecture by Tess Smidt)

RECAP LECTURE 2

RADIUS GRAPHS

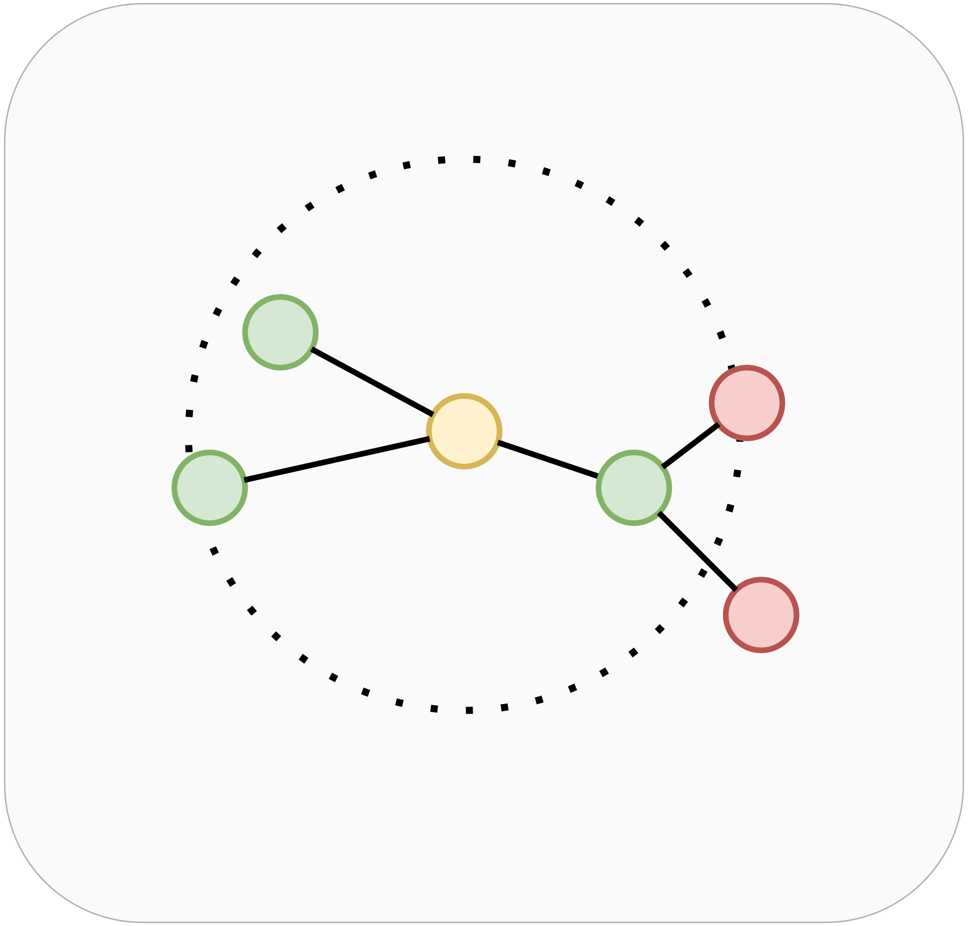

ROTATION INVARIANCE AND EQUIVARIANCE

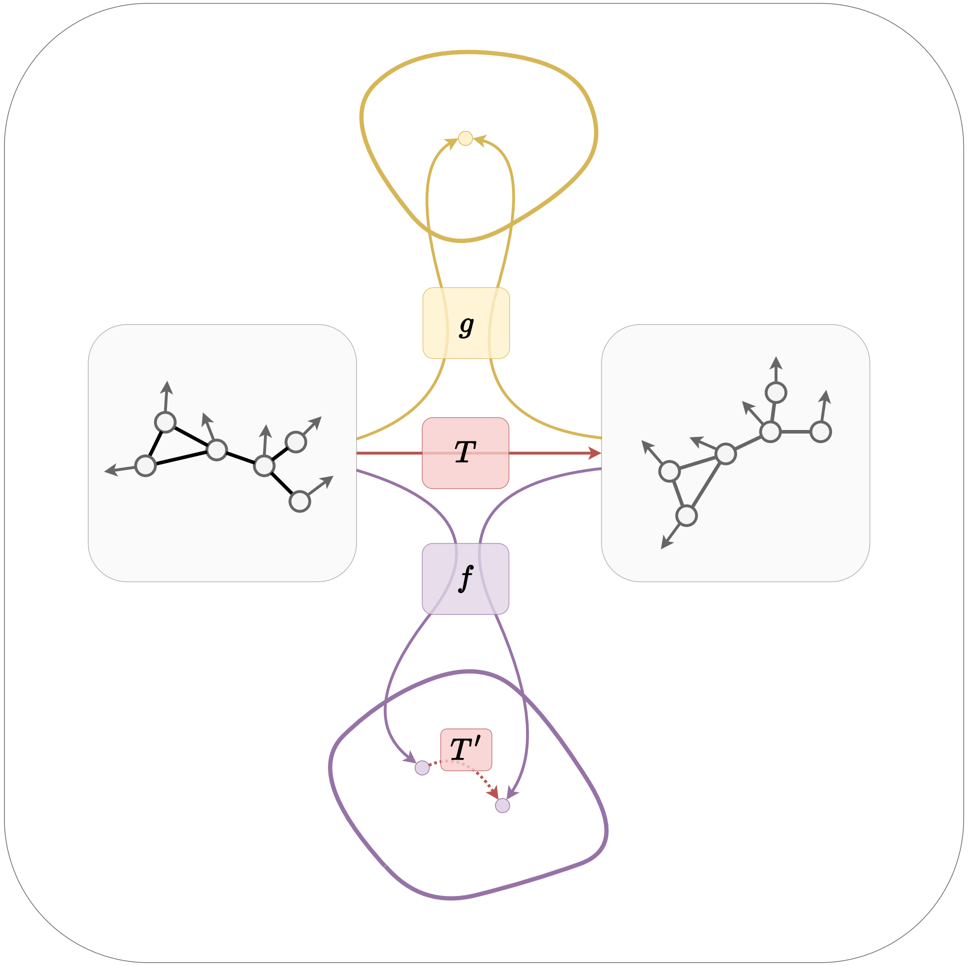

AUGMENTATION

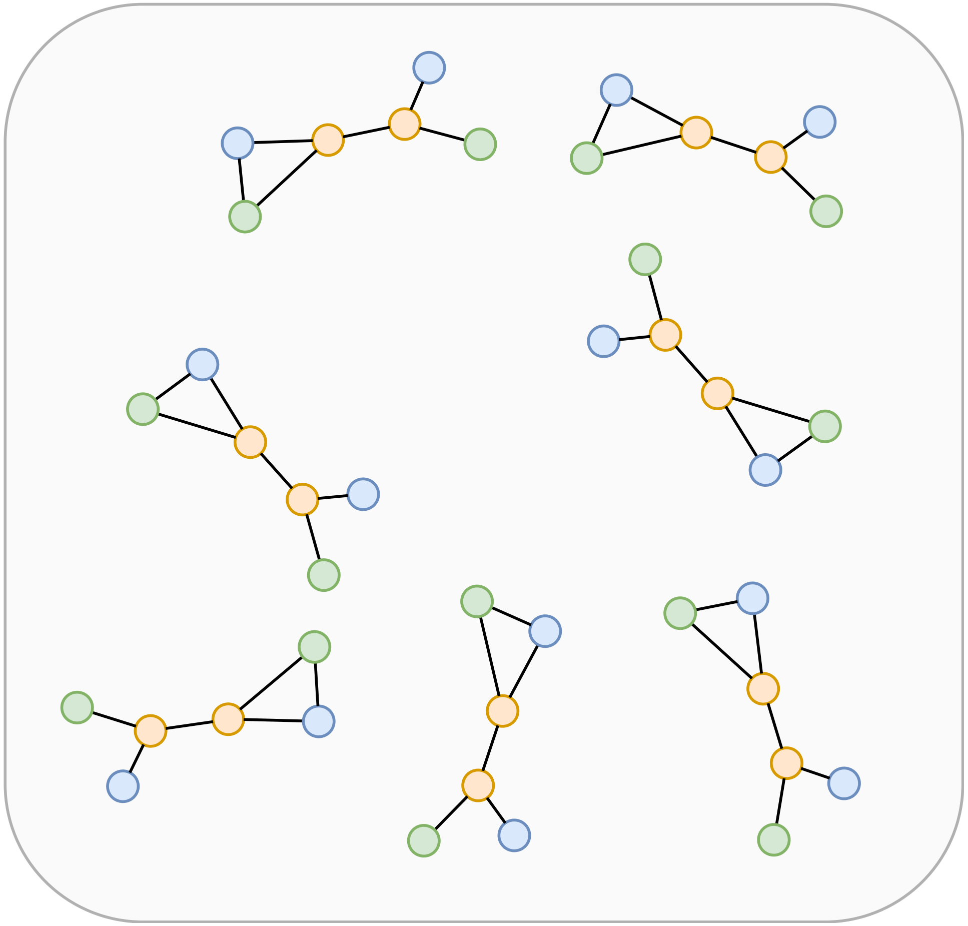

CANONICALIZATION

