
Learning Discrete Structures for Graph Neural Networks

Luca Franceschi^{1 2} Mathias Niepert³ Massimiliano Pontil^{1 2} Xiao He³

Abstract

Graph neural networks (GNNs) are a popular class of machine learning models whose major advantage is their ability to incorporate a sparse and discrete dependency structure between data points. Unfortunately, GNNs can only be used when such a graph-structure is available. In practice, however, real-world graphs are often noisy and incomplete or might not be available at all. With this work, we propose to jointly learn the graph structure and the parameters of graph convolutional networks (GCNs) by approximately solving a bilevel program that learns a discrete probability distribution on the edges of the graph. This allows one to apply GCNs not only in scenarios where the given graph is incomplete or corrupted but also in those where a graph is not available. We conduct a series of experiments that analyze the behavior of the proposed method and demonstrate that it outperforms related methods by a significant margin.

1. Introduction

Relational learning is concerned with methods that cannot only leverage the attributes of data points but also their relationships. Diagnosing a patient, for example, not only depends on the patient’s vitals and demographic information but also on the same information about their relatives, the information about the hospitals they have visited, and so on. Relational learning, therefore, does not make the assumption of independence between data points but models their dependency explicitly. Graphs are a natural way to represent relational information and there is a large number of learning algorithms leveraging graph structure. Graph neural networks (GNNs) (Scarselli et al., 2009) are one such

class of algorithms that are able to incorporate sparse and discrete dependency structures between data points.

While a graph structure is available in some domains, in others it has to be inferred or constructed. A possible approach is to first create a k -nearest neighbor (k NN) graph based on some measure of similarity between data points. This is a common strategy used by several learning methods such as LLE (Roweis & Saul, 2000) and Isomap (Tenenbaum et al., 2000). A major shortcoming of this approach, however, is that the efficacy of the resulting models hinges on the choice of k and, more importantly, on the choice of a suitable similarity measure over the input features. In any case, the graph creation and parameter learning steps are independent and require heuristics and trial and error. Alternatively, one could simply use a kernel matrix to model the similarity of examples implicitly at the cost of introducing a dense dependency structure.

With this paper, we follow a different route with the aim of learning *discrete* and *sparse* dependencies between data points while simultaneously training the parameters of graph convolutional networks (GCN), a class of GNNs. Intuitively, GCNs learn node representations by passing and aggregating messages between neighboring nodes (Kipf & Welling, 2017; Monti et al., 2017; Gilmer et al., 2017; Hamilton et al., 2017; Duran & Niepert, 2017; Velickovic et al., 2018). We propose to learn a generative probabilistic model for graphs, samples from which are used both during training and at prediction time. Edges are modelled with random variables whose parameters are treated as hyperparameters in a bilevel learning framework (Franceschi et al., 2018). We iteratively sample the structure while minimizing an inner objective (a training error) and optimize the edge distribution parameters by minimizing an outer objective (a validation error).

To the best of our knowledge, this is the first method that simultaneously learns the graph and the parameters of a GNN for semi-supervised classification. Moreover, and this might be of independent interest, we adapt gradient-based hyperparameter optimization to work for a class of discrete hyperparameters (edges, in this work). We conduct a series of experiments and show that the proposed method is competitive with and often outperforms existing approaches. We also verify that the resulting graph generative models have meaningful edge probabilities.

¹CSML, Istituto Italiano di Tecnologia, Genoa, Italy.

²University College London, London, UK. ³NEC Labs EU, Heidelberg, Germany. Work done in part as Luca Franceschi was visiting researcher at NEC. Correspondence to: Luca Franceschi <luca.franceschi@iit.it>, Xiao He <xiao.he@neclab.eu>.

2. Background

We first provide some background on graph theory, graph neural networks, and bilevel programming.

2.1. Graph Theory Basics

A graph G is a pair (V, E) with $V = \{v_1, \dots, v_N\}$ the set of vertices and $E \subseteq V \times V$ the set of edges. Let N and M be the number of vertices and edges, respectively. Each graph can be represented by an adjacency matrix A of size $N \times N$: $A_{i,j} = 1$ if there is an edge from vertex v_i to vertex v_j , and $A_{i,j} = 0$ otherwise. The graph Laplacian is defined by $L = D - A$ where $D_{i,i} = \sum_j A_{i,j}$ and $D_{i,j} = 0$ if $i \neq j$. We denote the set of all $N \times N$ adjacency matrices by \mathcal{H}_N .

