
Weisfeiler and Leman go sparse: Towards scalable higher-order graph embeddings

Christopher Morris*

Gaurav Rattan†

Petra Mutzel‡

Abstract

Graph kernels based on the 1-dimensional Weisfeiler-Leman algorithm and corresponding neural architectures recently emerged as powerful tools for (supervised) learning with graphs. However, due to the purely local nature of the algorithms, they might miss essential patterns in the given data and can only handle binary relations. The k -dimensional Weisfeiler-Leman algorithm addresses this by considering k -tuples, defined over the set of vertices, and defines a suitable notion of adjacency between these vertex tuples. Hence, it accounts for the higher-order interactions between vertices. However, it does not scale and may suffer from overfitting when used in a machine learning setting. Hence, it remains an important open problem to design WL-based graph learning methods that are simultaneously expressive, scalable, and non-overfitting. Here, we propose *local* variants and corresponding neural architectures, which consider a subset of the original neighborhood, making them more scalable, and less prone to overfitting. The expressive power of (one of) our algorithms is strictly higher than the original algorithm, in terms of ability to distinguish non-isomorphic graphs. Our experimental study confirms that the local algorithms, both kernel and neural architectures, lead to vastly reduced computation times, and prevent overfitting. The kernel version establishes a new state-of-the-art for graph classification on a wide range of benchmark datasets, while the neural version shows promising performance on large-scale molecular regression tasks.

1 Introduction

Graph-structured data is ubiquitous across application domains ranging from chemo- and bioinformatics [10, 103] to image [101] and social network analysis [27]. To develop successful machine learning models in these domains, we need techniques that can exploit the rich information inherent in the graph structure, as well as the feature information contained within nodes and edges. In recent years, numerous approaches have been proposed for machine learning with graphs—most notably, approaches based on *graph kernels* [71] or using *graph neural networks* (GNNs) [19, 45, 47]. Here, graph kernels based on the *1-dimensional Weisfeiler-Leman algorithm* (1-WL) [46, 111], and corresponding GNNs [83, 115] have recently advanced the state-of-the-art in supervised node and graph learning. Since the 1-WL operates via simple neighborhood aggregation, the purely local nature of these approaches can miss important patterns in the given data. Moreover, they are only applicable to binary structures, and therefore cannot deal with general t -ary structures, e.g., hypergraphs [124] or subgraphs, in a straight-forward way. A provably more powerful algorithm (for graph isomorphism testing) is the *k -dimensional Weisfeiler-Leman algorithm* (k -WL) [15, 46, 77]. The algorithm can capture more global, higher-order patterns by iteratively computing a coloring (or discrete labeling) for k -tuples, instead of single vertices, based on an appropriately defined notion of adjacency between two k -tuples. However, it fixes the cardinality of this neighborhood to $k \cdot n$, where n denotes the number of

*CERC in Data Science for Real-Time Decision-Making, Polytechnique Montréal

†Department of Computer Science, RWTH Aachen University

‡Department of Computer Science, University of Bonn

vertices of a given graph. Hence, the running time of each iteration does not take the *sparsity* of a given graph into account. Further, new neural architectures [77, 78] that possess the same power as the k -WL in terms of separating non-isomorphic graphs suffer from the same drawbacks, i.e., they have to resort to dense matrix multiplications. Moreover, when used in a machine learning setting with real-world graphs, the k -WL may capture the isomorphism type, which is the complete structural information inherent in a graph, after only a couple of iterations, which may lead to overfitting, see [82], and the experimental section of the present work.

Present work To address this, we propose a *local* version of the k -WL, the *local δ - k -dimensional Weisfeiler-Leman algorithm* (δ - k -LWL), which considers a subset of the original neighborhood in each iteration. The cardinality of the *local neighborhood* depends on the sparsity of the graph, i.e., the degrees of the vertices of a given k -tuple. We theoretically analyze the strength of a variant of our local algorithm and prove that it is strictly more powerful in distinguishing non-isomorphic graphs compared to the k -WL. Moreover, we devise a hierarchy of pairs of non-isomorphic graphs that a variant of the δ - k -LWL can separate while the k -WL cannot. On the neural side, we devise a higher-order graph neural network architecture, the δ - k -LGNN, and show that it has the same expressive power as the δ - k -LWL. Moreover, we connect it to recent advancements in learning theory for GNNs [41], which show that the δ - k -LWL architecture has better generalization abilities compared to dense architectures based on the k -WL. See Figure 1 for an overview of the proposed algorithms.

Experimentally, we apply the discrete algorithms (or kernels) and the (local) neural architectures to supervised graph learning, and verify that both are several orders of magnitude faster than the global, discrete algorithms or dense, neural architectures, and prevent overfitting. The discrete algorithms establish a new state-of-the-art for graph classification on a wide range of small- and medium-scale classical datasets. The neural version shows promising performance on large-scale molecular regression tasks.

Related work In the following, we review related work from graph kernels and GNNs. We refer to Appendix A for an in-depth discussion of related work, as well as a discussion of theoretical results for the k -WL.

Historically, kernel methods—which implicitly or explicitly map graphs to elements of a Hilbert space—have been the dominant approach for supervised learning on graphs. Important early work in this area includes kernels based on random-walks [42, 60, 70], shortest paths [13], and kernels based on the 1-WL [100]. Morris et al. [82] devised a local, set-based variant of the k -WL. However, the approach is (provably) weaker than the tuple-based algorithm, and they do not prove convergence to the original algorithm. For a thorough survey of graph kernels, see [71]. Recently, graph neural networks (GNNs) [45, 97] emerged as an alternative to graph kernels. Notable instances of this architecture include, e.g., [33, 51, 105], and the spectral approaches proposed in, e.g., [14, 29, 64, 81]—all of which descend from early work in [65, 80, 102, 97]. A survey of recent advancements in GNN techniques can be found, e.g., in [19, 113, 125]. Recently, connections to Weisfeiler-Leman type algorithms have been shown [11, 24, 43, 44, 75, 77, 83, 115]. Specifically, the authors of [83, 115] showed that the expressive power of any possible GNN architecture is limited by the 1-WL in terms of distinguishing non-isomorphic graphs. Morris et al. [83] introduced *k -dimensional GNNs* (k -GNN) which rely on a message-passing scheme between subgraphs of cardinality k . Similar to [82], the paper employed a local, set-based (neural) variant of the k -WL, which is (provably) weaker than the variant considered here. Later, this was refined in [77] by introducing *k -order invariant graph networks* (k -IGN), based on Maron et al. [78], which are equivalent to the folklore variant of the k -WL [44, 46] in terms of distinguishing non-isomorphic graphs. However, k -IGN may not scale since they rely on dense linear algebra routines. Chen et al. [24] connect the theory of universal approximation of permutation-invariant functions and the graph isomorphism viewpoint and introduce a variation of the 2-WL, which is more powerful than the former. Our comprehensive treatment of higher-order, sparse, (graph) neural networks for arbitrary k subsumes all of the algorithms and neural architectures mentioned above.

2 Preliminaries

We briefly describe the Weisfeiler-Leman algorithm and, along the way, introduce our notation, see Appendix B for expanded preliminaries. As usual, let $[n] = \{1, \dots, n\} \subset \mathbb{N}$ for $n \geq 1$, and let $\{\dots\}$ denote a multiset. We also assume elementary definitions from graph theory (such as

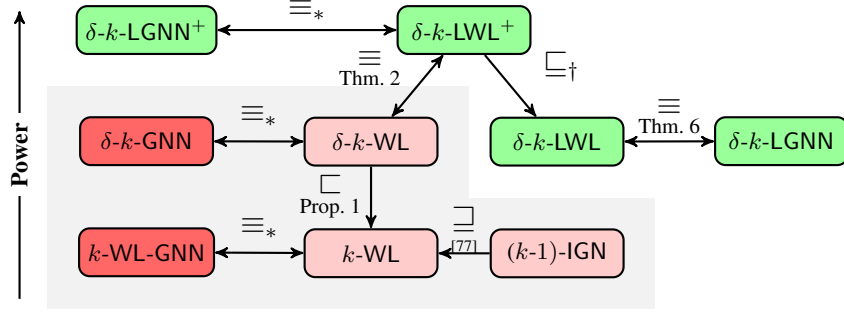


Figure 1: Overview of the power of proposed algorithms and neural architectures. The green and dark red nodes represent algorithms proposed in the present work. The grey region groups dense algorithms and neural architectures.

