

A Lagrangian Approach to Information Propagation in Graph Neural Networks

Matteo Tiezzi¹ and Giuseppe Marra²
and Stefano Melacci¹ and Marco Maggini¹ and Marco Gori¹

Abstract. In many real world applications, data are characterized by a complex structure, that can be naturally encoded as a graph. In the last years, the popularity of deep learning techniques has renewed the interest in neural models able to process complex patterns. In particular, inspired by the Graph Neural Network (GNN) model, different architectures have been proposed to extend the original GNN scheme. GNNs exploit a set of state variables, each assigned to a graph node, and a diffusion mechanism of the states among neighbor nodes, to implement an iterative procedure to compute the fixed point of the (learnable) state transition function. In this paper, we propose a novel approach to the state computation and the learning algorithm for GNNs, based on a constraint optimisation task solved in the Lagrangian framework. The state convergence procedure is implicitly expressed by the constraint satisfaction mechanism and does not require a separate iterative phase for each epoch of the learning procedure. In fact, the computational structure is based on the search for saddle points of the Lagrangian in the adjoint space composed of weights, neural outputs (node states), and Lagrange multipliers. The proposed approach is compared experimentally with other popular models for processing graphs.

1 Introduction

Due to their flexibility and approximation capabilities, the original processing and learning schemata of Neural Networks have been extended in order to deal with structured inputs. Based on the original feedforward model, able to process vectors of features as inputs, different architectures have been proposed to process sequences (Recurrent Neural Networks [23]), rasters of pixels (Convolutional Neural Networks [15]), directed acyclic graphs (Recursive Neural Networks [10, 9]), and general graph structures (Graph Neural Networks [20]). All these models generally share the same learning mechanism based on the error BackPropagation (BP) through the network architecture, that allows the computation of the loss gradient with respect to the connection weights. When processing structured data the original BP schema is straightforwardly extended by the process of *unfolding* that generates a network topology based on the current input structure by replicating a base neural network module (e.g. BP Through Time, BP Through Structure).

However, recently, some works [5] proposed a different approach to learning neural networks, where neural computations are ex-

pressed as constraints and the optimization is framed into the Lagrangian framework. These algorithms are naturally local and allow the learning of any computational structure, both acyclical or cyclical. The main drawback of these methods is that they are quite memory inefficient; in particular, they need to keep an extra-variable for each hidden neuron and for each example. This makes them inapplicable to large problems where BP is still the only viable option.

Graph Neural Networks (GNNs) [20] exploit neural networks to learn how to encode nodes of a graph for a given task taking into account both the information local to each node and the whole graph topology. The learning process requires, for each epoch, an iterative diffusion mechanism up to convergence to a stable fixed point, that is computationally heavy. A maximum number of iterations can be defined but this limits the local encoding to a maximum depth of the neighborhood of each node. In this paper, we propose a new learning mechanism for GNNs based on a Lagrangian formulation that allows the embedding of the fixed point computation into the problem constraints. In the proposed scheme the network state representations and the weights are jointly optimized without the need of applying the fixed point relaxation procedure at each weight update epoch.

The paper is organized as follows. The next section reviews the main developments in both the Neural Network models for processing graphs and learning methods based on the Lagrangian approach. Section 3 introduces the basics of the GNN model, whereas in Section 4 the Lagrangian formulation of GNNs is described. Section 5 reports the experimental evaluation of the proposed constraint-based learning for GNNs. Finally, the conclusions are drawn in Section 6.

2 Related Works

In many applications data are characterized by an underlying structure that lays on a non-Euclidean domain, i.e. graphs and manifolds. Whilst commonly addressed in relational learning, such domains have been initially not taken into account by popular machine learning techniques, that have been mostly devised for grid-like and Euclidean structured data [3]. Early machine learning approaches for structured data were designed for directed acyclic graphs [21, 9], while a more general framework was introduced in [20]. GNNs are able to directly deal with directed, undirected and cyclic graphs. The core idea is based on an iterative scheme of information diffusion among neighboring nodes, involving a propagation process aimed at reaching an equilibrium of the node states that represent a local encoding of the graph for a given task. The encoding is a computationally expensive process being based on the computation of the fixed point of the state transition function. Some proposals were aimed at simplifying this step, such as the scheme proposed in [17] that ex-

¹ Department of Information Engineering and Mathematical Sciences, University of Siena, Italy {mtiezzi, melacci, maggini, marco}@diism.unisi.it

² Department of Information Engineering, University of Florence, Italy, g.marra@unifi.it

plots gated recurrent units.

