Orbit-Equivariant Graph Neural Networks

ICLR 2024

Matthew Morris, Bernardo Cuenca Grau, & Ian Horrocks

Computer Science Department
University of Oxford

May 2024

Some Graph Theory

Similarity

An automorphism of G is a permutation σ on G satisfying $\sigma \cdot G = G$. Nodes $v, w \in V(G)$ are *similar* if there is an automorphism σ of G such that $\sigma(v) = w$. Similar nodes form equivalence classes called *orbits*.

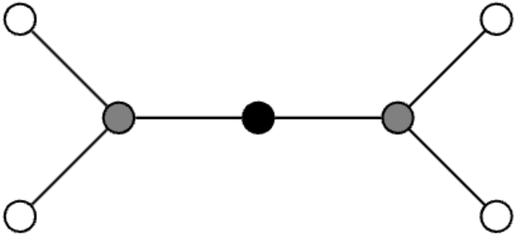

Figure: Graph with nodes colored by orbit.

Structural Inductive Biases

Equivariance

Node-labelling function f on a domain D closed under permutation is *equivariant* if $f(\sigma \cdot G) = \sigma \cdot f(G)$ holds for all $G \in D$ and permutations σ on $V(G)$.

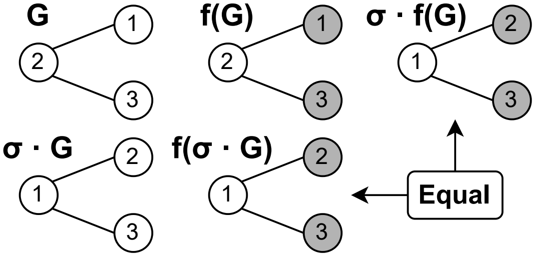

Figure: Example of an equivariant function f.

The Limits of Equivariance

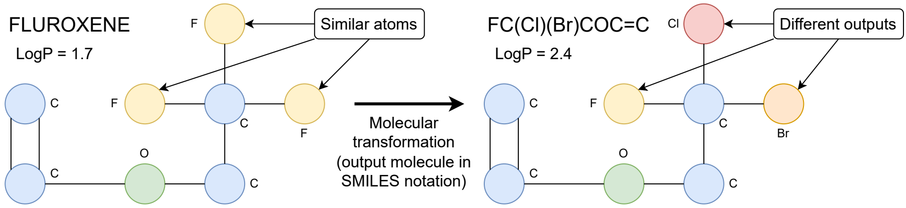

Figure: Non-equivariant molecular transformation that increases lipophilicity (LogP), where the nodes are labelled with the atom type, without positional information.

The Limits of Equivariance

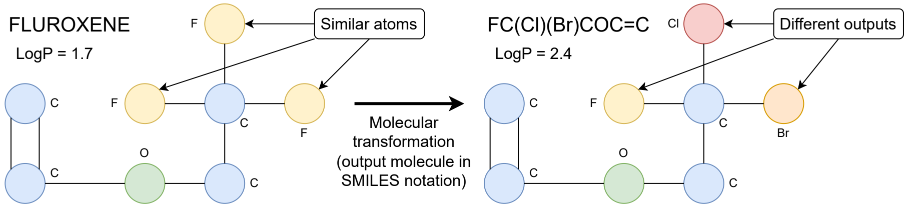

Figure: Non-equivariant molecular transformation that increases lipophilicity (LogP), where the nodes are labelled with the atom type, without positional information.

Proposition 1

Let f be an equivariant node-labelling function and let G be a labelled graph in its domain. If $v, w \in V(G)$ are similar, then $f(G)_v = f(G)_w$.

Orbit-Equivariance

Definition 1

A node-labelling function f on domain D closed under permutation is **orbit-equivariant** if, for all labelled graphs $G \in D$, permutations σ on $V(G)$, and orbits $r \in R(G)$, it holds that
$\{\!\{ f(\sigma \cdot G)_{\sigma(v)} \mid v \in r \}\!\} = \{\!\{ f(G)_v \mid v \in r \}\!\}$.

