E(n) Equivariant Topological Neural Networks

Claudio Battiloro[*,1]

Joint work with: E. Karaismailoğlu[*,2] M. Tec[*,1] G. Dasoulas[*,1]
M. Audirac[1] F. Dominici[1]

[1] *Harvard University* [2] *ETH Zurich*
**Equal Contribution*

The Thirteenth International Conference on Learning Representations
ICLR 2025

Topological Deep Learning

Goals and Motivation

- Graph-based representation: data are associated with the vertices to capture pairwise relations encoded by the edges

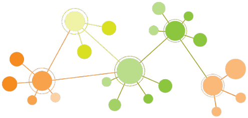

- In many systems, complex interactions cannot be reduced to dyadic relationships

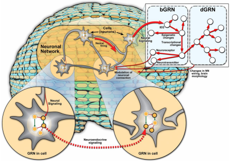

(a) In GRNs, some reactions occur when a set of genes interact

(b) In SNs, agents can interact in a group without being connected

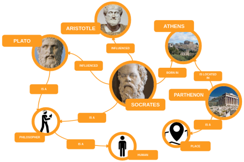

(c) In KGs, higher-order relationship could provide further insight and analysis

Topological Deep Learning
Combinatorial Topological Spaces

What topological descriptors
do we need to incorporate higher-order relationships?

Go beyond graphs: Simplicial Complexes, Cell Complexes, Hypergraphs, Network Sheaves, ...

- Each of these objects comes with a strong mathematical characterization
- And each of them is useful to model different types of higher-order interactions
- Several architectures operating on these spaces have been proposed
- Most of them share the common trait of being message-passing networks (MPNs)
  - Entities exchange messages whose structure and scheme are induced by the considered space

What topological descriptors
do we need to incorporate higher-order relationships?

Go beyond graphs: Simplicial Complexes, Cell Complexes, Hypergraphs, Network Sheaves, ...

- Each of these objects comes with a strong mathematical characterization
- And each of them is useful to model different types of higher-order interactions
- Several architectures operating on these spaces have been proposed
- Most of them share the common trait of being message-passing networks (MPNs)
  - Entities exchange messages whose structure and scheme are induced by the considered space

Topological Deep Learning
Combinatorial Topological Spaces

What topological descriptors
do we need to incorporate higher-order relationships?

Go beyond graphs: Simplicial Complexes, Cell Complexes, Hypergraphs, Network Sheaves, ...

- Each of these objects comes with a strong mathematical characterization
- And each of them is useful to model different types of higher-order interactions
- Several architectures operating on these spaces have been proposed
- Most of them share the common trait of being message-passing networks (MPNs)
  - Entities exchange messages whose structure and scheme are induced by the considered space

Combinatorial Complexes

- This common trait led to the general notion of Combinatorial Complex (CC)
- CCs are useful to design arbitrary MPNs on a variety of combinatorial objects
- A *combinatorial complex* (CC) [1] is a triple $(\mathcal{S}, \mathcal{X}, \mathrm{rk})$ consisting of a set $\mathcal{S}$, a subset $\mathcal{X}$ of $\mathcal{P}(\mathcal{S})\setminus\{\emptyset\}$, and a function $\mathrm{rk} : \mathcal{X} \to \mathbb{Z}_{\geq 0}$ with the following properties:
  - for all $s \in \mathcal{S}$, $\{s\} \in \mathcal{X}$
  - the function rk is order-preserving, i.e. if $x, y \in \mathcal{X}$ satisfy $x \subseteq y$, then $\mathrm{rk}(x) \leq \mathrm{rk}(y)$

  The elements of $\mathcal{S}$ are nodes, the elements of $\mathcal{X}$ are cells, $\mathrm{rk}(\cdot)$ is the rank function

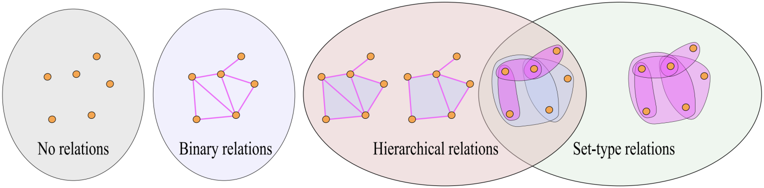

Combinatorial Complexes

- This common trait led to the general notion of Combinatorial Complex (CC)
- CCs are useful to design arbitrary MPNs on a variety of combinatorial objects
- A *combinatorial complex* (CC) [1] is a triple $(\mathcal{S}, \mathcal{X}, \mathrm{rk})$ consisting of a set $\mathcal{S}$, a subset $\mathcal{X}$ of $\mathcal{P}(S)\setminus\{\emptyset\}$, and a function $\mathrm{rk} : \mathcal{X} \to \mathbb{Z}_{\geq 0}$ with the following properties:
  - for all $s \in \mathcal{S}$, $\{s\} \in \mathcal{X}$
  - the function rk is order-preserving, i.e. if $x, y \in \mathcal{X}$ satisfy $x \subseteq y$, then $\mathrm{rk}(x) \leq \mathrm{rk}(y)$

  The elements of $\mathcal{S}$ are nodes, the elements of $\mathcal{X}$ are cells, $\mathrm{rk}(\cdot)$ is the rank function