Recent approaches differ in the choice of neighborhood aggregation method and graph level pooling scheme, and can be categorized in two main areas. *Spectral approaches* exploit particular embeddings of the graph and the convolution operation defined in the spectral domain [4]. However, they are characterized by computational drawbacks caused by the eigen-decomposition of the graph Laplacian. Simplified approaches are based on smooth reparametrization [12] or approximation of the spectral filters by a Chebyshev expansion [7]. Finally, in Graph Convolutional Networks (GCNs) [14], filters are restricted to operate in a 1-hop neighborhood of each node. *Spatial methods*, instead, exploit directly the graph topology, without the need of an intermediate representation. These approaches differ mainly in the definition of the aggregation operator used to compute the node states, that must be able to maintain weight sharing properties and to process nodes with different numbers of neighbors. The PATCHY-SAN [18] model converts graph-structured data into a grid-structured representation, extracting and normalizing neighborhoods containing a fixed number of nodes. In [8] the model exploits a weight matrix for each node degree, whereas DCNNs [1] compute the hidden node representation convolving inputs channels with power series of the transition probability matrix, learning weights for each neighborhood degree. GraphSAGE [11] exploits different aggregation functions to merge the node neighborhood information. Deep GNNs [2] stack layers of GNNs to obtain a deep architecture. In the context of graph classification tasks, SortPooling [26] uses a framework based on DGCNNs with a pooling strategy, that performs pooling by ordering vertices. Finally, the representational and discriminative power of GNN models were explored in [24], also introducing the novel GIN model.

A Lagrangian formulation of learning can be found in the seminal work of Yann LeCun [16], which studies a theoretical framework for Backpropagation. More recently, Carreira and Wang [5] introduced the idea of training networks, transformed into a constraints-based representation, though an extension of the learnable parameters space. Their optimization scheme was based on quadratic penalties, aiming at an approximate solution of the problem afterwards refined by a post-processing phase. Differently, [22] exploits closed-form solutions were most of the architectural constraints are softly enforced, and further additional variables are introduced to parametrize the neuron activations.

By framing the optimization of neural networks in the Lagrangian framework, where neural computations are expressed as constraints, their main goal is to obtain a local algorithm where computations of different layers can be carried out in parallel. On the contrary in the proposed approach, we use a novel mixed strategy. In particular, the majority of the computations still rely on Backpropagation while constraints are exploited only to express the diffusion mechanism. This allows to carry out both the optimization of the neural functions and the diffusion process at the same time, instead of alternating them into two distinct phases (as in [20]), with a theoretical framework supporting this approach (Lagrangian optimization).

It has already been shown that algorithms on graphs can be effectively learned exploiting a constrained fixed-point formulation. For example, SSE [6] exploits the Reinforcement Learning *policy iteration* algorithm for the interleaved evaluation of the fixed point equation and the improvement of the transition and output functions. Our approach, starting from similar assumptions, exploits the unifying Lagrangian framework for learning both the transition and the output functions. Thus, by framing the optimization algorithm into a standard gradient descent/ascent scheme, we are allowed to

use recent update rules (e.g. Adam) without the need to resort to ad-hoc moving average updates.

3 Graph Neural Networks

The term Graph Neural Network (GNN) refers to a general computational model, that exploits the processing and learning schemes of neural networks to process non Euclidean data, i.e. data organized as graphs.

Given an input graph $G = (V, E)$, where V is a finite set of *nodes* and $E \subseteq V \times V$ collects the *arcs*, GNNs apply a two-phase computation on G . In the *encoding* (or *aggregation*) phase the model computes a state vector for each node in V by (iteratively) combining the states of neighboring nodes (i.e. nodes $u, v \in V$ that are connected by an arc $(u, v) \in E$). In the second phase, usually referred to as *output* (or *readout*), the latent representations encoded by the states stored in each node are exploited to compute the model output. The GNN can implement either a *node-focused* function, where an output is produced for each node of the input graph, or a *graph-focused* function, where the representations of all the nodes are aggregated to yield a single output for the whole input graph.

The GNN is defined by a pair of (learnable) functions, that respectively implement the *state transition* function f_a required in the encoding phase and the *output* function f_r exploited in the output phase, as follows:

$$x_v^{(t)} = f_a(x_{\text{ne}[v]}^{(t-1)}, l_{\text{ne}[v]}, l_{(v, \text{ch}[v])}, l_{(\text{pa}[v], v)}, x_v^{(t-1)}, l_v | \theta_{f_a}), \quad (1)$$

$$y_v = f_r(x_v^{(T)} | \theta_{f_r}), \quad (2a)$$

$$y_G = f_r(\{x_v^{(T)}, v \in V\} | \theta_{f_r}), \quad (2b)$$

where $x_v^{(t)} \in \mathbb{R}^s$ is the state of the node v at iteration t , $\text{pa}[v] = \{u \in V : (u, v) \in E\}$ is the set of the *parents* of node v in G , $\text{ch}[v] = \{u \in V : (v, u) \in E\}$ are the *children* of v in G , $\text{ne}[v] = \text{pa}[v] \cup \text{ch}[v]$ are the *neighbors* of the node v in G , $l_u \in \mathbb{R}^m$ is the feature vector available for node $u \in V$, and $l_{(u, w)} \in \mathbb{R}^d$ is the feature vector available for the arc $(u, w) \in E^3$. The vectors θ_{f_a} and θ_{f_r} collect the model parameters (the neural network weights) to be adapted during the learning procedure. Equations (2a) and (2b) are the two variants of the output function for node-focused or graph-focused tasks, respectively.

