Orbit-Equivariant Graph Neural Networks

Matthew Morris, Bernardo Cuenca Grau, Ian Horrocks

University of Oxford

Equivariant Graph Functions

A node-labelling function f assigns output labels to each node. Node-labelling function f on domain D closed under permutation is *equivariant* if $f(\sigma \cdot G) = \sigma \cdot f(G)$ holds for all labelled graphs $G \in D$ and all permutations σ on $V(G)$. Graph Neural Networks (GNNs) are equivariant node-labelling functions.

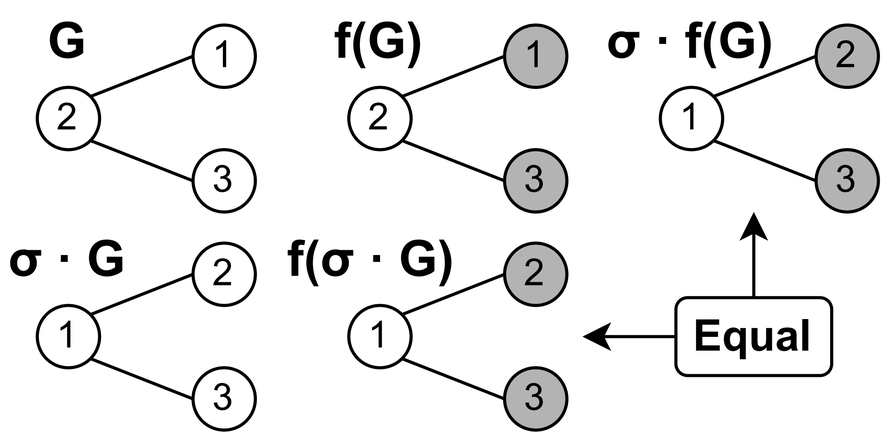

Figure 1: Example of an equivariant node-labelling function f.

The Limits of Equivariance

Nodes equivalent under automorphism are called *similar* and the equivalence classes *orbits*. If GNNs are used to transform (graphical representations of) molecules so as to optimize for some desired molecular property, equivariant functions cannot transform similar atoms in different ways.

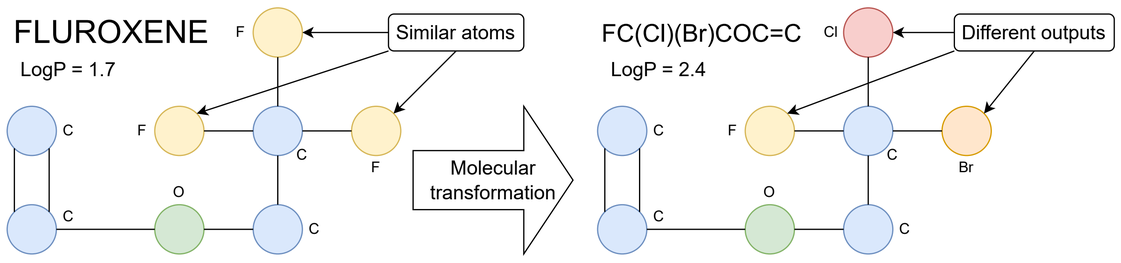

Figure 2: Non-equivariant molecular transformation that increases lipophilicity.

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 $f(G)_v = f(G)_w$.*

We require a function that yields a multiset of output labels for each orbit, so a less-restrictive property than equivariance is required.

Orbit-Equivariance

To retain the desirable inductive bias of equivariant functions, we require that no matter how the input graph is permuted, the multiset of outputs for each particular graph orbit should not change.