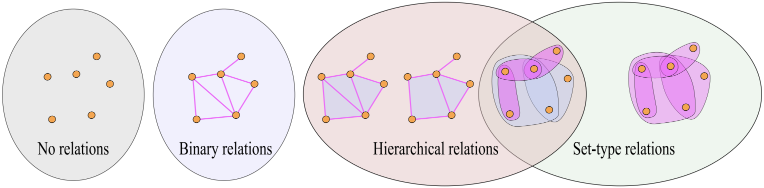

Combinatorial Complexes

Neighborhood Functions

- A neighborhood function $\mathcal{N} : \mathcal{X} \to \mathcal{P}(\mathcal{X})$ on $\mathcal{X}$ assigns to each cell x in $\mathcal{X}$ a nonempty collection of "neighbor cells" $\mathcal{N}(x) \subset \mathcal{X}$
- The up/down incidences $\mathcal{N}_{I,\uparrow}$ and $\mathcal{N}_{I,\downarrow}$ are defined by inclusion as

$$\mathcal{N}_{I,\uparrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x) + 1, x \subset y\}, \; \mathcal{N}_{I,\downarrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x) - 1, y \subset x\}$$

- The up/down adjacencies $\mathcal{N}_{A,\uparrow}$ and $\mathcal{N}_{A,\uparrow}$ are defined by common neighbor as

$$\mathcal{N}_{A,\uparrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x), \exists z \in \mathcal{X} : \mathrm{rk}(z) = \mathrm{rk}(x) - 1, z \subset y, \text{ and } z \subset x\},$$
$$\mathcal{N}_{A,\downarrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x), \exists z \in \mathcal{X} : \mathrm{rk}(z) = \mathrm{rk}(x) + 1, y \subset z, \text{ and } x \subset z\}$$

Combinatorial Complexes

Neighborhood Functions

- A neighborhood function $\mathcal{N} : \mathcal{X} \to \mathcal{P}(\mathcal{X})$ on $\mathcal{X}$ assigns to each cell x in $\mathcal{X}$ a nonempty collection of "neighbor cells" $\mathcal{N}(x) \subset \mathcal{X}$

- The up/down incidences $\mathcal{N}_{I,\uparrow}$ and $\mathcal{N}_{I,\downarrow}$ are defined by inclusion as

$$\mathcal{N}_{I,\uparrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x) + 1, x \subset y\},\ \mathcal{N}_{I,\downarrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x) - 1, y \subset x\}$$

- The up/down adjacencies $\mathcal{N}_{A,\uparrow}$ and $\mathcal{N}_{A,\uparrow}$ are defined by common neighbor as

$$\mathcal{N}_{A,\uparrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x), \exists z \in \mathcal{X} : \mathrm{rk}(z) = \mathrm{rk}(x) - 1, z \subset y, \text{ and } z \subset x\},$$
$$\mathcal{N}_{A,\downarrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x), \exists z \in \mathcal{X} : \mathrm{rk}(z) = \mathrm{rk}(x) + 1, y \subset z, \text{ and } x \subset z\}$$

Combinatorial Complexes

Neighborhood Functions

- The up/down adjacencies $\mathcal{N}_{A,\uparrow}$ and $\mathcal{N}_{A,\uparrow}$ are defined by common neighbor as

$$\mathcal{N}_{A,\uparrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x), \exists z \in \mathcal{X} : \mathrm{rk}(z) = \mathrm{rk}(x) - 1, z \subset y, \text{ and } z \subset x\},$$
$$\mathcal{N}_{A,\downarrow}(x) = \{y \in \mathcal{X} | \mathrm{rk}(y) = \mathrm{rk}(x) + 1, \exists z \in \mathcal{X} : \mathrm{rk}(z) > \mathrm{rk}(x), y \subset z, \text{ and } x \subset z\}$$

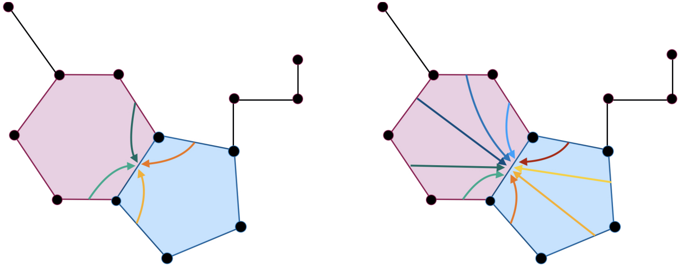

A CC with nodes as 0-cells, edges as 1-cells, and cycles as 2-cells
i.e. a polygonal cell complex

Combinatorial Complexes
Signals and MPNs

- A topological signal over $\mathcal{X}$ is a mapping $f : \mathcal{X} \to \mathbb{R}$
  - The feature vectors $\mathbf{h}_x \in \mathbb{R}^F$ and $\mathbf{h}_y \in \mathbb{R}^F$ of cells x and y are a collection of F topological signals, i.e. $\mathbf{h}_x = [f_1(x), \dots, f_F(x)]$ and $\mathbf{h}_y = [f_1(y), \dots, f_F(y)]$
- The l-th layer of a CC Message Passing Network (CCMPN) updates the embedding $\mathbf{h}_x^l$ as

$$\mathbf{h}_x^{l+1} = \beta\left(\mathbf{h}_x^l, \bigotimes_{\mathcal{N} \in \mathcal{CN}} \bigoplus_{y \in \mathcal{N}(x)} \psi_{\mathcal{N},\mathrm{rk}(x)}\left(\mathbf{h}_x^l, \mathbf{h}_y^l\right)\right)$$

