

Learning Label Initialization for *Time-Dependent* Harmonic Extension

Amitoz Azad

University of Caen Normandy

amitoz.sudo@gmail.com

Abstract

Node classification on graphs can be formulated as the Dirichlet problem on graphs where the signal is given at the labeled nodes, and the harmonic extension is done on the unlabeled nodes. This paper considers a *time-dependent* version of the Dirichlet problem on graphs and shows how to improve its solution by learning the proper initialization vector on the unlabeled nodes. Further, we show that the improved solution is at par with state-of-the-art methods used for node classification. Finally, we conclude this paper by discussing the importance of parameter t , pros, and future directions.

1 Introduction

Node classification on graphs has various real-world applications [Song *et al.*, 2021] ranging from computer vision, natural language processing, social network, biomedical sciences. Often it is done in a semi-supervised setting in which the labeled data is significantly less than the unlabeled data, and the goal is to use both the labeled and unlabeled data to obtain a good classification score. This makes it an ideal approach for the problems where labeled data is not cheap or scarcely available.

Many state-of-the-art methods for node classification use machine learning on graphs [Kipf and Welling, 2017; Veličković *et al.*, 2018; Monti *et al.*, 2017]. Graph Neural Networks (GNNs) are emerging as powerful tools for graph representation learning. The earliest method for node classification on graphs used Laplacian regularization [Zhu *et al.*, 2003; Belkin *et al.*, 2004; Zhou and Schölkopf, 2005]. The Laplacian regularization can also be formulated as a PDE approach on graphs (Dirichlet problem (10)). Several non-trivial Laplacians have been proposed, and the Dirichlet problem associated with each one of them could be used for node classification [Calder, 2018; Kyng *et al.*, 2015; Calder and Slepčev, 2020; El Alaoui *et al.*, 2016]. In this work, we consider the Dirichlet problem associated with normalized Laplacian as defined in [Zhou and Schölkopf, 2005]. More specifically, we consider the time-dependent formulation of the Dirichlet problem (15) and argue about making its solution as competitive as state-of-the-art methods for transductive learning (Table 2).

The triggering point of this work was the comparison of GCN’s [Kipf and Welling, 2017] accuracy and the solution of the Dirichlet problem (15) on handcrafted *knn* graphs for MNIST [LeCun *et al.*, 1998], FMNIST [Xiao *et al.*, 2017], and the popular citation graphs Cora, Citeseer, Pubmed [Sen *et al.*, 2008]. We observed that while the solutions of the PDE and the GCN are comparable for node classification on the handcrafted *knn* graphs, the latter performed significantly better on the citation graphs (Table 1). For the handcrafted graphs, the nodes correspond to the images, and node features are the pixel values of the images. The plausible reason for the similar performances on handcrafted *knn* graphs would be that edges connect the similar nodes, as the euclidean distance between the images happens to be a good metric for the graph construction in this case. The citation graphs are given to us by default; the nodes represent documents, edges correspond to citation links between the documents rather than their similarity, and each node has a high dimensional bag-of-words feature vector. The apparent reason for the superior performance of the GCN on citation graphs would be that not only does it take the graph as an input, but it also takes the node features as input and utilizes their hidden representations.

Main contributions. Recent work has shown that GNNs are good at utilizing the node feature information [Faber *et al.*, 2021]. Since the GCN beats the PDE-based approach on citation graphs by employing the node features along with the graph. This led us to ask: *How to make the solution of the PDE (15) at par with the GCN by utilizing the node feature information?*

- We show that this is achieved by utilizing the node features to learn the initial condition (15c) of the PDE (Figure 1).
- We benchmark the improved solution against the several state-of-the-art methods for node classification. Our approach yields competitive results (Table 2, Row 08).
- Although we learn the initialization vector to boost the performance of the PDE, further marginal gains are also added by learning the graph weights (Table 2, Row 09).
- Unlike GNNs, which need to be retrained every time to incorporate new labels, this approach can easily incorporate new labels after training. We demonstrate this by incorporating validation labels (Table 2, Row 11). Importantly, we do not use the new labels to retrain.

2 Background

In this section, we review notation on graphs and formulate node-classification as the *time-dependent* Dirichlet problem on graphs.

2.1 Preliminaries

A weighted graph $G = (V, E, w)$ consists of a finite set V of N nodes and a finite set $E \subset V \times V$ of edges. Let (u, v) be an edge connecting the nodes u and v . A weighted graph is associated with a weight function $w : V \times V \rightarrow [0, 1]$. It represents a measure of similarity between two nodes. The set of edges according to the weight function is given as: $E = \{(u, v) | w(u, v) \neq 0\}$. The set of nodes in the neighborhood of node u is denoted as $N(u)$. The notation $v \in N(u)$ means node v is in the neighborhood of node u , i.e. $N(u) = \{v \in V | (u, v) \in E\}$. In this paper, we consider symmetric graphs i.e. $w(u, v) = w(v, u)$ and $(u, v) \in E \leftrightarrow (v, u) \in E$. The degree of a node u is given as: $\delta_w(u) = \sum_{v \in N(u)} w(u, v)$.

