

Graph-level representations using ensemble-based readout functions

Jakub Binkowski¹[0000-0001-7386-5150], Albert Sawczyn¹[0000-0003-3142-4028],
Denis Janiak¹[0000-0003-1859-9093], Piotr Bielak¹[0000-0002-1487-2569], and
Tomasz Kajdanowicz¹[0000-0002-8417-1012]

Department of Artificial Intelligence,
Wrocław University of Science and Technology,
Wrocław, Poland
`jakub.binkowski@pwr.edu.pl`

Abstract. Graph machine learning models have been successfully deployed in various application areas. One of the most prominent types of models – Graph Neural Networks (GNNs) – provides an elegant way of extracting expressive node-level representation vectors, which can be used to solve node-related problems, such as classifying users in a social network. However, many tasks require representations at the level of the whole graph, e.g., molecular applications. In order to convert node-level representations into a graph-level vector, a so-called *readout* function must be applied. In this work, we study existing readout methods, including simple non-trainable ones, as well as complex, parametrized models. We introduce a concept of ensemble-based readout functions that combine either representations or predictions. Our experiments show that such ensembles allow for better performance than simple single readouts or similar performance as the complex, parametrized ones, but at a fraction of the model complexity.

Keywords: Graph Neural Networks · Readout · Pooling · Graph Classification · Graph Regression · Machine Learning.

1 Introduction

In recent years, machine learning has seen dramatic progress in areas where data is more complex, irregular, or structured. The abundance of methods treating objects of various structured domains eventually led to the movement of Geometric Deep Learning [3], which tries to characterize neural network architectures through the lens of data geometry and symmetries. In particular, research advancements in graph neural networks resulted in multiple novel architectures being able to address symmetries present in graph objects. One of the most important milestones was the development of Graph Neural Networks (GNNs), which rely on the convolution operation adapted to the irregular graph structure. GNNs can learn node representations by transforming initial input features throughout several message-passing layers. The final node representations might

be further utilized in a task of interest and are often learned end-to-end using a task-related objective.

In the realm of Graph Machine Learning, one can distinguish several types of tasks. Among them, the most prevalent are link prediction, node-level prediction (classification or regression) and graph-level prediction (classification or regression). It turns out that for node-level tasks, using node representations obtained from message-passing GNNs is straightforward as those representations could be immediately passed to the layer producing final predictions per each node. However, in the case of graph-level tasks, before the classification stage, one has to obtain a global graph-level representation vector, which summarizes the entire object. The common practice is to aggregate node representations from the last or all GNN layers through a so-called **readout function** [12]. Since graph-level representations must be invariant to node permutations, a readout function has to aggregate node embeddings regardless of their ordering, i.e., for each possible node permutation, it must return the same value. This enforces a significant constraint and poses a real challenge. Since graph-level tasks are of great importance, e.g., in fields like chemistry or biology, and their performance strongly depends on the readout function, this work focuses on the evaluation of various approaches to perform readout. Additionally, we include proposed ensemble methods and discuss obtained results to give an intuition for choosing the proper readout for a problem of interest.

In the simplest scenario, one might leverage permutation-invariant functions, such as: **sum**, **mean**, or **max**. However, as shown by Xu et al. [28], these approaches might be suboptimal and lead to significant data loss. Along with the advancements in various aggregation schemes for GNN message-passing, many works attempted to refine the readout function for graph-level tasks. Recently, Buterez et al. [5] performed a large-scale evaluation of various readouts, giving several unexpected insights. In this work, we aim to perform a follow-up study of selected readouts. In particular, we introduce an ensemble approach to readouts, in which we first perform several readout functions in parallel and then aggregate results in two scenarios: (1) representations of all readouts are aggregated, and the result is passed for prediction, (2) prediction is performed over embeddings from each readout individually, and then outcomes are aggregated.

Contributions This paper is aimed at bringing the following contributions:

1. We introduce ensemble-based readout models at representation and prediction levels, which help obtain rich graph representation from nodes' representations and outperform non-invariant SOTA results obtained with MLP and GRU models.
2. We perform a comparative study of readouts on datasets with varying sizes, characteristics, and target tasks.
3. We analyse the obtained results concerning the computational burden of considered readout functions, giving guidelines for readout function selection.

2 Related work

2.1 Graph Neural Networks

The early approaches, which enabled to perform tasks, like graph classification, involved graph kernels [16]. However, due to their limitations, e.g., the necessity of manually designing combinatorial features of a graph, as well as big advancements in deep learning, Graph Neural Networks were proposed. GNNs arose as a generalization of convolution to irregular graph structures. The groundbreaking works in graph machine learning proposed to perform graph convolution in a spectral domain, leveraging eigendecomposition of the graph Laplacian [4], or by employing polynomial spectral filters [8]. Gilmer et al. [12] proposed to unify GNNs through the message-passing framework, showing that previous convolution operators are its special cases. The rapid growth in the field led to many variants of Message Passing Neural Networks (MPNNs), among which the most renowned ones are GCN [15], GAT [25], GIN [28]. Such MPNNs were widely used in tasks like graph property prediction (classification or regression), node property prediction (node classification or regression), or link prediction.