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 $\{\!\{f(\sigma \cdot G)_{\sigma(v)} \mid v \in r\}\!\} = \{\!\{f(G)_v \mid v \in r\}\!\}$.*

Not all node-labelling functions are orbit-equivariant:

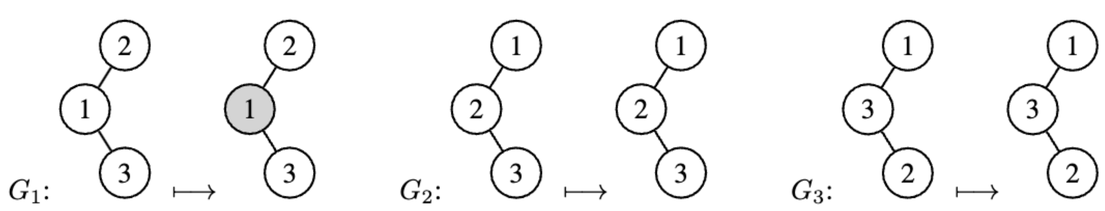

Figure 3: Non-orbit-equivariant node-labelling function.

Proposition 2. *All equivariant functions are orbit-equivariant, but there are orbit-equivariant functions which are not equivariant.*

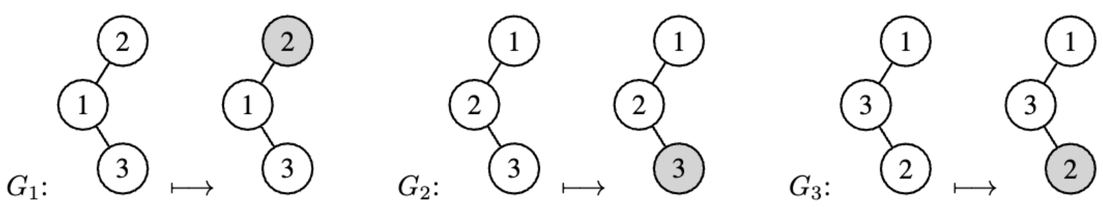

Figure 4: Non-equivariant but orbit-equivariant node-labelling function.

In summary: If $\mathcal{E}$, $\mathcal{O}$, and $\mathcal{U}$ are, respectively, the sets of all equivariant, orbit-equivariant, and universal node-labelling functions, then we have $\mathcal{E} \subset \mathcal{O} \subset \mathcal{U}$.

Classifying Orbit-Equivariant Functions

Definition 2. *For a node-labelling function f on domain D, we define $\texttt{max-orbit}(f)$ to be the maximum across all $G \in D$ and orbits $r \in R(G)$ of the number of unique values in $\{\!\{f(G)_v \mid v \in r\}\!\}$.*

For any equivariant function f, $\texttt{max-orbit}(f) = 1$. For any $m \in \mathbb{Z}^+$, there is an orbit-equivariant function f such that $\texttt{max-orbit}(f) = m$, since there are graphs (e.g. cycle graphs) with arbitrarily large orbits.

Proposition 3. *If f is orbit-equivariant and $\texttt{max-orbit}(f) = 1$, then f is equivariant.*

The 1-WL algorithm (computation equivalent to GNNs) can be used to compute graph orbit classes by computing iterative recolorings in the usual way, and using the final stable colorings to partition the nodes into their orbits. When 1-WL is used in this manner, we refer to it as *orbit-1-WL*. Orbit-1-WL is sound but not complete.

Theorem 1. *If orbit-1-WL colors $u, v \in V(G)$ differently, then u, v are not similar.*

Orbit-Equivariant GNNs

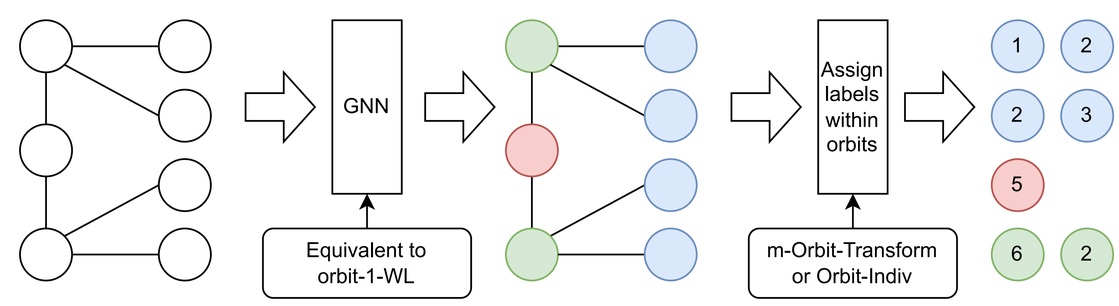

Figure 5: Architecture for constructing orbit-equivariant GNNs.