Let $H(V)$ be a Hilbert space of the real-valued function on the nodes of the graph. A function $f : V \rightarrow R$ of $H(V)$, represents a signal on a node and assigns a real value $f(u)$ to each node $u \in V$. Similarly, let $H(E)$ be a Hilbert space of the real-valued function defined on the edges. These two spaces are equipped with the following inner products:

$$\begin{aligned} \langle f, g \rangle_{H(V)} &= \sum_{u \in V} f(u)g(u), \quad \forall f, g \in H(V) \\ \langle F, G \rangle_{H(E)} &= \sum_{(u, v) \in E} F(u, v)G(u, v), \quad \forall F, G \in H(E) \end{aligned} \quad (1)$$

The weighted graph gradient (or difference) operator $\nabla_w : H(V) \rightarrow H(E)$ is defined as:

$$\nabla_w(f)(u, v) := \sqrt{\frac{w(u, v)}{\delta_w(v)}} f(v) - \sqrt{\frac{w(u, v)}{\delta_w(u)}} f(u) \quad (2)$$

The norm of the graph gradient on each node is given as:

$$\|\nabla_w(f)(u)\| = \left(\sum_{v \in N(u)} (\nabla_w(f))^2(u, v) \right)^{\frac{1}{2}} \quad (3)$$

The *Dirichlet* energy associated with the signal $f(u)$ is defined as :

$$J(f) := \frac{1}{2} \sum_{v \in V} \|\nabla_w(f)(u)\|^2 \quad (4)$$

Often it is also known as *Tikhonov regularization* [Belkin et al., 2004]. The weighted graph divergence operator $div_w : H(E) \rightarrow H(V)$ satisfies the discrete version of stokes law:

$$\begin{aligned} \langle \nabla_w(f), F \rangle_{H(E)} &= \langle f, -div_w(F) \rangle_{H(V)} \\ &\text{where } f \in H(V) \text{ and } F \in H(E) \end{aligned} \quad (5)$$

Using (1), (5) and let $f = \mathbf{1}_{\{u\}}$, it can be shown that divergence operator is given as:

$$div_w(F)(u) = \sum_{v \in N(u)} \sqrt{\frac{w(u, v)}{\delta_w(u)}} (F(u, v) - F(v, u)) \quad (6)$$

The graph Laplacian is defined as $\Delta_w : H(V) \rightarrow H(V)$:

$$\Delta_w(f)(u) := \frac{1}{2} div_w(\nabla_w(f))(u) \quad (7)$$

Using the definitions of ∇_w and div_w , the Laplacian is given as:

$$\Delta_w(f)(u) = \sum_{v \in N(u)} w(u, v) \left(\frac{f(v)}{\sqrt{\delta_w(u)\delta_w(v)}} - \frac{f(u)}{\delta_w(u)} \right) \quad (8)$$

2.2 Node Classification

Let us consider the problem of node classification on graphs in transductive settings for a multi-class problem. Let $V_0 = \{u_1, u_2, \dots, u_m\}$ be a subset of nodes in V over which the labels are given as $y_1, y_2, \dots, y_m \in \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_k\}$, where $\mathbf{e}_i \in R^k$ represents the i^{th} class out of k classes (one-hot vector). The goal of transductive learning is to extend these labels to the unlabeled nodes. One way to achieve this is via solving the following inverse problem:

$$\arg \min_{f(u)} \left\{ \sum_{u \in V} \frac{1}{2} \|\nabla_w(f)(u)\|^2 : f(u) = g(u) \quad \forall u \in V_0 \right\} \quad (9)$$

Here, $g : V_0 \rightarrow R^k$ is the label function on the nodes, such that $g(u) \in \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_k\}$. The other way to extend the known labels is via solving the following PDE (*a.k.a the Dirichlet problem*) on the graph, which is obtained through the E-L equation of the above optimization:

$$\begin{aligned} \Delta_w(f)(u) &= 0 \quad \forall u \in V \setminus V_0 \\ f(u) &= g(u) \quad \forall u \in V_0 \end{aligned} \quad (10)$$

It is worth remembering that the above equation is vector-valued PDE, with k components. The final label decision for the unlabeled node u is determined by the largest component of $f(u)$:

$$l(u) = \arg \max_{j \in \{1, \dots, k\}} \{f_j(u)\} \quad (11)$$

2.3 Connection With Label Propagation

We now digress briefly to mention the connection with the widely-used label propagation algorithm [Zhu, 2005]. For a symmetric graph and under somewhat different definitions of gradient and divergence

$$\begin{aligned} \nabla_w(f)(u, v) &= \sqrt{w(u, v)}(f(v) - f(u)) \\ div_w(F)(u) &= \sum_{v \in V} \sqrt{w(u, v)}(F(u, v) - F(v, u)), \end{aligned} \quad (12)$$

the Laplacian on the left-hand side of (10) yields the classic unnormalized (*a.k.a combinatorial*) Laplacian [Chung and Graham, 1997]. The Dirichlet problem then becomes:

$$\begin{aligned} \sum_{v \in N(u)} w(u, v)(f(v) - f(u)) &= 0, \quad \forall u \in V \setminus V_0 \\ f(u) &= g(u), \quad \forall u \in V_0 \end{aligned} \quad (13)$$