Orbit-Equivariance in the Hierarchy of Graph Functions

Not all orbit-equivariant functions are equivariant:

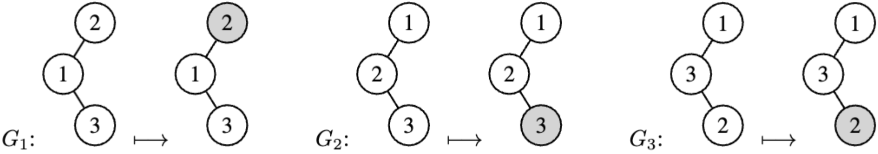

Figure: An orbit-equivariant function that is not equivariant.

Proposition 2

All equivariant functions are orbit-equivariant, but not vice-versa. There exist node-labelling functions which are not orbit-equivariant.

Max-Orbit and Orbit-Equivariance

Proposition 3

If f is orbit-equivariant and $\texttt{max-orbit}(f) = 1$, then f is equivariant.

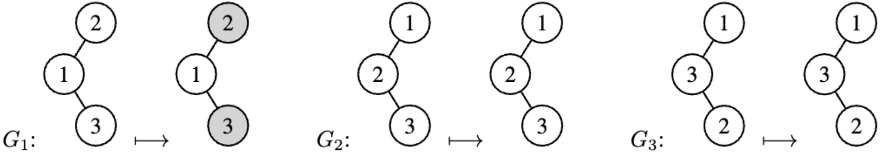

Figure: A node-labelling function with max-orbit $= 1$ that is not equivariant.

Non-equivariant GNN Architectures

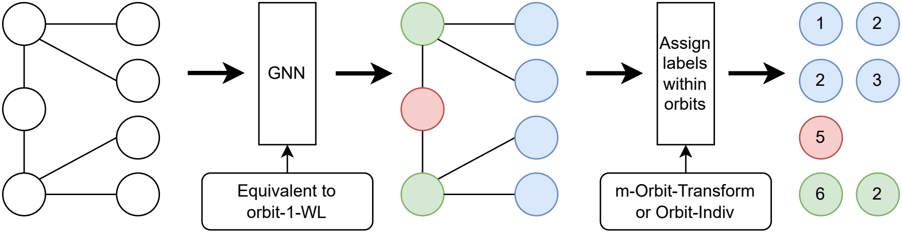

Figure: Architecture for constructing orbit-equivariant GNNs.

Model Expressivity

Theorem 2

The expressivity of our proposed models is as follows, where for a class of models X, the claim "X is not equivariant" means that "there exists $f \in X$ such that f is not equivariant":

1. Unique-ID-GNNs are not orbit-equivariant and, for any fixed n, can approximate any node-labelling function $f : G_n \to \mathbb{R}^n$, where G_n is the set of graphs with $\leq n$ nodes.
2. RNI-GNNs are equivariant in expectation and, for any fixed n, can approximate any equivariant function $f : G_n \to \mathbb{R}^n$ with probability arbitrarily close to 1. They can approximate some non-equivariant and orbit-equivariant functions with probability arbitrarily close to 1, but there exist RNI-GNNs which, with probability arbitrarily close to 1, are not orbit-equivariant.

Model Expressivity

Theorem 2

The expressivity of our proposed models is as follows, where for a class of models X, the claim "X is not equivariant" means that "there exists $f \in X$ such that f is not equivariant":

3. Orbit-Indiv-GNNs are not equivariant but are orbit-equivariant on graphs whose orbits are distinguishable by orbit-1-WL. For any $m \in \mathbb{Z}^+$, there exist Orbit-Indiv-GNNs f with $\texttt{max-orbit}(f) > m$.
4. m-Orbit-Transform-GNNs f are not equivariant but are orbit-equivariant on graphs whose orbits are distinguishable by orbit-1-WL. They have $\texttt{max-orbit}(f) \leq m$ and there exist m-Orbit-Transform-GNNs f with $\texttt{max-orbit}(f) = m$.

