DIRECTIONAL GRAPH NETWORKS

Dominique Beaini^{*1} Saro Passaro^{*2} Vincent Létourneau¹ William L. Hamilton³ Gabriele Corso² Pietro Liò²

ABSTRACT

The lack of anisotropic kernels in graph neural networks (GNNs) strongly limits their expressiveness, contributing to well-known issues such as over-smoothing. To overcome this limitation, we propose the first globally consistent anisotropic kernels for GNNs, allowing for graph convolutions that are defined according to topologicalyderived directional flows. First, by defining a vector field in the graph, we develop a method of applying directional derivatives and smoothing by projecting node-specific messages into the field. Then, we propose the use of the Laplacian eigenvectors as such vector field. We show that the method generalizes CNNs on an n-dimensional grid and is provably more discriminative than standard GNNs regarding the Weisfeiler-Lehman 1-WL test. We evaluate our method on different standard benchmarks and see a relative error reduction of 8% on the CIFAR10 graph dataset and 11% to 32% on the molecular ZINC dataset, and precision of 1.6% on the MolPCBA dataset.

1 INTRODUCTION

One of the most important distinctions between convolutional neural networks (CNNs) and graph neural networks (GNNs) is that CNNs allow for any convolutional kernel, while most GNN methods are limited to symmetric kernels (also called isotropic kernels) (Kipf & Welling, 2016; Gilmer et al., 2017). There are some implementations of asymmetric kernels using gated mechanisms (Bresson & Laurent, 2017; Veličković et al., 2017), motif attention (Peng et al., 2019), edge features (Gilmer et al., 2017), port numbering (Sato et al., 2019) or the 3D structure of molecules (Klicpera et al., 2019).

However, to the best of our knowledge, there are currently no methods that allow asymmetric graph kernels that are dependent on the full graph structure or directional flows. They either depend on local structures or local features. This is in opposition to images, which exhibit canonical directions: the horizontal and vertical axes. The absence of an analogous concept in graphs makes it difficult to define directional message passing and to produce an analogue of the directional frequency filters (or Gabor filters) widely present in image processing (Olah et al., 2020).

We propose a novel idea for GNNs: use vector fields in the

graph to define directions for the propagation of information. An overview of this framework is presented in figure 1. Using this approach, the usual message-passing structure of a GNN is projected onto globally-defined directions so that the contribution of each neighbouring node n_v is weighted by its alignment with the vector fields at the receiving node n_u . This allows to propagate information via directional derivatives or smoothing of the features.

In order to define globally consistent directional fields over general graphs, we propose to use the gradients of the lowfrequency eigenvectors ϕ_k of the graph Laplacian, since they are known to capture key information about the global structure of graphs (Bronstein et al., 2017; Chung et al., 1997). In particular, these eigenvectors can be used to define optimal partitions of the nodes in a graph, to give a natural ordering (Levy, 2006), and to find the dominant directions of the graph diffusion process (Chung & Yau, 2000). Further, we show that they generalize the horizontal and vertical directional flows in a grid (see figure 2), allowing them to guide the aggregation and mimic the asymmetric and directional kernels present in computer vision. In fact, we demonstrate mathematically that our work generalizes CNNs, by reproducing all convolutional kernels of radius R in an n-dimensional grid, while also bringing the powerful data augmentation capabilities of reflection, rotation or distortion of the directions. Additionally, we also prove that our directional graph networks (DGNs) are more discriminative than standard GNNs in regards to the Weisfeiler-Lehman 1-WL test, confirming an increase of expressiveness.

We tested our method on 5 standard datasets from (Dwivedi

^{*}Equal contribution ¹Valence Discovery, Montreal, Canada ²Department of Computer Science and Technology, University of Cambridge, Cambridge, UK ³MILA, McGill University, Montreal, Canada. Correspondence to: Dominique Beaini <dominique@valencediscovery.com>, Saro Passaro <sp976@cam.ac.uk>.

Proceedings of the 4th MLSys Conference, San Jose, CA, USA, 2021. Copyright 2021 by the author(s).

Directional Graph Networks



Figure 1. Overview of the steps required to aggregate messages in the direction of the eigenvectors.

et al., 2020) and (Hu et al., 2020), using two types of architectures, and either using or ignoring edge features, where observed state-of-the-art results from the proposed DGN.

2 THEORETICAL DEVELOPMENT

2.1 Intuitive overview

One of the biggest limitations of current GNN methods compared to CNNs is the inability to do message passing in a specific direction such as the horizontal one in a grid graph. In fact, it is difficult to define directions or coordinates based solely on the shape of the graph.

The lack of directions strongly limits the discriminative abilities of GNNs to understand local structures and simple feature transformations. Most GNNs are invariant to the permutation of the neighbours' features, so the nodes' received signal is not influenced by swapping the features of two neighbours. Therefore, several layers in a deep network will be employed to understand these simple changes instead of being used for higher level features, leading to problematic phenomena such as a over-squashing (Alon & Yahav, 2020).

In this work, one of the main contributions is the realisation that low-frequency eigenvectors of the Laplacian can overcome this limitation by providing intuitive directional flows over graphs. As a first example, taking a grid-shaped graph of size $N \times M$ with $\frac{N}{2} < M < N$, we find that the eigenvector associated to the smallest non-zero eigenvalue increases in the direction of the width N and the second one increases in the direction of the height M. This property generalizes to n-dimensional grids and motivates the use of gradients of Laplacian eigenvectors in general graphs.

We empirically validated this intuition by looking at the flow of the gradient of the eigenvectors for a variety of graphs, as shown in figure 2. For example, in the Minnesota map, the first 3 non-constant eigenvectors produce logical directions, namely South/North, suburb/city, and West/East.

Another important contribution—also noted in figure 2—is the ability to define any kind of directional flow based on prior knowledge of the problem. Hence, instead of relying on eigenvectors to find directions in a map, we can simply use the cardinal directions or the rush-hour traffic flow.

2.2 Vector fields in a graph

Based on a recent review from (Bronstein et al., 2017), this section presents the ideas of differential geometry applied to graphs, with the goal of finding proper definitions of scalar products, gradients and directional derivatives.

Let G = (V, E) be a graph with V the set of vertices and $E \subset V \times V$ the set of edges. The graph is undirected meaning that $(i, j) \in E$ iff $(j, i) \in E$. Define the vector spaces $L^2(V)$ and $L^2(E)$ as the set of maps $V \to \mathbb{R}$ and $E \to \mathbb{R}$ with $\boldsymbol{x}, \boldsymbol{y} \in L^2(V)$ and $\boldsymbol{F}, \boldsymbol{H} \in L^2(E)$ and scalar



Figure 2. Possible directional flows in different types of graphs. The node coloring is a potential map and the edges represent the gradient of the potential with the arrows in the direction of the flow. The first 3 columns present the arcosine of the normalized eigenvectors (acos $\hat{\phi}$) as node coloring, and their gradients represented as edge intensity. The last column presents examples of inductive bias introduced in the choice of direction. (a) The eigenvectors 1 and 2 are the horizontal and vertical flows of the grid. (b) The eigenvectors 1 and 2 are the flow in the longest and second-longest directions. (c) The eigenvectors 1, 2 and 3 flow respectively in the South-North, suburbs to the city center and West-East directions. We ignore ϕ_0 since it is constant and has no direction.

products

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle_{L^{2}(V)} := \sum_{i \in V} \boldsymbol{x}_{i} \boldsymbol{y}_{i}$$
$$\langle \boldsymbol{F}, \boldsymbol{H} \rangle_{L^{2}(E)} := \sum_{(i,j) \in E} \boldsymbol{F}_{(i,j)} \boldsymbol{H}_{(i,j)}$$
(1)

Think of E as the "tangent space" to V and of $L^2(E)$ as the set of "vector fields" on the space V with each row $F_{i,:}$ representing a vector at the *i*-th node.

Define the pointwise scalar product as the map $L^2(E) \times L^2(E) \to L^2(V)$ taking 2 vector fields and returning their inner product at each point of V, at the node *i* is defined by equation 2.

$$\langle \boldsymbol{F}, \boldsymbol{H} \rangle_i := \sum_{j:(i,j) \in E} \boldsymbol{F}_{i,j} \boldsymbol{H}_{i,j}$$
 (2)

In equation 3, we define the gradient ∇ as a mapping $L^2(V) \rightarrow L^2(E)$ and the divergence div as a mapping $L^2(E) \rightarrow L^2(V)$, thus leading to an analogue of the directional derivative in equation 4.

$$(\nabla \boldsymbol{x})_{(i,j)} := \boldsymbol{x}(j) - \boldsymbol{x}(i)$$

$$(\operatorname{div} \boldsymbol{F})_i := \sum_{j:(i,j)\in E} \boldsymbol{F}_{(i,j)}$$
(3)

Definition 1. The directional derivative of the function x on the graph G in the direction of the vector field \hat{F} where each vector is of unit-norm is

$$D_{\hat{F}}\boldsymbol{x}(i) := \langle \nabla \boldsymbol{x}, \hat{F} \rangle_i = \sum_{j:(i,j) \in E} (\boldsymbol{x}(j) - \boldsymbol{x}(i)) \hat{F}_{i,j} \quad (4)$$

|F| will denote the absolute value of F and $||F_{i,:}||_{L^p}$ the L^p -norm of the *i*-th row of F. We also define the forward/backward directions as the positive/negative parts of the field F^{\pm} .

2.3 Directional smoothing and derivatives

Next, we show how the vector field F is used to *guide* the graph aggregation by projecting the incoming messages. Specifically, we define the weighted aggregation matrices B_{av} and B_{dx} that allow to compute the directional smoothing and directional derivative of the node features, as presented visually in figure 1-d.