2.2. Graph Neural Networks

Graph neural networks are a popular class of machine learning models for graph-structured data. We focus specifically on graph convolutional networks (GCNs) and their application to semi-supervised learning. All GNNs have the same two inputs. First, a feature matrix $X \in \mathcal{X}_N \subset \mathbb{R}^{N \times n}$ where n is the number of different node features, second, a graph $G = (V, E)$ with adjacency matrix $A \in \mathcal{H}_N$. Given a set of class labels \mathcal{Y} and a labeling function $y : V \rightarrow \mathcal{Y}$ that maps (a subset of) the nodes to their true class label, the objective is, given a set of training nodes V_{Train} , to learn a function

$$f_w : \mathcal{X}_N \times \mathcal{H}_N \rightarrow \mathcal{Y}^N$$

by minimizing some regularized empirical loss

$$L(w, A) = \sum_{v \in V_{\text{Train}}} \ell(f_w(X, A)_v, y_v) + \Omega(w), \quad (1)$$

where $w \in \mathbb{R}^d$ are the parameters of f_w , $f_w(X, A)_v$ is the output of f_w for node v , $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^+$ is a point-wise loss function, and Ω is a regularizer. An example of the function f_w proposed by Kipf & Welling (2017) is the following two hidden layer GCN that computes the class probabilities as

$$f_w(X, A) = \text{Softmax}(\hat{A} \text{ReLU}(\hat{A} X W_1) W_2), \quad (2)$$

where $w = (W_1, W_2)$ are the parameters of the GCN and \hat{A} is the normalized adjacency matrix, given by $\hat{A} = \tilde{D}^{-1/2}(A + I)\tilde{D}^{-1/2}$, with diagonal, $\tilde{D}_{ii} = 1 + \sum_j A_{ij}$.

2.3. Bilevel Programming in Machine Learning

Bilevel programs are optimization problems where a set of variables occurring in the objective function are constrained to be an optimal solution of another optimization problem (see Colson et al., 2007, for an overview). Formally given two objective functions F and L , the outer and inner objectives, and two sets of variables, $\theta \in \mathbb{R}^m$ and $w \in \mathbb{R}^d$, the outer and inner variables, a bilevel program is given by

$$\min_{\theta, w_\theta} F(w_\theta, \theta) \quad \text{such that} \quad w_\theta \in \arg \min_w L(w, \theta). \quad (3)$$

Bilevel programs arise in numerous situations such as hyperparameter optimization, adversarial, multi-task, and meta-learning (Bennett et al., 2006; Flamary et al., 2014; Muñoz-González et al., 2017; Franceschi et al., 2018).

Solving Problem (3) is challenging since the solution sets of the inner problem are usually not available in closed-form. A standard approach involves replacing the minimization of L with the repeated application of an iterative optimization dynamics Φ such as (stochastic) gradient descent (Domke, 2012; Maclaurin et al., 2015; Franceschi et al., 2017). Let $w_{\theta, T}$ denote the inner variables after T iterations of the dynamics Φ , that is, $w_{\theta, T} = \Phi(w_{\theta, T-1}, \theta) = \Phi(\Phi(w_{\theta, T-2}, \theta), \theta)$, and so on. Now, if θ and w are real-valued and the objectives and dynamics smooth, we can compute the gradient of the function $F(w_{\theta, T}, \theta)$ w.r.t. θ , denoted throughout as the *hypergradient* $\nabla_\theta F(w_{\theta, T}, \theta)$, as

$$\partial_w F(w_{\theta, T}, \theta) \nabla_\theta w_{\theta, T} + \partial_\theta F(w_{\theta, T}, \theta), \quad (4)$$

where the symbol ∂ denotes the partial derivative (the Jacobian) and ∇ either the gradient (for scalar functions) or the total derivative. The first term can be computed efficiently in time $O(T(d + m))$ with reverse-mode algorithmic differentiation (Griewank & Walther, 2008) by unrolling the optimization dynamics, repeatedly substituting $w_{\Phi, t} = \Phi(w_{\theta, t-1}, \theta)$ and applying the chain rule. This technique allows to optimize a number of hyperparameters several orders of magnitude greater than classic methods for hyperparameter optimization (Feurer & Hutter, 2018).