In Table 1, we show some possible choices of the function f_a . It should be noted that this function may depend on a variable number of inputs, given that the nodes $v \in V$ may have different degrees $\text{de}[v] = |\text{ne}[v]|$. Moreover, in general, the proposed implementations are invariant with respect to permutations of the nodes in $\text{ne}[v]$, unless some predefined ordering is given for the neighbors of each node.

T is the number of iterations of the state transition function applied before computing the output. The recursive application of the state transition function f_a on the graph nodes yields a diffusion mechanism, whose range depends on T . In fact, by stacking t times the aggregation of 1-hop neighborhoods by f_a , information of one node can be transferred to the nodes that are distant at most t -hops. The number t may be seen as the *depth* of the GNN and thus each iteration can be considered a different *layer* of the GNN. A sufficient

³ With abuse of notation, we denote the set $\{x_u^{(t-1)} : u \in \text{ne}[v]\}$ by $x_{\text{ne}[v]}^{(t-1)}$. Similar definitions apply for $l_{\text{ne}[v]}$, $l_{(v, \text{ch}[v])}$, and $l_{(\text{pa}[v], v)}$.

Method: Function	Reference	Implementation of f_a
GNN: Sum	Scarselli et al. [20]	$\sum_{u \in \text{ne}[v]} h(x_u, l_u, l_{(v,u)}, l_{(u,v)}, x_v, l_v \theta_h)$
GIN: Sum	Xu et al. [24]	$h(x_v + \sum_{u \in \text{ne}[v]} x_u)$
GCN: Mean	Kipf and Welling [14]	$h\left(\frac{1}{ \text{ne}[v] +1}(x_v + \sum_{u \in \text{ne}[v]} x_u)\right)$
GraphSAGE: Max	Hamilton et al. [11]	$\max_{u \in \text{ne}[v]} h(x_u)$

Table 1. Simplified implementations of the state transition function f_a . The function $h()$ is implemented by a feedforward neural network with s outputs, whose input is the concatenation of its arguments (f.i. in the first case the input consists of a vector of $2s + 2m + 2d$ entries, with $l_{(u,v)} \in \mathbb{R}^d$ and $l_u \in \mathbb{R}^m$). For the sake of clarity, some of these formulas are reported in a simplified way w.r.t. the original proposal. For example, the "mean" function in [14] is a weighted mean, where the weights come from the normalized graph adjacency matrix, or the "max" function in [11] is followed by a concatenation.

number of layers is key to achieve a useful encoding of the input graph for the task at hand and, hence, the choice is problem-specific.

In the original GNN model [20], eq. (1) is executed until convergence of the state representation, i.e. until $x_v^{(t)} \simeq x_v^{(t-1)}$, $v \in V$. This scheme corresponds to the computation of the *fixed point* of the state transition function f_a on the input graph. In order to guarantee the convergence of this phase, the transition function is required to be a *contraction map*.

Henceforth, for compactness, we denote the state transition function, applied to a node $v \in V$, with:

$$f_{a,v} = f_a(x_{\text{ne}[v]}, l_{\text{ne}[v]}, l_{(v,\text{ch}[v])}, l_{(\text{pa}[v],v)}, x_v, l_v | \theta_{f_a}). \quad (3)$$

Basically, the encoding phase, through the iteration of f_a , finds a solution to the fixed point problem defined by the constraint

$$\forall v \in V, x_v = f_{a,v} \quad (4)$$

In this case, the states encode the information contained in the whole graph. This diffusion mechanism is more general than executing only a fixed number of iterations (i.e. stacking a fixed number of layers). However, it can be computationally heavy and, hence, many recent GNN architectures apply only a fixed number of iterations for all nodes.

4 A constraint-based formulation of Graph Neural Networks

Neural network learning can be cast as a Lagrangian optimization problem by a formulation that requires the minimization of the classical data fitting loss (and eventually a regularization term) and the satisfaction of a set of *architectural* constraints that describe the computation performed on the data. Given this formulation, the solution can be computed by finding the *saddle points* of the associated Lagrangian in the space defined by the original network parameters and the *Lagrange multipliers*. The constraints can be exploited to enforce the computational structure that characterizes the GNN models.

The computation of Graph Neural Networks is driven by the input graph topology that defines the constraints among the computed state variables $x_v, v \in V$. In particular, the fixed point computation aims to solving eq. (4), that imposes a constraint between the node states and the way they are computed by the state transition function.

In the original GNN learning algorithm, the computation of the fixed point is required at each epoch of the learning procedure, as implemented by the iterative application of the transition function. Moreover, also the gradient computation requires us to take into account the relaxation procedure, by a backpropagation schema through the

replicas of the state transition network exploited during the iterations for the fixed point computation. This procedure may be time consuming when the number of iterations T for convergence to the fixed point is high (for instance in the case of large graphs).