The directional average matrix B_{av} is the weighted aggregation matrix such that all weights are positives and all rows have an L^1 -norm equal to 1, as shown in equation 5 and theorem 2.1, with a proof in the appendix D.1.

$$\boldsymbol{B}_{av}(\boldsymbol{F})_{i,:} = \frac{|\boldsymbol{F}_{i,:}|}{||\boldsymbol{F}_{i,:}||_{L^1} + \epsilon}$$
(5)

The variable ϵ is an arbitrarily small positive number used to avoid floating-point errors. The L^1 -norm denominator is a local row-wise normalization. The aggregator works by assigning a large weight to the elements in the forward or backward direction of the field, while assigning a small weight to the other elements, with a total weight of 1.

Theorem 2.1 (Directional smoothing). The operation $y = B_{av}x$ is the directional average of x, in the sense that y_u is the mean of x_v , weighted by the direction and amplitude of F.

The directional derivative matrix B_{dx} is defined in (6) and theorem 2.2, with the proof in appendix D.2. Again, the denominator is a local row-wise normalization but can be replaced by a global normalization. diag(*a*) is a square, diagonal matrix with diagonal entries given by *a*. The aggregator works by subtracting the projected forward message by the backward message (similar to a center derivative), with an additional diagonal term to balance both directions.

$$\boldsymbol{B}_{dx}(\boldsymbol{F})_{i,:} = \hat{\boldsymbol{F}}_{i,:} - \operatorname{diag}\left(\sum_{j} \hat{\boldsymbol{F}}_{:,j}\right)_{i,:}$$

$$\hat{\boldsymbol{F}}_{i,:} = \left(\frac{\boldsymbol{F}_{i,:}}{||\boldsymbol{F}_{i,:}||_{L^{1}} + \epsilon}\right)$$
(6)

Theorem 2.2 (Directional derivative). Suppose \hat{F} have rows of unit L^1 norm. The operation $y = B_{dx}(\hat{F})x$ is the centered directional derivative of x in the direction of F, in the sense of equation 4, i.e.

$$oldsymbol{y} = D_{oldsymbol{\hat{F}}}oldsymbol{x} = \Big(oldsymbol{\hat{F}} - ext{diag}\Big(\sum_j oldsymbol{\hat{F}}_{:,j}\Big)\Big)oldsymbol{x}$$

These aggregators are directional, interpretable and complementary, making them ideal choices for GNNs. We discuss the choice of aggregators in more details in appendix A, while also providing alternative aggregation matrices such as the center-balanced smoothing, the forward-copy, the phantom zero-padding, and the hardening of the aggregators using softmax/argmax on the field.

2.4 Gradient of the Laplacian eigenvectors as interpretable vector fields

In this section we give theoretical support for the choice of gradients of the eigenfunctions of the Laplacian as sensible vectors along which to do directional message passing since they are interpretable and allow to reduce the oversmoothing. This section gives a theoretical ground to the intuitive directions presented in figure 2, and is the motivation behind steps (b-c) in figure 1.

As usual the combinatorial, degree-normalized and symmetric normalized Laplacian are defined as

$$L = D - A, \quad L_{\text{norm}} = D^{-1}L, \quad L_{\text{sym}} = D^{-\frac{1}{2}}LD^{-\frac{1}{2}}$$
(7)

The eigenvectors of these matrices are known to capture many essential properties of graphs, making them a natural foundation for directional message passing. Indeed, they hold such rich information about graph structure that their study is the focus of the mathematical subfield of spectral graph theory (Chung et al., 1997) and can be used to study the graph structure and Fourier modes (Hamilton, 2020). In the next paragraphs, we will prove that following the gradient of the eigenvectors allows to effectively reduce the heat-kernel distance between pairs of nodes.

Given a Markov process of k steps defined by the transition matrices W^k , j = 1, ..., k with $W = D^{-1}A$, we can define a continuous time random walk on the same graph below. Here, t is the continuous time and $p_k(x, y)$ is the probability to transition from x to y in k steps.

In (Barlow, 2017), the following identity is shown

$$q_t(x,y) = \sum_{n=0}^{\infty} \frac{e^{-t}t^k}{k!} p_k(x,y)$$

Or in matrix form $q_t = e^{t(W-I)} = e^{-tL_{norm}}$. This transition probability is also called the continuous time heat kernel because it satisfies the continuous time heat equation on graphs $\frac{d}{dt}q_t = -L_{norm}q_t$. In (Coifman & Lafon, 2006) the following distance is defined

Definition 2 (Diffusion distance). *The diffusion distance at time t between the nodes* x, y *is*

$$d_t(x,y) := \left(\sum_{z \in V} \left(q_t(x,z) - q_t(y,z)\right)^2\right)^{\frac{1}{2}}$$
(8)

Definition 3 (Gradient step). Suppose the two neighboring nodes x and z are such that $\phi(z) - \phi(x)$ is maximal among the neighbors of x, then we will say z is obtained from x by taking a step in the direction of the gradient $\nabla \phi$.

Theorem 2.3 (Gradient steps reduce diffusion distance). Let x, y be nodes such that $\phi_1(x) < \phi_1(y)$. Let x' be the node obtained from x by taking one step in the direction of $\nabla \phi_1$, then there is a constant C such that for $C \leq t$ we have

$$d_t(x', y) < d_t(x, y).$$

With the reduction in distance being proportional to $e^{-\lambda_1}$.

Hence, we see that moving from node x to node x' by following the gradient of the eigenvector ϕ_1 is guaranteed to reduce the heat kernel distance with a destination node y.

In the context of GNNs, Theorem 2.3 also has implications for the well-known problems of *over-smoothing* and *oversquashing* (Alon & Yahav, 2020; Hamilton, 2020), with the message reaching the stationary distribution of a random walk (Hamilton, 2020). Theorem 2.3 highlights how the DGN approach can alleviate these issues, since the Laplacian eigenfunctions reveal directions that allow efficient propagation of information between distant nodes instead of following a diffusion process.

Finally it is interesting to note that by selecting different eigenvectors as basis of directions, our method further aligns with a theorem that multiple independent aggregators are needed to distinguish neighbourhoods of nodes with continuous features (Corso et al., 2020).

2.5 Generalization of the convolution on a grid

In this section we show that our method generalizes CNNs by allowing to define any radius-R convolutional kernels in grid-shaped graphs. The radius-R kernel at node u is a convolutional kernel that takes the weighted sum of all nodes v at a distance $d(u, v) \leq R$.

Consider the lattice graph Γ of size $N_1 \times N_2 \times ... \times N_n$ where each vertices are connected to their direct nondiagonal neighbour. We know from Lemma D.1 that, for each dimension, there is an eigenvector that is only a function of this specific dimension. For example, the lowest frequency eigenvector ϕ_1 always flows in the direction of the longest length. Hence, the Laplacian eigenvectors of the grid can play a role analogous to the axes in Euclidean space, as shown in figure 2.

With this knowledge, we show in theorem 2.4 (proven in D.6), that we can generalize all convolutional kernels in an n-dimensional grid. This is a strong result since it demonstrates that our DGN framework generalizes CNNs when applied on a grid, thus closing the gap between GNNs and the highly successful CNNs on image tasks.

Theorem 2.4 (Generalization radius-R convolutional kernel in a lattice). For an n-dimensional lattice, any convolutional kernel of radius R can be realized by a linear combination of directional aggregation matrices and their compositions.

2.6 Comparison with Weisfeiler-Lehman (WL) test

We also compare the expressiveness of the Directional Graph Networks with the classical WL graph isomorphism test which is often used to classify the expressivity of graph neural networks (Xu et al., 2018). In theorem 2.5 (proven in appendix D.7) we show that DGNs are capable of distinguishing pairs of graphs that the 1-WL test (and so ordinary GNNs) cannot differentiate.

Theorem 2.5 (Comparison with 1-WL test). *DGNs using the mean aggregator, any directional aggregator of the first Laplacian eigenvector and injective degree-scalers are strictly more powerful than the 1-WL test.*

3 IMPLEMENTATION

We implemented the models using the DGL and Py-Torch libraries and we provide the code at the address https://github.com/Saro00/DGN. We test our method on standard benchmarks from (Dwivedi et al., 2020) and (Hu et al., 2020), namely ZINC, CIFAR10, PATTERN, Mol-HIV and MolPCBA with more details on the datasets and how we enforce a fair comparison in appendix C.2.

For the empirical experiments we inserted our proposed aggregation method in two different type of message passing architectures used in the literature: a *simple* convolutional architecture similar to the one present in GCN (equation 9a) (Kipf & Welling, 2016) and a more *complex* and general one typical of MPNNs (9b) (Gilmer et al., 2017) with or without edge features e_{ji} . The time complexity of our approach is O(Em), which is identical to PNA (Corso et al., 2020), where E is the number of edges and m the number of aggregators, with an additional O(Ek) to pre-compute the k-first eigenvectors, as explained in the appendix C.3.

$$X_i^{(t+1)} = U\left(\bigoplus_{(j,i)\in E} X_j^{(t)}\right)$$
(9a)

$$X_i^{(t+1)} = U\left(X_i^{(t)}, \bigoplus_{(j,i)\in E} M\left(X_i^{(t)}, X_j^{(t)}, \underbrace{e_{ji}}_{\text{optional}}\right)\right)$$
(9b)

Here, \bigoplus is an operator which concatenates the results of multiple aggregators, X is the node features, M is a linear transformation and U a multiple layer perceptron (MLP). This *simple* architecture of equation 9a is observed visually in steps (f-g) of figure 1.

We tested the directional aggregators across the datasets using the gradient of the first k eigenvectors $\nabla \phi_{1,...,k}$ as the underlying vector fields. Here, k is a hyperparameter, usually 1 or 2, but could be bigger for high-dimensional graphs. To deal with the arbitrary sign of the eigenvectors, we take the absolute value of the result of equation 6, making it invariant to a reflection of the field. In case of a disconnected graph, ϕ_i is the *i*-th eigenvector of each connected component. Despite the numerous aggregators proposed in appendix A, only B_{dx} and B_{av} are tested empirically.