Dataset	Eq15	GCN
MNIST	93.2 ± 0.0	91.3 ± 0.2
FMNIST	76.0 ± 0.0	77.4 ± 0.2
Cora	72.5 ± 0.0	81.5 ± — [†]
Citeseer	49.7 ± 0.0	70.3 ± — [†]
Pubmed	72.5 ± 0.0	79.0 ± — [†]

Table 1: Performances of (15) and the GCN on handcrafted and citation graphs. Observe the significant differences in scores over the latter. [†]Values taken from the GCN paper.

The above equation can be solved by the Jacobi’s iterative method. Which yields the classic label propagation:

$$f^{n+1}(u) = \frac{\sum_{v \in N(u)} w(u, v) f^n(v)}{\sum_{v \in N(u)} w(u, v)}, \quad \forall u \in V \setminus V_0$$

$$f^{n+1}(u) = g(u), \quad \forall u \in V_0 \quad (14)$$

2.4 Time-Dependent Dirichlet Problem

A traditional approach to solve E-L equation of a variational problem is to make it *time-dependent*. A celebrated example of this is *R.O.F.* PDE, which is used in image processing for anisotropic denoising [Rudin *et al.*, 1992]. Formulating the Dirichlet problem in (10) as time-dependent, it becomes:

$$\frac{\partial f(u, t)}{\partial t} = \Delta_w(f)(u, t) \quad \forall u \in V \setminus V_0 \quad (15a)$$

$$f(u, t) = g(u) \quad \forall u \in V_0 \quad (15b)$$

$$f(u, 0) = \psi_0(u) \quad \forall u \in V \quad (15c)$$

At steady-state, when $\frac{\partial f}{\partial t} \rightarrow 0$, the solution of above equation is equivalent to that of (10). Basically, it is a gradient flow in the direction of minimizing the Dirichlet energy while satisfying the boundary constraint. The above equation can also be viewed as heat diffusion with non-homogeneous boundary condition on graphs. Notice that the Dirichlet problem in (10) had only a boundary condition, but now it also has an initial condition, denoted as ψ_0 .

A common choice of setting ψ_0 is to let the signal be $g(u) \in R^k$ (one-hot vectors) on the labeled nodes and zero vectors elsewhere. It can be seen as an initial unnormalized probability vector on nodes. In this paper, we choose to address it as ‘front ψ_0 ’ since the approach in (15) is similar to level-set (front-propagation) methods in computer vision [Osher, 1993; Osher and Paragios, 2003].¹

3 Experiments and Results

This section walks through the research questions (RQs) and discusses the experiments done to answer them.

Code. The scripts used for the experiments and the instructions to run are available at Github: <https://github.com/aGIToz/Learning-Label-Initialization>.

¹Certain Dirichlet problems can be formulated as *time-dependent* PDE with no boundary condition.

RQ1. *How does the solution of (15) compare with the GCN [Kipf and Welling, 2017] for semi-supervised node classification?*

To answer this question, we tested them on handcrafted (*knn*) graphs and non-handcrafted citation graphs.

Handcrafted graphs. We created *knn* graphs for two datasets: MNIST [LeCun *et al.*, 1998] and F(fashion)-MNIST [Xiao *et al.*, 2017]. MNIST is a popular dataset of handwritten digits and FMNIST is a dataset from Zalando’s clothing images; both consist of 70k examples of sizes 28 by 28 and 10 classes. For the graph construction k was set to 10 and we used RBF kernel $w(u, v) = e^{-d(u, v)^2 / \sigma^2}$ for weights.