*—Follows directly from the proof of Theorem 6. $A \sqsubseteq B$ ($A \sqsubset B$, $A \equiv B$): algorithm A is more powerful (strictly more powerful, equally powerful) than B , \dagger —Follows by definition, strictness open.

graphs, directed graphs, vertices, edges, neighbors, trees, and so on). The vertex and the edge set of a graph G are denoted by $V(G)$ and $E(G)$ respectively. The *neighborhood* of v in $V(G)$ is denoted by $\delta(v) = N(v) = \{u \in V(G) \mid (v, u) \in E(G)\}$. Moreover, its complement $\bar{\delta}(v) = \{u \in V(G) \mid (v, u) \notin E(G)\}$. We say that two graphs G and H are *isomorphic* ($G \simeq H$) if there exists an adjacency preserving bijection $\varphi: V(G) \rightarrow V(H)$, i.e., (u, v) is in $E(G)$ if and only if $(\varphi(u), \varphi(v))$ is in $E(H)$, call φ an *isomorphism* from G to H . If the graphs have vertex/edges labels, the isomorphism is additionally required to match these labels. A *rooted tree* is a tree with a designated vertex called *root* in which the edges are directed in such a way that they point away from the root. Let p be a vertex in a directed tree then we call its out-neighbors *children* with parent p . Given a k -tuple of vertices $\mathbf{v} = (v_1, \dots, v_k)$, let $G[\mathbf{v}]$ denote the subgraph induced on the set $\{v_1, \dots, v_k\}$, where, the vertex v_i is labeled with i , for i in $[k]$.

Vertex refinement algorithms For a fixed positive integer k , let $V(G)^k$ denote the set of k -tuples of vertices of G . A *coloring* of $V(G)^k$ is a mapping $C: V(G)^k \rightarrow \mathbb{N}$, i.e., we assign a number (or color) to every tuple in $V(G)^k$. The *initial coloring* C_0 of $V(G)^k$ is specified by the isomorphism types of the tuples, i.e., two tuples \mathbf{v} and \mathbf{w} in $V(G)^k$ get a common color iff the mapping $v_i \rightarrow w_i$ induces an isomorphism between the labeled subgraphs $G[\mathbf{v}]$ and $G[\mathbf{w}]$. A *color class* corresponding to a color c is the set of all tuples colored c , i.e., the set $C^{-1}(c)$. For j in $[k]$, let $\phi_j(\mathbf{v}, w)$ be the k -tuple obtained by replacing the j^{th} component of \mathbf{v} with the vertex w . That is, $\phi_j(\mathbf{v}, w) = (v_1, \dots, v_{j-1}, w, v_{j+1}, \dots, v_k)$. If $\mathbf{w} = \phi_j(\mathbf{v}, w)$ for some w in $V(G)$, call \mathbf{w} a *j -neighbor* of \mathbf{v} (and vice-versa). The *neighborhood* of \mathbf{v} is then defined as the set of all tuples \mathbf{w} such that $\mathbf{w} = \phi_j(\mathbf{v}, w)$ for some j in $[k]$ and w in $V(G)$. The *refinement* of a coloring $C: V(G)^k \rightarrow \mathbb{N}$, denoted by \hat{C} , is a coloring $\hat{C}: V(G)^k \rightarrow \mathbb{N}$ defined as follows. For each j in $[k]$, collect the colors of the j -neighbors of \mathbf{v} as a multiset $S_j = \{\{C(\phi_j(\mathbf{v}, w)) \mid w \in V(G)\}\}$. Then, for a tuple \mathbf{v} , define $\hat{C}(\mathbf{v}) = (C(\mathbf{v}), M(\mathbf{v}))$, where $M(\mathbf{v})$ is the k -tuple (S_1, \dots, S_k) . For consistency, the strings $\hat{C}(\mathbf{v})$ thus obtained are lexicographically sorted and renamed as integers. Observe that the new color $\hat{C}(\mathbf{v})$ of \mathbf{v} is solely dictated by the color histogram of its neighborhood and the previous color of \mathbf{v} . In general, a different mapping $M(\cdot)$ could be used, depending on the neighborhood information that we would like to aggregate.

The k -dim. Weisfeiler-Leman For $k \geq 2$, the k -WL computes a coloring $C_\infty: V(G)^k \rightarrow \mathbb{N}$ of a given graph G , as follows.⁴ To begin with, the initial coloring C_0 is computed. Then, starting with C_0 , successive refinements $C_{i+1} = \hat{C}_i$ are computed until convergence. That is,

$$C_{i+1}(\mathbf{v}) = (C_i(\mathbf{v}), M_i(\mathbf{v})),$$

where

$$M_i(\mathbf{v}) = (\{\{C_i(\phi_1(\mathbf{v}, w)) \mid w \in V(G)\}\}, \dots, \{\{C_i(\phi_k(\mathbf{v}, w)) \mid w \in V(G)\}\}). \quad (1)$$

The successive refinement steps are also called *rounds* or *iterations*. Since the disjoint union of the color classes form a partition of $V(G)^k$, there must exist a finite $\ell \leq |V(G)|^k$ such that $C_\ell = \hat{C}_\ell$. In

⁴We define the 1-WL in the next subsection.

the end, the k -WL outputs C_ℓ as the *stable coloring* C_∞ . The k -WL *distinguishes* two graphs G and H if, upon running the k -WL on their disjoint union $G \cup H$, there exists a color c in \mathbb{N} in the stable coloring such that the corresponding color class S_c satisfies $|V(G)^k \cap S_c| \neq |V(H)^k \cap S_c|$, i.e., there exist an unequal number of c -colored tuples in $V(G)^k$ and $V(H)^k$. Hence, two graphs distinguished by the k -WL must be non-isomorphic. See Appendix C for its relation to the folklore k -WL.

The δ - k -dim. Weisfeiler-Leman Let $\mathbf{w} = \phi_j(\mathbf{v}, w)$ be a j -neighbor of \mathbf{v} . Call \mathbf{v} a *local* j -neighbor of \mathbf{w} if w is adjacent to the replaced vertex v_j . Otherwise, call \mathbf{v} a *global* j -neighbor of \mathbf{w} . For tuples \mathbf{v} and \mathbf{w} in $V(G)^k$, let the function $\text{adj}((\mathbf{v}, \mathbf{w}))$ evaluate to L or G, depending on whether \mathbf{w} is a local or a global neighbor, respectively, of \mathbf{v} . The δ - k -dimensional *Weisfeiler-Leman algorithm*, denoted by δ - k -WL, is a variant of the classic k -WL which *differentiates* between the local and the global neighbors during neighborhood aggregation [76]. Formally, the δ - k -WL algorithm refines a coloring C_i (obtained after i rounds) via the aggregation function

$$M_i^{\delta, \bar{\delta}}(\mathbf{v}) = (\{\{C_i(\phi_1(\mathbf{v}, w), \text{adj}(\mathbf{v}, \phi_1(\mathbf{v}, w))) \mid w \in V(G)\}\}, \dots, \{\{C_i(\phi_k(\mathbf{v}, w), \text{adj}(\mathbf{v}, \phi_k(\mathbf{v}, w))) \mid w \in V(G)\}\}), \quad (2)$$

instead of the k -WL aggregation specified by Equation (1). We define the 1-WL to be the δ -1-WL, which is commonly known as color refinement or naive vertex classification.

Comparison of k -WL variants Let A_1 and A_2 denote two vertex refinement algorithms, we write $A_1 \sqsubseteq A_2$ if A_1 distinguishes between all non-isomorphic pairs A_2 does, and $A_1 \equiv A_2$ if both directions hold. The corresponding strict relation is denoted by \sqsubset . The following result shows that the δ - k -WL is strictly more powerful than the k -WL for $k \geq 2$ (see Appendix C.1.1 for the proof).