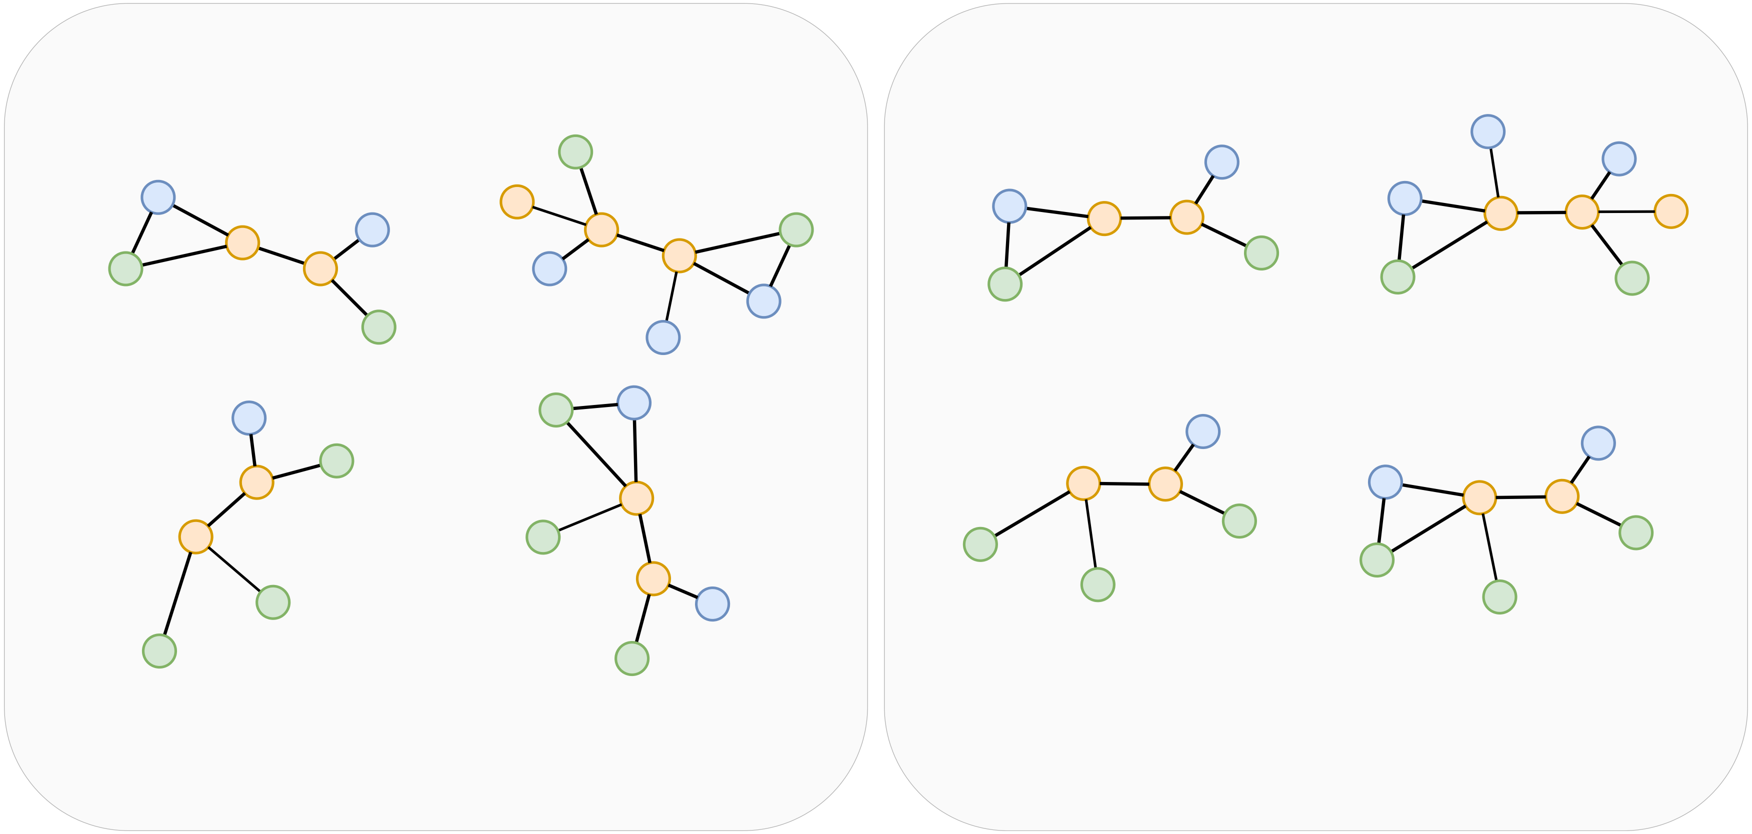

EDGES BETWEEN SPATIAL POINTS

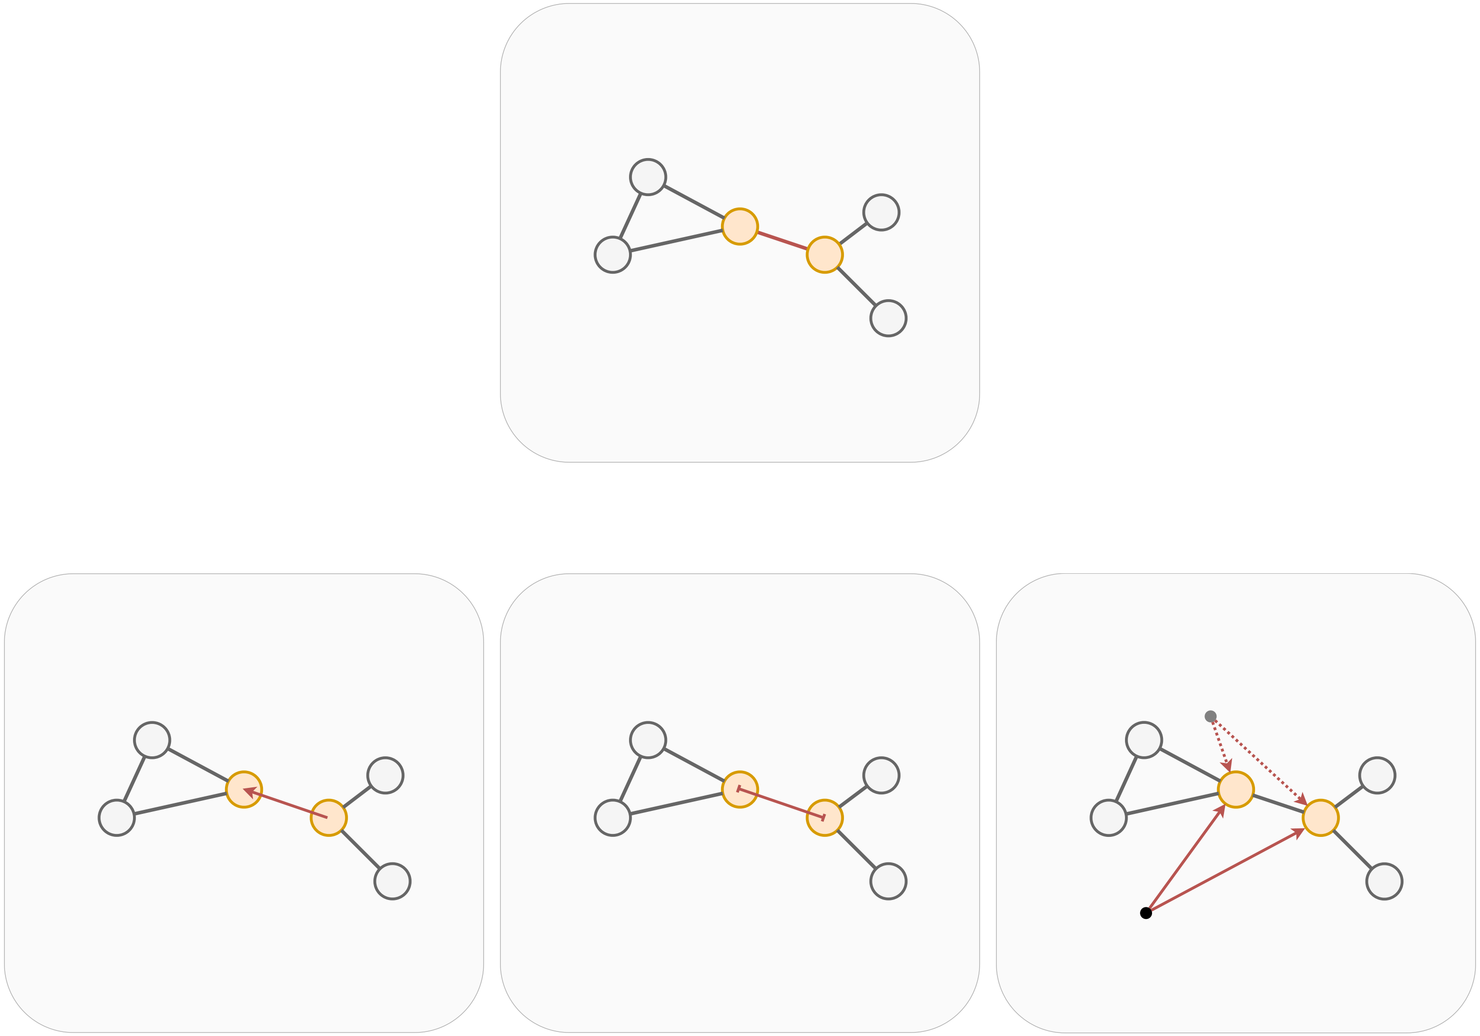

EQUIVARIANT MESSAGE PASSING

$$\mathbf{f}'_u = \phi\left(\mathbf{f}_u, \bigoplus_{v\in\mathcal{N}(u)} \psi_f\left(\mathbf{f}_u, \mathbf{f}_v, ||\mathbf{x}_u - \mathbf{x}_v||^2\right)\right)$$