The metrics used to measure the performance of a model depend are enforced for each dataset and provided by (Dwivedi et al., 2020) and (Hu et al., 2020). In particular, we use the mean absolute error (MAE), the accuracy (acc), the area under the receiver operating curve (ROC-AUC), and the average precision (AP).

4 RESULTS AND DISCUSSION

Directional aggregation Using the benchmarks introduced in section 3, we present in figure 3 a fair comparison of various aggregation strategies using the same parameter budget and hyperparameters. We see a consistent boost in the performance for *simple*, *complex* and *complex with edges* models using directional aggregators compared to the *mean-aggregator* baseline.

With our theoretical analysis in mind, we expected to perform well on PATTERN since the flow of the first eigenvectors are meaningful directions in a stochastic block model (i.e., these eigenvectors tend to correlate with community

Directional Graph Networks

	ZINC			PATTERN		CIFAR10		MolHIV	MolPCBA	
Aggregators	Simple	Complex	Complex-E	Simple	Complex	Simple	Complex	Simple	Complex	Complex-E
	MAE	MAE	MAE	% acc	% acc	% acc	% acc	% ROC-AUC	% AP	% AP
mean	0.316	0.353	0.262	80.77	83.34	55.9	62.8	75.1	26.04	26.38
mean pos ₁	0.349	0.332	0.297	80.76	83.74			75.8	26.97	27.50
mean pos ₁ pos ₂	0.344	0.330	0.284	84.51	81.25			76.1	26.03	25.65
mean dx ₁	0.296	0.233	0.191	84.22	83.44			78.0	26.79	27.91
mean dx ₁ dx ₂	0.337	0.271	0.205	81.61	86.62	52.9	69.8	76.5	27.16	26.55
mean av ₁	0.317	0.332	0.276	84.54	83.21			78.4	25.97	26.66
mean av ₁ av ₂	0.367	0.332	0.260	85.12	85.38	60.6	65.1	77.1	25.61	26.67
mean dx ₁ av ₁	0.290	0.245	0.192	85.17	86.68			79.0	26.40	27.47

Figure 3. Test set results using a parameter budget of ~ 100k with the same hyperparameters as (Corso et al., 2020), except MolPCBA with a budget of ~ 7*M*. The low-frequency Laplacian eigenvectors are used to define the directions, except for CIFAR10 that uses the coordinates of the image. For brevity, we denote dx_i and av_i as the directional derivative B_{dx}^i and smoothing B_{av}^i aggregators of the *i*-th direction. We also denote *pos*_i as the *i*-th eigenvector used as positional encoding for the mean aggregator.

	ZIM	١C	PATTERN	CIFAR10		MolHIV	MolPCBA
B 4edel	No edge features	Edge features	No edge features	No edge features	Edge features	No edge features	All models
woder	MAE	MAE	% acc	% acc	% acc	% ROC-AUC	% AP
GCN	0.469±0.002		65.880±0.074	54.46±0.10		76.06±0.97 *	20.20±0.24 *
GIN	0.408±0.008		85.590±0.011	53.28±3.70		75.58±1.40*	22.66±0.28 *
GraphSage	0.410±0.005		50.516±0.001	66.08±0.24			
GAT	0.463±0.002		75.824±1.823	65.48±0.33			
MoNet	0.407±0.007		85.482±0.037	53.42±0.43			
GatedGCN	0.422±0.006	0.363±0.009	84.480±0.122	69.19±0.28	69.37±0.48]	
PNA	0.320±0.032	0.188±0.004	86.567±0.075	70.46±0.44	70.47±0.72	79.05±1.32*	28.38±0.35 *
DGN	0.219±0.010	0.168±0.003	86.680±0.034	72.70±0.54	72.84±0.42	79.70±0.97	28.85±0.30 *

Figure 4. Fine-tuned results of the DGN model against models from (Dwivedi et al., 2020) and (Hu et al., 2020): GCN (Kipf & Welling, 2016), GraphSage (Hamilton et al., 2017), GIN (Xu et al., 2018), GAT (Veličković et al., 2017), MoNet (Monti et al., 2017), GatedGCN (Bresson & Laurent, 2017) and PNA (Corso et al., 2020). All the models use $\sim 100k$ parameters, except those with * who use 300k to 6.5*M*. In ZINC the DGN aggregators are {*mean*, *dx*₁, *max*, *min*}, in PATTERN {*mean*, *dx*₁, *av*₁}, in CIFAR10 {*mean*, *dx*₁, *dx*₂, *max*}, in MolHIV {*mean*, *dx*₁, *av*₁, *max*, *min*}, in MolPCBA {*mean*, *sum*, *max*, *dx*₁}.

membership). The results match our expectations, outperforming all the previous models.

In particular, we see a significant improvement in the molecular datasets (ZINC, MolHIV and MolPCBA) when using the directional aggregators, especially for the derivative aggregation B_{dx}^1 (noted dx_1 in figure 3). We believe this is due to the capacity to efficiently move messages across opposite parts of the molecule and to better understand the role of atom pairs. We further believe that the derivative aggregator is better able to capture high-frequency directional signals, similarly to the Gabor filters in computer vision.

Further, the thesis that DGNs can bridge the gap between CNNs and GNNs is supported by the clear improvements on CIFAR10 over the baselines.

Improvements over positional embeddings. In the work by (Dwivedi et al., 2020), they proposed the use of positional encoding of the eigenvectors. However, our experiments with the positional encoding of the first 2 non-trivial eigenvectors, noted pos_1 , pos_2 in figure 3, showed no clear improvement on most datasets. In fact, Dwivedi et al. noted that many eigenvectors and high network depths are required for improvements, yet we outperform their results with fewer parameters, less depth, and only 1-2 eigenvectors, further motivating their use as directional flows instead of positional encoding.

Comparison to the literature. In order to compare our model with the literature, we fine-tuned it on the various datasets and we report its performance in figure 4. We observe that DGN provides significant improvement across all benchmarks, highlighting the importance of anisotropic kernels that are dependent on the graph topology.

5 CONCLUSION

The proposed DGN method allows to address many problems of GNNs, including the lack of anisotropy, the low expressiveness, the over-smoothing and over-squashing. For the first time in graph networks, we generalize the directional properties of CNNs and their data augmentation capabilities. Based on the intuitive idea that the lowfrequency eigenvectors of the graph Laplacian gives an interpretable directional flow, we backed our work by a set of strong theoretical results showing that these eigenvectors are important in connecting nodes that are far away and improving the expressiveness in regards to the WL-test.

We further discuss avenues for future work, including the hardening of the aggregators A.5, the zero-padding at the boundaries A.7, the implementation of radius-R kernels A.8, and the study of directional data augmentation B.1.

REFERENCES

- Alon, U. and Yahav, E. On the bottleneck of graph neural networks and its practical implications. *arXiv:2006.05205 [cs, stat]*, 2020. URL http://arxiv.org/abs/2006.05205.
- Barlow, M. T. Random Walks and Heat Kernels on Graphs. London Mathematical Society Lecture Note Series. Cambridge University Press, 2017.
- Bresson, X. and Laurent, T. Residual gated graph convnets. arXiv preprint arXiv:1711.07553, 2017.
- Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., and Vandergheynst, P. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017. ISSN 1053-5888, 1558-0792. doi: 10.1109/MSP.2017.2693418. URL http://arxiv. org/abs/1611.08097.
- Chung, F. and Yau, S. T. Discrete green's functions. Journal of Combinatorial Theory, Series A, 91(1):191– 214, 2000. ISSN 0097-3165. doi: 10.1006/jcta.2000. 3094. URL http://www.sciencedirect.com/ science/article/pii/S0097316500930942.
- Chung, F., Graham, F., on Recent Advances in Spectral Graph Theory, C. C., (U.S.), N. S. F., Society, A. M., and of the Mathematical Sciences, C. B. Spectral Graph Theory. CBMS Regional Conference Series. Conference Board of the mathematical sciences, 1997. ISBN 9780821803158. URL https://books.google. ca/books?id=41K8DqAAQBAJ.
- Coifman, R. R. and Lafon, S. Diffusion maps. Applied and Computational Harmonic Analysis, 21(1):5-30, 2006. ISSN 1063-5203. doi: https://doi.org/10.1016/j.acha.2006.04.006. URL https://www.sciencedirect.com/ science/article/pii/S1063520306000546. Special Issue: Diffusion Maps and Wavelets.
- Corso, G., Cavalleri, L., Beaini, D., Liò, P., and Veličković,
 P. Principal neighbourhood aggregation for graph nets. arXiv preprint arXiv:2004.05718, 2020.
- Doshi, V. and Eun, D. Y. Fiedler vector approximation via interacting random walks. arXiv:2002.00283 [math], 2000. doi: 10.1145/3379487. URL http://arxiv. org/abs/2002.00283.
- Dwivedi, V. P., Joshi, C. K., Laurent, T., Bengio, Y., and Bresson, X. Benchmarking graph neural networks. arXiv preprint arXiv:2003.00982, 2020.
- Fiedler, M. Algebraic connectivity of graphs. *Czechoslo-vak Mathematical Journal*, 23:298–305, 01 1973. doi: 10.21136/CMJ.1973.101168.

- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1263–1272. JMLR. org, 2017.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, pp. 1024–1034, 2017.
- Hamilton, W. L. Graph Representation Learning. Morgan and Claypool, 2020.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE* conference on computer vision and pattern recognition, pp. 770–778, 2016.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*, 2020.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167, 2015.
- Islam, M. A., Jia, S., and Bruce, N. D. B. How much position information do convolutional neural networks encode? *arXiv:2001.08248 [cs]*, 2020. URL http: //arxiv.org/abs/2001.08248.
- Jin, W., Barzilay, R., and Jaakkola, T. Junction tree variational autoencoder for molecular graph generation. *arXiv:1802.04364 [cs, stat]*, 2018. URL http:// arxiv.org/abs/1802.04364.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907, 2016.
- Klicpera, J., Groß, J., and Günnemann, S. Directional message passing for molecular graphs. ICLR2020, 2019. URL https://openreview.net/forum? id=BleWbxStPH.
- Knyazev, B., Taylor, G. W., and Amer, M. Understanding attention and generalization in graph neural networks. In *Advances in Neural Information Processing Systems*, pp. 4204–4214, 2019.
- Kondor, R., Son, H. T., Pan, H., Anderson, B., and Trivedi, S. Covariant compositional networks for learning graphs. arXiv preprint arXiv:1801.02144, 2018.