Proposition 1. For $k \geq 2$, the following holds:

$$\delta\text{-}k\text{-WL} \sqsubset k\text{-WL}.$$

3 Local δ - k -dimensional Weisfeiler-Leman algorithm

In this section, we define the new *local δ - k -dimensional Weisfeiler-Leman algorithm* (δ - k -LWL). This variant of the δ - k -WL considers only local neighbors during the neighborhood aggregation process, and discards any information about the global neighbors. Formally, the δ - k -LWL algorithm refines a coloring C_i (obtained after i rounds) via the aggregation function,

$$M_i^\delta(\mathbf{v}) = (\{\{C_i(\phi_1(\mathbf{v}, w)) \mid w \in N(v_1)\}\}, \dots, \{\{C_i(\phi_k(\mathbf{v}, w)) \mid w \in N(v_k)\}\}), \quad (3)$$

instead of Equation (2), hence considering only the local j -neighbors of the tuple \mathbf{v} in each iteration. The indicator function adj used in Equation (2) is trivially equal to L here, and is thus omitted. The coloring function for the δ - k -LWL is then defined by

$$C_{i+1}^{k, \delta}(\mathbf{v}) = (C_i^{k, \delta}(\mathbf{v}), M_i^\delta(\mathbf{v})).$$

We also define the δ - k -LWL⁺, a minor variation of the δ - k -LWL. Later, we will show that the δ - k -LWL⁺ is equivalent in power to the δ - k -WL (Theorem 2). Formally, the δ - k -LWL⁺ algorithm refines a coloring C_i (obtained after i rounds) via the aggregation function,

$$M^{\delta, +}(\mathbf{v}) = (\{\{C_i(\phi_1(\mathbf{v}, w)), \#_i^1(\mathbf{v}, \phi_1(\mathbf{v}, w))\} \mid w \in N(v_1)\}\}, \dots, \{\{C_i(\phi_k(\mathbf{v}, w)), \#_i^k(\mathbf{v}, \phi_k(\mathbf{v}, w))\} \mid w \in N(v_k)\}\}), \quad (4)$$

instead of δ - k -LWL aggregation defined in Equation (3). Here, the function

$$\#_i^j(\mathbf{v}, \mathbf{x}) = |\{\mathbf{w} : \mathbf{w} \sim_j \mathbf{v}, C_i(\mathbf{w}) = C_i(\mathbf{x})\}|, \quad (5)$$

where $\mathbf{w} \sim_j \mathbf{v}$ denotes that \mathbf{w} is j -neighbor of \mathbf{v} , for j in $[k]$. Essentially, $\#_i^j(\mathbf{v}, \mathbf{x})$ counts the number of j -neighbors (local or global) of \mathbf{v} which have the same color as \mathbf{x} under the coloring C_i (i.e., after i rounds). For a fixed \mathbf{v} , the function $\#_i^j(\mathbf{v}, \cdot)$ is uniform over the set $S \cap N_j$, where S is a color class obtained after i iterations of the δ - k -LWL⁺ and N_j denotes the set of j -neighbors of \mathbf{v} . Note that after the stable partition has been reached $\#_i^j(\mathbf{v})$ will not change anymore. Intuitively, this variant captures local and to some extent global information, while still taking the sparsity of the underlying graph into

account. Moreover, observe that each iteration of the δ - k -LWL⁺ has the same asymptotic running time as an iteration of the δ - k -LWL, and that the information of the # function is already implicitly contained in Equation (2).

The following theorem shows that the local variant δ - k -LWL⁺ is at least as powerful as the δ - k -WL when restricted to the class of connected graphs. The possibly slower convergence leads to advantages in a machine learning setting, see Sections 4 and 6, and also Section 5 for a discussion of practicality, running times, and remaining challenges.

Theorem 2. For the class of connected graphs, the following holds for all $k \geq 1$:

$$\delta$$
- k -LWL⁺ \equiv δ - k -WL.

Along with Proposition 1, this establishes the superiority of the δ - k -LWL⁺ over the k -WL.

Corollary 3. For the class of connected graphs, the following holds for all $k \geq 2$:

$$\delta$$
- k -LWL⁺ \sqsubseteq k -WL.

In fact, the proof of Proposition 1 shows that the infinite family of graphs G_k, H_k witnessing the strictness condition can even be distinguished by the δ - k -LWL, for each corresponding $k \geq 2$. We note here that the restriction to connected graphs can easily be circumvented by adding a specially marked vertex, which is connected to every other vertex in the graph.

Kernels based on vertex refinement algorithms After running the δ - k -LWL (and the other vertex refinements algorithms), the concatenation of the histogram of colors in each iteration can be used as a feature vector in a kernel computation. Specifically, in the histogram for every color c in \mathbb{N} there is an entry containing the number of nodes or k -tuples that are colored with c .

Local converges to global: proof of Theorem 2 The main technique behind the proof is to construct tree-representations of the colors assigned by the k -WL (or its variants). Given a graph G , a tuple \mathbf{v} , and an integer $\ell \geq 0$, the *unrolling tree* of the graph G at \mathbf{v} of depth ℓ is a rooted directed tree UNR $[G, \mathbf{s}, \ell]$ (with vertex and edge labels) which encodes the color assigned by k -WL to the tuple \mathbf{v} after ℓ rounds, see Appendix D.2 for a formal definition and Figure 4 for an illustration. The usefulness of these tree representations is established by the following lemma. Formally, let \mathbf{s} and \mathbf{t} be two k -vertex-tuples in $V(G)^k$.

Lemma 4. The colors of \mathbf{s} and \mathbf{t} after ℓ rounds of k -WL are identical if and only if the unrolling tree UNR $[G, \mathbf{s}, \ell]$ is isomorphic to the unrolling tree UNR $[G, \mathbf{t}, \ell]$.

For different k -WL variants, the construction of these unrollings are slightly different, since an unrolling tree needs to faithfully represent the corresponding aggregation process for generating new colors. For the variants δ - k -WL, δ - k -LWL, and δ - k -LWL⁺, we define respective unrolling trees δ -UNR $[G, \mathbf{s}, \ell]$, L-UNR $[G, \mathbf{s}, \ell]$, and L⁺-UNR $[G, \mathbf{s}, \ell]$ along with analogous lemmas, as above, stating their correctness/usefulness. Finally, we show that for *connected graphs*, the δ -UNR unrolling trees (of sufficiently large depth) at two tuples \mathbf{s} and \mathbf{t} are identical only if the respective δ - k -LWL⁺ unrolling trees (of sufficiently larger depth) are identical, as shown in the following lemma.

Lemma 5. Let G be a connected graph, and let \mathbf{s} and \mathbf{t} in $V(G)^k$. If the stable colorings of \mathbf{s} and \mathbf{t} under δ - k -LWL⁺ are identical, then the stable colorings of \mathbf{s} and \mathbf{t} under δ - k -WL are also identical.

Hence, the local algorithm δ - k -LWL⁺ is at least as powerful as the global δ - k -WL, for connected graphs, i.e., δ - k -LWL⁺ \sqsubseteq δ - k -WL. The exact details and parameters of this proof can be found in the Appendix.

4 Higher-order neural architectures

Although the discrete kernels defined in the previous section are quite powerful, they are limited due to their fixed feature construction scheme, hence suffering from poor adaption to the learning task at hand and the inability to handle continuous node and edge labels in a meaningful way. Moreover, they often result in high-dimensional embeddings forcing one to resort to non-scalable, kernelized optimization procedures. This motivates our definition of a new neural architecture, called *local δ - k -GNN* (δ - k -LGNN). Given a labeled graph G , let each tuple \mathbf{v} in $V(G)^k$ be annotated with an

initial feature $f^{(0)}(\mathbf{v})$ determined by its isomorphism type. In each layer $t > 0$, we compute a new feature $f^{(t)}(\mathbf{v})$ as

$$f_{\text{mrg}}^{W_1} \left(f^{(t-1)}(\mathbf{v}), f_{\text{agg}}^{W_2} \left(\{ \{ f^{(t-1)}(\phi_1(\mathbf{v}, w)) \mid w \in \delta(v_1) \} \}, \dots, \{ \{ f^{(t-1)}(\phi_k(\mathbf{v}, w)) \mid w \in \delta(v_k) \} \} \right) \right),$$

in $\mathbb{R}^{1 \times e}$ for a tuple \mathbf{v} , where $W_1^{(t)}$ and $W_2^{(t)}$ are learnable parameter matrices from $\mathbb{R}^{d \times e}$.⁵ Moreover, $f_{\text{mrg}}^{W_2}$ and the permutation-invariant $f_{\text{agg}}^{W_1}$ can be arbitrary (permutation-invariant) differentiable functions, responsible for merging and aggregating the relevant feature information, respectively. Initially, we set $f^{(0)}(\mathbf{v})$ to a one-hot encoding of the (labeled) isomorphism type of $G[\mathbf{v}]$. Note that we can naturally handle discrete node and edge labels as well as directed graphs, see Section 4 on how to deal with continuous information. The following result demonstrates the expressive power of the δ - k -GNN, in terms of distinguishing non-isomorphic graphs.