We consider a Lagrangian formulation of the problem by adding free variables corresponding to the node states x_v , such that the fixed point is directly defined by the constraints themselves, as

$$\forall v \in V, \mathcal{G}(x_v - f_{a,v}) = 0 \quad (5)$$

where $\mathcal{G}(x)$ is a function characterized by $\mathcal{G}(0) = 0$, such that the satisfaction of the constraints implies the solution of eq. (4). Apart from classical choices, like $\mathcal{G}(x) = x$ or $\mathcal{G}(x) = x^2$, we can design different function shape (see Section 5.1), with desired properties. For instance, a possible implementation is $\mathcal{G}(x) = \max(\|x\|_1 - \epsilon, 0)$, where $\epsilon \geq 0$ is a parameter that can be used to allow tolerance in the satisfaction of the constraint. The hard formulation of the problem requires $\epsilon = 0$, but by setting ϵ to a small positive value it is possible to obtain a better generalization and tolerance to noise.

In the following, for simplicity, we will refer to a node-focused task, such that for some (or all) nodes $v \in S \subseteq V$ of the input graph G , a target output y_v is provided as a supervision⁴. If $L(f_r(x_v | \theta_{f_r}), y_v)$ is the loss function used to measure the target fitting approximation for node $v \in S$, the formulation of the learning task is:

$$\begin{aligned} \min_{\theta_{f_a}, \theta_{f_r}, X} \quad & \sum_{v \in S} L(f_r(x_v | \theta_{f_r}), y_v) \\ \text{subject to} \quad & \mathcal{G}(x_v - f_{a,v}) = 0, \quad \forall v \in V \end{aligned} \quad (6)$$

where θ_{f_a} and θ_{f_r} are the weights of the MLPs implementing the state transition function and the output function, respectively, and $X = \{x_v : v \in V\}$ is the set of the introduced free state variables.

This problem statement implicitly includes the definition of the fixed point of the state transition function in the optimal solution, since for any solution the constraints are satisfied and hence the computed optimal x_v are solutions of eq. (4). As shown in the previous subsection, the constrained optimization problem of eq. (6) can be faced in the Lagrangian framework by introducing for each constraint a Lagrange multiplier λ_v , to define the Lagrangian function

⁴ For the sake of simplicity we consider only the case when a single graph is provided for learning. The extension for more graphs is straightforward for node-focused tasks, since they can be considered as a single graph composed by the given graphs as disconnected components.

$\mathcal{L}(\theta_{f_a}, \theta_{f_r}, X, \Lambda)$ as:

$$\mathcal{L}(\theta_{f_a}, \theta_{f_r}, X, \Lambda) = \sum_{v \in S} [L(f_r(x_v | \theta_{f_r}), y_v) + \lambda_v \mathcal{G}(x_v - f_{a,v})], \quad (7)$$

where Λ is the set of the $|V|$ Lagrangian multipliers. Finally, we can define the unconstrained optimization problem as the search for saddle points in the adjoint space $(\theta_{f_a}, \theta_{f_r}, X, \Lambda)$ as:

$$\min_{\theta_{f_a}, \theta_{f_r}, X} \max_{\Lambda} \mathcal{L}(\theta_{f_a}, \theta_{f_r}, X, \Lambda) \quad (8)$$

that can be solved by gradient descent with respect to the variables $\theta_{f_a}, \theta_{f_r}, X$ and gradient ascent with respect to the Lagrange multipliers Λ , exploiting the Basic Differential Multiplier Method, introduced in [19] in the context of neural networks. We are interested in having a strong enforcement of the *diffusion constraints*, and common penalty-based methods are hard to tune and not always guaranteed to converge to the constraint satisfaction. BDMM could be seen as a simplified procedure that implements the principles behind the common Multiplier Methods, in order to enforce the hard fulfilment of the given constraints.

The gradient can be computed locally to each node, given the local variables and those of the neighboring nodes. In fact, the derivatives of the Lagrangian⁵ with respect to the considered parameters are:

$$\frac{\partial \mathcal{L}}{\partial x_v} = L' f'_{r,v} + \lambda_v \mathcal{G}'_v (1 - f'_{a,v}) - \sum_{w:v \in ne[w]} \lambda_w \mathcal{G}'_w f'_{a,w} \quad (9)$$

$$\frac{\partial \mathcal{L}}{\partial \theta_{f_a}} = - \sum_{v \in S} \lambda_v \mathcal{G}'_v f'_{a,v} \quad (10)$$

$$\frac{\partial \mathcal{L}}{\partial \theta_{f_r}} = \sum_{v \in S} L' f'_{r,v} \quad (11)$$

$$\frac{\partial \mathcal{L}}{\partial \lambda_v} = \mathcal{G}_v \quad (12)$$

where, $f_{a,v} = f_a(x_{ne[v]}, l_{ne[v]}, l_{(v, ch[v])}, l_{(pa[v], v)}, x_v, l_v | \theta_{f_a})$, $f'_{a,v}$ is its first derivative⁶, $f_{r,v} = f_r(x_v | \theta_{f_r})$, $f'_{r,v}$ is its first derivative, $\mathcal{G}_v = \mathcal{G}(x_v - f_{a,v})$ and \mathcal{G}'_v is its first derivative, and, finally, L' is the first derivative of L . Being f_a and f_r implemented by feed-forward neural networks, their derivatives are obtained easily by applying a classical backpropagation scheme, in order to optimize the Lagrangian function in the descent-ascent scheme, aiming at the saddle point, following [19].