Krizhevsky, A., 2009.

- Lanczos, C. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. United States Governm. Press Office Los Angeles, CA, 1950.
- Levy, B. Laplace-beltrami eigenfunctions towards an algorithm that "understands" geometry. In *IEEE International Conference on Shape Modeling and Applications* 2006 (SMI'06), pp. 13–13, 2006. doi: 10.1109/SMI. 2006.21.
- Maron, H., Ben-Hamu, H., Shamir, N., and Lipman, Y. Invariant and equivariant graph networks. arXiv preprint arXiv:1812.09902, 2018.
- Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., and Bronstein, M. M. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5115–5124, 2017.
- O'Gara, S. and McGuinness, K. Comparing data augmentation strategies for deep image classification. *Session 2: Deep Learning for Computer Vision*, 2019. doi: http://doi.org10.21427/148b-ar75. URL https: //arrow.tudublin.ie/impstwo/7.
- Olah, C., Cammarata, N., Schubert, L., Goh, G., Petrov, M., and Carter, S. An overview of early vision in InceptionV1. *Distill*, 5(4):e00024.002, 2020. ISSN 2476-0757. doi: 10.23915/distill. 00024.002. URL https://distill.pub/2020/ circuits/early-vision.
- Peng, H., Li, J., Gong, Q., Wang, S., Ning, Y., and Yu, P. S. Graph convolutional neural networks via motifbased attention. arXiv:1811.08270 [cs], 2019. URL http://arxiv.org/abs/1811.08270.
- Sato, R., Yamada, M., and Kashima, H. Approximation ratios of graph neural networks for combinatorial problems. arXiv preprint arXiv:1905.10261, 2019.
- Shorten, C. and Khoshgoftaar, T. M. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):60, 2019. ISSN 2196-1115. doi: 10.1186/s40537-019-0197-0. URL https://doi.org/10.1186/s40537-019-0197-0.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. arXiv preprint arXiv:1710.10903, 2017.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.

A APPENDIX - CHOICES OF DIRECTIONAL AGGREGATORS

This appendix helps understand the choice of B_{av} and B_{dx} in section 2.3 and presents different directional aggregators that can be used as an alternative to the ones proposed.

A simple alternative to the directional smoothing and directional derivative operator is to simply take the *forward/backward* values according to the underlying positive/negative parts of the field F, since it can effectively replicate them. However, there are many advantage of using $B_{av,dx}$. First, one can decide to use either of them and still have an interpretable aggregation with half the parameters. Then, we also notice that $B_{av,dx}$ regularize the parameter by forcing the network to take both forward and backward neighbours into account at each time, and avoids one of the neighbours becoming too important. Lastly, they are robust to a change of sign of the eigenvectors since B_{av} is sign invariant and B_{dx} will only change the sign of the results, which is not the case for *forward/backward* aggregations.

A.1 Visual interpretation of the aggregators

We further provide a visual interpretation of the B_{av} and B_{dx} aggregators in figure 5. Interestingly, we also note in appendix A.2 that B_{av} and B_{dx} yield respectively the mean and Laplacian aggregations when F is a vector field such that all entries are constant $F_{ij} = \pm C$.



Figure 5. Illustration of how the directional aggregation works at a node n_v , with the arrows representing the direction and intensity of the field F.

A.2 Retrieving the mean and Laplacian aggregations

It is interesting to note that we can recover simple aggregators from the aggregation matrices $B_{av}(F)$ and $B_{dx}(F)$. Let F be a vector field such that all edges are equally weighted $F_{ij} = \pm C$ for all edges (i, j). Then, the aggregator B_{av} is equivalent to a mean aggregation:

$$\boldsymbol{B}_{av}(\boldsymbol{F})\boldsymbol{x} = \boldsymbol{D}^{-1}\boldsymbol{A}\boldsymbol{x}$$

Under the condition $F_{ij} = C$, the differential aggregator is equivalent to a Laplacian operator L normalized using the degree D

$$\boldsymbol{B}_{dx}(C\boldsymbol{A})\boldsymbol{x} = \boldsymbol{D}^{-1}(\boldsymbol{A} - \boldsymbol{D})\boldsymbol{x} = -\boldsymbol{D}^{-1}\boldsymbol{L}\boldsymbol{x}$$

A.3 Global field normalization

The proposed aggregators are defined with a row-wise normalized field

$$\hat{oldsymbol{F}_{i,:}} = rac{oldsymbol{F}_{i,:}}{||oldsymbol{F}_{i,:}||_{L^P}}$$

meaning that all the vectors are of unit-norm and the aggregation/message passing is done only according to the direction of the vectors, not their amplitude. However, it is also possible to do a global normalization of the field F by taking a matrix-norm instead of a vector-norm. Doing so will modulate the aggregation by the amplitude of the field at each node. One needs to be careful since a global normalization might be very sensitive to the number of nodes in the graph.

A.4 Center-balanced aggregators

A problem arises in the aggregators B_{dx} and B_{av} proposed in equations 5 and 6 when there is an imbalance between the positive and negative terms of F^{\pm} . In that case, one of the directions overtakes the other in terms of associated weights.

An alternative is also to normalize the forward and backward directions separately, to avoid having either the backward or forward direction dominating the message.

$$\boldsymbol{B}_{av-center}(\boldsymbol{F})_{i,:} = \frac{\boldsymbol{F}_{i,:}^{\prime +} + \boldsymbol{F}_{i,:}^{\prime -}}{||\boldsymbol{F}_{i,j}^{\prime +} + \boldsymbol{F}_{i,j}^{\prime -}||_{L_1}} \quad , \qquad \boldsymbol{F}_{i,:}^{\prime \pm} = \frac{|\boldsymbol{F}_{i,:}^{\pm}|}{||\boldsymbol{F}_{i,:}^{\pm}||_{L^1} + \epsilon} \tag{10}$$

The same idea can be applied to the derivative aggregator equation 11 where the positive and negative parts of the field F^{\pm} are normalized separately to allow to project both the *forward* and *backward* messages into a vector field of unit-norm. F^+ is the out-going field at each node and is used for the *forward* direction, while F^- is the in-going field used for the *backward* direction. By averaging the *forward* and *backward* derivatives, the proposed matrix $B_{dx-center}$ represents the centered derivative matrix.

$$\boldsymbol{B}_{dx\text{-center}}(\boldsymbol{F})_{i,:} = \boldsymbol{F}'_{i,:} - \operatorname{diag}\left(\sum_{j} \boldsymbol{F}'_{:,j}\right)_{i,:}, \quad \boldsymbol{F}'_{i,:} = \frac{1}{2}\left(\underbrace{\frac{\boldsymbol{F}^{+}_{i,:}}_{||\boldsymbol{F}^{+}_{i,:}||_{L^{1}} + \epsilon}_{\text{forward field}} + \underbrace{\frac{\boldsymbol{F}^{-}_{i,:}}_{\text{backward field}}}_{\text{backward field}}\right)$$
(11)

A.5 Hardening the aggregators

The aggregation matrices that we proposed, mainly B_{dx} and B_{av} depend on a smooth vector field F. At any given node, the aggregation will take a weighted sum of the neighbours in relation to the direction of F. Hence, if the field F_v at a node v is *diagonal* in the sense that it gives a non-zero weight to many neighbours, then the aggregator will compute a weighted average of the neighbours.

Although there are clearly good reasons to have this weighted-average behaviour, it is not necessarily desired in every problem. For example, if we want to move a single node across the graph, this behaviour will smooth the node at every step. Instead, we propose below to soften and harden the aggregations by forcing the field into making a decision on the direction it takes.

Soft hardening the aggregation is possible by using a softmax with a temperature T on each row to obtain the field F_{softhard} .

$$(\boldsymbol{F}_{\text{softhard}})_{i,:} = \operatorname{sign}(\boldsymbol{F}_{i,:})\operatorname{softmax}(T|\boldsymbol{F}_{i,:}|)$$
(12)

Hardening the aggregation is possible by using an infinite temperature, which changes the softmax functions into argmax. In this specific case, the node with the highest component of the field will be copied, while all other nodes will be ignored.

$$(\boldsymbol{F}_{\text{hard}})_{i,:} = \operatorname{sign}(\boldsymbol{F}_{i,:})\operatorname{argmax}(|\boldsymbol{F}_{i,:}|)$$
(13)

An alternative to the aggregators above is to take the *softmin/argmin* of the negative part and the *softmax/argmax* of the positive part.

A.6 Forward and backward copy

The aggregation matrices B_{av} and B_{dx} have the nice property that if the field is flipped (change of sign), the aggregation gives the same result, except for the sign of B_{dx} . However, there are cases where we want to propagate information in the forward direction of the field, without smoothing it with the backward direction. In this case, we can define the strictly forward and strictly backward fields below, and use them directly with the aggregation matrices.

$$F_{\text{forward}} = F^+$$
, $F_{\text{backward}} = F^-$ (14)

Further, we can use the hardened fields in order to define a forward copy and backward copy, which will simply copy the node in the direction of the highest field component.

$$F_{\text{forward copy}} = F_{\text{hard}}^+$$
, $F_{\text{backward copy}} = F_{\text{hard}}^-$ (15)

A.7 Phantom zero-padding

Some recent work in computer vision has shown the importance of zero-padding to improve CNNs by allowing the network to understand it's position relative to the border (Islam et al., 2020). In contrast, using boundary conditions or reflection padding makes the network completely blind to positional information. In this section, we show that we can mimic the zero-padding in the direction of the field F for both aggregation matrices B_{av} and B_{dx} .