Theorem 6. Let (G, l) be a labeled graph. Then for all $t \geq 0$ there exists a sequence of weights $\mathbf{W}^{(t)}$ such that

$$C_t^{k, \delta}(\mathbf{v}) = C_t^{k, \delta}(\mathbf{w}) \iff f^{(t)}(\mathbf{v}) = f^{(t)}(\mathbf{w}).$$

Hence, for all graphs, the following holds for all $k \geq 1$:

$$\delta\text{-}k\text{-LGNN} \equiv \delta\text{-}k\text{-LWL}.$$

Moreover, the δ - k -GNN inherits the main strength of the δ - k -LWL, i.e., it can be implemented using sparse matrix multiplication. Note that it is not possible to come up with an architecture, i.e., instantiations of $f_{\text{mrg}}^{W_1}$ and $f_{\text{agg}}^{W_2}$, such that it becomes more powerful than the δ - k -LWL, see [83]. However, all results from the previous section can be lifted to the neural setting. That is, one can derive neural architectures based on the δ - k -LWL⁺, δ - k -WL, and k -WL, called δ - k -LGNN⁺, δ - k -GNN, and k -WL-GNN, respectively, and prove results analogous to Theorem 6.

Incorporating continuous information Since many real-world graphs, e.g., molecules, have continuous features (real-valued vectors) attached to vertices and edges, using a one-hot encoding of the (labeled) isomorphism type is not a sensible choice. Let $a: V(G) \rightarrow \mathbb{R}^{1 \times d}$ be a function such that each vertex v is annotated with a feature $a(v)$ in $\mathbb{R}^{1 \times d}$, and let $\mathbf{v} = (v_1, \dots, v_k)$ be a k -tuple of vertices. Then we can compute an initial feature

$$f^{(0)}(\mathbf{v}) = f_{\text{enc}}^{W_3}((a(v_1), \dots, a(v_k))), \quad (6)$$

for the tuple \mathbf{v} . Here, $f_{\text{enc}}: (\mathbb{R}^{1 \times d})^k \rightarrow \mathbb{R}^{1 \times e}$ is an arbitrary differentiable, parameterized function, e.g., a multi-layer perceptron or a standard GNN aggregation function, that computes a joint representation of the k node features $a(v_1), \dots, a(v_k)$. Moreover, it is also straightforward to incorporate the labeled isomorphism type and continuous edge label information. We further explore this in the experimental section.

Generalization abilities of the neural architecture Garg et al. [41], studied the generalization abilities of a standard GNN architecture for binary classification using a margin loss. Under mild conditions, they bounded the empirical Rademacher complexity as $\tilde{O}(rdL/\sqrt{m}\gamma)$, where d is the maximum degree of the employed graphs, r is the number of components of the node features, L is the number of layers, and γ is a parameter of the loss function. It is straightforward to transfer the above bound to the higher-order (local) layer from above. Hence, this shows that local, sparsity-aware, higher-order variants, e.g., δ - k -LGNN, exhibit a smaller generalization error compared to dense, global variants like the k -WL-GNN.

5 Practicality, barriers ahead, and possible road maps

As Theorem 2 shows, the δ - k -LWL⁺ and its corresponding neural architecture, the δ - k -LGNN⁺, have the same power in distinguishing non-isomorphic graphs as δ - k -WL. Although for dense graphs, the local algorithms will have the same running time, for sparse graphs, the running time for each iteration can be upper-bounded by $|n^k| \cdot kd$, where d denotes the maximum or average degree of the graph. Hence, the local algorithm takes the sparsity of the underlying graph into account, resulting in improved computation times compared to the non-local δ - k -WL and the k -WL (for the same number

⁵For clarity of presentation we omit biases.

of iterations). These observations also translate into practice, see Section 6. The same arguments can be used in favor of the δ - k -LWL and δ - k -LGNN, which lead to even sparser algorithms.

Obstacles The biggest obstacle in applying the algorithms to truly large graphs is the fact that the algorithm considers all possible k -tuples leading to a lower bound on the running time of $\Omega(n^k)$. Lifting the results to the folklore k -WL, e.g., [77], only “shaves off one dimension”. Moreover, applying higher-order algorithms for large k might lead to overfitting issues, see also Section 6.

Possible solutions Recent sampling-based approaches for graph kernels or GNNs, see, e.g., [22, 23, 51, 54, 82] address the dependence on n^k , while appropriate pooling methods along the lines of Equation (6) address the overfitting issue. Finally, new directions from the theory community, e.g., [50] point further directions, which might result in more scalable algorithms.

6 Experimental evaluation

Our intention here is to investigate the benefits of the local, sparse algorithms, both kernel and neural architectures, compared to the global, dense algorithms, and standard kernel and GNN baselines. More precisely, we address the following questions:

Q1 Do the local algorithms, both kernel and neural architectures, lead to improved classification and regression scores on real-world benchmark datasets compared to global, dense algorithms and standard baselines?

Q2 Does the δ - k -LWL⁺ lead to improved classification accuracies compared to the δ - k -LWL? Does it lead to higher computation times?

Q3 Do the local algorithms prevent overfitting to the training set?

Q4 How much do the local algorithms speed up the computation time compared to the non-local algorithms or dense neural architectures?

The source code of all methods and evaluation procedures is available at <https://www.github.com/chrsmrts/sparsewl>.

Datasets To evaluate kernels, we use the following, well-known, small-scale datasets: ENZYMES [98, 13], IMDB-BINARY, IMDB-MULTI [119], NCI1, NCI109 [109], PTC_FM [53], PROTEINS [31, 13], and REDDIT-BINARY [119]. To show that our kernels also scale to larger datasets, we additionally used the mid-scale datasets: YEAST, YEASTH, UACC257, UACC257H, OVCAR-8, OVCAR-8H [117]. For the neural architectures, we used the large-scale molecular regression datasets ZINC [34, 57] and ALCHEMY [21]. To further compare to the (hierarchical) k -GNN [83] and k -IGN [77], and show the benefits of our architecture in presence of continuous features, we used the QM9 [91, 112] regression dataset.⁶ All datasets can be obtained from <http://www.graphlearning.io> [84]. See Appendix E.1 for further details.

Kernels We implemented the δ - k -LWL, δ - k -LWL⁺, δ - k -WL, and k -WL kernel for k in $\{2, 3\}$. We compare our kernels to the Weisfeiler-Leman subtree kernel (1-WL) [100], the Weisfeiler-Leman Optimal Assignment kernel (WLOA) [68], the graphlet kernel (GR) [99], and the shortest-path kernel [13] (SP). All kernels were (re-)implemented in C++11. For the graphlet kernel, we counted (labeled) connected subgraphs of size three. We followed the evaluation guidelines outlined in [84]. We also provide precomputed Gram matrices for easier reproducibility.

Neural architectures We used the GIN and GIN- ϵ architecture [115] as neural baselines. For data with (continuous) edge features, we used a 2-layer MLP to map them to the same number of components as the node features and combined them using summation (GINE and GINE- ϵ). For the evaluation of the neural architectures of Section 4, δ - k -LGNN, δ - k -GNN, and k -WL-GNN, we implemented them using PYTORCH GEOMETRIC [36], using a Python-wrapped C++11 preprocessing routine to compute the computational graphs for the higher-order GNNs.⁷ We used the GIN- ϵ layer to express $f_{\text{msg}}^{W_1}$ and $f_{\text{aggr}}^{W_2}$ of Section 4.

See Appendix E.2 for a detailed description of all evaluation protocols and hyperparameter selection routines.

Results and discussion In the following we answer questions **Q1** to **Q4**.

⁶We opted for comparing on the QM9 dataset to ensure a fair comparison concerning hyperparameter selection.

⁷We opted for not implementing the δ - k -LGNN⁺ as it would involve precomputing #.