Datasets

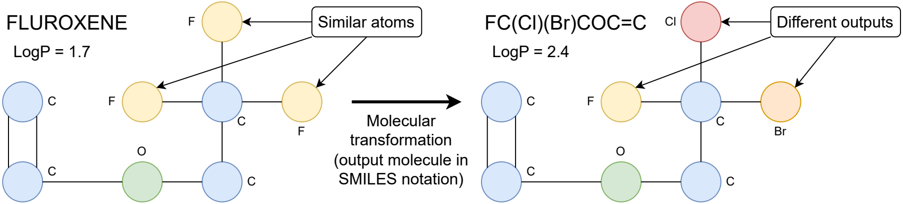

Figure: Molecular transformation.

- **Bioisostere** - swap out some atoms in a given molecule to achieve minimal lipophilicity. This is an important factor in drug design.
- **Alchemy-Max-Orbit-**m - based off Alchemy dataset, designed to test model performance for orbits of increasing sizes.

Results

Table: Mean and standard deviation of final model accuracy percentage on the test datasets.

Dataset	Model	Graph Accuracy	Orbit Accuracy	Node Accuracy
Bioisostere (cross-entropy loss)	GCN	52.4 ± 6.37	92.9 ± 1.14	94.4 ± 0.79
	Unique-ID-GCN	66.1 ± 5.13	94.5 ± 0.97	95.6 ± 0.62
	RNI-GCN	63.6 ± 4.29	93.9 ± 0.86	95.1 ± 0.76
	Orbit-Indiv-GCN	$\mathbf{69.9 \pm 4.68}$	$\mathbf{95.4 \pm 0.63}$	$\mathbf{96.3 \pm 0.49}$
	2-Orbit-Transform	57.1 ± 6.43	93 ± 0.99	94.1 ± 0.79
Alchemy-Max-Orbit-2 (orbit-sorting cross-entropy)	Unique-ID-GCN	20 ± 4.57	79.9 ± 2.01	77 ± 1.52
	RNI-GCN	0 ± 0	74.5 ± 1.7	75.2 ± 1.66
	Orbit-Indiv-GCN	$\mathbf{51.9 \pm 4.38}$	$\mathbf{87.5 \pm 1.78}$	$\mathbf{90.6 \pm 1.12}$
	2-Orbit-Transform	47.9 ± 6.45	86.8 ± 1.53	85.1 ± 1.66
Alchemy-Max-Orbit-6 (orbit-sorting cross-entropy)	Unique-ID-GCN	66.8 ± 7.15	84.8 ± 2.97	95.4 ± 1.07
	RNI-GCN	44.9 ± 7.19	78.5 ± 3.39	91.4 ± 1.47
	Orbit-Indiv-GCN	$\mathbf{83.4 \pm 4.22}$	$\mathbf{88.9 \pm 2.71}$	$\mathbf{97.1 \pm 1.46}$
	6-Orbit-Transform	10.6 ± 4.14	71.2 ± 2.47	87.6 ± 1.08

Results

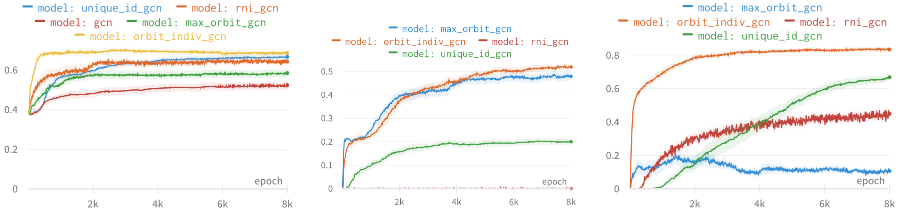

Figure: Graph accuracy with standard error on the test datasets across all models using orbit-sorting cross-entropy: Bioisostere (left), Alchemy-Max-Orbit-2 (center), and Alchemy-Max-Orbit-6 (right).

Future Work

- Orbit-equivariance can be generalized to other data structures besides graphs.
- Designing optimal bioisosteres is complex and deserves further investigation.
- Identification of other problems that require non-equivariant and orbit-equivariant models to solve.
- Finally, design better orbit-equivariant GNNs to solve such problems.