Theorem 2. *The expressivity of our proposed GNN models is as follows:*

1. *Unique-ID-GNNs are not orbit-equivariant and can approximate any node-labelling function $f : G_n \to \mathbb{R}^n$, where domain G_n is the set of all graphs with $\leq n$ nodes.*
2. *RNI-GNNs are equivariant in expectation and 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.*
3. *Orbit-Indiv-GNNs are not equivariant but are orbit-equivariant on graphs whose orbits are distinguishable by orbit-1-WL.*
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 $\texttt{max-orbit}(f) = m$.*

Experiments

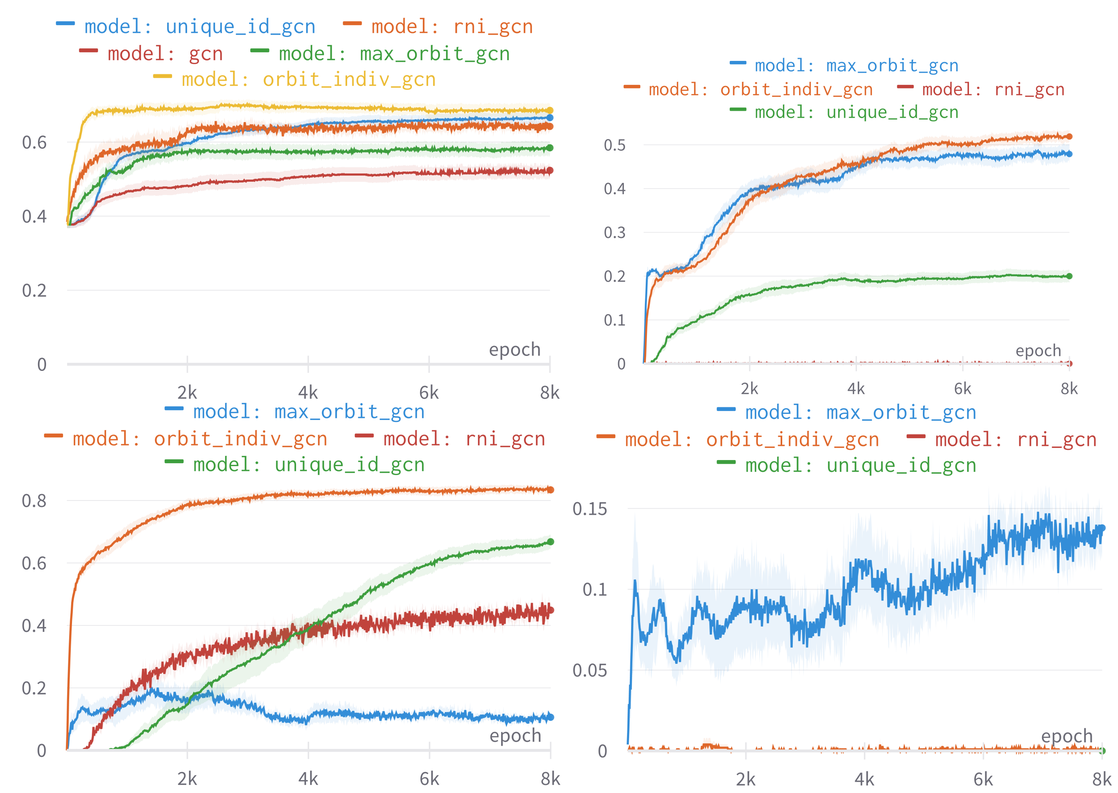

Figure 6: Graph accuracy with standard error on the test datasets across all models using orbit-sorting cross-entropy: Bioisostere (top left), Alchemy-Max-Orbit-2 (top right), Alchemy-Max-Orbit-6 (bottom left), and Alchemy-Max-Orbit-6 using cross-entropy loss (bottom right).