$$\mathbf{x}'_u = \mathbf{x}_u + \sum_{v\neq u} (\mathbf{x}_u - \mathbf{x}_v)\, \psi_c\left(\mathbf{f}_u, \mathbf{f}_v, ||\mathbf{x}_u - \mathbf{x}_v||^2\right)$$

Based on Chap 5. of: Bronstein, Michael M., et al. "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges." arXiv preprint arXiv:2104.13478 (2021).

OBSERVED AND EXPECTED SHIFTS

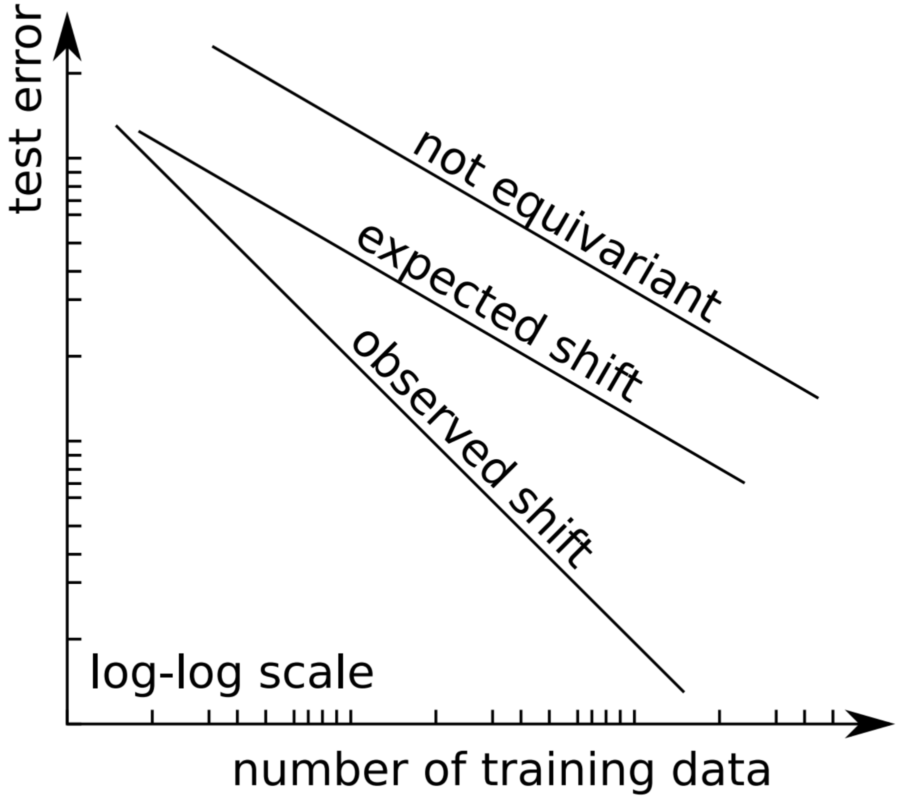

(Geiger, Mario, and Tess Smidt. "e3nn: Euclidean neural networks." arXiv preprint arXiv:2207.09453 (2022))