Starting with the B_{av} matrix, in the case of a missing neighbour in the forward/backward direction, the matrix will compensate by adding more weights to the other direction, due to the denominator which performs a normalization. Instead, we would need the matrix to consider both directions separately so that a missing direction would result in zero padding. Hence, we define $B_{av,0pad}$ below, where either the F^+ or F^- will be 0 on a boundary with strictly in-going/out-going field.

$$(\boldsymbol{B}_{av,0pad})_{i,:} = \frac{1}{2} \left(\frac{|\boldsymbol{F}_{i,:}^{+}|}{||\boldsymbol{F}_{i,:}^{+}||_{L^{1}} + \epsilon} + \frac{|\boldsymbol{F}_{i,:}^{-}|}{||\boldsymbol{F}_{i,:}^{-}||_{L^{1}} + \epsilon} \right)$$
(16)

Following the same argument, we define $B_{dx,0pad}$ below, where either the forward or backward term is ignored. The diagonal term is also removed at the boundary so that the result is a center derivative equal to the subtraction of the forward term with the 0-term on the back (or vice-versa), instead of a forward derivative.

$$B_{dx-0pad}(\mathbf{F})_{i,:} = \begin{cases} \mathbf{F}_{i,:}^{\prime+} & \text{if } \sum_{j} \mathbf{F}_{i,j}^{\prime-} = 0\\ \mathbf{F}_{i,:}^{\prime-} & \text{if } \sum_{j} \mathbf{F}_{i,j}^{\prime+} = 0\\ \frac{1}{2} \left(\mathbf{F}_{i,:}^{\prime+} + \mathbf{F}_{i,:}^{\prime-} - \text{diag} \left(\sum_{j} \mathbf{F}_{:,j}^{\prime+} + \mathbf{F}_{:,j}^{\prime-} \right)_{i,:} \right), & \text{otherwise} \end{cases}$$
(17)
$$\mathbf{F}_{i,:}^{\prime+} = \frac{\mathbf{F}_{i,:}^{+}}{||\mathbf{F}_{i,:}^{+}||_{L^{1}} + \epsilon} & \mathbf{F}_{i,:}^{\prime-} = \frac{\mathbf{F}_{i,:}^{-}}{||\mathbf{F}_{i,:}^{-}||_{L^{1}} + \epsilon}$$

A.8 Extending the radius of the aggregation kernel

We aim at providing a general radius-R kernel B_R that assigns different weights to different subsets of nodes n_u at a distance R from the center node n_v .

First, we decompose the matrix B(F) into positive and negative parts $B^{\pm}(F)$ representing the forward and backward steps aggregation in the field F.

$$B(F) = B^{+}(F) - B^{-}(F)$$
 (18)

Thus, defining $B_{fb}^{\pm}(F)_{i,:} = \frac{F_{i,:}^{\pm}}{||F_{i,:}||_{L^{p}}}$, we can find different aggregation matrices by using different combinations of walks of radius R. First demonstrated for a grid in theorem 2.4, we generalize it in equation 19 for any graph G.

Definition 4 (General radius R n-directional kernel). Let S_n be the group of permutations over n elements with a set of directional fields F_i .

$$\boldsymbol{B}_{R} := \sum_{\substack{V = \{v_{1}, v_{2}, \dots, v_{n}\} \in \mathbb{N}^{n} \\ ||V||_{L^{1}} \leq R, \quad -R \leq v_{i} \leq R \\ \text{Any choice of walk V with at most R steps}} \sum_{\substack{\sigma \in S_{n} \\ \text{optional} \\ permutations}} a_{V} \prod_{\substack{j=1 \\ Aggregator following the steps V, permuted by S_{n}}}^{N} (\boldsymbol{B}_{fb}^{sgn(v_{\sigma(j)})}(\boldsymbol{F}_{\sigma(j)}))^{|v_{\sigma(j)}|}$$
(19)

In this equation, n is the number of directional fields and R is the desired radius. V represents all the choices of walk $\{v_1, v_2, ..., v_n\}$ in the direction of the fields $\{F_1, F_2, ..., F_n\}$. For example, $V = \{3, 1, 0, -2\}$ has a radius R = 6, with 3 steps *forward* of F_1 , 1 step *forward* of F_2 , and 2 steps *backward* of F_4 . The sign of each B_{fb}^{\pm} is dependent to the sign of

 $v_{\sigma(j)}$, and the power $|v_{\sigma(j)}|$ is the number of aggregation steps in the directional field $F_{\sigma(j)}$. The full equation is thus the combination of all possible choices of paths across the set of fields F_i , with all possible permutations. Note that we are restricting the sum to v_i having only a possible sign; although matrices don't commute, we avoid choosing different signs since it will likely self-intersect a lower radius walk. The permutations σ are required since, for example, the path $up \rightarrow left$ is different (in a general graph) than the path $left \rightarrow up$.

This matrix B_R has a total of $\sum_{r=0}^{R} (2n)^r = \frac{(2n)^{R+1}-1}{2n-1}$ parameters, with a high redundancy since some permutations might be very similar, e.g. for a grid graph we have that $up \to left$ is identical to $left \to up$. Hence, we can replace the permutation S_n by a reverse ordering, meaning that $\prod_j^N B_j = B_N \dots B_2 B_1$. Doing so does not perfectly generalize the radius-R kernel for all graphs, but it generalizes it on a grid and significantly reduces the number of parameters to $\sum_{r=0}^{R} \sum_{l=1}^{min(n,r)} 2^r {n \choose l} {r-1 \choose l-1}$.

A.9 Arcsine of the eigenvectors

Since the eigenvectors ϕ_i are equivalent to the Fourier basis and represent the waves in the graphs, then it is expected that they behave similarity to sine/cosine waves when the graph is similar to a grid. This is further highlighted by the proof that the eigenvectors of a grid are all sines/cosines in appendix D.4.

Hence, when we define the field F as $F^i = \nabla \phi_i$, we must realize that the gradient will be lower near the minima/maxima of the eigenvector, as it is the case with sine/cosine waves. In the paper, we cope with this problem by dividing by the norm of the field $||F||_{L^1}$ in equations 5 and 6.

Another solution is to use the arcsine of the eigenvectors so that the function eigenvectors become similar to triangle functions and the gradient is almost uniform. However, since the arcsine function works only in the range [-1, 1], then we must first normalize the eigenvector by it's maximum, as given by equation 20.

$$\boldsymbol{F}_{asin}^{i} = \nabla \arcsin\left(\frac{\phi_{i}}{\max(|\phi_{i}|)}\right) \tag{20}$$

B APPENDIX - DATA AUGMENTATION

B.1 Generalizing image augmentation to graphs

The simplest augmentation is the vector field flipping, which is done changing the sign of the field F, as stated in definition 5. This changes the sign of B_{dx} , but leaves B_{av} unchanged.

Definition 5 (Reflection of the vector field). For a vector field F, the reflected field is -F.

Let F_1, F_2 be vector fields in a graph, with \hat{F}_1 and \hat{F}_2 being the field normalized such that each row has a unitary L^2 norm. Define the angle vector $\boldsymbol{\alpha}$ by $\langle (\hat{F}_1)_{i,:}, (\hat{F}_2)_{i,:} \rangle = \cos(\alpha_i)$. The vector field \hat{F}_2^{\perp} is the normalized component of \hat{F}_2 perpendicular to \hat{F}_1 . The equation below defines \hat{F}_2^{\perp} . The next equation defines the angle

$$(\hat{F}_{2}^{\perp})_{i,:} = rac{(\hat{F}_{2} - \langle \hat{F}_{1}, \hat{F}_{2} \rangle \hat{F}_{1})_{i,:}}{||(\hat{F}_{2} - \langle \hat{F}_{1}, \hat{F}_{2} \rangle \hat{F}_{1})_{i,:}||}$$

Notice that we then have the decomposition $(\hat{F}_2)_{i,:} = \cos(\alpha_i)(\hat{F}_1)_{i,:} + \sin(\alpha_i)(\hat{F}_2)_{i,:}$

Definition 6 (Rotation of the vector fields). For \hat{F}_1 and \hat{F}_2 non-colinear vector fields with each vector of unitary length, their rotation by the angle θ in the plane formed by $\{\hat{F}_1, \hat{F}_2\}$ is

$$\hat{F}_{1}^{\theta} = \hat{F}_{1} \operatorname{diag}(\cos \theta) + \hat{F}_{2}^{\perp} \operatorname{diag}(\sin \theta)
\hat{F}_{2}^{\theta} = \hat{F}_{1} \operatorname{diag}(\cos(\theta + \alpha)) + \hat{F}_{2}^{\perp} \operatorname{diag}(\sin(\theta + \alpha))$$
(21)

Finally, the following augmentation has a similar effect to a wave distortion applied on images.

Definition 7 (Random distortion of the vector field). For vector field \mathbf{F} and anti-symmetric random noise matrix \mathbf{R} , its randomly distorted field is $\mathbf{F}' = \mathbf{F} + \mathbf{R} \circ \mathbf{A}$.

B.2 Preliminary results of data augmentation

To evaluate the effectiveness of the proposed augmentation, we trained the models on a reduced version of the CIFAR10 dataset. The results in figure 6 show clearly a higher expressive power of the dx aggregator, enabling it to fit well the training data. For a small dataset, this comes at the cost of overfitting and a reduced test-set performance, but we observe that randomly rotating or distorting the kernels counteracts the overfitting and improves the generalization.