We initialize the variables in \mathcal{X} and Λ to zero, while the neural weights $\theta_{f_a}, \theta_{f_r}$ are randomly chosen. In particular, this *differential optimization* process consists of a gradient-descent step to update $\theta_{f_a}, \theta_{f_r}, X$, and a gradient-ascent step to update Λ , until we converge to the desired stationary point. Hence, the redefined differential equation system gradually fulfills the constraints, undergoing oscillations along the constraint subspace. To ease this procedure, we add the $\mathcal{G}()$ function, with the purpose of obtaining a more stable learning process.

Even if the proposed formulation adds the free state variables x_v and the Lagrange multipliers $\lambda_v, v \in V$, there is no significant increase in the memory requirements since the state variables are also

⁵ When parameters are vectors, the reported gradients should be considered element-wise.

⁶ The derivative is computed with respect to the same argument as in the partial derivative on the left side.

required in the original formulation and there is just a Lagrange multiplier for each node.

The diffusion mechanism of the state computation is enforced by means of the constraints. The learning algorithm is based on a mixed strategy where (i) Backpropagation is used to efficiently update the weights of the neural networks that implement the state transition and output functions, and, (ii) the diffusion mechanism evolves gradually by enforcing the convergence of the state transition function to a fixed point by virtue of the constraints. This last point is a novel approach in training Graph Neural Networks. In fact, in classical approaches, the encoding phase (see Section 3) is completely executed during the forward pass to compute the node states and, only after this phase is completed, the backward step is applied to update the weights of f_a and f_r . In the proposed scheme, both the neural network weights and the node state variables are simultaneously updated, forcing the state representation function towards a fixed point of f_a in order to satisfy the constraints. In other words, the learning proceeds by jointly updating the function weights and by diffusing information among nodes, through their state, up to a stationary condition where both the objective function is minimized and the state transition function has reached a fixed point.

In our proposed algorithm, the diffusion process is turned itself into an optimization process that must be carried out both when learning and when making predictions. As a matter of fact, inference itself requires the diffusion of information through the graph, that, in our case, corresponds with satisfying the constraints of Eq. (5). For this reason, the testing phase requires a (short) optimization routine to be carried out, that simply looks for the satisfaction of Eq. (5) for test nodes, and it is implemented using the same code that is used to optimize Eq.(8), avoiding to update the previously learned state transition and output functions.

4.1 Complexity analysis

Common graph models exploit synchronous updates among all nodes and multiple iterations for the node state embedding, with a computational complexity for each parameter update $\mathcal{O}(T(|V| + |E|))$, where T is the number of iterations, $|V|$ the number of nodes and $|E|$ the number of edges. By simultaneously carrying on the optimization of neural models and the diffusion process, our scheme relies only on 1-hop neighbors for each parameter update, hence showing a computational cost of $\mathcal{O}(|V| + |E|)$. From the memory cost viewpoint, the persistent state variable matrix requires $\mathcal{O}(|V|)$ space. However, it represents a much cheaper cost than most of GNN models, usually requiring $\mathcal{O}(T|V|)$ space. In fact, those methods need to store all the intermediate state values of all the iterations, for a latter use in back-propagation.

5 Experiments

The evaluation was carried out on two classes of tasks. Artificial tasks (Subgraph matching and Clique detection) are commonly exploited as benchmarks for GNNs, thus, allowing a direct comparison of the proposed constraint based optimization algorithm with respect to the original GNN learning scheme, on the same architecture. The second class of tasks consists of graph classification in the domains of social networks and bioinformatics. The goal is to compare the performances of the proposed approach, hereafter referred to as Lagrangian Propagation GNN (LP-GNN), that is based on a simpler model, with respect to deeper architectures such as Graph Convolutional Neural Networks.

With reference to Table 1, in our experiments we validated two formulations of the state transition function $f_{a,v}$, with two different aggregation scheme. In particular:

$$f_{a,v}^{(\text{SUM})} = \sum_{u \in ne[v]} h(x_u, l_u, l_{(v,u)}, l_{(u,v)}, x_v, l_v | \theta_h) \quad (13)$$

$$f_{a,v}^{(\text{AVG})} = \frac{1}{|ne[v]|} \sum_{u \in ne[v]} h(x_u, l_u, l_{(v,u)}, l_{(u,v)}, x_v, l_v | \theta_h) \quad (14)$$

5.1 Artificial Tasks

Subgraph Matching Given a graph G and a graph S such that $|S| \leq |G|$, the subgraph matching problem consists in finding the nodes of a subgraph $\hat{S} \subset G$ which is isomorphic to S . The task is that of learning a function τ , such that $\tau_S(G, n) = 1, n \in V$, when the node n belongs to the given subgraph S , otherwise $\tau_S(G, n) = 0$. It is designed to identify the nodes in the input graph that belong to a single subgraph given a priori during learning. The problem of finding a given subgraph is common in many practical problems and corresponds, for instance, to finding a particular small molecule inside a greater compound. An example of a subgraph structure is shown in Fig. 1. Our dataset is composed of 100 different graphs, each one having 7 nodes. The number of nodes of the target subgraph S is instead 3.