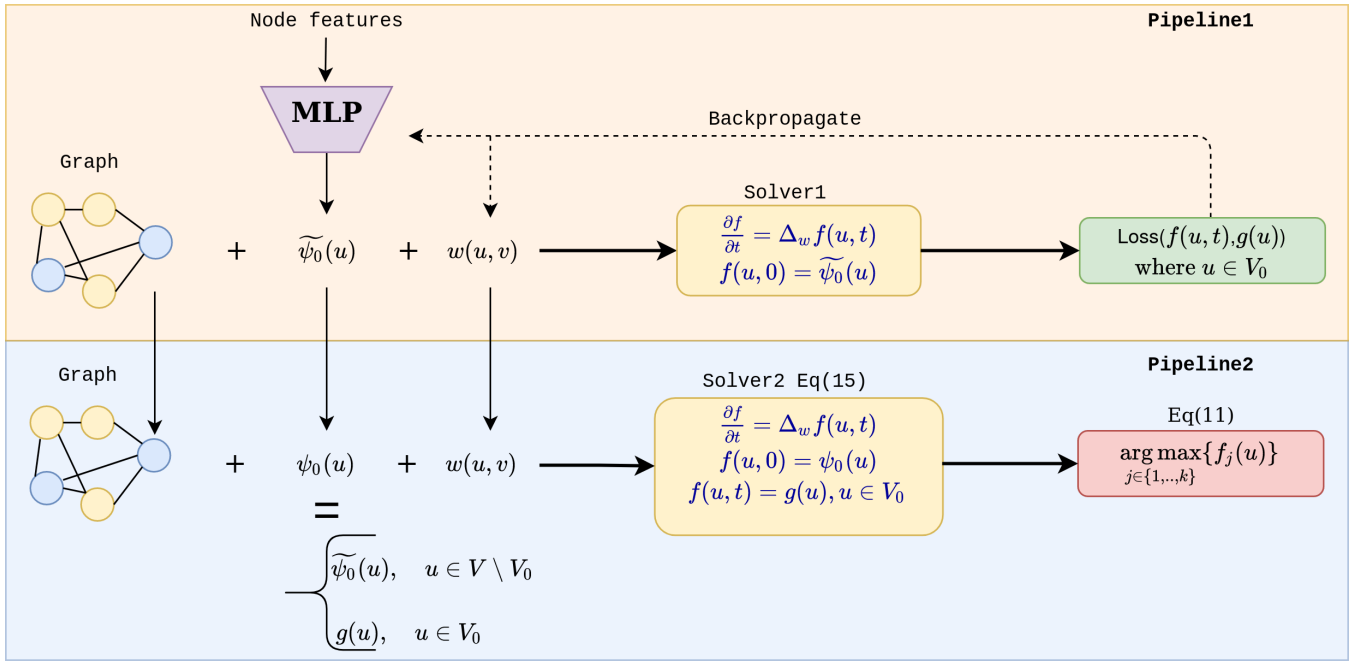
Non-handcrafted graphs. We used the popular citation graphs, which have been widely used for benchmarking Graph nets. These are Cora, Citeseer, Pubmed [Sen *et al.*, 2008]. Here the nodes represent documents; edges correspond to citation links between them, and each node has a sparse bag-of-words feature vector and a class label.

Software. We used PyTorch framework [Paszke *et al.*, 2019]. To run the geometric PDE (15), we used torch-geometric [Fey and Lenssen, 2019] and torchdiffeq [Chen *et al.*, 2018]. Torchdiffeq is a popular ODE solver written in PyTorch. It allows to backpropagate via adjoint sensitivity method [Pontryagin *et al.*, 1962] and comes with different numerical schemes. For all the experiments, we used the Dormand-Prince (dopri5) method. It is also recommended by the authors of the library, and it is a common choice in NeuralODE literature (Sec 4.1).

Split setting. For the MNIST and FMNIST datasets, we kept the seed size (labeled nodes) 20 nodes per class, for validation 500 nodes per class, and the rest for the test set. The citation graphs come with a prebuilt split (*a.k.a. planetoid split*), which is very often used for benchmarking node classification with GNNs. These datasets come with default training, validation, and test sets. The training sets of the citation graphs were used as seeds (labeled nodes). The validation sets of all the datasets were used to tune the hyperparameters of the GCN and the ODE solver. It was also used to tune σ parameter in the RBF kernel in the weights of handcrafted *knn* graphs. The initialization of the front ψ_0 was done as a one-hot vector on the labeled nodes and zero vector on the unlabelled nodes.

Evaluation. We report the accuracy on the test nodes for 100 random iterations (Table 1). The GCN’s results on citation graphs are taken from their paper [Kipf and Welling, 2017]. For every dataset, the same graph structure was given to the ODE solver and the GCN. Note that the performance of GCN is sensitive to weight initialization (random iterations) and hence comes with variations, but the solution of (15) is deterministic and does not fluctuate for different iterations (Table 1).

Observations. We observe a notable discrepancy between the performances of the GCN and (15) for citation graphs compared to the handcrafted *knn* graphs (Table 1). The GCN performs significantly better on the citation graphs. The likely reason for the poor performance of the PDE on the citation graphs could be attributed to the fact that the edges


 Figure 1: Proposed architecture to learn the front ψ_0 on the unlabeled nodes.

do not explicitly translate to the similarity between the nodes on these graphs. In fact, they correspond to citation links between the documents. This breaks the premise of LP (label propagation) like algorithms, which assume that edges and weights correspond to the similarity between the nodes. Often LP like algorithms are evaluated on the graphs constructed using node features, but in this case, the graphs correspond to citation *networks*, and they are already made for us. The GCN, on the other hand, is able to perform well on citation networks as it uses the graph structure and node features, whereas (15) only takes the graph structure as an input. This leads us to our next research question.