Method	Dataset								
	ENZYMES	IMDB-BINARY	IMDB-MULTI	NCI1	NCI109	PTC_FM	PROTEINS	REDDIT-BINARY	
Baseline	GR	29.7 \pm 0.6	58.9 \pm 1.0	39.0 \pm 0.8	66.1 \pm 0.4	66.3 \pm 0.2	61.3 \pm 1.1	71.2 \pm 0.6	60.0 \pm 0.2
	SP	40.7 \pm 0.9	58.5 \pm 0.4	39.4 \pm 0.3	74.0 \pm 0.3	73.0 \pm 0.4	61.3 \pm 1.3	75.6 \pm 0.5	84.6 \pm 0.3
	1-WL	50.7 \pm 1.2	72.5 \pm 0.5	50.0 \pm 0.5	84.2 \pm 0.3	84.3 \pm 0.3	62.6 \pm 2.0	72.6 \pm 1.2	72.8 \pm 0.5
	WLOA	56.8 \pm 1.6	72.7 \pm 0.9	50.1 \pm 0.7	84.9 \pm 0.3	85.2 \pm 0.3	61.8 \pm 1.5	73.2 \pm 0.6	88.1 \pm 0.4
Neural	Gin-0	38.8 \pm 1.7	72.7 \pm 0.9	49.9 \pm 0.8	78.5 \pm 0.5	76.7 \pm 0.8	58.2 \pm 3.3	71.3 \pm 0.9	89.8 \pm 0.6
	Gin- ϵ	39.4 \pm 1.7	72.9 \pm 0.6	49.6 \pm 0.9	78.6 \pm 0.3	77.0 \pm 0.5	57.7 \pm 2.0	71.1 \pm 0.8	90.3 \pm 0.3
Global	2-WL	36.7 \pm 1.7	68.2 \pm 1.1	48.1 \pm 0.5	67.1 \pm 0.3	67.5 \pm 0.2	62.3 \pm 1.6	75.0 \pm 0.8	OOM
	3-WL	42.3 \pm 1.1	67.8 \pm 0.8	47.0 \pm 0.7	OOT	OOT	61.5 \pm 1.7	OOM	OOM
	δ -2-WL	37.5 \pm 1.2	68.1 \pm 1.1	47.9 \pm 0.7	67.0 \pm 0.5	67.2 \pm 0.4	61.9 \pm 0.9	75.0 \pm 0.4	OOM
	δ -3-WL	43.0 \pm 1.4	67.5 \pm 1.0	47.3 \pm 0.9	OOT	OOT	61.2 \pm 2.0	OOM	OOM
	δ -2-LWL	56.6 \pm 1.2	73.3 \pm 0.5	50.2 \pm 0.6	84.7 \pm 0.3	84.2 \pm 0.4	60.3 \pm 3.2	75.1 \pm 0.3	89.7 \pm 0.4
	δ -2-LWL ⁺	52.9 \pm 1.4	75.7 \pm 0.7	62.5 \pm 1.0	91.4 \pm 0.2	89.3 \pm 0.2	62.6 \pm 1.6	79.3 \pm 1.1	91.1 \pm 0.5
Local	δ -3-LWL	57.6 \pm 1.2	72.8 \pm 1.2	49.3 \pm 1.0	83.4 \pm 0.2	82.4 \pm 0.4	61.3 \pm 1.6	OOM	OOM
	δ -3-LWL ⁺	56.8 \pm 1.2	76.2 \pm 0.8	64.2 \pm 0.9	82.7 \pm 0.5	81.9 \pm 0.4	61.3 \pm 2.0	OOM	OOM

Table 1: Classification accuracies in percent and standard deviations, OOT— Computation did not finish within one day, OOM— Out of memory.

(a) Training versus test accuracy of local and global kernels for a subset of the datasets.

Set		Dataset		
		ENZYMES	IMDB-BINARY	IMDB-MULTI
δ -2-WL	Train	91.2	83.8	57.6
	Test	37.5	68.1	47.9
δ -2-LWL	Train	98.8	83.5	59.9
	Test	56.6	73.3	50.2
δ -2-LWL ⁺	Train	99.5	95.1	86.5
	Test	52.9	75.7	62.5

(b) Mean MAE (mean std. MAE, logMAE) on large-scale (multi-target) molecular regression tasks.

Method	Dataset	
	ZINC (FULL)	ALCHEMY (FULL)
Baseline		
GINE- ϵ	0.084 \pm 0.004	0.103 \pm 0.001 -2.956 \pm 0.029
2-WL-GNN	0.133 \pm 0.013	0.093 \pm 0.001 -3.394 \pm 0.035
δ -2-GNN	0.042 \pm 0.003	0.080 \pm 0.001 -3.516 \pm 0.021
δ -2-LGNN	0.045 \pm 0.006	0.083 \pm 0.001 -3.476 \pm 0.025

Table 2: Additional results for kernel and neural approaches.

A1 Kernels See Table 1. The local algorithm, for $k = 2$ and 3, severely improves the classification accuracy compared to the k -WL and the δ - k -WL. For example, on the ENZYMES dataset the δ -2-LWL achieves an improvement of almost 20%, and the δ -3-LWL achieves the best accuracies over all employed kernels, improving over the 3-WL and the δ -3-WL by more than 13%. This observation holds over all datasets. Our algorithms also perform better than neural baselines. See Table 5 in the appendix for additional results on the mid-scale datasets. However, it has to be noted that increasing k does not always result in increased accuracies. For example, on all datasets (excluding ENZYMES), the performance of the δ -2-LWL is better or on par with the δ -3-LWL. Hence, with increasing k the local algorithm is more prone to overfitting.


Neural architectures See Table 2b and Figure 2. On the ZINC and ALCHEMY datasets, the δ -2-LGNN is on par or slightly worse than the δ -2-GNN. Hence, this is in contrast to the kernel variant. We assume that this is due to the δ -2-GNN being more flexible than its kernel variant in weighing the importance of global and local neighbors. This is further highlighted by the worse performance of the 2-WL-GNN, which even performs worse than the (1-dimensional) GINE- ϵ on the ZINC dataset. On the QM9 dataset, see Figure 2a, the δ -2-LGNN performs better than the higher-order methods from [77, 83] while being on par with the MPNN architecture. We note here that the MPNN was specifically tuned to the QM9 dataset, which is not the case for the δ -2-LGNN (and the other higher-order architectures).

A2 See Table 1. The δ -2-LWL⁺ improves over the δ -2-LWL on all datasets excluding ENZYMES. For example, on IMDB-MULTI, NCI1, NCI109, and PROTEINS the algorithm achieves an improvement over of 4%, respectively, achieving a new state-of-the-art. The computation times are only increased slightly, see Table 8 in the appendix. Similar results can be observed on the mid-scale datasets, see Tables 5 and 9 in the appendix.


A3 Kernels As Table 2a (Table 6 for all datasets) shows the δ -2-WL reaches slightly higher training accuracies over all datasets compared to the δ -2-LWL, while the testing accuracies are much lower (excluding PTC_FM and PROTEINS). This indicates that the δ -2-WL overfits on the training set. The higher test accuracies of the local algorithm are likely due to the smaller neighborhood, which promotes that the number of colors grow slower compared to the global algorithm. The δ - k -LWL⁺

	Method	QM9
Baseline	GINE- ϵ	0.081 \pm 0.003
	MPNN	0.034 \pm 0.001
	1-2-GNN	0.068 \pm 0.001
	1-3-GNN	0.088 \pm 0.007
	1-2-3-GNN	0.062 \pm 0.001
	3-IGN	0.046 \pm 0.001
	δ -2-LGNN	0.029 \pm 0.001

(a) Mean std. MAE compared to [45, 77, 83].



(b) ZINC



(c) ALCHEMY

Figure 2: Additional results for neural architectures.

(a) Speed up ratios of local kernel computations for a subset of the datasets.

Method	Dataset		
	ENZYMES	IMDB-BINARY	NC11
Global 2-WL	10.4	3.6	14.3
δ -2-WL	10.1	3.6	14.5
δ -2-LWL ⁺	1.2	1.2	1.3
δ -2-LWL	1.0	1.0	1.0

(b) Average speed up ratios over all epochs (training and testing).

Method	Dataset	
	ZINC	ALCHEMY
Dense 2-WL-GNN	2.2	1.1
δ -2-GNN	2.5	1.7
GINE- ϵ	0.2	0.4
δ -2-LGNN	1.0	1.0

Table 3: Speed up ratios of local over global algorithms.

inherits the strengths of both algorithms, i.e., achieving the overall best training accuracies while achieving state-of-the-art testing accuracies.

Neural architectures See Figure 2. In contrast to the kernel variants, the 2-WL and the δ -2-WL, the corresponding neural architectures, the 2-WL-GNN and the δ -2-GNN, seem less prone to overfitting. However, especially on the ALCHEMY dataset, the δ -2-LGNN overfits less.