Figure 1. An example of a subgraph matching problem, where the graph with the blue nodes is matched against the bigger graph.

Clique localization A clique is a complete graph, i.e. a graph in which each node is connected with all the others. In a network, overlapping cliques (i.e. cliques that share some nodes) are admitted. Clique localization is a particular instance of the subgraph matching problem, with S being complete. However, the several symmetries contained in a clique makes the graph isomorphism test more difficult. Indeed, it is known that the graph isomorphism has polynomial time solutions only in absence of symmetries. A clique example is shown in Fig. 2. In the experiments, we consider a dataset composed by graphs having 7 nodes each, where the dimension of the maximal clique is 3 nodes.

We designed a batch of experiments on these two tasks aimed at validating our simple local optimization approach to constraint-based networks. In particular, we want to show that our optimization scheme can learn better transition and output functions than the corresponding GNN of [20]. Moreover, we want to investigate the behaviour of the algorithm for different choices of the function




Figure 2. An example of a graph containing a clique. The blue nodes represent a fully connected subgraph of dimension 4, whereas the red nodes do not belong to the clique.

$\mathcal{G}(x)$, i.e. when changing how we enforce the state convergence constraints. In particular, we tested functions with different properties: ϵ -insensitive functions, i.e. $\mathcal{G}(x) = 0, \forall x : -\epsilon \leq x \leq \epsilon$, unilateral functions, i.e. $\mathcal{G}(x) \in \mathbb{R}^+$, and bilateral functions, i.e. $\mathcal{G}(x) \in \mathbb{R}$ (a \mathcal{G} function is either unilateral or bilateral). The considered functions are shown in Table 2.

Following the experimental setting of [20], we exploited a training, validation and test set having the same size, i.e. 100 graphs each. We tuned the hyperparameters on the validation data, by selecting the node state dimension from the set $\{5, 10, 35\}$, the dropout drop-rate from the set $\{0., 0.7\}$, the state transition function from $\{f_{a,v}^{(\text{AVG})}, f_{a,v}^{(\text{SUM})}\}$ and their number of hidden units from $\{5, 20, 50\}$. We used the Adam optimizer (TensorFlow). Learning rate for parameters θ_{f_a} and θ_{f_r} is selected from the set $\{10^{-5}, 10^{-4}, 10^{-3}\}$, and the learning rate for the variables x_v and l_v from the set $\{10^{-4}, 10^{-3}, 10^{-2}\}$.

We compared our model with the equivalent GNN in [20], with the same number of hidden neurons of the f_a and f_r functions. Results are presented in Table 3.

Constraints characterized by *unilateral functions* usually offer better performances than equivalent bilateral constraints. This might be due to the fact that keeping constraints positive (as in unilateral constraints) provides a more stable learning process. Moreover, smoother constraints (i.e. *squared*) or eps-insensitive constraints tend to perform slightly better than the hard versions. This can be due to the fact that as the constraints move closer to 0 they tend to give a small or null contribution, for *squared* and *abs-eps* respectively, acting as regularizers.

5.2 Graph Classification

We used 6 graph classification benchmarks: 4 bioinformatics datasets (MUTAG, PTC, NCI1, PROTEINS) and 2 social network datasets (IMDB-BINARY, IMDB-MULTI) [25], which are becoming popular for benchmarking GNN models. In the bioinformatic graphs, the nodes have categorical input labels (e.g. atom symbol). In the social networks, there are no input node labels. In this case, we followed what has been recently proposed in [24], i.e. using one-hot encodings of node degrees. Dataset statistics are summarized in Table 4.

We compared the proposed Lagrangian Propagation GNN (LP-GNN) scheme with some of the state-of-the-art neural models for graph classification, such as Graph Convolutional Neural Networks. All the GNN-like models have a number of layers/iterations equal to 5. An important difference with these models is that, by using a different transition function at each iteration, at a cost of a much larger number of parameters, they have a much higher representational power. Even though our model could, in principle, stack mul-

	lin	$lin-\epsilon$	abs	$abs-\epsilon$	$squared$
$\mathcal{G}(x)$	x	$\max(x, \epsilon) - \max(-x, \epsilon)$	$ x $	$\max(x - \epsilon, 0)$	x^2
Unilateral	×	×	✓	✓	✓
ϵ -insensitive	×	✓	×	✓	×

Table 2. The considered variants of the \mathcal{G} function. By introducing ϵ -insensitive constraint satisfaction, we can inject into our hard-optimization scheme a controlled amount (i.e. ϵ) of unsatisfaction tolerance.