DISCUSSION: LIVE DEMO

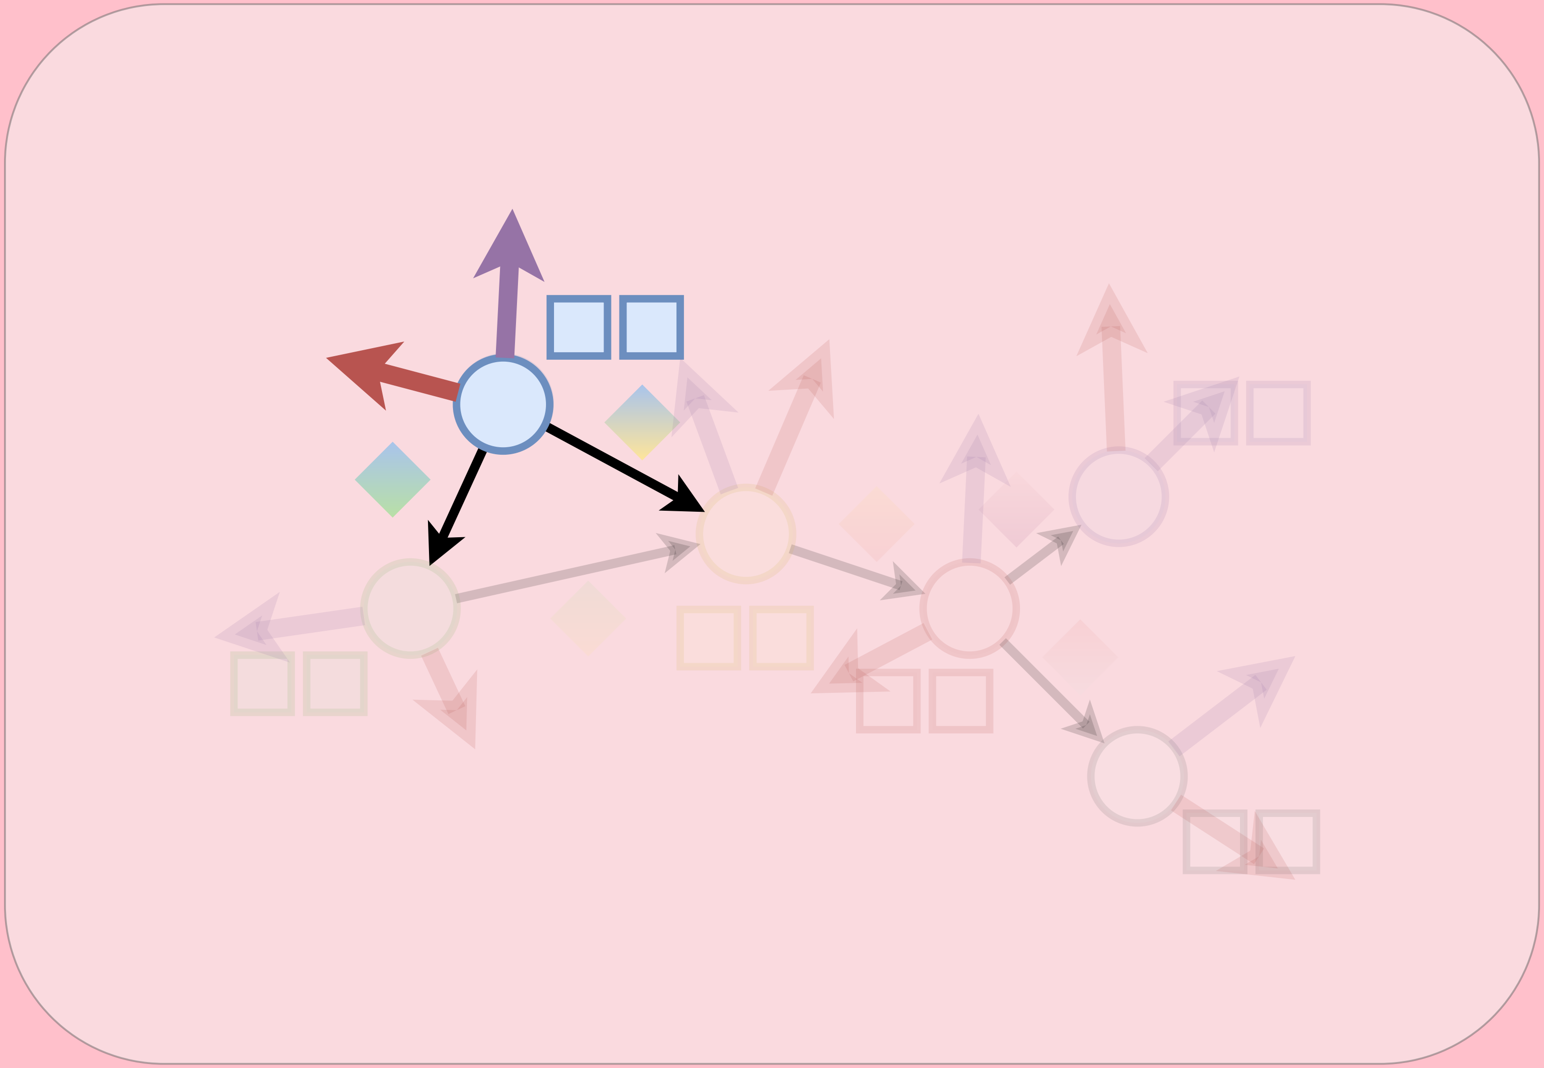

DISCUSSION: LIVE DEMO

gnn-live-demo.web.app

Session ID: "lecture2" (not Lecture2, Lecture 2, ..., etc.)