3. Learning Discrete Graph Structures

With this paper we address the challenging scenarios where a graph structure is either completely missing, incomplete, or noisy. To this end, we learn a *discrete* and *sparse* dependency structure between data points while simultaneously training the parameters of a GCN. We frame this as a bilevel programming problem whose outer variables are the parameters of a generative probabilistic model for graphs. The proposed approach, therefore, optimizes both the parameters of a GCN and the parameters of a graph generator so as to minimize the classification error on a given dataset. We developed a practical algorithm based on truncated reverse-mode algorithmic differentiation (Williams & Peng, 1990) and hypergradient estimation to approximately solve the bilevel problem. A schematic illustration of the resulting method is presented in Figure 1.

3.1. Jointly Learning the Structure and Parameters

Let us suppose that information about the true adjacency matrix A is missing or incomplete. Since, ultimately, we are interested in finding a model that minimizes the generalization error, we assume the existence of a second subset of instances with known target, V_{val} (the validation set), from

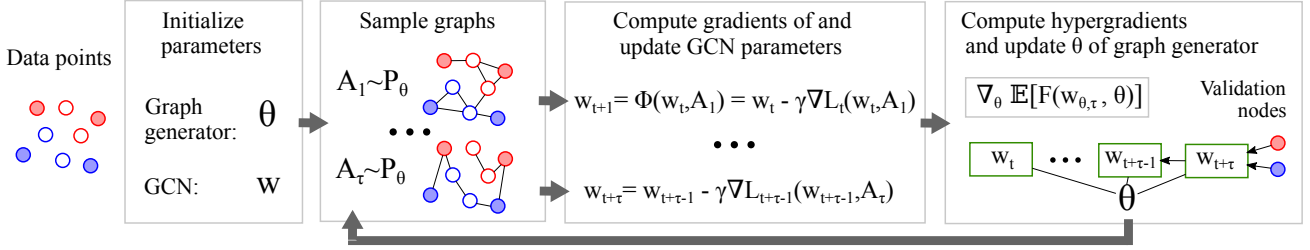


Figure 1. Schematic representation of our approach for learning discrete graph structures for GNNs.

which we can estimate the generalization error. Hence, we propose to find $A \in \mathcal{H}_N$ that minimizes the function

$$F(w_A, A) = \sum_{v \in V_{\text{val}}} \ell(f_{w_A}(X, A)_v, y_v), \quad (5)$$

where w_A is the minimizer, assumed unique, of L (see Eq. (1) and Sec. 2.3) for a *fixed adjacency matrix* A . We can then consider Equations (1) and (5) as the inner and outer objective of a mixed-integer bilevel programming problem where the outer objective aims to find an optimal discrete graph structure and the inner objective the optimal parameters of a GCN given a graph.

The resulting bilevel problem is intractable to solve exactly even for small graphs. Moreover, it contains both continuous and discrete-valued variables, which prevents us from directly applying Eq. (4). A possible solution is to construct a continuous relaxation (see e.g. Frecon et al., 2018), another is to work with parameters of a probability distribution over graphs. The latter is the approach we follow in this paper. We maintain a generative model for the graph structure and reformulate the bilevel program in terms of the (continuous) parameters of the resulting distribution over discrete graphs. Specifically, we propose to model each edge with a Bernoulli random variable. Let $\overline{\mathcal{H}}_N = \text{Conv}(\mathcal{H}_N)$ be the convex hull of the set of all adjacency matrices for N nodes. By modeling all possible edges as a set of mutually independent Bernoulli random variables with parameter matrix $\theta \in \overline{\mathcal{H}}_N$ we can sample graphs as $\mathcal{H}_N \ni A \sim \text{Ber}(\theta)$. Eqs. (1) and (5) can then be replaced, by using the expectation over graph structures. The resulting bilevel problem can be written as

$$\min_{\theta \in \overline{\mathcal{H}}_N} \mathbb{E}_{A \sim \text{Ber}(\theta)} [F(w_\theta, A)] \quad (6)$$

$$\text{such that } w_\theta = \arg \min_w \mathbb{E}_{A \sim \text{Ber}(\theta)} [L(w, A)]. \quad (7)$$

By taking the expectation, both the inner and the outer objectives become continuous (and possibly smooth) functions of the Bernoulli parameters. The bilevel problem given by Eqs. (6)-(7) is still challenging to solve efficiently. This is because the solution of the inner problem is not available in closed form for GCNs (the objective is non-convex); and the expectations are intractable to compute exactly¹.

¹This is different than e.g. (model free) reinforcement learning,

An efficient algorithm, therefore, will only be able to find approximate stochastic solutions, that is, $\theta_{i,j} \in (0, 1)$.