Model			Subgraph		Clique	
	\mathcal{G}	ϵ	$Acc(avg)$	$Acc(std)$	$Acc(avg)$	$Acc(std)$
LP-GNN	abs	0.00	96.25	0.96	88.80	4.82
		0.01	96.30	0.87	88.75	5.03
		0.10	95.80	0.85	85.88	4.13
	lin	0.00	95.94	0.91	84.61	2.49
		0.01	95.94	0.91	85.21	0.54
		0.10	95.80	0.85	85.14	2.17
	$squared$	-	96.17	1.01	93.07	2.18
GNN [20]	-	-	95.86	0.64	91.86	1.12

Table 3. Accuracies on the artificial datasets, for the proposed model (Lagrangian Propagation GNN - LP-GNN) and the standard GNN model for different settings.

Datasets	IMDB-B	IMDB-M	MUTAG	PROT.	PTC	NCI1
# graphs	1000	1500	188	1113	344	4110
# classes	2	3	2	2	2	2
Avg # nodes	19.8	13.0	17.9	39.1	25.5	29.8
DCNN	49.1	33.5	67.0	61.3	56.6	62.6
PATCHY-SAN	71.0 ± 2.2	45.2 ± 2.8	92.6 ± 4.2	75.9 ± 2.8	60.0 ± 4.8	78.6 ± 1.9
DGCNN	70.0	47.8	85.8	75.5	58.6	74.4
AWL	74.5 ± 5.9	51.5 ± 3.6	87.9 ± 9.8	-	-	-
GIN	75.1 ± 5.1	52.3 ± 2.8	89.4 ± 5.6	76.2 ± 2.8	64.6 ± 7.0	82.7 ± 1.7
GNN	60.9 ± 5.7	41.1 ± 3.8	88.8 ± 11.5	76.4 ± 4.4	61.2 ± 8.5	51.5 ± 2.6
LP-GNN*	71.2 ± 4.7	46.6 ± 3.7	90.5 ± 7.0	77.1 ± 4.3	64.4 ± 5.9	68.4 ± 2.1

Table 4. We report the average accuracies and standard deviations for the graph classification benchmarks, evaluated on the test set, and we compare multiple GNN models. The proposed model is denoted as LP-GNN and marked with a star. Even though it exploits only shallow representation of nodes, our model performs, on average, on-par to other top models, setting a new state-of-the-art for the Proteins dataset.

multiple diffusion processes at different levels (i.e. different latent representation of the nodes) and, then, have multiple transition functions, we have not explored this direction in this paper. In particular, the models used in the comparison are: Diffusion-Convolutional Neural Networks (DCNN) [1], PATCHY-SAN [18], Deep Graph CNN (DGCNN) [26], AWL [13], GIN-GNN [24], original GNN [20]. Apart from original GNN, we report the accuracy as reported in the referred papers.

We followed the evaluation settings in [18]. In particular, we performed 10-fold cross-validation and reported both the average and standard deviation of validation accuracies across the 10 folds within the cross-validation. The stopping epoch is selected as the epoch with the best cross-validation accuracy averaged over the 10 folds. We tuned the hyperparameters by searching: (1) the number of hidden units for both the f_a and f_r functions from the set $\{5, 20, 50, 70, 150\}$; (2) the state transition function from

$\{f_{a,v}^{(AVG)}, f_{a,v}^{(SUM)}\}$; (3) the dropout ratio from $\{0, 0.7\}$; (4) the size of the node state x_v from $\{10, 35, 50, 70, 150\}$; (5) learning rates for both the θ_{f_a} , θ_{f_r} , x_v and λ_v from $\{0.1, 0.01, 0.001\}$. Results are shown in Table 4.

As previously stated, differently from the baseline models, our approach does not rely on a deep stack of layers based on differently learnable filters. Despite of this fact, the simple GNN model trained by the proposed scheme offers performances that, on average, are preferable or on-par to the ones obtained by more complex models that exploit a larger amount of parameters.

Moreover, it is interesting to note that for current GNN models, the role of the architecture depth is twofold. First, as it is common in deep learning, depth is used to perform a multi-layer feature extraction of node inputs. Secondly, it allows node information to flow through the graph fostering the realisation of a diffusion mechanism. Conversely, our model strictly splits these two processes. We believe

this distinction to be a fundamental ingredient for a clearer understanding of which mechanism, between diffusion and node deep representation, is concurring in achieving specific performances. Indeed, in this paper, we show that the diffusion mechanism paired only with a simple shallow representation of nodes is sufficient to match performances of much deeper and complex networks.

6 Conclusions and Future Work

We showed that formulation of the GNN learning task as a constrained optimization problem allows us to avoid the explicit computation of the fixed point needed to encode the graph. The proposed framework defines how to jointly optimize the model weights and the state representation without the need of separate phases. This approach simplifies the computational scheme of GNNs and allows us to incorporate alternative strategies in the fixed point optimization by the choice of the constraint function $\mathcal{G}()$. As shown in the experimental evaluation, the appropriate functions may affect generalization and robustness to noise.

Future work will be devoted to explore systematically the properties of the proposed algorithm in terms of convergence and complexity. Moreover, we plan to extend the experimental evaluation to verify the algorithm behaviour with respect to either the characteristics of the input graphs, such as the graph diameter, the variability in the node degrees, the type of node and arc features or to the model architecture (f.i. type of the state transition function, of the constraint function, etc.). Furthermore, the proposed constraint-based scheme can be extended to all the other methods proposed in the literature that exploit more sophisticated architectures.