As expected, the performance decreases when the rotation or distortion is too high since the augmented graph changes too much. In computer vision images similar to CIFAR10 are usually rotated by less than 30° (Shorten & Khoshgoftaar, 2019; O'Gara & McGuinness, 2019). Further, due to the constant number of parameters across models, less parameters are attributed to the mean aggregation in the directional models, thus it cannot fit well the data when the rotation/distortion is too strong since the directions are less informative. We expect large models to perform better at high angles.



Figure 6. Accuracy of the various models using data augmentation with a *complex* architecture of ~ 100k parameters and trained on 10% of the CIFAR10 training set (4.5k images). An angle of x corresponds to a rotation of the kernel by a random angle sampled uniformly in $(-x^{\circ}, x^{\circ})$ using definition 6 with $F_{1,2}$ being the gradient of the horizontal/vertical coordinates. A noise of 100x% corresponds to a distortion of each eigenvector with a random noise uniformly sampled in $(-x \cdot m, x \cdot m)$ where m is the average absolute value of the eigenvector's components. The *mean* baseline model is not affected by the augmentation since it does not use the underlining vector field.

C APPENDIX - IMPLEMENTATION DETAILS

C.1 Choosing a basis of the Laplacian eigenspace

When using eigenvectors of the Laplacian ϕ_i to define directions in a graph, we need to keep in mind that there is never a single eigenvector associated to an eigenvalue, but a whole eigenspace. If an eigenvalue has multiplicity of k, the associated eigenspace has dimension k and any collection of k orthogonal vectors could be chosen as basis of that space and as vectors for the definitions of the aggregation matrices **B** defined in the previous sections.

Disconnected graphs. When a graph is disconnected, then the eigenfunctions will simply be the combination of the eigenfunctions of each connected components. Hence, one must consider ϕ_i as the *i*-th eigenvector of each component when taken separately.

Normalizing the eigenvectors. For an eigenvalue of multiplicity 1, there are always two unit norm eigenvectors of opposite sign, which poses a problem during the directional aggregation. We can make a choice of sign and later take the absolute value (i.e. B_{av} in equation 5). An alternative that applies to multiplicities higher than 1 is to take samples of orthonormal bases of the eigenspace and use each choice to augment the training (see section B.1).

Multiplicities greater than 1. Although multiplicities higher than one do happen for low-frequencies (square grids have a multiplicity 2 for λ_1) this is not common in "real-world graphs" since it suggests symmetries in the graph which are uncommon. Furthermore, we found no λ_1 multiplicity greater than 1 in the ZINC and PATTERN datasets. We further discuss these rare cases and how to deal with them in appendix C.5.

Orthogonal directions. Although all ϕ are orthogonal, their gradients, used to define directions, are not always *locally* orthogonal (e.g. there are many horizontal flows in the grid). This concern is left to be addressed in future work.

C.2 Benchmarks and datasets

We use a variety of benchmarks proposed by (Dwivedi et al., 2020) and (Hu et al., 2020) to test the empirical performance of our proposed methods. In particular, to have a wide variety of graphs and tasks we chose:

- 1. ZINC, a graph regression dataset from molecular chemistry. The task is to predict a score that is a subtraction of computed properties logP SA, with logP being the computed octanol-water partition coefficient, and SA being the synthetic accessibility score (Jin et al., 2018).
- 2. CIFAR10, a graph classification dataset from computer vision (Krizhevsky, 2009). The task is to classify the images into 10 different classes, with a total of 5000 training image per class and 1000 test image per class. Each image has 32×32 pixels, but the pixels have been clustered into a graph of ~ 100 super-pixels. Each super-pixel becomes a node in an *almost* grid-shaped graph, with 8 edges per node. The clustering uses the code from (Knyazev et al., 2019), and results in a different number of super-pixels per graph.
- 3. PATTERN, a node classification synthetic benchmark generated with Stochastic Block Models, which are widely used to model communities in social networks. The task is to classify the nodes into 2 communities and it tests the fundamental ability of recognizing specific predetermined subgraphs.
- 4. MolHIV, a graph classification benchmark from molecular chemistry. The task is to predict whether a molecule inhibits HIV virus replication or not. The molecules in the training, validation and test sets are divided using a scaffold splitting procedure that splits the molecules based on their two-dimensional structural frameworks.
- 5. MolPCBA, a graph classification benchmark from molecular chemistry. It consists of measured biological activities of small molecules generated by high-throughput screening. The dataset consists of a total of 437,929 molecules divided using a scaffold slitting procedure and a set of 128 properties to predict for each.

For the results in figure 3, our goal is to provide a fair comparison to demonstrate the capacity of our proposed aggregators. Therefore, we compare the various methods on both types of architectures using the same hyperparameters tuned in previous works (Corso et al., 2020) for similar networks. The models vary exclusively in the aggregation method and the width of the architectures to keep a set parameter budget. Following the indication of the benchmarks' authors, we averaged the performances of the models on 4 runs with different initialization seeds for the benchmarks from Dwivedi et al. (2020) (ZINC, PATTERN and CIFAR10) and 10 runs for the ones from Hu et al. (2020) (MolHIV and MolPCBA¹).

For the results in figure 4, we took the fine tuned results of other models from the corresponding public leaderboards by Dwivedi et al. (2020) and Hu et al. (2020). For the DGN results we fine tuned the model taking the lowest validation loss across runs with the following hyperparameters (you can also find the fine tuned commands in the documentation of the code repository):

- 1. ZINC: weight decay $\in \{1 \cdot 10^{-5}, 10^{-6}, 3 \cdot 10^{-7}\}$, aggregators $\in \{(mean, avg_1), (mean, dx_1), (mean, av_1, dx_1), (mean, min, max, av_1), (mean, min, max, dx_1)\}$
- 2. CIFAR10: weight decay $\in \{3 \cdot 10^{-6}\}$, dropout $\in \{0.1, 0.3\}$, aggregators $\in \{(mean, av_1, av_2), (mean, dx_1, dx_2), (mean, dx_1, dx_2, av_1, av_2), (mean, max, min, dx_1, dx_2), (mean, max, min, av_1, av_2)\}$
- 3. PATTERN: weight decay $\in \{0, 10^{-8}\}$, architecture $\in \{simple, complex\}$, aggregators $\in \{(mean, av_1), (mean, av_1, dx_1), (mean, av_1, dx_1)\}$
- 4. MolHIV: aggregators $\in \{(mean, dx_1), (mean, av_1), (mean, dx_1, av_1), (mean, max, dx_1), (mean, max, dx_1), (mean, max, min, av_1, dx_1)\}, dropout \in \{0.1, 0.3, 0.5\}, L \in \{4, 6\}$
- 5. for MolPCBA, given we did not start from any previously tuned architecture, we performed a line search with the following hyperparameters: mix of aggregators $\in \{mean, max, min, sum, dx_1, dx_2, av_1, av_2\}$, dropout $\in \{0.1, 0.2, 0.3, 0.4\}$, $L \in \{4, 6, 8\}$, weight decay $\in \{10^{-7}, 10^{-6}, 3 \cdot 10^{-6}, 10^{-5}, 3 \cdot 10^{-5}\}$, batch size $\in \{128.512.2048, 3072\}$, learning rate $\in \{10^{-2}, 10^{-3}, 5 \cdot 10^{-4}, 2 \cdot 10^{-4}\}$, learning rate patience $\in \{4, 6, 8\}$, learning rate reduce factor $\in \{0.5, 0.8\}$, architecture type $\in \{simple, complex, towers\}$, edge features dimension $\in \{0, 8, 16, 32\}$

¹For MolPCBA, due to the computational cost of running models in the large dataset and the relatively low variance, we only used 1 run for the results in figure 3, but 10 runs in those for figure 4

In CIFAR10 it is impossible to numerically compute a deterministic vector field with eigenvectors due to the multiplicity of λ_1 being greater than 1. This is caused by the symmetry of the square image, and is extremely rare in real-world graphs. Therefore, we used as underlying vector field the gradient of the coordinates of the image. Note that these directions are provided in the nodes' features in the dataset and available to all models, that they are co-linear to the eigenvectors of the grid as per lemma D.1, and that they mimic the inductive bias in CNNs.

C.3 Implementation and computational complexity

Unlike several more expressive graph networks (Kondor et al., 2018; Maron et al., 2018), our method does not require a computational complexity superlinear with the size of the graph. The calculation of the first k eigenvectors during pretraining, done using Lanczos method (Lanczos, 1950) and the sparse module of Scipy, has a time complexity of O(Ek)where E is the number of edges. During training the complexity is equivalent to a m-aggregator GNN O(Em) (Corso et al., 2020) for the aggregation and O(Nm) for the MLP.

To all the architectures we added residual connections (He et al., 2016), batch normalization (Ioffe & Szegedy, 2015) and graph size normalization (Dwivedi et al., 2020).

For some of the datasets with non-regular graphs, we combine the various aggregators with logarithmic degree-scalers as in (Corso et al., 2020).

An important thing to note is that, for dynamic graphs, the eigenvectors need to be re-computed dynamically with the changing edges. Fortunately, there are random walk based algorithms that can estimate ϕ_1 quickly, especially for small changes to the graph (Doshi & Eun, 2000). In the current empirical results, we do not work with dynamic graphs.

C.4 Running time

The precomputation of the first four eigenvectors for all the graphs in the datasets takes 38s for ZINC, 96s for PATTERN and 120s for MolHIV on CPU. Table 1 shows the average running time on GPU for all the various model from figure 3. On average, the epoch running time is 15% slower for the DGN compared to the mean aggregation, but a faster convergence for DGN means that the total training time is on average 2% faster for DGN.

Table 1. Average running time for the non-fine tuned models from figure 3. Each entry represents average time per epoch / average total training time. For the first four datasets, each of the models has a parameter budget $\sim 100k$ and was run on a Tesla T4 (15GB GPU). The *avg increase* row is the average of the relative running time of all rows compared to the *mean* row, with a negative value meaning a faster running time.

