

1 **We thank all reviewers for their fair and constructive reviews.** 🙌

2 **General remarks** We will revise and restructure the appendix (and the main paper), especially concerning refer-
3 ences/theorem numbering and redundancy between the main text.

4 **Scalability** We agree that our methods are not ready for industrial-scale graphs. We view our contribution as method-
5 ological work that is the first step to make higher-order methods more practical by leveraging *sparsity*, which is perhaps
6 the most significant parameter associated with a class of graphs. Moreover, we want to stress that we empirically
7 verified that our method offers significant benefits (compared to standard GNNs) in the regime of small graphs such as
8 molecules. We will integrate the discussion of Appendix F into the main text (complexity, sampling, ...).

9 **R1 Computation of isomorphism type** We agree that the method is only of theoretical interest for a large choice of k .
10 Nonetheless, we believe that we provided sufficient empirical evidence showing that for small $k \in \{2, 3\}$, the method
11 provides benefits over standard GNN and other higher-order architectures. Note that strictly speaking, the k -WL does
12 not need to perform an isomorphism test. It merely performs an equality test. Lines 97-98 (in the "appendix.pdf")
13 describe the condition for two tuples to have the same initial labeling: it is only the identity mapping $\text{id}: [k] \mapsto [k]$
14 which we demand to be a (partial) graph isomorphism. Hence, we only need to check "equality" of the two tuples and
15 not their "isomorphism". If we had considered two *sets* of k vertices instead of two *tuples* of k vertices, then we would
16 have needed isomorphism testing (as the reviewer suggests).

17 **Subgraph learning** Yes, we agree that results on link and triadic prediction would further highlight the approach's
18 generality. We plan to include them in a revised version/future work and add a discussion to the main paper.

19 **Aggregation scheme** We considered all tuples. We will make the complexity, especially the dependence on k , clearer in
20 the main paper, e.g., by including the discussion of Appendix F in the main paper.

21 **Directed graphs** Thank you for bringing this to our attention. The applicability of our arguments to the direct case
22 depends on the way we define the k -WL for directed graphs. It is typical to consider both in-neighbours and out-
23 neighbours, but separately. Hence for a tuple T , we have "local-in"-neighbors and "local-out"-neighbors. The local
24 unfolding at the tuple T will have both in-neighbours and out-neighbours at depth-1, and so on. Therefore, in this kind
25 of WL, we are essentially dealing with the underlying undirected graph of the given directed graph. If the graph is
26 weakly connected, we will still see all 1-neighbours of a tuple (local-in, local-out, global) somewhere down the local
27 unfolding tree. In that case, we are fine with weak connectivity. Of course, we could choose to define a WL-variant
28 with only "local-out"-neighbors, where we agree that strong connectivity will be necessary. Since the typical notion of
29 sparsity for directed/undirected graphs is the same, i.e., low edge-density, we would argue that it is natural to stick with
30 the first formalism. In that case, weak connectivity will suffice. We will clarify this in the revised version.

31 **R2 Performance of δ -2-LWL⁺** The δ -2-LWL⁺ is as a middle ground between the purely local δ -2-LWL and the global
32 δ -2-WL. The +-Version achieves good generalization performance due to slower convergence while preserving certain
33 global information, which is needed to derive Theorem 2. We will expand on this in the revised version.

34 **Other GNN architectures** We agree and will include more recently-proposed GNNs as baselines.

35 **Lower results for GIN** We used the implementation available in Pytorch Geometric. The lower numbers are because
36 we did not use an initial one-hot degree feature for datasets that did not provide node labels. We will add a comment
37 and add a row for results using degree features. (Note that all kernels are computed without this information)

38 **Leman vs. Lehman** Leman/Lehman personally stated that he prefers the transcription Leman (through private commu-
39 nication with Russian graph theorist Ilia Ponomarenko).¹ Moreover, the spelling Leman is also used throughout the
40 theory community.

41 **Proof of Prop. 1** Thank you for bringing this to our attention. The missing part directly follows from the CFI
42 construction outlined in [15]. However, a formal proof is quite involved and would require repeating the reasoning of
43 the paper above. We will add a proof sketch to the revised version of the paper.

44 **R3 Scalability/limitations** See general remarks above. We will report absolute running times (<1h for the local
45 architecture on the 10k subset.).

46 **Readability/relevance to a broader audience** We will incorporate intuition in the revised paper and emphasize the
47 neural architecture. Moreover, we will try to make the main paper as self-contained as possible.

48 **R4 Notation** We agree and will make the main text more self-contained.

49 **δ - k -LGNN** We note that the δ - k -LGNN is slightly weaker than the δ - k -GNN as it does not use the $\#$ labeling. (It
50 could be included but would require preprocessing.). Moreover, observe that the δ - k -GNN may learn to combine the
51 local and global neighborhood information in a more fine-grained way compared to the kernel (due to its discrete
52 nature). Moreover, we stress here that the δ - k -LGNN still achieves good performance while being much faster than the
53 δ - k -GNN or k -WL-GNN. We will add an expanded discussion in the revised paper.

54 **Explanation on experiment part** Thank you for the good suggestion; we will incorporate it into a revised version.

¹<https://www.itu.zcu.cz/wl2018/pdf/leman.pdf>