  - $\mathbf{h}_x^0 := \mathbf{h}_x$ are the initial features
  - $\mathcal{CN}$ is a collection of neighborhood functions.
  - $\bigoplus$ is an intra-neighborhood permutation invariant aggregator
  - $\bigotimes$ is an inter-neighborhood (possibly) permutation invariant aggregator

Combinatorial Complexes

Signals and MPNs

- A topological signal over $\mathcal{X}$ is a mapping $f : \mathcal{X} \to \mathbb{R}$
  - The feature vectors $\mathbf{h}_x \in \mathbb{R}^F$ and $\mathbf{h}_y \in \mathbb{R}^F$ of cells x and y are a collection of F topological signals, i.e. $\mathbf{h}_x = [f_1(x), \dots, f_F(x)]$ and $\mathbf{h}_y = [f_1(y), \dots, f_F(y)]$
- The l-th layer of a CC Message Passing Network (CCMPN) updates the embedding $\mathbf{h}_x^l$ as

$$\mathbf{h}_x^{l+1} = \beta\left(\mathbf{h}_x^l, \bigotimes_{\mathcal{N}\in\mathcal{CN}} \bigoplus_{y\in\mathcal{N}(x)} \psi_{\mathcal{N},\mathrm{rk}(x)}\left(\mathbf{h}_x^l, \mathbf{h}_y^l\right)\right)$$

  - $\mathbf{h}_x^0 := \mathbf{h}_x$ are the initial features
  - $\mathcal{CN}$ is a collection of neighborhood functions.
  - $\bigoplus$ is an intra-neighborhood permutation invariant aggregator
  - $\bigotimes$ is an inter-neighborhood (possibly) permutation invariant aggregator

Equivariance

- Symmetries are formally framed as groups
- Let G be a group with an action $a : G \times \mathcal{Y} \to \mathcal{Y}$ on a set $\mathcal{Y}$
- A function $e : \mathcal{Y} \to \mathcal{Y}$ is equivariant w.r.t. the a if, for all $x \in \mathcal{Y}$ and $g \in G$, it holds

$$e(a(g, x)) = a(g, e(x))$$

  If $e(a(g, x)) = e(x)$, e is said to be invariant w.r.t. the action of G
- CCMPNs are permutation equivariant, i.e. they are equivariant to the relabeling of cells

Equivariance

- Symmetries are formally framed as groups
- Let G be a group with an action $a : G \times \mathcal{Y} \rightarrow \mathcal{Y}$ on a set $\mathcal{Y}$
- A function $e : \mathcal{Y} \rightarrow \mathcal{Y}$ is equivariant w.r.t. the a if, for all $x \in \mathcal{Y}$ and $g \in G$, it holds

  $$e(a(g, x)) = a(g, e(x))$$

  If $e(a(g, x)) = e(x)$, e is said to be invariant w.r.t. the action of G
- CCMPNs are permutation equivariant, i.e. they are equivariant to the relabeling of cells

E(n) Equivariance

- Many problems exhibit n-dimensional rotation and translation symmetries
- It is desirable to design models that are equivariant w.r.t. the action of the $E(n)$ group
- An element of the $E(n)$ group is a tuple $(\mathbf{O}, \mathbf{b})$ consisting of an orthogonal matrix $\mathbf{O} \in \mathbb{R}^{n \times n}$ (the rotation) and a vector $\mathbf{b} \in \mathbb{R}^n$ (the translation)
- The action a of $E(n)$ on $\mathbf{x} \in \mathbb{R}^n$ is

$$a((\mathbf{O}, \mathbf{b}), \mathbf{x}) = \mathbf{O}\mathbf{x} + \mathbf{b}$$

- Consider the setting in which nodes (0-cells) come with both non-geometric $\mathbf{h}_x \in \mathbb{R}^F$ and geometric $\mathbf{x}_x \in \mathbb{R}^n$ features, such as positions
- **Research Question**: How can we design $E(n)$ equivariant CCMPNs?

E(n) Equivariance

- Many problems exhibit n-dimensional rotation and translation symmetries
- It is desirable to design models that are equivariant w.r.t. the action of the $E(n)$ group
- An element of the $E(n)$ group is a tuple $(\mathbf{O}, \mathbf{b})$ consisting of an orthogonal matrix $\mathbf{O} \in \mathbb{R}^{n \times n}$ (the rotation) and a vector $\mathbf{b} \in \mathbb{R}^n$ (the translation)
- The action a of $E(n)$ on $\mathbf{x} \in \mathbb{R}^n$ is

$$a((\mathbf{O}, \mathbf{b}), \mathbf{x}) = \mathbf{O}\mathbf{x} + \mathbf{b}$$

- Consider the setting in which nodes (0-cells) come with both non-geometric $\mathbf{h}_x \in \mathbb{R}^F$ and geometric $\mathbf{x}_x \in \mathbb{R}^n$ features, such as positions
- Research Question: How can we design $E(n)$ equivariant CCMPNs?