Finally, LP-GNN can be extended allowing the diffusion mechanism to take place at multiple layers allowing a *controlled* integration of diffusion and deep feature extraction mechanisms.

Acknowledgments

This work was partly supported by the PRIN 2017 project RexLearn, funded by the Italian Ministry of Education, University and Research (grant no. 2017TWNMH2).

REFERENCES

- [1] James Atwood and Don Towsley, ‘Diffusion-convolutional neural networks’, in *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 1993–2001, (2016).
- [2] Monica Bianchini, Giovanna Maria Dimitri, Marco Maggini, and Franco Scarselli, *Deep Neural Networks for Structured Data*, 29–51, Springer International Publishing, Cham, 2018.
- [3] Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst, ‘Geometric deep learning: Going beyond euclidean data’, *IEEE Signal Process. Mag.*, **34**(4), 18–42, (2017).
- [4] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun, ‘Spectral networks and locally connected networks on graphs’, in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, (2014).
- [5] Miguel Carreira-Perpinan and Weiran Wang, ‘Distributed optimization of deeply nested systems’, in *Artificial Intelligence and Statistics*, pp. 10–19, (2014).
- [6] Hanjun Dai, Zornitsa Kozareva, Bo Dai, Alex Smola, and Le Song, ‘Learning steady-states of iterative algorithms over graphs’, in *International Conference on Machine Learning*, pp. 1114–1122, (2018).
- [7] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst, ‘Convolutional neural networks on graphs with fast localized spectral filtering’, in *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3837–3845, (2016).
- [8] David K. Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P. Adams, ‘Convolutional networks on graphs for learning molecular fingerprints’, in *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 2224–2232, (2015).
- [9] P. Frasconi, M. Gori, and A. Sperduti, ‘A general framework for adaptive processing of data structures’, *Trans. Neur. Netw.*, **9**(5), 768–786, (September 1998).
- [10] C. Goller and A. Kuchler, ‘Learning task-dependent distributed representations by backpropagation through structure’, in *Proceedings of International Conference on Neural Networks (ICNN’96)*, volume 1, pp. 347–352 vol.1, (June 1996).
- [11] William L. Hamilton, Rex Ying, and Jure Leskovec, ‘Inductive representation learning on large graphs’, in *NIPS*, (2017).
- [12] Mikael Henaff, Joan Bruna, and Yann LeCun, ‘Deep convolutional networks on graph-structured data’, *CoRR*, **abs/1506.05163**, (2015).
- [13] Sergey Ivanov and Evgeny Burnaev, ‘Anonymous walk embeddings’, *arXiv preprint arXiv:1805.11921*, (2018).
- [14] Thomas N. Kipf and Max Welling, ‘Semi-supervised classification with graph convolutional networks’, in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, (2017).
- [15] Yann LeCun and Yoshua Bengio, ‘The handbook of brain theory and neural networks’, in *The Handbook of Brain Theory and Neural Networks*, ed., Michael A. Arbib, chapter Convolutional Networks for Images, Speech, and Time Series, 255–258, MIT Press, Cambridge, MA, USA, (1998).
- [16] Yann LeCun, D Touresky, G Hinton, and T Sejnowski, ‘A theoretical framework for back-propagation’, in *Proceedings of the 1988 connectionist models summer school*, volume 1, pp. 21–28. CMU, Pittsburgh, Pa: Morgan Kaufmann, (1988).
- [17] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard S. Zemel, ‘Gated graph sequence neural networks’, in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, (2016).
- [18] Mathias Niepert, Mohamed Ahmed, and Konstantin Kutikov, ‘Learning convolutional neural networks for graphs’, in *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pp. 2014–2023, (2016).
- [19] John C Platt and Alan H Barr, ‘Constrained differential optimization’, in *Neural Information Processing Systems*, pp. 612–621, (1988).
- [20] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini, ‘The graph neural network model’, *IEEE Trans. Neural Networks*, **20**(1), 61–80, (2009).

- [21] A. Sperduti and A. Starita, 'Supervised neural networks for the classification of structures', *Trans. Neur. Netw.*, **8**(3), 714–735, (May 1997).
- [22] Gavin Taylor, Ryan Burmeister, Zheng Xu, Bharat Singh, Ankit Patel, and Tom Goldstein, 'Training neural networks without gradients: A scalable admm approach', in *International conference on machine learning*, pp. 2722–2731, (2016).
- [23] Ronald J. Williams and David Zipser, 'A learning algorithm for continually running fully recurrent neural networks', *Neural Computation*, **1**, 270–280, (1989).
- [24] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka, 'How powerful are graph neural networks?', *CoRR*, **abs/1810.00826**, (2018).
- [25] Pinar Yanardag and SVN Vishwanathan, 'Deep graph kernels', in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1365–1374. ACM, (2015).
- [26] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen, 'An end-to-end deep learning architecture for graph classification', in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pp. 4438–4445, (2018).