RQ2. *Knowing that the GCN beats the solution of (15) by utilizing the node features and the graph structure, could we use the node features and make the solution of (15) at par with the GCN on the citation graphs?*

We show that this is indeed possible by learning the front ψ_0 (15c) on the unlabeled nodes. Not only does this improve the solution, but the results are made as competitive as several state-of-the-art methods (Table 2).

Proposed model. Figure 1 shows the proposed architecture to enhance the solution of (15). It is composed of two Pipelines. Pipeline1 is a training pipeline; it is used for learning the ψ_0 on unlabeled nodes. Node features are plugged into an MLP to generate an initial estimate of the front ($\tilde{\psi}_0$), which is further given as an input to Solver1. Apart from $\tilde{\psi}_0$, the Solver1 also takes the graph structure and graph weights $w(u, v)$ as inputs. The graph weights are also made learnable. Although not shown in the Figure 1, the graph weights are learned in such a way to ensure they are symmetric $w(u, v) = w(v, u)$, and between zero and one. The loss function in Pipeline1 could be seen as a way to enforce

the boundary condition (15b); one could argue that such reasoning is motivated by the inspiring work in PINNs [Raissi *et al.*, 2019], where a boundary condition of a PDE is turned into a loss function. Pipeline2 corresponds to the evaluation pipeline. After every epoch, the values of $\tilde{\psi}_0$ and $w(u, v)$ are updated via full-batch gradient descent, and fed into Solver2 to evaluate the PDE on the validation set and for saving the front and the weights. Pipeline2 could be regarded as the final model as we use it in the end, to evaluate the performance on the test set, and it can be separately deployed once the front and the weights are learned. Note the relationship between ψ_0 and $\tilde{\psi}_0$. For the unlabeled nodes, ψ_0 vector is equal to $\tilde{\psi}_0$ vector, and for the labeled nodes, ψ_0 vector is equal to the $g(u)$ vector (one-hot). Unlike Solver1, the Solver2 keeps the signal on the labeled nodes as $g(u)$ throughout the evolution of (15) from 0 to t .²

Baselines. We compare the improved performances of (15) with three graph networks: Graph Convolutional Network (GCN) [Kipf and Welling, 2017], Graph Attention Network (GAT) [Veličković *et al.*, 2018], Mixture Model Networks (MoNet) [Monti *et al.*, 2017]. Additionally, we compare to a hybrid model Correct and Smooth (CS) [Huang *et al.*, 2021] and to Continuous Graph Neural Networks (CGNN) [Xhonneux *et al.*, 2020] with and without feature mixing. The latter two models are discussed briefly in Sec 4.1.

Datasets. Apart from the citation graphs, we include three more real-world datasets: Amazon Photo [Shchur *et al.*,

²The reader may be intrigued about the need for Solver2. Why not put the boundary condition (15b) in Solver1 and minimize the loss function? It is not possible as the loss will remain zero.

Model	Cora	Citeseer	Pubmed	Photo	CountyFB	FMAAsia
01 GCN	81.5 ± — [†]	70.3 ± — [†]	79.0 ± — [†]	93.4 ± 0.6	76.4 ± 1.4	85.9 ± 0.7
02 GAT	83.0 ± 0.7 [†]	72.5 ± 0.7 [†]	79.0 ± 0.3 [†]	94.1 ± 0.7	76.3 ± 1.6	87.0 ± 0.8
03 MoNet	81.7 ± 0.5 [†]	70.3 ± 0.9	78.8 ± 0.4 [†]	94.3 ± 0.3	79.4 ± 1.5	86.1 ± 1.0
04 C&S (MLP)	80.0 ± 0.9	67.3 ± 1.3	75.9 ± 1.9	94.2 ± 0.5	82.2 ± 1.2	87.3 ± 0.7
05 CGNN(w/o)	82.7 ± 0.7 [†]	71.5 ± 1.1 [†]	81.7 ± 0.5 [†]	93.3 ± 0.2	80.5 ± 1.4	85.3 ± 0.4
06 CGNN(w)	82.8 ± 0.6 [†]	72.1 ± 0.8 [†]	81.7 ± 0.8 [†]	92.2 ± 0.7	81.9 ± 0.9	85.4 ± 0.3
07 Eq15	72.5 ± 0.0	49.7 ± 0.0	72.5 ± 0.0	91.3 ± 0.5	81.4 ± 1.8	86.7 ± 0.8
08 Eq15 learned ψ_0	82.9 ± 0.7	72.2 ± 0.6	81.1 ± 0.9	94.3 ± 0.5	82.4 ± 1.5	87.0 ± 0.9
09 Eq15 learned ψ_0 & w	83.5 ± 0.7	72.4 ± 0.7	81.9 ± 0.9	94.4 ± 0.5	82.6 ± 1.8	87.1 ± 0.8
10 Eq15 inc. val	79.5 ± 0.0	60.4 ± 0.0	78.4 ± 0.0	92.4 ± 0.6	82.8 ± 1.4	87.2 ± 0.7
11 Eq15 inc. val, (ψ_0, w)	87.2 ± 0.0	73.3 ± 0.0	83.5 ± 0.0	95.5 ± 0.5	85.3 ± 1.4	88.6 ± 0.8

Table 2: Test accuracy over different datasets. [†]Values taken from the original papers. Coloring scheme: red-1st, blue-2nd, green-3rd.