E(n) Equivariance

- Many problems exhibit n-dimensional rotation and translation symmetries
- It is desirable to design models that are equivariant w.r.t. the action of the $E(n)$ group
- An element of the $E(n)$ group is a tuple $(\mathbf{O}, \mathbf{b})$ consisting of an orthogonal matrix $\mathbf{O} \in \mathbb{R}^{n \times n}$ (the rotation) and a vector $\mathbf{b} \in \mathbb{R}^n$ (the translation)
- The action a of $E(n)$ on $\mathbf{x} \in \mathbb{R}^n$ is

$$a((\mathbf{O}, \mathbf{b}), \mathbf{x}) = \mathbf{O}\mathbf{x} + \mathbf{b}$$

- Consider the setting in which nodes (0-cells) come with both non-geometric $\mathbf{h}_x \in \mathbb{R}^F$ and geometric $\mathbf{x}_x \in \mathbb{R}^n$ features, such as positions
- **Research Question**: How can we design $E(n)$ equivariant CCMPNs?

E(n) Equivariant Topological Neural Networks

- We represent equivariance by scalarization
  - Geometric features are transformed into invariant scalars, and then processed before being combined along the original directions
- The l-th layer of an $E(n)$ Equivariant Topological Neural Network (ETNN) updates the embeddings $\mathbf{h}_x^l$ of every cell $x \in \mathcal{X}$ and the positions $\mathbf{x}_z^l$ of every node $z \in \mathcal{S}$ as

$$\mathbf{h}_x^{l+1} = \beta\left(\mathbf{h}_x^l, \bigotimes_{\mathcal{N}\in\mathcal{CN}} \bigoplus_{y\in\mathcal{N}(x)} \underbrace{\psi_{\mathcal{N},\mathrm{rk}(x)}\left(\mathbf{h}_x^l, \mathbf{h}_y^l, \mathrm{Inv}\left(\{\mathbf{x}_z^l\}_{z\in x}, \{\mathbf{x}_z^l\}_{z\in y}\right)\right)}_{\mathbf{m}_{x,y}^{\mathcal{N}}}\right)$$

$$\mathbf{x}_z^{l+1} = \mathbf{x}_z^l + C \sum_{\mathcal{N}\in\mathcal{CN}} \sum_{t\in S:\{t\}\in\mathcal{N}(z)} \left(\mathbf{x}_z^l - \mathbf{x}_t^l\right) \xi\left(\mathbf{m}_{z,t}^{\mathcal{N}}\right)$$

  - Inv is an invariant function, and we refer to a specific Inv as a geometric invariant.

E(n) Equivariant Topological Neural Networks

- We represent equivariance by scalarization
  - Geometric features are transformed into invariant scalars, and then processed before being combined along the original directions

- The l-th layer of an $E(n)$ Equivariant Topological Neural Network (ETNN) updates the embeddings $\mathbf{h}_x^l$ of every cell $x \in \mathcal{X}$ and the positions $\mathbf{x}_z^l$ of every node $z \in \mathcal{S}$ as

$$\mathbf{h}_x^{l+1} = \beta\left(\mathbf{h}_x^l, \bigotimes_{\mathcal{N}\in\mathcal{CN}} \bigoplus_{y\in\mathcal{N}(x)} \underbrace{\psi_{\mathcal{N},\mathrm{rk}(x)}\left(\mathbf{h}_x^l, \mathbf{h}_y^l, \mathrm{Inv}\left(\{\mathbf{x}_z^l\}_{z\in x}, \{\mathbf{x}_z^l\}_{z\in y}\right)\right)}_{\mathbf{m}_{x,y}^{\mathcal{N}}}\right)$$

$$\mathbf{x}_z^{l+1} = \mathbf{x}_z^l + C \sum_{\mathcal{N}\in\mathcal{CN}} \sum_{t\in S:\{t\}\in\mathcal{N}(z)} \left(\mathbf{x}_z^l - \mathbf{x}_t^l\right)\xi\left(\mathbf{m}_{z,t}^{\mathcal{N}}\right)$$

  - Inv is an invariant function, and we refer to a specific Inv as a geometric invariant.

E(n) Equivariant Topological Neural Networks

- We represent equivariance by scalarization
  - Geometric features are transformed into invariant scalars, and then processed before being combined along the original directions

- The l-th layer of an $E(n)$ Equivariant Topological Neural Network (ETNN) updates the embeddings $\mathbf{h}_x^l$ of every cell $x \in \mathcal{X}$ and the positions $\mathbf{x}_z^l$ of every node $z \in \mathcal{S}$ as

$$\mathbf{h}_x^{l+1} = \beta\left(\mathbf{h}_x^l, \bigotimes_{\mathcal{N}\in\mathcal{CN}} \bigoplus_{y\in\mathcal{N}(x)} \underbrace{\psi_{\mathcal{N},\mathrm{rk}(x)}\left(\mathbf{h}_x^l, \mathbf{h}_y^l, \mathrm{Inv}\left(\{\mathbf{x}_z^l\}_{z\in x}, \{\mathbf{x}_z^l\}_{z\in y}\right)\right)}_{\mathbf{m}_{x,y}^{\mathcal{N}}}\right)$$

$$\mathbf{x}_z^{l+1} = \mathbf{x}_z^l + C \sum_{\mathcal{N}\in\mathcal{CN}} \sum_{t\in S:\{t\}\in\mathcal{N}(z)} \left(\mathbf{x}_z^l - \mathbf{x}_t^l\right) \xi\left(\mathbf{m}_{z,t}^{\mathcal{N}}\right)$$

  - Inv is an invariant function, and we refer to a specific Inv as a geometric invariant.