2.2 Graph-level prediction

While MPNNs can transform input features into nodes' representations, predicting at a graph level requires a readout function that summarises a graph's node embeddings into a single vector, further passed to a prediction model (classifier or regressor). Such function should be invariant to node order permutations [3], and it is desired to be injective for the sake of expressivity [28].

The primary approach for readouts leveraged simple permutation-invariant functions, including `sum`, `mean`, or `max`. Xu et al. [28] proved the properties of these functions, showing their advantages and limitations. Although the work considered these functions in the context of neighbourhood aggregation in graph convolution, the same applies to them being readouts. Also, the authors raised the issue of function injectivity and stated that in certain situations, the `sum` operator might satisfy this property.

Recently, researchers developed more advanced approaches considering permutation invariance with respect to inputs. Most of these works aimed to perform various tasks on sets, which resemble the scenario for readouts. Zaheer et al. [31] presented the DeepSets architecture as the main framework to obtain permutation invariant representations of sets. We adopted this approach in our experimental scenario, and we will discuss it in more detail in the next section. Moreover, the SetTransformer [17] model enabled the efficient application of Transformer [24] architecture to sets, which scales to large-sized inputs.

It is also worth noting that several works proposed to leverage local pooling, which attempts to group similar nodes together iteratively. For instance, DiffPool [30] clusters together similar nodes in several iterations, such that in each iteration number of clusters decreases up to only 1 at the last layer, which serves as a graph representation. Besides, there are other local pooling methods, like

Graph Memory Networks[14], or based on GRACLU [22]. However, as shown by Mesquita et al. [19], there are no significant performance benefits from using local pooling, hence in this work, we do not consider them in the experiments.

Recently, Buterez et al. [5] showed results of various readouts based on large-scale evaluation, stating that in some situations losing permutation-invariance and employing models, like MLP or GRU recurrent neural network [6], might bring superior performance. Moreover, they showed that general adaptive (i.e. parametrized and learnable) readouts often lead to better efficacy with sufficient data. Our work also employs MLP and GRU, however, as opposed to Buterez et al., we also propose to use node-shared parameters (which, in hand with a prediction layer, forms the DeepSets architecture [31]). Moreover, we propose leveraging an ensemble approach that utilizes several readouts in a single model, similar to the aggregation scheme in PNA graph convolution [7].

3 Graph Neural Network Readouts

Graph Neural Networks operate on graph input data, transforming initial node features into rich node representations. Formally, we denote a graph as $\mathcal{G} = (\mathbf{X}, \mathcal{A})$, where $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$ are node features, and \mathcal{A} is an adjacency matrix describing the connection between nodes, such that $a_{uv} = \mathbf{1}((u, v) \in \mathcal{E})$ for \mathcal{E} is a set of edges and $u, v \in \mathcal{V}$ are nodes. In this work, we consider unweighted edges. Further, a GNN layer is a function taking a graph as input and transforming it into a latent space – see: Eq. (1). Such formulation describes a generic framework of Message Passing Neural Network. First, it aggregates the neighborhood of a node through a permutation-invariant \oplus function to compute a message, and then it combines the message with the node’s representation through the Θ function. Both functions could be parameterized by neural networks. Particular instances of GNNs, like GCN or GIN, vary only in the Θ and \oplus functions. The whole network often comprises several such layers, and their number should be adjusted to a particular dataset.

$$\mathbf{h}_u^{(l+1)} = \Theta^{(l)} \left(\mathbf{h}_u^{(l)}, \oplus^{(l)} \left(\{\mathbf{h}_v^{(l)}, \forall v \in \mathcal{N}(u)\} \right) \right) \quad (1)$$

3.1 Problem statement

Formally, a readout function, denoted as \mathcal{R} , is an aggregation function that takes a matrix of stacked node embeddings (usually taken from the last GNN layer) at the input and returns a single vector at its output, i.e., $\mathcal{R}(\mathbf{H}^{(L)}) \rightarrow \mathbf{z}$, where $\mathbf{H}^{(L)} \in \mathbb{R}^{|\mathcal{V}| \times d_{\mathcal{V}}}$, and $\mathbf{z} \in \mathbb{R}^{d_{\mathcal{G}}}$, where L is the number of layers, $d_{\mathcal{V}}$ dimension of a hidden node representation, and $d_{\mathcal{G}}$ dimension of graph representation. In many scenarios, the dimension of the graph representation vector is equal to the dimension of the nodes’ representation, i.e., $d_{\mathcal{G}} = d_{\mathcal{V}}$. Fig. 1 visually depicts the concept of the readout function.