2018], CountyFB [Jia and Benson, 2021], FMAAsia [Rozemberczki and Sarkar, 2020]. In Amazon Photo nodes represent items, and edges represent if two items are frequently bought together. Node feature is product review as a bag-of-words vector, and the task is to map the nodes to their respective product category. In CountyFB dataset, the nodes are US Counties, and edges come from FacebookSocialConnectedness index. Nodes have demographic features (income, migration, education ...) and few features from Facebook friendships. The task is to predict the 2016 election outcome (Republican vs Democrat). The nodes in FMAAsia dataset represents users of streaming service LastFM in Asia, and edges are friendships between them. The task is to predict the home country of the users.

Evaluation setting. We use the standard planetoid split setting for the citation graphs and report the performance over 100 random iterations. The planetoid split amounts to just 5% of seeds (training set) per class. Small training splits in graph datasets could produce results sensitive to the choice of the split [Shchur *et al.*, 2018]. Therefore, for Photo, CountyFB, and FMAAsia, we keep the split size 40% (train), 40% (valid), 20% (test) per class and evaluate the performance for 10 different splits.³

Training setting. We used Adam or RMSprop optimizer in all our experiments. The optimization was done with full-batch gradient descent. The L2 weight regularization and dropouts were also used in all optimizations. Most of the MLP used in the architecture (Figure 1) are with one or two hidden channels with ReLU activations. The loss function was cross entropy. The hyperparameters of the Solver1 and Solver2 were kept same. The Dormand-Prince (dopri5) numerical scheme was used for all the experiments related to (15). The validation splits were used for tuning the hyperparams and for the final model selection (*saving the front and the weights*). Hyperparameter optimization was done using a large random search (coarse to fine). We used either categorical distribution or uniform distribution on the hyperparams while running the random search. The search was done using

NVIDIA 1080Ti and P100 GPUs.⁴

Observations. We see that the proposed approach significantly improves the performance of (15) and makes it competitive with several other methods (Table 2). Row 07 corresponds to the classical solution of (15) (with zeros and one-hot vectors initialization). Row 08 shows the improvement in the solution with the learned front ψ_0 . And Row 09 shows further slight gains by learning the weights along with the front ψ_0 . Notice that while (15) is inherently deterministic, there are variations from Row 07 onwards. These variations can be attributed to the weight initializations in MLP (Pipeline1), and the fact that evaluation is done for the multiple test splits for Photo, CountyFB, and FMAAsia datasets.

RQ3. *How to incorporate new labels in the final model once the training has been ended?*

This can be quickly done by including the new labels in the boundary condition (15b) and then updating a saved front ψ_0 ⁵ by assigning one-hot vectors to the new labels and finally running Pipeline2. Row 10 (Table 2) corresponds to the solution of (15), including the validation labels in (15b) and with the default initialization of the front. Row 11 corresponds to the solution of (15) by updating a saved front ψ_0 and using saved weights.

4 Discussion

In this section, we discuss some relevant works and conclude this work with key take-homes.

4.1 Related Work

Hybrid models. The Correct & Smooth (C&S) [Huang *et al.*, 2021] paper is relevant to our work. Their work focuses on first generating base predictions by training an MLP, which is then followed by a correction step (error propagation) and subsequently a label smoothing step. All three steps are done separately. The correction and smoothing steps are done using the iterative scheme in [Zhou *et al.*, 2004]. On the other hand, we learn the base predictions (the front ψ_0),

³Each split is generated by torch.random.seed from 0 to 9.

⁴The final evaluation was done on 1080Ti.

⁵Learned by training on the training set only.

but there is no error propagation step, and the final smoothing is done via the PDE (15) in Pipeline2. Also, unlike their work, learning the base prediction is followed by a smoothing step in Pipeline1 (Solver1) and the MLP is learned in an end-to-end manner. Interestingly, we found their approach to be significantly more effective for larger seed size (Table 2). The APPNP [Klicpera *et al.*, 2018] model was a precursor to C&S paper. It consists of using a neural network to generate base predictions and then smoothing using power iterations of topic-sensitive page rank algorithm [Haveliwala, 2003], and the whole model is trained in an end-to-end manner. It is similar to Pipeline1 of our model, which is also trained end-to-end, except that smoothing in Pipeline1 is done via a PDE. They had no counterpart of Pipeline2. Another interesting but distinct hybrid model is from [Wang and Leskovec, 2020], in which they propose to learn graph weights as a pre-processing step and then use those graph weights in a GCN for node classification.