Geometric Invariants

- Permutation invariant functions of pairwise distances: if x and y have the same cardinality,

$$\bigoplus_{z \in x, t \in y} \tau \left(\| \mathbf{x}_z - \mathbf{x}_t \| \right)$$

  is a geometric invariant.

- A basic instance is:
  - $\bigoplus = \sum$
  - $\tau = \mathrm{Id}$

  i.e. the sum of pairwise distances.

Geometric Invariants

- Distances of permutation invariant functions: regardless of the cardinalities of x and y,

$$\tau\left(\left\|\bigoplus_{z\in x}\mathbf{x}_z - \bigoplus_{t\in y}\mathbf{x}_t\right\|\right)$$

is a geometric invariant if $\bigoplus$ is linear

- A basic instance is:
  - $\bigoplus = \frac{1}{|\cdot|}\sum$
  - $\tau = \mathrm{Id}$

i.e. the distance between centroids

Geometric Invariants

- Hausdorff distance: Regardless of the cardinalities of x and y, the Hausdorff distance, defined as

$$\max\left\{\max_{z\in x}\min_{t\in y}\|\mathbf{x}_z-\mathbf{x}_t\|,\max_{t\in y}\min_{z\in x}\|\mathbf{x}_t-\mathbf{x}_z\|\right\}$$

  is a geometric invariant. The Hausdorff distance measures the mutual proximity of two sets whose elements live in a metric space.

- Volume of the convex hull: Regardless of the cardinalities of x and y, any function of the volumes of the convex hulls of $\{\mathbf{x}_z\}_{z\in x}$ and $\{\mathbf{x}_z\}_{z\in y}$ is a geometric invariant

Overview

- ETNNs generalize several scalarization-based architectures on combinatorial objects
- By tailoring the definition of cells and the choice of the neighborhood functions, $E(n)$ Equivariant Graph/Simplicial/Cellular/Hypergraph Networks can be derived

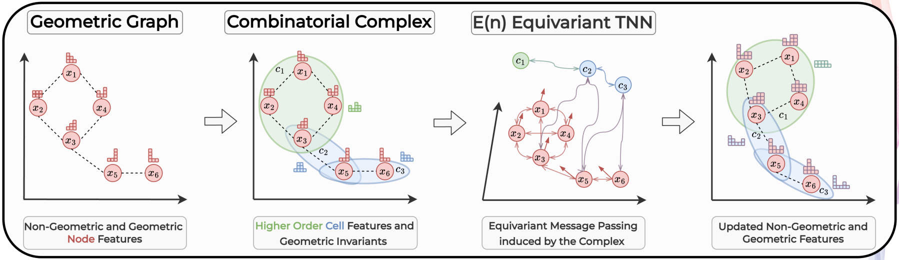

Real-World Combinatorial Complexes

Molecular Combinatorial Complexes

- In Molecular Combinatorial Complexes, the set of vertices $\mathcal{S}$ is the set of atoms, while the set of cells $\mathcal{X}$ and the rank function rk are such that:
  - Atoms are 0-cells
  - Bonds and functional groups made of two atoms are 1-cells
  - Rings and functional groups made of more than two atoms are 2-cells

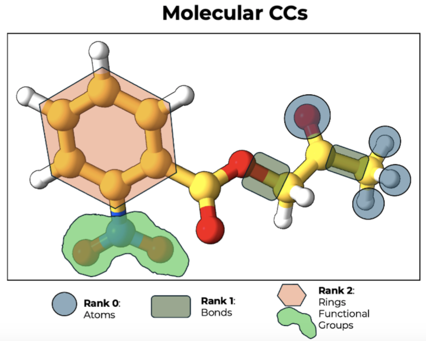

- This modeling cannot be framed neither as a graph nor as a cell complex nor a hypergraph!

Real-World Combinatorial Complexes

Geospatial Combinatorial Complexes

- Points, polylines, and polygons constitute some of the most basic geometric objects in geospatial data
- These objects can be modeled as cells of a CC, and their rank corresponds to their intrinsic dimensionality:
  - Points are 0-cells
  - traffic roads are 1-cells
  - Census tracts are 2-cells

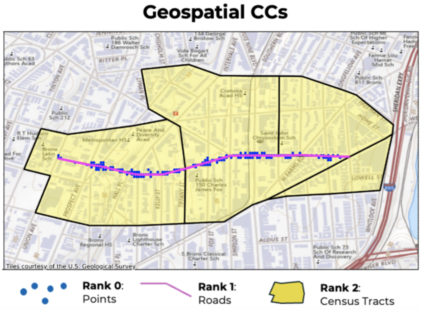

- This modeling cannot be framed neither as a graph nor as a cell complex nor a hypergraph!