		ZINC		PATTERN			
Aggregators	Simple	Complex	Complex-E		Simple	Complex	
mean	3.29s/1505s	3.58s/1584s	3.56s/1654s	1.	53.1s/10154s	117.8s/9031s	
mean dx_1	3.86s/1122s	3.77s/1278s	4.22s/1371s	1	44.9s/8109s	127.2s/8417s	
mean $dx_1 dx_2$	4.23s/1360s	4.55s/1560s	4.63s/1680s	1	53.3s/8057s	167.9s/9326s	
mean av ₁	3.68s/1297s	3.84s/1398s	3.92s/1272s	1	28.0s/8680s	88.1s/7456s	
mean $av_1 av_2$	3.95s/1432s	4.03s/1596s	4.07s/1721s	1	34.2s/8115s	170.4s/11114s	
mean $dx_1 av_1$	3.89s/1079s	4.09s/1242s	4.58s/1510s	1	18.6s/6221s	144.2s/9112s	
avg increase	+19%/-16% +13%/-11%		+20%/-9%		-11%/-23%	+18%/+1%	
	CIF	FAR10	MolHI	V	Mol	IPCBA	
Aggregators	Simple	Complex	Simple	e	Complex	Complex-E	
mean	83.6s/10526s	78.7s/10900	s 11.4s/21	39s	279s/30128s	356s/38126s	
mean dx_1			12.6s/234	48s	304s/34129s	461s/43419s	
mean $dx_1 dx_2$	98.4s/8405s	100.9s/5191	s 14.1s/234	45s	314s/36581s	334s/38363s	
mean av ₁			12.2s/21	77s	297s/30316s	436s/54545s	
mean $av_1 av_2$	117.1s/12834s	s 89.5s/14481	s 13.9s/21	50s	315s/42297s	333s/36641s	
mean $dx_1 av_1$			14.0s/20	70s	326s/37523s	461s/59109s	
avg increase	+29%/+1%	+21%/-10%	6 +17%/+1	%	+12%/+20%	+14%/+22%	

C.5 Eigenvector multiplicity

The possibility to define equivariant directions using the low-frequency Laplacian eigenvectors is subject to the uniqueness of those vectors. When the dimension of the eigenspaces associated with the lowest eigenvalues is 1, the eigenvectors are defined up to a constant factor. In section 2.4, we propose the use of unit vector normalization and an absolute value to eliminate the scale and sign ambiguity. When the dimension of those eigenspaces is greater than 1, it is not possible to define equivariant directions using the eigenvectors.

Fortunately, it is very rare for the Laplacian matrix to have repeated eigenvalues in real-world datasets. We validate this claim by looking at ZINC and PATTERN datasets where we found no graphs with repeated Fiedler vector and only one graph out of 26k with multiplicity of the second eigenvector greater than 1.

When facing a graph that presents repeated Laplacian eigenvalues, we propose to randomly shuffle, during training time, different eigenvectors randomly sampled in the eigenspace. This technique will act as a data augmentation of the graph during training time allowing the network to train with multiple directions at the same time.

D APPENDIX - MATHEMATICAL PROOFS

D.1 Proof for theorem 2.1 (Directional smoothing)

The operation $y = B_{av}x$ is the directional average of x, in the sense that y_u is the mean of x_v , weighted by the direction and amplitude of F.

Proof. This should be a simple proof, that if we want a weighted average of our neighbours, we simply need to multiply the weights by each neighbour, and divide by the sum of the weights. Of course, the weights should be positive.

D.2 Proof for theorem 2.2 (Directional derivative)

Suppose \hat{F} have rows of unit L^1 norm. The operation $y = B_{dx}(\hat{F})x$ is the centered directional derivative of x in the direction of F, in the sense of equation 4, i.e.

$$oldsymbol{y} = D_{oldsymbol{\hat{F}}}oldsymbol{x} = \Big(oldsymbol{\hat{F}} - ext{diag}\Big(\sum_j oldsymbol{\hat{F}}_{:,j}\Big)\Big)oldsymbol{x}$$

Proof. Since F rows have unit L^1 norm, $\hat{F} = F$. The *i*-th coordinate of the vector $\left(F - \operatorname{diag}\left(\sum_j F_{:,j}\right)\right) x$ is

$$\left(\mathbf{F}\mathbf{x} - \operatorname{diag}\left(\sum_{j} \mathbf{F}\right) \mathbf{x} \right)_{i} = \sum_{j} \mathbf{F}_{i,j} \mathbf{x}(j) - \left(\sum_{j} \mathbf{F}_{i,j}\right) \mathbf{x}(i)$$
$$= \sum_{j:(i,j)\in E} (\mathbf{x}(j) - \mathbf{x}(i)) \mathbf{F}_{i,j}$$
$$= D_{\mathbf{F}} \mathbf{x}(i)$$

D.3 Proof of theorem 2.3 (Gradient steps reduce diffusion distance)

Let x, y be nodes such that $\phi_1(x) < \phi_1(y)$. Let x' be the node obtained from x by taking one step in the direction of $\nabla \phi_1$, then there is a constant C such that for $C \le t$ we have

$$d_t(x', y) < d_t(x, y).$$

With the reduction in distance being proportional to $e^{-\lambda_1}$.

Proof. Recall that $p_k(x,y) = (D^{-1}A)_{x,y}^k$ is the discrete heat kernel at step k, $q_t(x,y) = \sum_{k\geq 0} \frac{e^{-t}t^k}{k!} p_k(x,y)$ is the continuous heat kernel at time t. In (Barlow, 2017), it is shown that the continuous heat kernel is computed by $q_t(x,y) = e^{-tL_{\text{norm}}}$. Following (Coifman & Lafon, 2006) we can diagonalise q_t to get the identity

$$d_t(x,y) = \left(\sum_{i=1}^{n-1} e^{-2t\lambda_i} \left(\phi_i(x) - \phi_i(y)\right)^2\right)^{\frac{1}{2}}$$
(22)

The inequality $d_t(x', y) < d_t(x, y)$ is equivalent to

$$\sum_{i=2}^{n-1} e^{-2t\lambda_i} \left(\left(\phi_i(x') - \phi_i(y) \right)^2 - \left(\phi_i(x) - \phi_i(y) \right)^2 \right) < e^{-2t\lambda_1} \left(\left(\phi_1(x) - \phi_1(y) \right)^2 - \left(\phi_1(x') - \phi_1(y) \right)^2 \right)$$
(23)

The term on the left is bounded above by

$$\sum_{i=2}^{n-1} e^{-2t\lambda_i} \left| \left(\phi_i(x') - \phi_i(y) \right)^2 - \left(\phi_i(x) - \phi_i(y) \right)^2 \right|$$

and this last term is in turn bounded above by

$$e^{-2t\lambda_2} \sum_{i=2}^{n-1} \left| \left(\phi_i(x') - \phi_i(y) \right)^2 - \left(\phi_i(x) - \phi_i(y) \right)^2 \right|$$

Inequality 23 will then hold if

$$e^{-2t\lambda_2} \sum_{i=2}^{n-1} \left| \left(\phi_i(x') - \phi_i(y) \right)^2 - \left(\phi_i(x) - \phi_i(y) \right)^2 \right| < e^{-2t\lambda_1} \left(\left(\phi_1(x) - \phi_1(y) \right)^2 - \left(\phi_1(x') - \phi_1(y) \right)^2 \right)$$

and this is equivalent to

$$\frac{1}{2(\lambda_1 - \lambda_2)} \log \left(\frac{\left(\left(\phi_1(x) - \phi_1(y) \right)^2 - \left(\phi_1(x') - \phi_1(y) \right)^2 \right)}{\sum_{i=2}^{n-1} \left| \left(\phi_i(x') - \phi_i(y) \right)^2 - \left(\phi_i(x) - \phi_i(y) \right)^2 \right|} \right) < t$$

if we take t to be larger than the term on the left the inequality we get $d_t(x', y) < d_t(x, y)$.

The constant C in the statement is the constant on the left side of the inequality. It is also interesting to note that C is expected to be positive since the term $\lambda_1 - \lambda_2$ is negative and the argument of the log will most likely be < 1.

D.4 Proof for Lemma D.1 (Cosine eigenvectors)

Consider the lattice graph Γ of size $N_1 \times N_2 \times ... \times N_n$, that has vertices $\prod_{i=1,...,n} \{1,...,N_i\}$ and the vertices $(x_i)_{i=1,...,n}$ and $(y_i)_{i=1,...,n}$ are connected by an edge iff $|x_i - y_i| = 1$ for one index *i* and 0 for all other indices. Note that there are no diagonal edges in the lattice. The eigenvector of the Laplacian of the grid $L(\Gamma)$ are given by ϕ_j .

Lemma D.1 (Cosine eigenvectors). The Laplacian of Γ has an eigenvalue $2 - 2\cos\left(\frac{\pi}{N_i}\right)$ with the associated eigenvector ϕ_j that depends only the variable in the *i*-th dimension and is constant in all others, with $\phi_j = \mathbf{1}_{N_1} \otimes \mathbf{1}_{N_2} \otimes ... \otimes \mathbf{x}_{1,N_i} \otimes ... \otimes \mathbf{1}_{N_n}$, and $\mathbf{x}_{1,N_i}(j) = \cos\left(\frac{\pi j}{n} - \frac{\pi}{2n}\right)$

Proof. First, recall the well known result that the path graph on N vertices P_N has eigenvalues

$$\lambda_k = 2 - 2\cos\left(\frac{\pi k}{n}\right)$$

with associated eigenvector \boldsymbol{x}_k with *i*-th coordinate

$$\boldsymbol{x}_k(i) = \cos\left(rac{\pi ki}{n} + rac{\pi k}{2n}
ight)$$

The Cartesian product of two graphs $G = (V_G, E_G)$ and $H = (V_H, E_H)$ is defined as $G \times H = (V_{G \times H}, E_{G \times H})$ with $V_{G \times H} = V_G \times V_H$ and $((u_1, u_2), ((v_1, v_2)) \in E_{G \times H}$ iff either $u_1 = v_1$ and $(u_2, v_2) \in E_H$ or $(u_1, v_1) \in V_G$ and $u_2 = v_2$. It is shown in (Fiedler, 1973) that if $(\mu_i)_{i=1,...,m}$ and $(\lambda_j)_{j=1,...,n}$ are the eigenvalues of G and H respectively, then the eigenvalues of the Cartesian product graph $G \times H$ are $\mu_i + \lambda_j$ for all possible eigenvalues μ_i and λ_j . Also, the eigenvectors associated to the eigenvalue $\mu_i + \lambda_j$ are $u_i \otimes v_j$ with u_i an eigenvector of the Laplacian of G associated to the eigenvalue μ_i and v_j an eigenvector of the Laplacian of H associated to the eigenvalue λ_j .