Readout function has to be invariant to the permutation of the node order in a graph, meaning that $\mathcal{R}(\mathbf{H}^{(L)}) = \mathcal{R}(\mathbf{P}\mathbf{H}^{(L)})$, where \mathbf{P} is a row-permutation




Fig. 1. A readout function in the whole graph-level prediction pipeline.

matrix. As showed by Zaheer et al. [31], permutation-invariant functions over the elements could be expressed as in Eq. (2), where ρ and ϕ could be functions parametrized by neural networks. Zaheer et al. considered such parametrization with MLP networks as the DeepSets architecture. However, it might be hard to train representations in a DeepSets setting, which was theoretically discussed by Wagstaff et al. [27]. Nonetheless, we include DeepSets approach in our experiments as well. Besides, there are several different approaches to performing prediction on sets, e.g., SetTransformer [17], or Set2Set [26].

$$f(X) = \rho \left(\sum_{x \in X} \phi(x) \right) \quad (2)$$

3.2 Non-parametrized readout functions

This work considers three basic non-parametrized functions: **sum**, **mean**, and **max**. These are the simplest realizations of readouts which satisfy the necessary invariance conditions. However, there are certain drawbacks to such aggregations. The concerns were raised by Xu et al. [28], who characterized each of the three aggregators with respect to their expressivity, stating that **sum** could capture the most information from the node embeddings, yet the effectiveness of such readout function might depend on the underlying data. However, as proved by Corso et al. [7], a single aggregation might not satisfy injectivity over real numbers, hence they propose to concatenate results of several aggregations at each convolutional layer. Motivated by these results, we propose an ensemble approach for building readout function. In particular, we evaluate the effectiveness of a readout function which composes **sum**, **mean**, and **max** aggregator into one vector.

3.3 Parametrized readout functions

In contrast to **sum**, **mean**, and **max**, parametrized readout functions could be optimized during training. In this work, we consider three such functions: DeepSets, Virtual Node, and an ensemble approach. Also, motivated by the surprising performance of permutation-sensitive methods [5], we include MLP- and GRU-based readouts. Now, let us discuss each parametrized readout in more detail.

DeepSets We use DeepSets architecture as in Eq. (2), where ϕ and ρ are parametrized by an MLP network. This approach is a special case of the encoder-decoder architecture, in which ϕ is the encoder with weights shared among all nodes, and ρ is the decoder transforming node representations summed over the embedding dimension. We employ two variants of DeepSets model, namely **DeepSets-Base** and **DeepSets-Large**, which differ by the number of layers in the encoder. Both versions have the same 2-layer MLP at the decoder, which predicts the final output.

Virtual Node In this approach, no explicit readout function is present in the architecture. Instead, the graph adjacency matrix is altered such that it includes one extra node connected with all other nodes in the graph [12]. Then, the graph-level representation is the embedding of this single node trained along with the remaining node embeddings (virtual node features are initialized with zeros). The aggregation function of the graph convolution might be seen as an implicit readout, and due to its permutation-invariance, this approach is also insensitive to the modification in the node ordering.

Dense and GRU In their recent work, Buterez et al. [5] empirically showed that on many datasets loosening permutation-invariance and leveraging MLP (Dense) or GRU architectures might lead to outperformance over other readouts, including the learnable ones. However, they argue that this approach is particularly suitable for data with consistent node ordering like canonical-SMILES [20] in molecular datasets.

3.4 Ensemble readouts

In this work, we propose to evaluate ensemble approaches to readout functions that combine results from many readouts. As described by Xu et al. [28], each aggregation from **sum**, **mean**, **max** has specific desirable properties, hence using three of them might bring a performance improvement. We propose to use these three readout functions and consider three approaches to aggregate representation computed with each of them.

1. First, we propose to simply concatenate the results of each aggregation, which are then forwarded to a target predictor. We denote that approach as **ConcatR**, where \parallel denotes the concatenation operator:

$$\mathbf{z} = \mathcal{R}_1(\mathbf{H}^{(L)}) \parallel \mathcal{R}_2(\mathbf{H}^{(L)}) \parallel \dots \parallel \mathcal{R}_N(\mathbf{H}^{(L)}), \quad (3)$$

2. In the second scenario, we build the final graph-level representation as a weighted sum of each readout’s output. This approach is formally expressed by Eq. (4), where $\{\mathcal{R}_1, \dots, \mathcal{R}_N\}$ is a set of N readout functions, $(\mathbf{W}_r^{(p)}, \mathbf{b}_r^{(p)})$ are weights and biases of a linear projection layer specific to the readout r , and $(w_r^{(c)}, b_r^{(c)})$ are weights and biases of a combination layer. Note that

representations from each readout might result in different scales, so we experiment with turning the projection layer on and off.