Real-World Combinatorial Complexes

Geospatial Combinatorial Complexes

- The same rationale can be applied to any irregular multi-resolution geospatial data
- Another example is the political division of spatial domains, e.g., zip codes, counties, and states. In this case:
  - Zip-codes are 0-cells
  - Counties are 1-cells
  - States are 2-cells

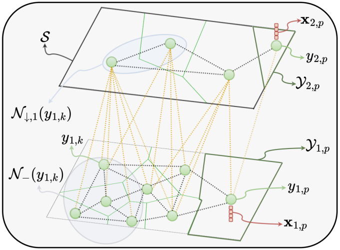

- Within our framework, multi-level features, associated with each rank, can be used without the need for any a priori aggregation (as it is usually done for geospatial data)!

Molecular Property Prediction on QM9

Supervised Benchmark

- ETNN-* is the best ETNN configuration we tried per each target
- ETNN-graph-W and EMPSN* are EGNN [2] and EMPSN [3] derived from the ETNN framework with our codebase
- We generalize the notion of a virtual node by introducing a virtual cell and a dedicated neighborhood function
- In this way, we also show how ETNN can handle heterogeneous interactions

Task Units (↓)	α bohr^3	$\Delta\varepsilon$ meV	$\varepsilon_{\text{HOMO}}$ meV	$\varepsilon_{\text{LUMO}}$ meV	μ D	C_ν cal/mol K
NMP	.092	69	43	38	.030	.040
Schnet	.235	63	41	34	.033	.033
Cormorant	.085	61	34	38	.038	.026
L1Net	.088	68	46	35	.043	.031
LieConv	.084	49	30	25	.032	.038
DimeNet++	**.044**	33	25	20	.030	.023
TFN	.223	58	40	38	.064	.101
SE(3)-Tr.	.142	53	35	33	.051	.054
Equiformer (SotA)	.046	**30**	**15**	**14**	**.011**	**.023**
TopNet†	.083	47	37	24	.035	.032
EGNN	.071	48	29	25	.029	.031
ETNN-*	.062	45	26	22	.022	.030
ETNN-graph-W	.067	46	27	25	.030	.036
ETNN-w/o VC	.161	73	49	48	.306	.051
EMPSN*	.061	42	29	25	.030	.028
ETNN-single-*	.069	46	26	23	.025	.033
Improvement over EGNN	**-13%**	**-6%**	**-10%**	**-12%**	**-26%**	**-3%**

Computational Advantages of ETNN

- ETNN-* is the best ETNN configuration we tried per each target
- ETNN-graph-W and EMPSN* are EGNN [2] and EMPSN [3] derived from the ETNN framework with our codebase

	Avg. runtime (seconds/epoch)
EGNN-graph-W	60.34
ETNN	193.16
EMPSN*	251.3

	Avg. memory consumption (GB)
EGNN-graph-W	13.63
ETNN	12.05
EMPSN*	27.7

Hyperlocal Air Pollution Downscaling

A New Semisupervised Benchmark for TDL

- The goal is predicting $\mathrm{PM}_{2.5}$ at a high resolution of $\approx 22m$ (point level, 0-cells)
- Thus, this is a node regression task
- To obtain train/val/test split being mindful of spatial correlation, we randomly sample at the census tract level (2-cells)
  - In this way, points (0-cells) belonging to the same census tract are in the same set

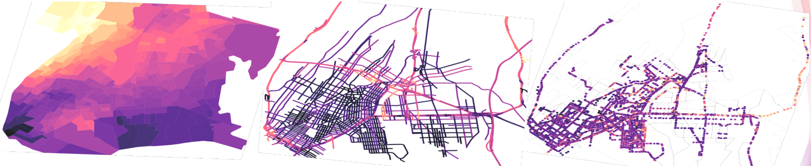

2-cell features (census tracts): coarse $\mathrm{PM}_{2.5}$ estimates, demographics, land use (e.g., residential, commercial, parks), etc.

1-cell features (roads): annual average daily traffic and traffic composition, etc.

0-cell features (points): distance to the nearest school and intersection, etc.

Baseline	AEV(R^2)	Std. Err (R^2)	MSE
Linear	0.51%	0.95%	1.106
MLP	2.35%	1.61%	1.022
GNN	2.44%	0.99%	0.987
EGNN	1.43%	1.2%	1.041
ETNN	9.34%	2.05%	0.935

(a) Baseline model comparison

Baseline	DEV (R^2)	Std. Err (R^2)	MSE
no virtual node	-1.80%	2.32%	0.957
no position update: invariant ETNN	-1.37%	2.52%	0.956
no geometric features: CCMPN	-1.08%	2.05%	0.946

(b) Ablation study results

Expressivity of ETNNs

- The ability of GNNs to differentiate between non-isomorphic graphs is typically employed as an expressivity metric
- Vanilla GNNs are at maximum as powerful as the Weisfeiler-Leman (WL) test
- To include geometric graphs, the Geometric WL test (GWL) has been introduced [4]
- Two geometric graphs are called geometrically isomorphic if there exists an edge-preserving bijection b between their nodes s.t. their positions are equivalent up to the $E(n)$ action:

$$\mathbf{h}^{\mathcal{G}_2}_{x_{b(i)}}, \mathbf{x}^{\mathcal{G}_2}_{x_{b(i)}} = \mathbf{h}^{\mathcal{G}_1}_{x_i}, \mathbf{O}\mathbf{x}^{\mathcal{G}_1}_{x_i} + \mathbf{b}$$