A4 Kernels See Table 3a (Tables 8 and 9 for all datasets). The local algorithm severely speeds up the computation time compared to the δ - k -WL and the k -WL for $k = 2$ and 3. For example, on the ENZYMES dataset the δ -2-LWL is over ten times faster than the δ -2-WL. The improvement of the computation times can be observed across all datasets. For some datasets, the $\{2, 3\}$ -WL and the δ - $\{2, 3\}$ -WL did not finish within the given time limit or went out of memory. For example, on four out of eight datasets, the δ -3-WL is out of time or out of memory. In contrast, for the corresponding local algorithm, this happens only two out of eight times. Hence, the local algorithm is more suitable for practical applications.

Neural architectures See Table 3b. The local algorithm severely speeds up the computation time of training and testing. Especially, on the ZINC dataset, which has larger graphs compared to the ALCHEMY dataset, the δ -2-LGNN achieves a computation time that is more than two times lower compared to the δ -2-GNN and the 2-WL-GNN.

7 Conclusion

We introduced local variants of the k -dimensional Weisfeiler-Leman algorithm. We showed that one variant and its corresponding neural architecture are strictly more powerful than the k -WL while taking the underlying graph’s sparsity into account. To demonstrate the practical utility of our findings, we applied them to graph classification and regression. We verified that our local, sparse algorithms lead to vastly reduced computation times compared to their global, dense counterparts while establishing new state-of-the-art results on a wide range of benchmark datasets. *We believe that our local, higher-order kernels and GNN architectures should become a standard approach in the regime of supervised learning with small graphs, e.g., molecular learning.*

Future work includes a more fine-grained analysis of the proposed algorithm, e.g., moving away from the restrictive graph isomorphism objective and deriving a deeper understanding of the neural architecture’s capabilities when optimized with stochastic gradient descent.

Broader impact

We view our work mainly as a methodological contribution. It studies the limits of current (supervised) graph embeddings methods, commonly used in chemoinformatics [103], bioinformatics [10], or network science [27]. Currently, methods used in practice, such as GNNs or extended-connectivity fingerprints [93] have severe limitations and might miss crucial patterns in today’s complex, inter-connected data. We investigate how to scale up graph embeddings that can deal with higher-order interactions of vertices (or atom of molecules, users in social networks, variables in optimization, . . .) to larger graphs or networks. Hence, our method paves the way for more resource-efficient and expressive graph embeddings.

We envision that our (methodological) contributions enable the design of more expressive and scalable graph embeddings in fields such as quantum chemistry, drug-drug interaction prediction, in-silicio, data-driven drug design/generation, and network analysis for social good. However, progress in graph embeddings might also trigger further advancements in hostile social network analysis, e.g., extracting more fine-grained user interactions for social tracking.

Example impact We are actively cooperating with chemists on drug design to evaluate further our approach to new databases for small molecules. Here, the development of new databases is quite tedious, and graph embeddings can provide hints to the wet lab researcher where to start their search. However, still, humans need to do much of the intuition-driven, manual wet lab work. Hence, we do not believe that our methods will result in job losses in the life sciences in the foreseeable future.

Acknowledgments and disclosure of funding

We thank Matthias Fey for answering our questions with regard to PYTORCH GEOMETRIC and Xavier Bresson for providing the ZINC dataset. This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2047/1 – 390685813 and under DFG Research Grants Program–RA 3242/1-1–411032549.

References

- [1] S. Abu-El-Haija, B. Perozzi, A. Kapoor, N. Alipourfard, K. Lerman, H. Harutyunyan, G. Ver Steeg, and A. Galstyan. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *International Conference on Machine Learning*, pages 21–29, 2019.
- [2] B. M. Anderson, T.-S. Hy, and R. Kondor. Cormorant: Covariant molecular neural networks. In *Advances in Neural Information Processing Systems 32*, pages 14510–14519, 2019.
- [3] S. Arora, S. S. Du, W. Hu, Z. Li, R. Salakhutdinov, and R. Wang. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, pages 8139–8148, 2019.
- [4] V. Arvind, J. Köbler, G. Rattan, and O. Verbitsky. On the power of color refinement. In *International Symposium on Fundamentals of Computation Theory*, pages 339–350, 2015.
- [5] V. Arvind, F. Fuhlbrück, J. Köbler, and O. Verbitsky. On Weisfeiler-Leman invariance: Subgraph counts and related graph properties. In *International Symposium on Fundamentals of Computation Theory*, pages 111–125, 2019.
- [6] A. Atserias and E. N. Maneva. Sherali-adams relaxations and indistinguishability in counting logics. *SIAM Journal on Computing*, 42(1):112–137, 2013.
- [7] A. Atserias, L Mancinska, D. E. Roberson, R. Sámal, S. Severini, and A. Varvitsiotis. Quantum and non-signalling graph isomorphisms. *Journal of Combinatorial Theory, Series B*, 136: 289–328, 2019.
- [8] L. Babai. Graph isomorphism in quasipolynomial time. In *ACM SIGACT Symposium on Theory of Computing*, pages 684–697, 2016.
- [9] S. Bai, F. Zhang, and P. H. S. Torr. Hypergraph convolution and hypergraph attention. *CoRR*, abs/1901.08150, 2019.

- [10] A.-L. Barabasi and Z. N. Oltvai. Network biology: Understanding the cell’s functional organization. *Nature Reviews Genetics*, 5(2):101–113, 2004.
- [11] P. Barceló, E. V. Kostylev, M. Monet, J. Pérez, J. L. Reutter, and J. Pablo Silva. The logical expressiveness of graph neural networks. In *International Conference on Learning Representations*, 2020.
- [12] C. Berkholz, P. S. Bonsma, and M. Grohe. Tight lower and upper bounds for the complexity of canonical colour refinement. In *Annual European Symposium on Algorithms*, pages 145–156. Springer, 2013.
- [13] K. M. Borgwardt and H.-P. Kriegel. Shortest-path kernels on graphs. In *IEEE International Conference on Data Mining*, pages 74–81, 2005.
- [14] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and deep locally connected networks on graphs. In *International Conference on Learning Representation*, 2014.
- [15] J. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identifications. *Combinatorica*, 12(4):389–410, 1992.
- [16] C. Cangea, P. Velickovic, N. Jovanovic, T. Kipf, and P. Liò. Towards sparse hierarchical graph classifiers. *CoRR*, abs/1811.01287, 2018.
- [17] S. Cao, W. Lu, and Q. Xu. GraRep: Learning graph representations with global structural information. In *ACM International Conference on Information and Knowledge Management*, pages 891–900, 2015.
- [18] I. Chami, Z. Ying, C. Ré, and J. Leskovec. Hyperbolic graph convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 4869–4880, 2019.
- [19] I. Chami, S. Abu-El-Haija, B. Perozzi, C. Ré, and K. Murphy. Machine learning on graphs: A model and comprehensive taxonomy. *CoRR*, abs/2005.03675, 2020.
- [20] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011.
- [21] G. Chen, P. Chen, C.-Y. Hsieh, C.-K. Lee, B. Liao, R. Liao, W. Liu, J. Qiu, Q. Sun, J. Tang, R. S. Zemel, and S. Zhang. Alchemy: A quantum chemistry dataset for benchmarking AI models. *CoRR*, abs/1906.09427, 2019.
- [22] J. Chen, T. Ma, and C. Xiao. Fastgcn: Fast learning with graph convolutional networks via importance sampling. In *International Conference on Learning Representation*, 2018.
- [23] J. Chen, J. Zhu, and L. Song. Stochastic training of graph convolutional networks with variance reduction. In *International Conference on Machine Learning*, pages 941–949, 2018.
- [24] Z. Chen, S. Villar, L. Chen, and J. Bruna. On the equivalence between graph isomorphism testing and function approximation with GNNs. In *Advances in Neural Information Processing Systems*, pages 15868–15876, 2019.
- [25] Z. Chen, L. Chen, S. Villar, and J. Bruna. Can graph neural networks count substructures? *CoRR*, abs/2002.04025, 2020.
- [26] G. Corso, L. Cavalleri, D. Beaini, P. Liò, and P. Velickovic. Principal neighbourhood aggregation for graph nets. *CoRR*, abs/2004.05718, 2020.
- [27] Easley D. and Kleinberg J. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.
- [28] G. Dasoulas, L. Dos Santos, K. Scaman, and A. Virmaux. Coloring graph neural networks for node disambiguation. In *International Joint Conference on Artificial Intelligence*, pages 2126–2132, 2020.