$$\mathbf{z} = \sum_{r \in \{\mathcal{R}_1, \dots, \mathcal{R}_N\}} w_r^{(c)} \left(\mathbf{W}_r^{(p)} \mathbf{z}_r + \mathbf{b}_r^{(p)} \right) + b_r^{(c)} \quad (4)$$

We propose two variants of such representation-level ensembles:

- **WMeanR** (*Weighted Mean of Readouts*) – the representation from each readout is multiplied by a learnable weight and summed together (the projection layer is replaced with an identity map: $\mathbf{W}_r^{(p)} = \mathbb{I}$, $\mathbf{b}_r^{(p)} = \mathbf{0}$),
 - **WMeanR+Proj** (*Weighted Mean of Readouts with Projection layer*) – the representation from each readout is first transformed by the projection layer, and then weighted and summed together (as in the previous model).
3. The last scenario involves the aggregation of results from the three readouts at the prediction level, as shown in Eq. (5). The symbols are the same as in Eq. (4), and $\psi(\cdot)$ denotes a predictor model. Again, we include models with and without a projection layer in our experiments.

$$\hat{\mathbf{y}} = \sum_{r \in \{\mathcal{R}_1, \dots, \mathcal{R}_N\}} \left(w_r^{(c)} \psi(\mathbf{W}_r^{(p)} \mathbf{z}_r + \mathbf{b}_r^{(p)}) + b_r^{(c)} \right) \quad (5)$$

- **MeanPred** (*Mean of readouts' Predictions*) – predictions are computed for each readout separately and then averaged over all readouts (the projection layer is replaced with an identity map: $\mathbf{W}_r^{(p)} = \mathbb{I}$, $\mathbf{b}_r^{(p)} = \mathbf{0}$; we also used fixed values for the combination parameters: $w_r^{(c)} = 1$, $b_r^{(c)} = 0$),
- **WMeanPred** (*Weighted Mean of readouts' Predictions*) – predictions are computed for each readout separately, then aggregated by learnable weights (the projection layer is replaced with an identity map: $\mathbf{W}_r^{(p)} = \mathbb{I}$, $\mathbf{b}_r^{(p)} = \mathbf{0}$),
- **WMeanPred+Proj** (*Weighted Mean of readouts' Predictions with Projection layer*) – projection layer is applied to representations from each readout, then predictions are computed and aggregated by learnable weights

Altogether, we introduce **6 different ensemble-based variants of readout functions**.

4 Experiments

Here, we describe experiments and obtained results in detail. First, we elaborate on the datasets and splits, further, we specify the experimental setting, and finally, we discuss the results.

4.1 Datasets

In our experiments, we utilized four datasets, which are MUTAG [16], ENZYMES [1], ZINC [2] (12K subset of the dataset, as proposed in [9]), and REDDIT-MULTI-12K [29]. Datasets were selected to cover various tasks, domains, and graph sizes. Table 1 summarizes datasets statistics. For each dataset from the TUD repository (MUTAG, ENZYMES, REDDIT-MULTI-12K), we performed random splits to *train/val/test* subsets in proportions 80%/10%/10%. In the case of the ZINC dataset, we leverage pre-defined random splits without any further modifications. In the case of binary classification on the MUTAG dataset, we measure performance with the *F1* score, for multi-class classification on ENZYMES and REDDIT-MULTI-12K datasets, we leveraged the macro-averaged *F1* score, and for ZINC the R^2 metric was used.

Table 1. Tasks and statistics for each of four datasets used in the experiments.

Dataset	Task	# graphs	Avg. nodes	Avg. edges
ENZYMES	multi-class	600	32.9	125.4
MUTAG	binary	188	18.0	39.9
ZINC	regression	12 000	23.2	49.8
REDDIT-MULTI-12K	multi-class	11 929	391.0	913.3

4.2 Evaluation protocol

We evaluated three different and most prevalent graph convolutions, i.e., GCN, GAT and GIN, for each readout described in section 3. We set the hyperparameters of each model based on the ones found in [9] with slight modifications when necessary. All hyperparameters specific to each dataset are presented in Table 2. For **Dense** and **GRU** readout, we adopted the architectures from Buterez et al. [5]. For **DeepSets-Base**, we used MLP with two layers of size 128 followed by batch normalization [13] and ReLU activation, and the last layer additionally contains dropout [23] with probability set to 0.4. For **DeepSets-Large**, we leveraged 6 such layers, followed by dropout with probability set to 0.4. The two considered **DeepSets** architectures were arranged to resemble **Dense** approach. It is worth noting that the number of trainable parameters of **Dense** method is relatively large, even when compared to **DeepSets-Large**. For the final predictor, i.e., classifier or regressor, the resultant representation of the readout function is always passed to an MLP with one hidden layer of size 128 (three layers including input and output).