NeuralODEs. There has been some prior work of bringing NeuralODEs [Chen *et al.*, 2018] to do transductive learning on Graphs. The earlier attempts using NeuralODE framework parameterized the derivative function with a few layers of GNN [Poli *et al.*, 2020; Zhuang *et al.*, 2019]. The first attempt to do the transductive learning using a diffusion based method without parameterizing the derivative was in CGNN [Xhonneux *et al.*, 2020]. GRAND [Chamberlain *et al.*, 2021] went steps further and added the attention mechanism [Vaswani *et al.*, 2017] to learn the graph weights, graph rewiring after each backward pass, and proposed a non-linear version. It is worth noting that in all these models, the evolution of a PDE is done in a latent space. That is, the evolution is stacked between an encoder and decoder using MLPs. Whereas, in our case, the evolution occurs in the actual space of the label vector \mathbf{R}^k as required by (15); also, the Dirichlet boundary condition (15b) is respected during the evaluation on the validation and test sets. Another important difference is that our model can incorporate new labels without retraining the model (Table 2).

4.2 Conclusion

Importance of parameter t . The parameter t at which the evolution of (15) is stopped is quite important. It acts like a scaling parameter of a variational problem. It was carefully tuned to yield the best accuracy and avoid reaching steady-state. This is because at steady-state, one gets the harmonic extension, and the solution of the Dirichlet problem is unique⁶ and independent of the initialization on the unlabeled nodes.

Pros. An essential advantage of our model is the ability to incorporate new labels without retraining the whole model (Table 2). It has a practical application for a huge graph where retraining would require considerable time. It must be emphasized that GNNs and NeuralODEs do not have this advantage, although this advantage was also a highlight of C&S paper [Huang *et al.*, 2021]. Deeper GCNs are known to suffer from oversmoothing and cannot propagate information to longer range [Li *et al.*, 2018]. However, the proposed

approach avoids this limiting behavior for t tends to infinity.⁷ Finally, we believe the most significant advantage of this approach is simplicity, which makes it easier to debug. For example, in the case of absurd predictions over some nodes, one could quickly check the learned front ψ_0 and interpret the absurdity in terms of initial probability. This is not possible in the case of GNNs, or if the evolution of a PDE is stacked between encoder-decoder.

Societal impact. Trustworthiness in A.I. is a growing concern. The key to using A.I. more responsibly to serve society is to use effective yet straightforward ML models. Table 2 shows that our model is effective. It is simple as it basically corresponds to the PDE (15), and in this work, we show how to ameliorate its solution by learning the initial front on unlabeled nodes.

Future directions. *Machine learning perspective:* for a very large and dense graph, learning graph weights directly may not be efficient. One way to circumvent this issue is to bring the attention mechanisms [Vaswani *et al.*, 2017; Veličković *et al.*, 2018] to learn similarities between the features. *PDE perspective:* the Dirichlet energy in (9) corresponds to $p = 2$ case of p -Dirichlet energy, given as $J_p(f) = \frac{1}{p} \sum_{u \in V} \|\nabla_w(f)(u)\|^p$. It would be interesting to consider the variational problem (9) for $p = 1$ because the PDE corresponding to that would be anisotropic diffusion (edge-preserving), which yields piecewise constant results. Hence it is desirable as intuitively, the vector f should be composed of zeros and ones. Another interesting direction would be to consider the Dirichlet problem with novel Laplacians [Calder and Slepčev, 2020; Calder, 2018; El Alaoui *et al.*, 2016], which effectively do semi-supervised learning for very small seed sizes.

Acknowledgements

This work was done with the support of the French National Research Agency through the project ANR-17-CE38-0004. The author thanks Julien Rabin for insightful discussions at the start of this work.

References

- [Belkin *et al.*, 2004] Mikhail Belkin, Irina Matveeva, and Partha Niyogi. Regularization and semi-supervised learning on large graphs. In *COLT*, 2004.
- [Calder and Slepčev, 2020] Jeff Calder and Dejan Slepčev. Properly-weighted graph laplacian for semi-supervised learning. *Applied mathematics & optimization*, 82(3):1111–1159, 2020.
- [Calder, 2018] Jeff Calder. The game theoretic p -laplacian and semi-supervised learning with few labels. *Nonlinearity*, 32(1):301, 2018.
- [Chamberlain *et al.*, 2021] Benjamin Paul Chamberlain, James Rowbottom, Maria Gorinova, Stefan Webb, Emanuele Rossi, and Michael M Bronstein. Grand: Graph neural diffusion. In *ICML*, 2021.

⁶Assuming the graph is connected.

⁷Remark: In NeuralODE literature the *Number Of Function Evaluations* is considered as a proxy for the depth.