- [29] M. Defferrard, Bresson X., and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3844–3852, 2016.
- [30] H. Dell, M. Grohe, and G. Rattan. Lovász meets Weisfeiler and Leman. In *International Colloquium on Automata, Languages, and Programming*, pages 40:1–40:14, 2018.
- [31] P. D. Dobson and A. J. Doig. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of Molecular Biology*, 330(4):771 – 783, 2003.
- [32] S. S. Du, K. Hou, R. R. Salakhutdinov, B. Póczos, R. Wang, and K. Xu. Graph Neural Tangent Kernel: Fusing graph neural networks with graph kernels. In *Advances in Neural Information Processing Systems*, pages 5723–5733, 2019.
- [33] D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*, pages 2224–2232, 2015.
- [34] V. P. Dwivedi, C. K. Joshi, T. Laurent, Y. Bengio, and X. Bresson. Benchmarking graph neural networks. *CoRR*, abs/2003.00982, 2020.
- [35] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.
- [36] M. Fey and J. E. Lenssen. Fast graph representation learning with PyTorch Geometric. *CoRR*, abs/1903.02428, 2019.
- [37] M. Fey, J. E. Lenssen, F. Weichert, and H. Müller. SplineCNN: Fast geometric deep learning with continuous B-spline kernels. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 869–877, 2018.
- [38] D. Flam-Shepherd, T. Wu, P. Friederich, and A. Aspuru-Guzik. Neural message passing on high order paths. *CoRR*, abs/2002.10413, 2020.
- [39] Martin Fürer. On the combinatorial power of the Weisfeiler-Lehman algorithm. In *International Conference on Algorithms and Complexity*, pages 260–271, 2017.
- [40] H. Gao and S. Ji. Graph U-Nets. In *International Conference on Machine Learning*, pages 2083–2092, 2019.
- [41] V. K. Garg, S. Jegelka, and T. S. Jaakkola. Generalization and representational limits of graph neural networks. *CoRR*, abs/2002.06157, 2020.
- [42] T. Gärtner, P. Flach, and S. Wrobel. On graph kernels: Hardness results and efficient alternatives. In *Learning Theory and Kernel Machines*, pages 129–143. 2003.
- [43] F. Geerts, F. Mazowiecki, and G. A. Pérez. Let’s agree to degree: Comparing graph convolutional networks in the message-passing framework. *CoRR*, abs/2004.02593, 2020.
- [44] Floris Geerts. The expressive power of kth-order invariant graph networks. *CoRR*, abs/2007.12035, 2020.
- [45] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, 2017.
- [46] M. Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*. Lecture Notes in Logic. Cambridge University Press, 2017.
- [47] M. Grohe. Word2vec, Node2vec, Graph2vec, X2vec: Towards a theory of vector embeddings of structured data. *CoRR*, abs/2003.12590, 2020.
- [48] M. Grohe and M. Otto. Pebble games and linear equations. *Journal of Symbolic Logic*, 80(3): 797–844, 2015.

- [49] M. Grohe, K. Kersting, M. Mladenov, and E. Selman. Dimension reduction via colour refinement. In *European Symposium on Algorithms*, pages 505–516, 2014.
- [50] M. Grohe, P. Schweitzer, and Wiebking D. Deep Weisfeiler Leman. *CoRR*, abs/2003.10935, 2020.
- [51] W. L. Hamilton, R. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1025–1035, 2017.
- [52] M. Heimann, T. Safavi, and D. Koutra. Distribution of node embeddings as multiresolution features for graphs. In *IEEE International Conference on Data Mining*, pages 289–298, 2019.
- [53] C. Helma, R. D. King, S. Kramer, and A. Srinivasan. The Predictive Toxicology Challenge 2000–2001. *Bioinformatics*, 17(1):107–108, 01 2001.
- [54] W. Huang, T. Zhang, Y. Rong, and J. Huang. Adaptive sampling towards fast graph representation learning. In *Advances in Neural Information Processing Systems*, pages 4563–4572, 2018.
- [55] N. Immerman and E. Lander. *Describing Graphs: A First-Order Approach to Graph Canonization*, pages 59–81. Springer, 1990.
- [56] A. Jacot, C. Hongler, and F. Gabriel. Neural Tangent kernel: convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, pages 8580–8589, 2018.
- [57] W. Jin, R. Barzilay, and T. S. Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International Conference on Machine Learning*, pages 2328–2337, 2018.
- [58] Y. Jin, G. Song, and C. Shi. GraLSP: Graph neural networks with local structural patterns. *CoRR*, abs/1911.07675, 2019.
- [59] F. D. Johansson and D. Dubhashi. Learning with similarity functions on graphs using matchings of geometric embeddings. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 467–476, 2015.
- [60] H. Kashima, K. Tsuda, and A. Inokuchi. Marginalized kernels between labeled graphs. In *International Conference on Machine Learning*, pages 321–328, 2003.
- [61] S. Kiefer and B. D. McKay. The iteration number of colour refinement. *CoRR*, abs/2005.10182, 2020.
- [62] S. Kiefer and P. Schweitzer. Upper bounds on the quantifier depth for graph differentiation in first order logic. In *ACM/IEEE Symposium on Logic in Computer Science*, pages 287–296, 2016.
- [63] S. Kiefer, P. Schweitzer, and E. Selman. Graphs identified by logics with counting. In *International Symposium on Mathematical Foundations of Computer Science*, pages 319–330, 2015.
- [64] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- [65] D. B. Kireev. Chemnet: A novel neural network based method for graph/property mapping. *Journal of Chemical Information and Computer Sciences*, 35(2):175–180, 1995.
- [66] J. Klicpera, J. Groß, and S. Günnemann. Directional message passing for molecular graphs. In *International Conference on Learning Representations*, 2020.
- [67] R. Kondor and H. Pan. The multiscale Laplacian graph kernel. In *Advances in Neural Information Processing Systems*, pages 2982–2990, 2016.
- [68] N. M. Kriege, P.-L. Giscard, and R. C. Wilson. On valid optimal assignment kernels and applications to graph classification. In *Advances in Neural Information Processing Systems*, pages 1615–1623, 2016.

- [69] N. M. Kriege, C. Morris, A. Rey, and C. Sohler. A property testing framework for the theoretical expressivity of graph kernels. In *International Joint Conference on Artificial Intelligence*, pages 2348–2354, 2018.
- [70] N. M. Kriege, M. Neumann, C. Morris, K. Kersting, and P. Mutzel. A unifying view of explicit and implicit feature maps of graph kernels. *Data Mining and Knowledge Discovery*, 33(6): 1505–1547, 2019.
- [71] N. M. Kriege, F. D. Johansson, and C. Morris. A survey on graph kernels. *Applied Network Science*, 5(1):6, 2020.
- [72] J. B. Lee, R. A. Rossi, X. Kong, S. Kim, E. Koh, and A. Rao. Graph convolutional networks with motif-based attention. In *28th ACM International Conference on Information*, pages 499–508, 2019.
- [73] M. Lichter, I. Ponomarenko, and P. Schweitzer. Walk refinement, walk logic, and the iteration number of the Weisfeiler-Leman algorithm. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 1–13, 2019.
- [74] Andreas Loukas. What graph neural networks cannot learn: depth vs width. In *International Conference on Learning Representations*, 2020.
- [75] T. Maehara and H. NT. A simple proof of the universality of invariant/equivariant graph neural networks. *CoRR*, abs/1910.03802, 2019.
- [76] P. N. Malkin. Sherali–adams relaxations of graph isomorphism polytopes. *Discrete Optimization*, 12:73 – 97, 2014.
- [77] H. Maron, H. Ben-Hamu, H. Serviansky, and Y. Lipman. Provably powerful graph networks. In *Advances in Neural Information Processing Systems*, pages 2153–2164, 2019.
- [78] H. Maron, H. Ben-Hamu, N. Shamir, and Y. Lipman. Invariant and equivariant graph networks. In *International Conference on Learning Representations*, 2019.
- [79] C. Meng, S. C. Mouli, B. Ribeiro, and J. Neville. Subgraph pattern neural networks for high-order graph evolution prediction. In *AAAI Conference on Artificial Intelligence*, pages 3778–3787, 2018.
- [80] C. Merkwirth and T. Lengauer. Automatic generation of complementary descriptors with molecular graph networks. *Journal of Chemical Information and Modeling*, 45(5):1159–1168, 2005.
- [81] F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model CNNs. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5425–5434, 2017.
- [82] C. Morris, K. Kersting, and P. Mutzel. Glocalized Weisfeiler-Lehman kernels: Global-local feature maps of graphs. In *IEEE International Conference on Data Mining*, pages 327–336. IEEE, 2017.
- [83] C. Morris, M. Ritzert, M. Fey, W. L. Hamilton, Jan Eric Lenssen, G. Rattan, and M. Grohe. Weisfeiler and Leman go neural: Higher-order graph neural networks. In *AAAI Conference on Artificial Intelligence*, pages 4602–4609, 2019.
- [84] C. Morris, N. M. Kriege, F. Bause, K. Kersting, P. Mutzel, and M. Neumann. TUDataset: A collection of benchmark datasets for learning with graphs. *CoRR*, abs/2007.08663, 2020.
- [85] R. L. Murphy, B. Srinivasan, V. A. Rao, and B. Ribeiro. Janossy pooling: Learning deep permutation-invariant functions for variable-size inputs. In *International Conference on Learning Representations*, 2019.
- [86] Ryan L. Murphy, B. Srinivasan, V. A. Rao, and B. Ribeiro. Relational pooling for graph representations. In *International Conference on Machine Learning*, pages 4663–4673, 2019.