Expressivity of ETNNs

- The ability of GNNs to differentiate between non-isomorphic graphs is typically employed as an expressivity metric
- Vanilla GNNs are at maximum as powerful as the Weisfeiler-Leman (WL) test
- To include geometric graphs, the Geometric WL test (GWL) has been introduced [4]
- Two geometric graphs are called geometrically isomorphic if there exists an edge-preserving bijection b between their nodes s.t. their positions are equivalent up to the $E(n)$ action:

$$\mathbf{h}^{\mathcal{G}_2}_{x_{b(i)}}, \mathbf{x}^{\mathcal{G}_2}_{x_{b(i)}} = \mathbf{h}^{\mathcal{G}_1}_{x_i}, \mathbf{O}\mathbf{x}^{\mathcal{G}_1}_{x_i} + \mathbf{b}$$

Expressivity of ETNNs

- Two geometric graphs are said to be k-hop distinct if for all graph isomorphisms, there is some node such that the corresponding k-hop subgraphs are distinct

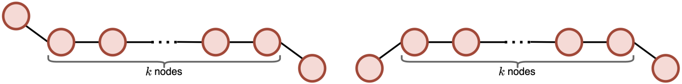

k-chain geometric graphs are $\left(\left\lfloor \frac{k}{2} \right\rfloor + 1\right)$-hop distinct

- How good are ETNNs in distinguishing k-hop distinct graph?

Expressivity of ETNNs

- Two geometric graphs are said to be k-hop distinct if for all graph isomorphisms, there is some node such that the corresponding k-hop subgraphs are distinct

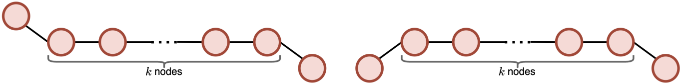

k-chain geometric graphs are $\left(\left\lfloor \frac{k}{2} \right\rfloor + 1\right)$-hop distinct

- How good are ETNNs in distinguishing k-hop distinct graph?

Expressivity of ETNNs

- The Geometric Augmented Hasse graph $\mathcal{G}_{\mathcal{X}}$ of $(\mathcal{S}, \mathcal{X}, \mathrm{rk}) + \mathcal{CN}$ is a (possibly) directed geometric graph $\mathcal{G}_{\mathcal{X}} = (\mathcal{X}, \mathcal{E})$ with cells as nodes, edges given by

$$\mathcal{E} = \{(x, y) | x \in \mathcal{X}, y \in \mathcal{X}, \exists \mathcal{N} \in \mathcal{CN} : x \in \mathcal{N}(y) \text{ or } y \in \mathcal{N}(x)\},$$

and positions of cell x being a linear permutation invariant function of its node positions

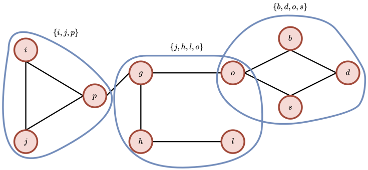

(a)

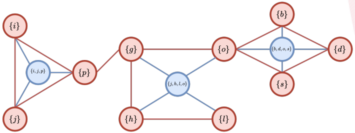

(b)

- (c) A CC with nodes of the underlying graph as cells of rank 0, and with 3 arbitrary cells of rank 1
- (d) The corresponding geometric augmented Hasse graph if $\mathcal{CN}$ = up/down incidences + graph adjacency, and $\bigoplus$ is the mean
- This is also an example of a skeleton-preserving lift

Expressivity of ETNNs

- The Geometric Augmented Hasse graph $\mathcal{G}_{\mathcal{X}}$ of $(\mathcal{S}, \mathcal{X}, \mathrm{rk}) + \mathcal{CN}$ is a (possibly) directed geometric graph $\mathcal{G}_{\mathcal{X}} = (\mathcal{X}, \mathcal{E})$ with cells as nodes, edges given by

$$\mathcal{E} = \{(x, y) | x \in \mathcal{X}, y \in \mathcal{X}, \exists \mathcal{N} \in \mathcal{CN} : x \in \mathcal{N}(y) \text{ or } y \in \mathcal{N}(x)\},$$

and positions of cell x being a linear permutation invariant function of its node positions

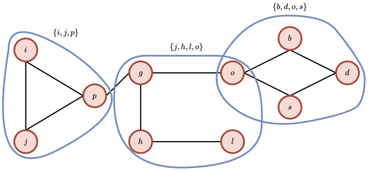

(c)

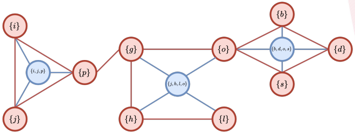

(d)

  - (c) A CC with nodes of the underlying graph as cells of rank 0, and with 3 arbitrary cells of rank 1
  - (d) The corresponding geometric augmented Hasse graph if $\mathcal{CN}$ = up/down incidences + graph adjacency, and $\bigoplus$ is the mean
  - This is also an example of a skeleton-preserving lift