Based on the well-established protocol schemes, we apply five times repeated bootstrap evaluation of each combination of graph convolution and readout. The models for each dataset were trained for a minimum of 10 epochs with early stopping on validation set loss and the patience set to 25 epochs. We leverage

dynamic learning rate starting from a value of 1×10^{-3} , which was multiplied by 0.5 on each plateau with patience set to 10 epochs (minimum learning rate was set to 1×10^{-6}). Models were optimized with AdamW optimizer [18], following default PyTorch parameters (except for learning rate). In the case of classification, we used cross-entropy loss, and for regression - mean-squared error. Experiments were implemented with, PyTorch [21], `torch-geometric` [11] and `pytorch-lightning` [10] libraries. To ensure reproducibility, we provide the full implementation of our experiments, containing entire hyperparameters configurations, in our public git repository <https://github.com/graphml-lab-pwr/ensemble-readouts>.

Table 2. Hyperparameters specification for each dataset, where L is number of layers (last layer index), d_V is the hidden graph representation dimension, and d_G is the graph representation dimension.

Dataset	L	d_V	d_G
ENZYMES	3	128	128
MUTAG	3	64	128
ZINC	3	128	128
REDDIT-MULTI-12K	3	128	128

4.3 Results

Following the experimental setup described in the previous section, we obtain results presented in Table 3, which shows the achieved efficacy for each dataset, split by graph convolution and readout (for all metrics *greater=better*). We observe that in most cases the best results are obtained either by *parametrized* or *ensemble* methods. Among all parametrized readouts, non-invariant methods turned out to perform best only on MUTAG, which is the smallest dataset used in our experiments (containing only 188 graphs overall, and only *19 graphs* in the test set). We hypothesize that such results might be caused by overfitting due to a relatively large number of trainable parameters in these models and a fixed node ordering. Therefore, we consider such methods as the worst choice for the readout function. On the other hand, the `DeepSets` architecture outperformed other models for ENZYMES using all types of GNNs, and for ZINC using GCN and GAT. The usage of a `Virtual Node` did not lead to any improvement.

Furthermore, the results showed that **ensemble** approaches outperformed other readouts for REDDIT-MULTI-12K dataset using GCN and GIN layers, as well as using the GIN layer on the ZINC dataset. It is worth noting that when using GIN convolution, the best ensemble approach is consistently the best overall or nearly the best. We did not observe such consistent results for GCN and GAT, yet the **ensemble** approaches provide better or comparative metric values.

Whether to use ensembles over representations or predictions depends on the dataset and graph convolution. Each ensemble enables an additional interpretability property, since the trained weights might be further analyzed to assess the contribution of each readout method.

Regarding the projection layer considered in these models, we observe that projection layers in ensembles do not guarantee better performance and could be omitted for the sake of decreasing model size. Nonetheless, when using a combination of various readouts from non-aligned embedding spaces, one should always verify the necessity of additional projection layers. We also tested a combination of DeepSets with *non-parameterized* methods in an *ensemble* setting, but we did not observe an increase in performance, hence we omitted results for brevity.

Table 3. Results of the conducted experiments. Each cell contains mean and standard deviation in percent of the metric over 5 runs. Metric $F1$ score was used for ENZYMES, MUTAG, REDDIT-MULTI-12K, and R^2 for ZINC (for all metrics *greater = better*). Best results are presented in **bold**, best within a class (*non-parametrized*, *parametrized*, *ensemble*) are presented with underline. *NON-PAR*, *PAR* and *ENS* denote non-parametrized, parametrized and ensemble-based readout functions, respectively.