FINAL RECAP

REPRESENTING A GRAPH

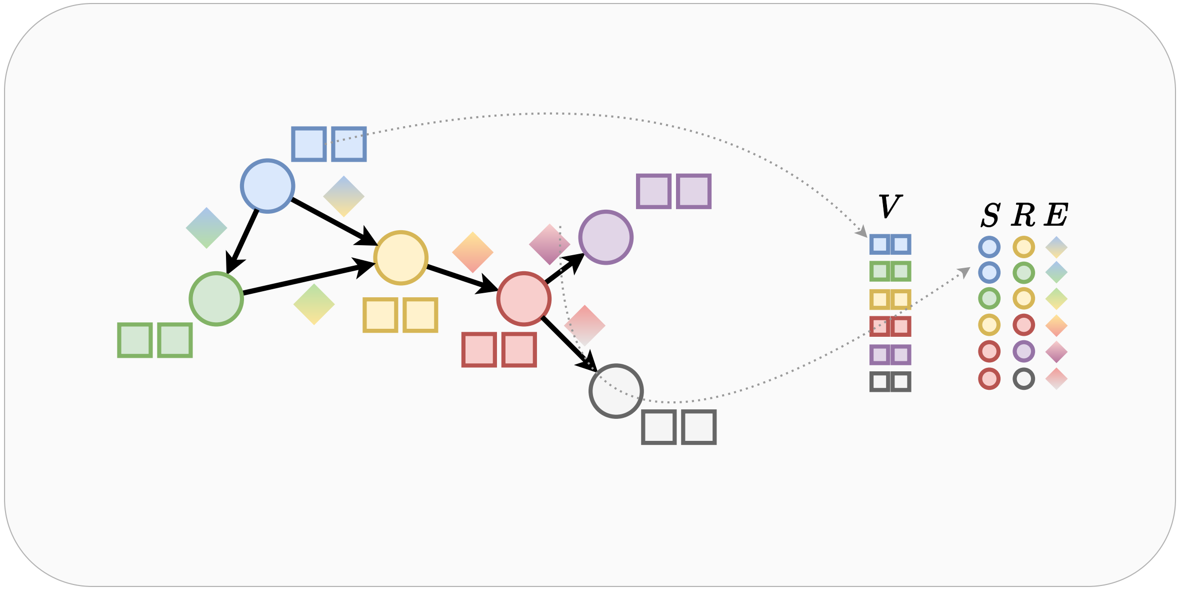

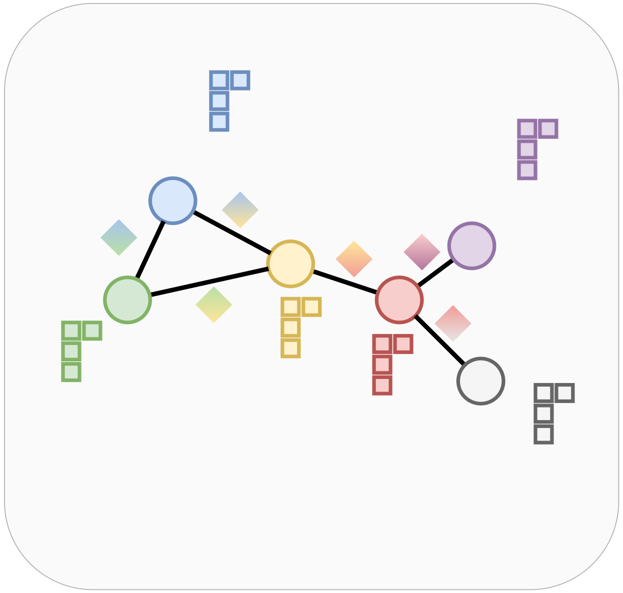

INVARIANCE AND EQUIVARIANCE

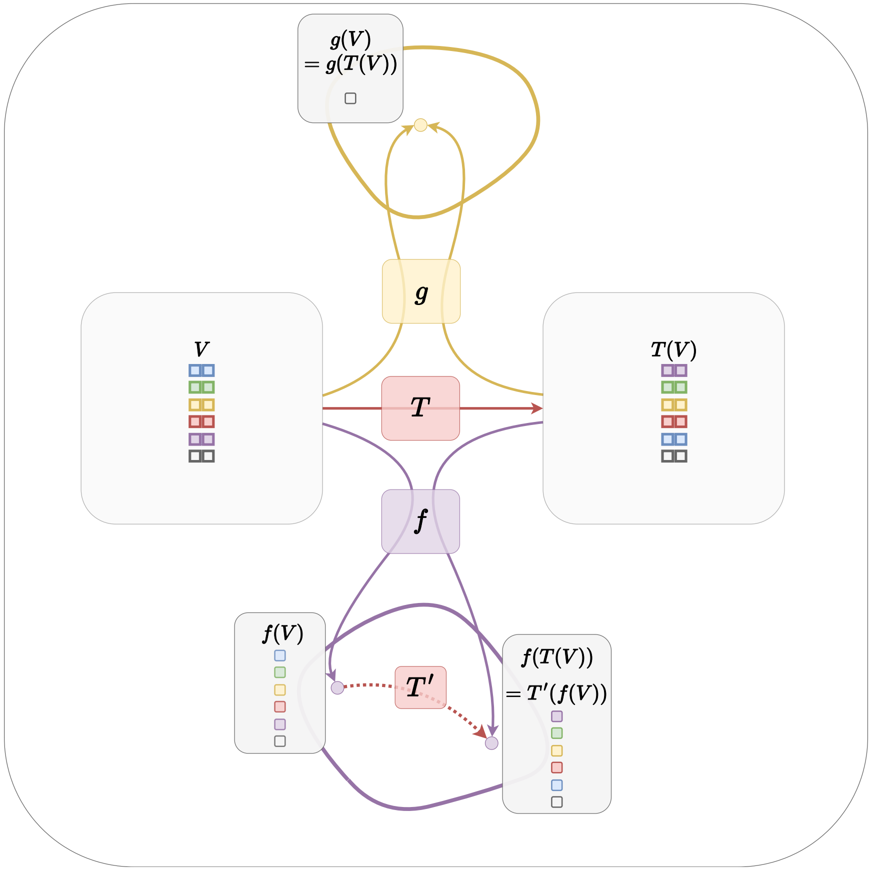

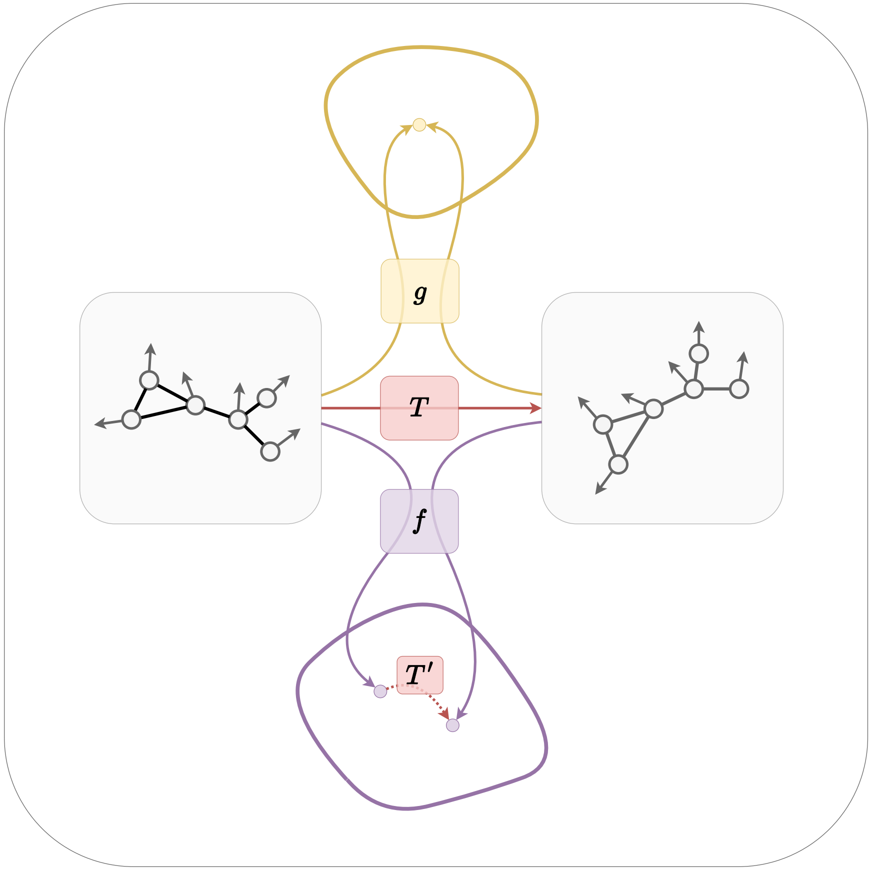

MESSAGE PASSING ON GRAPH

Algorithm 4 Basic graph message passing

Input: Weight matrices, $\mathbf{W}_{\text{self}}$, $\mathbf{W}_{\text{neigh}}$, and bias, $\mathbf{b}$, neighborhood function, $\mathcal{N}$.

Input: Graph, $\mathcal{G}$ with nodes $\mathcal{V} = \{v_i\}_{i=0}^{V}$ and edges $\mathcal{E} = \{e_{u \to v} | u, v \in \mathcal{V}\}$, and a specified K number of rounds of message passing.

Output: Updated node features $\mathbf{h}_u^{(K+1)}$ for all nodes u

Initialize $\mathbf{h}_u^{(0)}$ as v_u for all nodes u

for $k \in [0, 1, \ldots, K]$ **do**

  for $u \in \mathcal{V}$ **do**

    for $v \in \mathcal{N}(u)$ **do**

      Compute messages : $\mathbf{M}_{v \to u} = \mathbf{W}_{\text{neighbors}} \mathbf{h}_v^{(k)} + \mathbf{b}$