Expressivity of ETNNs

- We show that a subclass of ETNNs on a CC $\mathcal{X}$ are equivalent to EGNNs over $\mathcal{G}_{\mathcal{X}}$
- Thus, we can study the expressivity of ETNNs by studying how the GWL behaves on Geometric Augmented Hasse Graphs
- Proposition. Assume to have:
  - Two k-hop distinct geometric graphs $\mathcal{G}_1 = (\mathcal{S}_1, \mathcal{E}_1)$ and $\mathcal{G}_2 = (\mathcal{S}_2, \mathcal{E}_2)$
  - Two CCs $\mathcal{X}_{\mathcal{G}_1}$ and $\mathcal{X}_{\mathcal{G}_2}$ and a collection $\mathcal{CN}$ of neighborhoods obtained via a skeleton preserving lift and leading to undirected geometric Hasse graphs $\mathcal{G}_{\mathcal{X}_{\mathcal{G}_1}}$ and $\mathcal{G}_{\mathcal{X}_{\mathcal{G}_2}}$

  An ETNN operating over $\mathcal{X}_{\mathcal{G}_1}$ and $\mathcal{X}_{\mathcal{G}_2}$ can distinguish $\mathcal{G}_1$ and $\mathcal{G}_2$ in M layers, where M is the number of layers required to have at least one cell in $\mathcal{X}_{\mathcal{G}_1}/\mathcal{X}_{\mathcal{G}_2}$ whose receptive field is the whole set of nodes in $\mathcal{G}_1/\mathcal{G}_2$

Expressivity of ETNNs

- We show that a subclass of ETNNs on a CC $\mathcal{X}$ are equivalent to EGNNs over $\mathcal{G}_{\mathcal{X}}$
- Thus, we can study the expressivity of ETNNs by studying how the GWL behaves on Geometric Augmented Hasse Graphs
- Proposition. Assume to have:
  - Two k-hop distinct geometric graphs $\mathcal{G}_1 = (\mathcal{S}_1, \mathcal{E}_1)$ and $\mathcal{G}_2 = (\mathcal{S}_2, \mathcal{E}_2)$
  - Two CCs $\mathcal{X}_{\mathcal{G}_1}$ and $\mathcal{X}_{\mathcal{G}_2}$ and a collection $\mathcal{CN}$ of neighborhoods obtained via a skeleton preserving lift and leading to undirected geometric Hasse graphs $\mathcal{G}_{\mathcal{X}_{\mathcal{G}_1}}$ and $\mathcal{G}_{\mathcal{X}_{\mathcal{G}_2}}$

  An ETNN operating over $\mathcal{X}_{\mathcal{G}_1}$ and $\mathcal{X}_{\mathcal{G}_2}$ can distinguish $\mathcal{G}_1$ and $\mathcal{G}_2$ in M layers, where M is the number of layers required to have at least one cell in $\mathcal{X}_{\mathcal{G}_1}/\mathcal{X}_{\mathcal{G}_2}$ whose receptive field is the whole set of nodes in $\mathcal{G}_1/\mathcal{G}_2$

Expressivity of ETNNs

- The expressiveness of ETNNs depends on the cells and the neighborhood functions

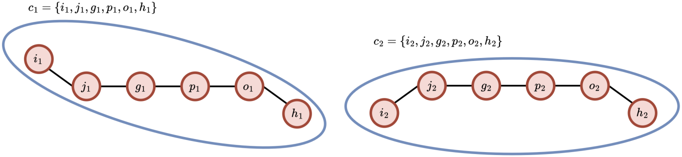

(e)

(f)

(e) Lift 1. (f) The corresponding geometric augmented Hasse graphs.

- 1 layer is enough for ETNNs to distinguish the graphs; EGNNs need 3 layers

Expressivity of ETNNs

- The expressiveness of ETNNs depends on the cells and the neighborhood functions

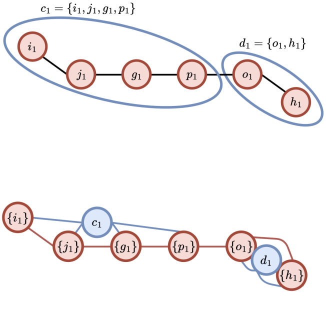

(a) Lift 2. (b) The corresponding geometric augmented Hasse graphs.

- 2 layers are enough for ETNNs to distinguish the graphs; EGNNs need 3 layers

Conclusions

- ETNN is the first unifying framework for scalarization-based $E(n)$ equivariant networks over combinatorial topological spaces or, more in general, over combinatorial objects used to model interactions among entities

- ETNNs are simple, expressive, and extremely flexible

- Many open venues: tyme-varying scenarios, additional geometric invariants, beyond-scalarization approaches,...

(a) Paper

(b) Claudio's X

(c) Claudio's Linkedin

References

- M. Hajij, G. Zamzmi, T. Papamarkou, N. Miolane, A. Guzmán-Sáenz, K. N. Ramamurthy, T. Birdal, T. K. Dey, S. Mukherjee, S. N. Samaga *et al.*, "Topological deep learning: Going beyond graph data," *arXiv preprint arXiv:2206.00606*, 2022.
- V. G. Satorras, E. Hoogeboom, and M. Welling, "E (n) equivariant graph neural networks," in *International conference on machine learning*. PMLR, 2021, pp. 9323–9332.
- F. Eijkelboom, R. Hesselink, and E. J. Bekkers, "E (n) equivariant message passing simplicial networks," in *International Conference on Machine Learning*. PMLR, 2023, pp. 9071–9081.
- C. K. Joshi, C. Bodnar, S. V. Mathis, T. Cohen, and P. Lio, "On the expressive power of geometric graph neural networks," in *International Conference on Machine Learning*. PMLR, 2023, pp. 15 330–15 355.