- [Chen *et al.*, 2018] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *NeurIPS*, 2018.
- [Chung and Graham, 1997] Fan RK Chung and Fan Chung Graham. *Spectral graph theory*. AMS, 1997.
- [El Alaoui *et al.*, 2016] Ahmed El Alaoui, Xiang Cheng, Aaditya Ramdas, Martin J Wainwright, and Michael I Jordan. Asymptotic behavior of ℓ_p -based laplacian regularization in semi-supervised learning. In *COLT*, 2016.
- [Faber *et al.*, 2021] Lukas Faber, Yifan Lu, and Roger Wattenhofer. Should graph neural networks use features, edges, or both? *arXiv:2103.06857*, 2021.
- [Fey and Lenssen, 2019] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. In *ICLR Workshop*, 2019.
- [Haveliwala, 2003] Taher H Haveliwala. Topic-sensitive pagerank: A context-sensitive ranking algorithm for web search. *IEEE transactions on knowledge and data engineering*, 2003.
- [Huang *et al.*, 2021] Qian Huang, Horace He, Abhay Singh, Ser-Nam Lim, and Austin R Benson. Combining label propagation and simple models out-performs graph neural networks. In *ICLR*, 2021.
- [Jia and Benson, 2021] Junteng Jia and Austin R Benson. A unifying generative model for graph learning algorithms: Label propagation, graph convolutions, and combinations. *arXiv:2101.07730*, 2021.
- [Kipf and Welling, 2017] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- [Klicpera *et al.*, 2018] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate: Graph neural networks meet personalized pagerank. In *ICLR*, 2018.
- [Kyng *et al.*, 2015] Rasmus Kyng, Anup Rao, Sushant Sachdeva, and Daniel A Spielman. Algorithms for lipschitz learning on graphs. In *COLT*, 2015.
- [LeCun *et al.*, 1998] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *IEEE*, 1998.
- [Li *et al.*, 2018] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *AAAI*, 2018.
- [Monti *et al.*, 2017] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *CVPR*, 2017.
- [Osher and Paragios, 2003] Stanley Osher and Nikos Paragios. *Geometric level set methods in imaging, vision, and graphics*. Springer Science & Business Media, 2003.
- [Osher, 1993] Stanley Osher. A level set formulation for the solution of the dirichlet problem for hamilton-jacobi equations. *SIAM Journal on Mathematical Analysis*, 1993.
- [Paszke *et al.*, 2019] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- [Poli *et al.*, 2020] Michael Poli, Stefano Massaroli, Junyoung Park, Atsushi Yamashita, Hajime Asama, and Jinkyoo Park. Graph neural ordinary differential equations. In *AAAI-DLG*, 2020.
- [Pontryagin *et al.*, 1962] LS Pontryagin, EF Mishchenko, VG Boltyanskiy, and RV Gamkrelidze. Mathematical theory of optimal processes. 1962.
- [Raissi *et al.*, 2019] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 2019.
- [Rozemberczki and Sarkar, 2020] Benedek Rozemberczki and Rik Sarkar. Characteristic functions on graphs: Birds of a feather, from statistical descriptors to parametric models. In *ACM ICIKM*, 2020.
- [Rudin *et al.*, 1992] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D : Nonlinear phenomena*, 1992.
- [Sen *et al.*, 2008] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 2008.
- [Shchur *et al.*, 2018] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. Pitfalls of graph neural network evaluation. *arXiv:1811.05868*, 2018.
- [Song *et al.*, 2021] Zixing Song, Xiangli Yang, Zenglin Xu, and Irwin King. Graph-based semi-supervised learning: A comprehensive review. *arXiv:2102.13303*, 2021.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [Veličković *et al.*, 2018] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2018.
- [Wang and Leskovec, 2020] Hongwei Wang and Jure Leskovec. Unifying graph convolutional neural networks and label propagation. *arXiv:2002.06755*, 2020.
- [Xhonneux *et al.*, 2020] Louis-Pascal Xhonneux, Meng Qu, and Jian Tang. Continuous graph neural networks. In *ICML*, 2020.
- [Xiao *et al.*, 2017] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv:1708.07747*, 2017.
- [Zhou and Schölkopf, 2005] Dengyong Zhou and Bernhard Schölkopf. Regularization on discrete spaces. In *Joint Pattern Recognition Symposium*, 2005.
- [Zhou *et al.*, 2004] Dengyong Zhou, Olivier Bousquet, Thomas N Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. In *NeurIPS*, 2004.
- [Zhu *et al.*, 2003] Xiaojin Zhu, Zoubin Ghahramani, and John D Lafferty. Semi-supervised learning using gaussian fields and harmonic functions. In *ICML*, 2003.
- [Zhu, 2005] Xiaojin Zhu. *Semi-supervised learning with graphs*. PhD thesis, Carnegie Mellon University, 2005.
- [Zhuang *et al.*, 2019] Juntang Zhuang, Nicha Dvornek, Xiaoxiao Li, and James S Duncan. Ordinary differential equations on graph networks. Technical report, 2019.