		ENZYMES (F1 \uparrow)			MUTAG (F1 \uparrow)		
		GAT	GCN	GIN	GAT	GCN	GIN
NON-PAR	max	56.68 \pm 7.43	54.63 \pm 4.86	48.29 \pm 4.33	77.55 \pm 7.40	78.57 \pm 6.94	82.61 \pm 9.39
	mean	57.60 \pm 7.31	45.62 \pm 6.74	50.46 \pm 2.77	75.99 \pm 9.52	76.60 \pm 8.59	77.70 \pm 8.66
	sum	51.61 \pm 12.71	49.23 \pm 8.08	<u>51.19 \pm 9.65</u>	<u>84.65 \pm 6.87</u>	<u>82.62 \pm 6.28</u>	81.47 \pm 8.83
PAR	DeepSets-Base	62.32 \pm 10.09	64.39 \pm 6.26	55.66 \pm 9.77	66.63 \pm 28.74	75.42 \pm 13.24	84.68 \pm 6.86
	DeepSets-Large	63.65 \pm 7.01	58.08 \pm 11.49	42.44 \pm 11.31	78.00 \pm 3.93	72.59 \pm 25.46	79.66 \pm 5.97
	Dense	59.70 \pm 7.58	54.24 \pm 5.32	40.43 \pm 4.15	83.91 \pm 3.89	88.61 \pm 8.58	85.77 \pm 3.26
	GRU	43.72 \pm 7.39	47.56 \pm 3.75	27.96 \pm 6.42	86.41 \pm 5.74	85.90 \pm 5.30	86.76 \pm 4.02
	Virtual Node	55.49 \pm 10.04	45.54 \pm 3.99	51.37 \pm 8.42	<u>75.77 \pm 8.74</u>	78.41 \pm 6.73	84.46 \pm 7.14
ENS	ConcatR	56.38 \pm 8.12	44.76 \pm 10.68	46.06 \pm 9.24	<u>85.65 \pm 7.03</u>	85.63 \pm 8.36	85.63 \pm 8.36
	WMeanR	50.53 \pm 8.03	44.82 \pm 7.43	49.04 \pm 3.65	<u>85.65 \pm 7.03</u>	83.32 \pm 8.73	85.26 \pm 8.27
	WMeanR+Proj	57.90 \pm 5.75	47.09 \pm 11.09	49.77 \pm 9.11	84.09 \pm 6.93	<u>85.65 \pm 7.03</u>	<u>86.31 \pm 6.08</u>
	MeanPred	50.95 \pm 9.59	47.56 \pm 11.81	<u>54.49 \pm 5.51</u>	83.21 \pm 5.61	80.18 \pm 9.37	86.20 \pm 8.39
	WMeanPred	51.89 \pm 10.17	50.06 \pm 9.43	48.52 \pm 8.90	82.10 \pm 7.16	79.62 \pm 10.43	84.58 \pm 8.13
	WMeanPred+Proj	<u>58.35 \pm 9.43</u>	<u>52.59 \pm 7.55</u>	50.67 \pm 10.82	84.18 \pm 8.35	84.95 \pm 8.26	86.25 \pm 6.95
		REDDIT-MULTI-12K (F1 \uparrow)			ZINC (R^2 \uparrow)		
		GAT	GCN	GIN	GAT	GCN	GIN
NON-PAR	max	37.66 \pm 1.56	38.15 \pm 2.95	39.09 \pm 2.12	65.74 \pm 0.76	65.88 \pm 1.43	71.62 \pm 3.41
	mean	34.89 \pm 2.73	37.77 \pm 1.77	36.53 \pm 0.55	<u>68.10 \pm 0.57</u>	67.84 \pm 0.79	78.29 \pm 2.85
	sum	39.02 \pm 1.16	<u>40.38 \pm 0.97</u>	<u>39.86 \pm 0.92</u>	67.88 \pm 0.89	<u>68.21 \pm 0.26</u>	<u>81.68 \pm 1.65</u>
PAR	DeepSets-Base	36.59 \pm 2.35	35.42 \pm 2.75	39.08 \pm 1.82	64.42 \pm 9.36	73.82 \pm 2.39	75.40 \pm 2.17
	DeepSets-Large	37.61 \pm 1.71	37.29 \pm 2.06	16.90 \pm 10.11	70.09 \pm 2.77	83.69 \pm 3.26	81.96 \pm 1.24
	Dense	32.12 \pm 1.67	32.65 \pm 1.10	34.91 \pm 1.36	61.93 \pm 1.73	62.57 \pm 1.31	69.65 \pm 1.39
	GRU	27.49 \pm 1.38	26.62 \pm 1.49	29.21 \pm 2.40	58.23 \pm 8.00	65.26 \pm 1.20	63.81 \pm 2.91
	Virtual Node	37.26 \pm 1.60	<u>40.39 \pm 0.44</u>	<u>40.37 \pm 1.21</u>	67.43 \pm 1.44	69.16 \pm 0.31	76.67 \pm 3.01
ENS	ConcatR	38.62 \pm 1.90	41.54 \pm 0.75	41.06 \pm 1.05	68.08 \pm 0.56	67.24 \pm 1.33	83.23 \pm 2.56
	WMeanR	<u>38.83 \pm 0.94</u>	41.02 \pm 0.75	39.22 \pm 2.58	67.66 \pm 0.88	67.08 \pm 0.82	81.04 \pm 2.00
	WMeanR+Proj	37.70 \pm 1.04	40.65 \pm 1.17	40.37 \pm 1.15	67.97 \pm 0.63	67.80 \pm 0.26	82.18 \pm 2.31
	MeanPred	38.46 \pm 1.52	40.59 \pm 1.79	40.20 \pm 1.25	<u>68.22 \pm 0.58</u>	67.90 \pm 0.47	78.14 \pm 2.24
	WMeanPred	37.56 \pm 1.30	39.91 \pm 0.98	38.62 \pm 2.02	67.28 \pm 1.90	<u>68.22 \pm 0.37</u>	80.06 \pm 3.41
	WMeanPred+Proj	38.18 \pm 1.11	40.65 \pm 1.95	38.92 \pm 1.34	67.77 \pm 1.19	66.98 \pm 2.22	85.66 \pm 2.70

In order to discuss the results from the perspective of model size, we plotted the efficacy achieved by models on each dataset against the number of trainable parameters, presented in Fig. 2. Without accounting for a specific GNN type, the plots reveal that we can often benefit from **ensemble** or **parametrized** approaches which are comparable in size and efficacy. The exception is the dense model represented by the rightmost outlier points, which in connection with a lack of permutation-invariance position it as a poor choice for a readout function.