Finally, noticing that a lattice of shape $N_1 \times N_2 \times ... \times N_n$ is really the Cartesian product of path graphs of length N_1 up to N_n , we conclude that there are eigenvalues $2 - 2 \cos\left(\frac{\pi}{N_i}\right)$. Denoting by $\mathbf{1}_{N_j}$ the vector in \mathbf{R}^{N_j} with only ones as coordinates, then the eigenvector associated to the eigenvalue $2 - 2 \cos\left(\frac{\pi}{N_i}\right)$ is

$$\mathbf{1}_{N_1}\otimes \mathbf{1}_{N_2}\otimes ...\otimes x_{1,N_i}\otimes ...\otimes \mathbf{1}_{N_n}$$

where x_{1,N_i} is the eigenvector of the Laplacian of P_{N_i} associated to its first non-zero eigenvalue. $2 - 2\cos\left(\frac{\pi}{N_i}\right)$.

D.5 Radius 1 convolution kernels in a grid

In this section we show any radius 1 convolution kernel can be obtained as a linear combination of the $B_{dx}(\nabla \phi_i)$ and $B_{av}(\nabla \phi_i)$ matrices for the right choice of Laplacian eigenvectors ϕ_i . First we show this can be done for 1-d convolution kernels.

Theorem D.2. On a path graph, any 1D convolution kernel of size 3 k is a linear combination of the aggregators B_{av} , B_{dx} and the identity I.

Proof. Recall from the previous proof that the first non zero eigenvalue of the path graph P_N has associated eigenvector $\phi_1(i) = \cos(\frac{\pi i}{N} - \frac{\pi}{2N})$. Since this is a monotone decreasing function in *i*, the *i*-th row of $\nabla \phi_1$ will be

$$(0, ..., 0, s_{i-1}, 0, -s_{i+1}, 0, ..., 0)$$

with s_{i-1} and $s_{i+1} > 0$. We are trying to solve

$$(a\mathbf{B}_{av} + b\mathbf{B}_{dx} + c\mathbf{Id})_{i,:} = (0, ..., 0, x, y, z, 0, ..., 0)$$

with x, y, z, in positions i - 1, i and i + 1. This simplifies to solving

$$a\frac{1}{\|s\|_{L^1}}|s| + b\frac{1}{\|s\|_{L^2}}s + c(0,1,0) = (x,y,z)$$

with $s = (s_{i-1}, 0, -s_{i+1})$, which always has a solution because $s_{i-1}, s_{i+1} > 0$.

Theorem D.3 (Generalization radius-1 convolutional kernel in a grid). Let Γ be the *n*-dimensional lattice as above and let ϕ_j be the eigenvectors of the Laplacian of the lattice as in theorem D.1. Then any radius 1 kernel k on Γ is a linear combination of the aggregators $B_{av}(\phi_i)$, $B_{dx}(\phi_i)$ and I.

Proof. This is a direct consequence of D.2 obtained by adding n 1-dimensional kernels, with each kernel being in a different axis of the grid as per Lemma D.1. See figure 7 for a visual example in 2D.



Figure 7. Realization of a radius-1 convolution using the proposed aggregators. I_x is the input feature map, * the convolutional operator, I_y the convolution result, and $B^i = B(\nabla \phi_i)$.

D.6 Proof for theorem 2.4 (Generalization radius-*R* convolutional kernel in a lattice)

For an n-dimensional lattice, any convolutional kernel of radius R can be realized by a linear combination of directional aggregation matrices and their compositions.

Proof. For clarity, we first do the 2 dimensional case for a radius 2, then extended to the general case. Let k be the radius 2 kernel on a grid represented by the matrix

$$\boldsymbol{a}_{5\times5} = \begin{pmatrix} 0 & 0 & a_{-2,0} & 0 & 0 \\ 0 & a_{-1,-1} & a_{-1,0} & a_{-1,1} & 0 \\ a_{0,-2} & a_{0,-1} & a_{0,0} & a_{0,1} & a_{0,2} \\ 0 & a_{1,-1} & a_{1,0} & a_{1,1} & 0 \\ 0 & 0 & a_{2,0} & 0 & 0 \end{pmatrix}$$

since we supposed the $N_1 \times N_2$ grid was such that $N_1 > N_2$, by theorem D.1, we have that ϕ_1 is depending only in the first variable x_1 and is monotone in x_1 . Recall from D.1 that

$$\phi_1(i) = \cos\left(\frac{\pi i}{N_1} + \frac{\pi}{2N_1}\right)$$

The vector $\frac{N_1}{\pi} \nabla \arccos(\phi_1)$ will be denoted by F_1 in the rest. Notice all entries of F_1 are 0 or ± 1 . Denote by F_2 the gradient vector $\frac{N_2}{\pi} \nabla \arccos(\phi_k)$ where ϕ_k is the eigenvector given by theorem D.1 that is depending only in the second variable x_2 and is monotone in x_1 and recall

$$\phi_k(i) = \cos\left(\frac{\pi i}{N_2} + \frac{\pi}{2N_2}\right)$$

For a matrix B, let B^{\pm} the positive/negative parts of B, is matrices with positive entries such that $B = B^{+} - B^{-}$. Let B_{r1} be a matrix representing the radius 1 kernel with weights

$$\boldsymbol{a}_{3\times3} = \begin{pmatrix} 0 & a_{-1,0} & 0 \\ a_{0,-1} & a_{0,0} & a_{0,1} \\ 0 & a_{1,0} & 0 \end{pmatrix}$$

The matrix B_{r1} can be obtained by theorem D.3. Then the radius 2 kernel k is defined by all the possible combinations of 2 positive/negative steps, plus the initial radius-1 kernel.

$$\boldsymbol{B}_{r2} = \sum_{\substack{-2 \leq i, j \leq 2\\|i|+|j|=2}} \underbrace{\left(a_{i,j}(\boldsymbol{F}_1^{sgn(i)})^{|i|}(\boldsymbol{F}_2^{sgn(j)})^{|j|}\right)}_{\text{Any combination of 2 steps}} + \underbrace{\boldsymbol{B}_{r1}}_{\text{all possible single-steps}}$$

with sgn the sign function sgn(i) = + if $i \ge 0$ and - if i < 0. The matrix B_{r2} then realises the kernel $a_{5\times 5}$.

We can further extend the above construction to ${\cal N}$ dimension grids and radius ${\cal R}$ kernels k

$$\underbrace{\sum_{\substack{V = \{v_1, v_2, \dots, v_N\} \in \mathbb{N}^n \\ ||V||_{L^1} \leq R \\ -R \leq v_i \leq R}}_{V \in \mathcal{I}_i \leq R} a_V \underbrace{\prod_{j=1}^{i^*} (F_j^{sgn(v_j)})^{|v_j|}}_{\text{Aggregator following the steps defined in } V}$$

N

Any choice of walk V with at most R-steps

with $F_j = \frac{N_j}{\pi} \nabla \arccos \phi_j$, ϕ_j the eigenvector with lowest eigenvalue only dependent on the *j*-th variable and given in theorem D.1 and \prod is the matrix multiplication. V represents all the choices of walk $\{v_1, v_2, ..., v_n\}$ in the direction of the fields $\{F_1, F_2, ..., F_n\}$. For example, $V = \{3, 1, 0, -2\}$ has a radius R = 6, with 3 steps *forward* of F_1 , 1 step *forward* of F_2 , and 2 steps *backward* of F_4 .

D.7 Proof for theorem 2.5 (Comparison with 1-WL test)

DGNs using the *mean* aggregator, any directional aggregator of the first Laplacian eigenvector and injective degree-scalers are strictly more powerful than the 1-WL test.

Proof. We will show that (1) DGNs are at least as powerful as the 1-WL test and (2) there is a pair of graphs which are not distinguishable by the 1-WL test which DGNs can discriminate.

Since the DGNs include the mean aggregator combined with at least an injective degree-scaler, (Corso et al., 2020) show that the resulting architecture is at least as powerful as the 1-WL test.



Figure 8. Illustration of an example pair of graphs which the 1-WL test cannot distinguish but DGNs can. The table shows the node feature updates done at every layer. MPNN with mean/sum aggregators and the 1-WL test only use the updates in the first row and therefore cannot distinguish between the nodes in the two graphs. DGNs also use directional aggregators that, with the vector field given by the first eigenvector of the Laplacian matrix, provides different updates to the nodes in the two graphs.

Then, to show that the DGNs are strictly more powerful than the 1-WL test it suffices to provide an example of a pair of graphs that DGNs can differentiate and 1-WL cannot. Such a pair of graphs is illustrated in figure 8.

The 1-WL test (as any MPNN with, for example, sum aggregator) will always have the same features for all the nodes labelled with a and for all the nodes labelled with b and, therefore, will classify the graphs as isomorphic. DGNs, via the directional smoothing or directional derivative aggregators based on the first eigenvector of the Laplacian matrix, will update the features of the a nodes differently in the two graphs (figure 8 presents also the aggregation functions) and will, therefore, be capable of distinguishing them.