    end for

    Compute self message: $\mathbf{M}_{\text{self}} = \mathbf{W}_{\text{self}} \mathbf{h}_u^{k}$

    Compute total message: $\mathbf{M}_u = \mathbf{M}_{\text{self}} + \sum_{v \in \mathcal{N}(u)} \mathbf{M}_{v \to u}$

    Update node: $\mathbf{h}_u^{(k+1)} \leftarrow \sigma(\mathbf{M}_u)$

  end for

end for

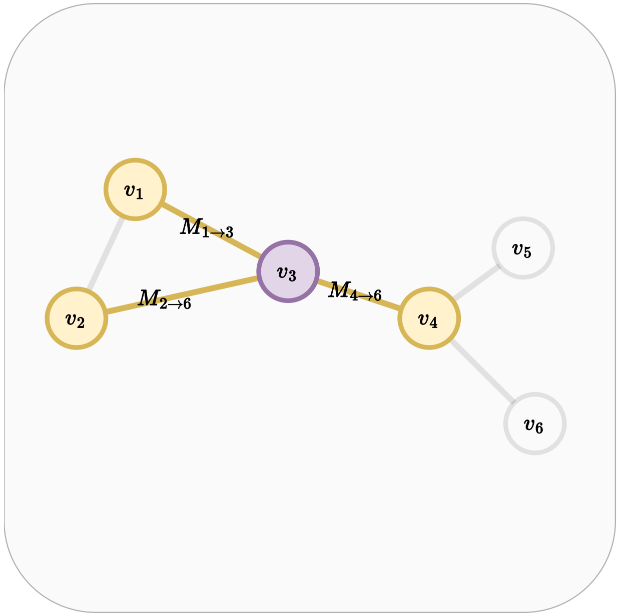

EQUIVARIANT MESSAGE PASSING

$$\mathbf{f}_u' = \phi\left(\mathbf{f}_u, \bigoplus_{v \in \mathcal{N}(u)} \psi_f\left(\mathbf{f}_u, \mathbf{f}_v, ||\mathbf{x}_u - \mathbf{x}_v||^2\right)\right)$$

$$\mathbf{x}_u' = \mathbf{x}_u + \sum_{v \neq u} (\mathbf{x}_u - \mathbf{x}_v)\, \psi_c\left(\mathbf{f}_u, \mathbf{f}_v, ||\mathbf{x}_u - \mathbf{x}_v||^2\right)$$

Based on Chap 5. of: Bronstein, Michael M., et al. "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges." arXiv preprint arXiv:2104.13478 (2021).