Fig. 2. Comparison of the number of parameters with achieved efficacy for the four considered datasets.

5 Conclusion

In this work, we evaluated various neural network readouts on several datasets covering different domains, cardinalities, graph properties, and tasks. We introduced **ensemble**-based approach that combines basic *non-parametrized* readout functions in several ways and evaluated them against other *parametrized* methods known in the literature, including DeepSets, Virtual Node, and non-invariant MLP and GRU. We showed that DeepSets and **ensemble** approaches outperform other methods in the majority of experiments and hence should be a first choice when designing models for graph-level predictions. We also do not recommend using non-invariant readouts, as they introduce larger models, might lead to overfitting, and not necessarily lead to better generalization. For future work, we suggest investigating probabilistic modelling of node representations,

including Bayesian approach, to prevent information loss, which might happen while using deterministic readouts presented in this work.

Acknowledgments This work was financed by (1) the Polish Ministry of Education and Science, CLARIN-PL; (2) the European Regional Development Fund as a part of the 2014–2020 Smart Growth Operational Programme, CLARIN — Common Language Resources and Technology Infrastructure, project no. POIR.04.02.00-00C002/19; (3) the statutory funds of the Department of Artificial Intelligence, Wrocław University of Science and Technology, Poland; (4) Horizon Europe Framework Programme MSCA Staff Exchanges grant no. 101086321 (OMINO).

References

1. Borgwardt, K.M., Ong, C.S., Schönauer, S., Vishwanathan, S.V.N., Smola, A.J., Kriegel, H.P.: Protein function prediction via graph kernels. *Bioinformatics* (Oxford, England) **21 Suppl 1**, i47–56 (Jun 2005). <https://doi.org/10.1093/bioinformatics/bti1007>
2. Bresson, X., Laurent, T.: A Two-Step Graph Convolutional Decoder for Molecule Generation. arXiv, Vancouver, Canada. (Jun 2019). <https://doi.org/10.48550/arXiv.1906.03412>, <http://arxiv.org/abs/1906.03412>, arXiv:1906.03412 [cs, stat]
3. Bronstein, M.M., Bruna, J., Cohen, T., Veličković, P.: Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges (May 2021). <https://doi.org/10.48550/arXiv.2104.13478>, <http://arxiv.org/abs/2104.13478>, arXiv:2104.13478 [cs, stat]
4. Bruna, J., Zaremba, W., Szlam, A., LeCun, Y.: Spectral Networks and Locally Connected Networks on Graphs. arXiv (May 2014), <http://arxiv.org/abs/1312.6203>, arXiv:1312.6203 [cs]
5. Buterez, D., Janet, J.P., Kiddle, S.J., Oglic, D., Liò, P.: Graph Neural Networks with Adaptive Readouts (Oct 2022), <https://openreview.net/forum?id=yts7fLpWY9G>
6. Cho, K., van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the Properties of Neural Machine Translation: Encoder-Decoder Approaches (Oct 2014), <http://arxiv.org/abs/1409.1259>, arXiv:1409.1259 [cs, stat]
7. Corso, G., Cavalleri, L., Beaini, D., Liò, P., Veličković, P.: Principal Neighbourhood Aggregation for Graph Nets
8. Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. arXiv (Feb 2017), <http://arxiv.org/abs/1606.09375>, arXiv:1606.09375 [cs, stat]
9. Dwivedi, V.P., Joshi, C.K., Luu, A.T., Laurent, T., Bengio, Y., Bresson, X.: Benchmarking Graph Neural Networks. *Journal of Machine Learning Research* **23** (Dec 2022), <http://arxiv.org/abs/2003.00982>, arXiv:2003.00982 [cs, stat]
10. Falcon, W., team, T.P.L.: PyTorch Lightning (Mar 2019). <https://doi.org/10.5281/zenodo.7545285>, <https://zenodo.org/record/7545285>
11. Fey, M., Lenssen, J.E.: Fast Graph Representation Learning with PyTorch Geometric. arXiv (Apr 2019), <http://arxiv.org/abs/1903.02428>, arXiv:1903.02428 [cs, stat]

12. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural Message Passing for Quantum Chemistry. In: Proceedings of the 34th International Conference on Machine Learning. vol. 70. arXiv, Sydney NSW Australia (Jun 2017). <https://doi.org/10.48550/arXiv.1704.01212>, <http://arxiv.org/abs/1704.01212>, arXiv:1704.01212 [cs]
13. Ioffe, S., Szegedy, C.: Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift (Mar 2015). <https://doi.org/10.48550/arXiv.1502.03167>, <http://arxiv.org/abs/1502.03167>, arXiv:1502.03167 [cs]
14. Khasahmadi, A.H., Hassani, K., Moradi, P., Lee, L., Morris, Q.: Memory-Based Graph Networks. arXiv (Jun 2020). <https://doi.org/10.48550/arXiv.2002.09518>, <http://arxiv.org/abs/2002.09518>, arXiv:2002.09518 [cs, stat]
15. Kipf, T.N., Welling, M.: Semi-Supervised Classification with Graph Convolutional Networks. arXiv (Feb 2017). <https://doi.org/10.48550/arXiv.1609.02907>, <http://arxiv.org/abs/1609.02907>, arXiv:1609.02907 [cs, stat]
16. Kriege, N., Mutzel, P.: Subgraph Matching Kernels for Attributed Graphs. In: Proceedings of the 29th International Conference on International Conference on Machine Learning. Omnipress, Edinburgh Scotland
17. Lee, J., Lee, Y., Kim, J., Kosiorek, A.R., Choi, S., Teh, Y.W.: Set Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks. In: Proceedings of the 36th International Conference on Machine Learning. pp. 3744–3753. arXiv (May 2019). <https://doi.org/10.48550/arXiv.1810.00825>, <http://arxiv.org/abs/1810.00825>, arXiv:1810.00825 [cs, stat]
18. Loshchilov, I., Hutter, F.: Decoupled Weight Decay Regularization (Jan 2019). <https://doi.org/10.48550/arXiv.1711.05101>, <http://arxiv.org/abs/1711.05101>, arXiv:1711.05101 [cs, math]
19. Mesquita, D., Souza, A.H., Kaski, S.: Rethinking pooling in graph neural networks (Oct 2020), <http://arxiv.org/abs/2010.11418>, arXiv:2010.11418 [cs]
20. O’Boyle, N.M.: Towards a Universal SMILES representation - A standard method to generate canonical SMILES based on the InChI. *Journal of Cheminformatics* 4(1), 22 (Sep 2012). <https://doi.org/10.1186/1758-2946-4-22>, <https://doi.org/10.1186/1758-2946-4-22>
21. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library (Dec 2019). <https://doi.org/10.48550/arXiv.1912.01703>, <http://arxiv.org/abs/1912.01703>, arXiv:1912.01703 [cs, stat]
22. Rhee, S., Seo, S., Kim, S.: Hybrid Approach of Relation Network and Localized Graph Convolutional Filtering for Breast Cancer Subtype Classification (Jun 2018). <https://doi.org/10.48550/arXiv.1711.05859>, <http://arxiv.org/abs/1711.05859>, arXiv:1711.05859 [cs]
23. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15(56), 1929–1958 (2014), <http://jmlr.org/papers/v15/srivastava14a.html>
24. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention Is All You Need (Dec 2017), <http://arxiv.org/abs/1706.03762>, arXiv:1706.03762 [cs]
25. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph Attention Networks (Feb 2018), <http://arxiv.org/abs/1710.10903>, arXiv:1710.10903 [cs, stat]

26. Vinyals, O., Bengio, S., Kudlur, M.: Order Matters: Sequence to sequence for sets (Feb 2016). <https://doi.org/10.48550/arXiv.1511.06391>, <http://arxiv.org/abs/1511.06391>, arXiv:1511.06391 [cs, stat]
27. Wagstaff, E., Fuchs, F.B., Engelcke, M., Posner, I., Osborne, M.: On the Limitations of Representing Functions on Sets (Oct 2019), <http://arxiv.org/abs/1901.09006>, arXiv:1901.09006 [cs, stat]
28. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How Powerful are Graph Neural Networks? (Feb 2019), <http://arxiv.org/abs/1810.00826>, arXiv:1810.00826 [cs, stat]
29. Yanardag, P., Vishwanathan, S.: Deep graph kernels. In: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. p. 1365–1374. KDD '15, Association for Computing Machinery, New York, NY, USA (2015). <https://doi.org/10.1145/2783258.2783417>, <https://doi.org/10.1145/2783258.2783417>
30. Ying, R., You, J., Morris, C., Ren, X., Hamilton, W.L., Leskovec, J.: Hierarchical Graph Representation Learning with Differentiable Pooling (Feb 2019). <https://doi.org/10.48550/arXiv.1806.08804>, <http://arxiv.org/abs/1806.08804>, arXiv:1806.08804 [cs, stat]
31. Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R., Smola, A.: Deep Sets (Apr 2018). <https://doi.org/10.48550/arXiv.1703.06114>, <http://arxiv.org/abs/1703.06114>, arXiv:1703.06114 [cs, stat]