- [87] M. Niepert, M. Ahmed, and K. Kutzkov. Learning convolutional neural networks for graphs. In *International Conference on Machine Learning*, pages 2014–2023, 2016.
- [88] G. Nikolentzos, P. Meladianos, and M. Vazirgiannis. Matching node embeddings for graph similarity. In *AAAI Conference on Artificial Intelligence*, pages 2429–2435, 2017.
- [89] G. Nikolentzos, P. Meladianos, S. Limnios, and M. Vazirgiannis. A degeneracy framework for graph similarity. In *International Joint Conference on Artificial Intelligence*, pages 2595–2601, 2018.
- [90] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019.
- [91] R. Ramakrishnan, O. Dral, P., M. Rupp, and O. A. von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1, 2014.
- [92] B. Rieck, C. Bock, and K. M. Borgwardt. A persistent Weisfeiler-Lehman procedure for graph classification. In *International Conference on Machine Learning*, pages 5448–5458, 2019.
- [93] D. Rogers and M. Hahn. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754, May 2010.
- [94] Y. Rong, W. Huang, T. Xu, and J. Huang. DropEdge: Towards deep graph convolutional networks on node classification. In *International Conference on Learning Representations*, 2020.
- [95] R. Sato, M. Yamada, and H. Kashima. Approximation ratios of graph neural networks for combinatorial problems. In *Neural Information Processing Systems*, pages 4083–4092, 2019.
- [96] R. Sato, M. Yamada, and H. Kashima. Random features strengthen graph neural networks. *CoRR*, abs/2002.03155, 2020.
- [97] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- [98] I. Schomburg, A. Chang, C. Ebeling, M. Gremse, C. Heldt, G. Huhn, and D. Schomburg. BRENDA, the enzyme database: updates and major new developments. *Nucleic acids research*, 32(Database issue):D431–3, January 2004.
- [99] N. Shervashidze, S. V. N. Vishwanathan, T. H. Petri, K. Mehlhorn, and K. M. Borgwardt. Efficient graphlet kernels for large graph comparison. In *International Conference on Artificial Intelligence and Statistics*, pages 488–495, 2009.
- [100] N. Shervashidze, P. Schweitzer, E. J. van Leeuwen, K. Mehlhorn, and K. M. Borgwardt. Weisfeiler-Lehman graph kernels. *Journal of Machine Learning Research*, 12:2539–2561, 2011.
- [101] M. Simonovsky and N. Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 29–38, 2017.
- [102] A. Sperduti and A. Starita. Supervised neural networks for the classification of structures. *IEEE Transactions on Neural Networks*, 8(2):714–35, 1997.
- [103] J. Stokes, K. Yang, K. Swanson, W. Jin, A. Cubillos-Ruiz, N. Donghia, C. MacNair, S. French, L. Carfrae, Z. Bloom-Ackerman, V. Tran, A. Chiappino-Pepe, A. Badran, I. Andrews, E. Chory, G. Church, E. Brown, T. Jaakkola, R. Barzilay, and J. Collins. A deep learning approach to antibiotic discovery. *Cell*, 180:688–702.e13, 02 2020.
- [104] M. Togninalli, E. Ghisu, F. Llinares-López, B. Rieck, and K. M. Borgwardt. Wasserstein Weisfeiler-Lehman graph kernels. In *Advances in Neural Information Processing Systems*, pages 6436–6446, 2019.

- [105] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- [106] S. Verma and Z.-L. Zhang. Hunt for the unique, stable, sparse and fast feature learning on graphs. In *Advances in Neural Information Processing Systems*, pages 88–98, 2017.
- [107] S. Verma and Z.-L. Zhang. Stability and generalization of graph convolutional neural networks. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1539–1548, 2019.
- [108] O. Vinyals, S. Bengio, and M. Kudlur. Order matters: Sequence to sequence for sets. In *International Conference on Learning Representations*, 2016.
- [109] N. Wale, I. A. Watson, and G. Karypis. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowledge and Information Systems*, 14(3):347–375, 2008.
- [110] B. Weisfeiler. *On Construction and Identification of Graphs*. Lecture Notes in Mathematics, Vol. 558. Springer, 1976.
- [111] B. Weisfeiler and A. Leman. The reduction of a graph to canonical form and the algebra which appears therein. *Nauchno-Tekhnicheskaya Informatsia*, 2(9):12–16, 1968. English translation by G. Ryabov is available at https://www.iti.zcu.cz/wl2018/pdf/wl_paper_translation.pdf.
- [112] Z. Wu, B. Ramsundar, E. N. Feinberg, J. Gomes, C. Geniesse, A. S. Pappu, K. Leswing, and V. Pande. MoleculeNet: A benchmark for molecular machine learning. *Chemical Science*, 9: 513–530, 2018.
- [113] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A comprehensive survey on graph neural networks. *CoRR*, abs/1901.00596, 2019.
- [114] K. Xu, C. Li, Y. Tian, T. Sonobe, K. Kawarabayashi, and S. Jegelka. Representation learning on graphs with jumping knowledge networks. In *International Conference on Machine Learning*, pages 5453–5462, 2018.
- [115] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- [116] N. Yadati, M. Nimishakavi, P. Yadav, V. Nitin, A. Louis, and P. P. Talukdar. HyperGCN: A new method for training graph convolutional networks on hypergraphs. In *Advances in Neural Information Processing Systems*, pages 1509–1520, 2019.
- [117] X. Yan, H. Cheng, J. Han, and P. S. Yu. Mining significant graph patterns by leap search. In *ACM SIGMOD International Conference on Management of Data*, pages 433–444, 2008.
- [118] P. Yanardag and S. V. N. Vishwanathan. A structural smoothing framework for robust graph comparison. In *Advances in Neural Information Processing Systems*, pages 2125–2133, 2015.
- [119] P. Yanardag and S. V. N. Vishwanathan. Deep graph kernels. In *ACM SIGKDD International Conference on Knowledge Discovery and Data*, pages 1365–1374. ACM, 2015.
- [120] R. Ying, J. You, C. Morris, X. Ren, W. L. Hamilton, and J. Leskovec. Hierarchical graph representation learning with differentiable pooling. In *Advances in Neural Information Processing Systems*, pages 4800–4810, 2018.
- [121] J. You, R. Ying, and J. Leskovec. Position-aware graph neural networks. In *International Conference on Machine Learning*, pages 7134–7143, 2019.
- [122] M. Zhang, Z. Cui, M. Neumann, and C. Yixin. An end-to-end deep learning architecture for graph classification. In *AAAI Conference on Artificial Intelligence*, pages 4428–4435, 2018.
- [123] R. Zhang, Y. Zou, and J. Ma. Hyper-SAGNN: A self-attention based graph neural network for hypergraphs. In *International Conference on Learning Representations*, 2020.

- [124] D. Zhou, J. Huang, and B. Schölkopf. Learning with hypergraphs: Clustering, classification, and embedding. In *Advances in Neural Information Processing Systems*, pages 1601–1608, 2006.
- [125] J. Zhou, G. Cui, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun. Graph neural networks: A review of methods and applications. *CoRR*, abs/1812.08434, 2018.