GRAPH MESSAGE PASSING NETWORKS

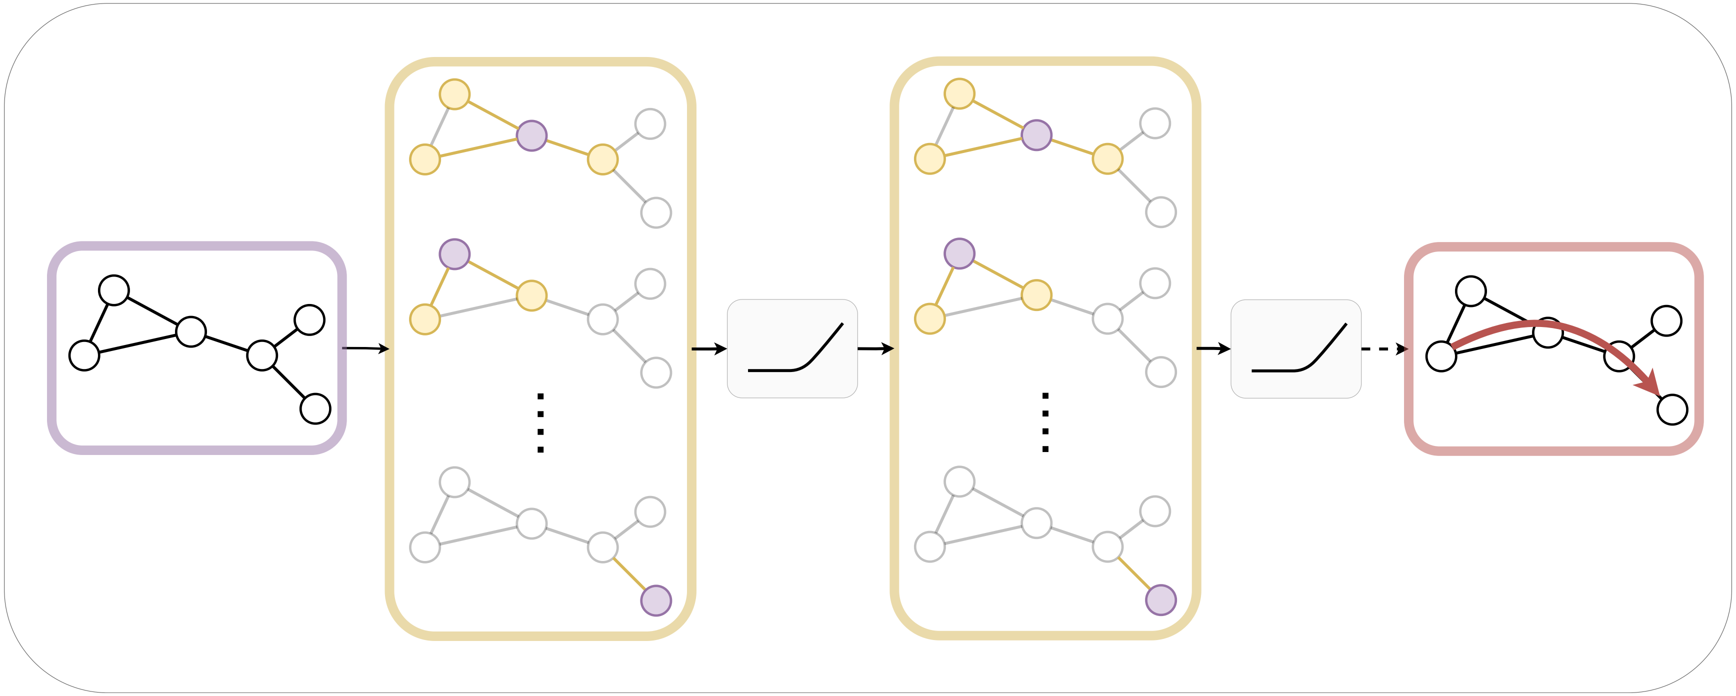

(Adapted from Thomas Kipf, https://tkipf.github.io/graph-convolutional-networks/)

RESOURCES

- Graph neural networks
  - A Gentle Introduction to Graph Neural Networks: https://distill.pub/2021/gnn-intro/
  - Graph Convolutional Networks, blog by Thomas Kipf: https://tkipf.github.io/graph-convolutional-networks/
  - ... and the paper: Kipf, Thomas N., and Max Welling. "Semi-supervised classification with graph convolutional networks." arXiv preprint arXiv:1609.02907 (2016).
  - Understanding Convolutions on Graphs: https://distill.pub/2021/understanding-gnns/
  - Hamilton, William L. Graph Representation Learning. Morgan & Claypool Publishers, 2020.
  - Battaglia, Peter W., et al. "Relational inductive biases, deep learning, and graph networks." arXiv preprint arXiv:1806.01261 (2018).
  - CS224W at Stanford: Machine Learning with Graphs
  - Theoretical Foundations of Graph Neural Networks by Petar Veličković: https://www.youtube.com/watch?v=uF53xsT7mjc
- Geometric deep learning
  - ICLR 2021 Keynote Talk by Michael Bronstein: Geometric Deep Learning: The Erlangen Programme of ML
  - e3nn.ogr: a modular PyTorch framework for Euclidean neural networks
  - Max Welling's talk "Learning equivariant and hybrid message passing on graphs": https://www.youtube.com/watch?v=hUrbS1BhBWc
  - The Geometric Deep Learning web page: https://geometricdeeplearning.com/
  - Jing, Bowen, et al. "Learning from protein structure with geometric vector perceptrons." arXiv preprint arXiv:2009.01411 (2020).
  - Bronstein, Michael M., et al. "Geometric deep learning: Grids, Groups, Graphs, Geodesics, and Gauges." arXiv preprint arXiv:2104.13478 (2021).
  - Deep learning for molecules and materials: https://dmol.pub/index.html
- Graphs in general
  - http://networksciencebook.com/
  - https://www.cs.cornell.edu/home/kleinber/networks-book/