Before describing a method to solve the optimization problem given by Eqs. (6)-(7) approximately with hypergradient descent, we first turn to the question of obtaining a final GCN model that we can use for prediction. For a given distribution P_θ over graphs with N nodes and with parameters θ , the expected output of a GCN is

$$f_w^{\text{exp}}(X) = \mathbb{E}_A [f_w(X, A)] = \sum_{A \in \mathcal{H}_N} P_\theta(A) f_w(X, A). \quad (8)$$

Unfortunately, computing this expectation is intractable even for small graphs; we can, however, compute an empirical estimate of the output as

$$\hat{f}_w(X) = \frac{1}{S} \sum_{i=1}^S f_w(X, A_i), \quad (9)$$

where $S > 0$ is the number of samples we wish to draw. Note that \hat{f} is an unbiased estimator of f_w^{exp} . Hence, to use a GCN f_w learned with the bilevel formulation for prediction, we sample S graphs from the distribution P_θ and compute the prediction as the empirical mean of the values of f_w .

Given the parametrization of the graph generator with Bernoulli variables ($P_\theta = \text{Ber}(\theta)$), one can sample a new graph in $O(N^2)$. Sampling from a large number of Bernoulli variables, however, is highly efficient, trivially parallelizable, and possible at a rate of millions per second. Other sampling strategies such as MCMC sampling are possible in constant time. Given a set of sampled graphs, it is more efficient to evaluate a sparse GCN S times than to use the Bernoulli parameters as weights of the GCN's adjacency matrix². Indeed, for GCN models, computing \hat{f}_w has a cost of $O(SCd)$, rather than $O(N^2d)$ for a fully connected graph, where $C = \sum_{i,j} \theta_{i,j}$ is the expected number of edges, and d is the dimension of the weights. Another advantage of using a graph-generative model is that we can interpret it probabilistically which is not the case when learning a dense adjacency matrix.

where the objective function is usually unknown, depending in an unknown way from the action and the environment.

²Note also that $\mathbb{E} f_w(X, A) \neq f_w(X, \mathbb{E}A) = f_w(X, \theta)$, as the model f_w is, in general, nonlinear.

3.2. Structure Learning via Hypergradient Descent

The bilevel programming formalism is a natural fit for the problem of learning both a graph generative model and the parameters of a GNN for a specific downstream task. Here, the outer variables θ are the parameters of the graph generative model and the inner variables w are the parameters of the GCN.

We now discuss a practical algorithm to approach the bilevel problem defined by Eqs. (6) and (7). Regarding the inner problem, we note that the expectation

$$\mathbb{E}_{A \sim \text{Ber}(\theta)} [L(w, A)] = \sum_{A \in \mathcal{H}_N} P_\theta(A) L(w, A) \quad (10)$$

is composed of a sum of 2^{N^2} terms, which is intractable even for relatively small graphs. We can, however, choose a tractable approximate learning dynamics Φ such as stochastic gradient descent (SGD),

$$w_{\theta, t+1} = \Phi(w_{\theta, t}, A_t) = w_{\theta, t} - \gamma_t \nabla L(w_{\theta, t}, A_t), \quad (11)$$

where γ_t is a learning rate and $A_t \sim \text{Ber}(\theta)$ is drawn at each iteration. Under appropriate assumptions and for $t \rightarrow \infty$, SGD converges to a weight vector w_θ that depends on the edges' probability distribution (Bottou, 2010).

Let $w_{\theta, T}$ be an approximate minimizer of $\mathbb{E}[L]$ (where T may depend on θ). We now need to compute an estimator for the hypergradient $\nabla_\theta \mathbb{E}_{A \sim \text{Ber}(\theta)} [F]$. Recalling Eq. (4), we have

$$\begin{aligned} \nabla_\theta \mathbb{E} [F(w_{\theta, T}, A)] &= \mathbb{E} [\nabla_\theta F(w_{\theta, T}, A)] = \\ &\mathbb{E} [\partial_w F(w_{\theta, T}, A) \nabla_\theta w_{\theta, T} + \partial_A F(w_{\theta, T}, A) \nabla_\theta A], \end{aligned} \quad (12)$$

where we can swap the gradient and expectation operators since the expectation is over a finite random variable, assuming that the loss function F bounded. We use the so-called straight-through estimator (Bengio et al., 2013) and set $\nabla_\theta A := I$ (which would be 0 a.e. otherwise); $\nabla_\theta A$ appears both explicitly in (12) and in the computation of $\nabla_\theta w_{\theta, T}$, through $\partial_\theta \Phi(w_{\theta, t}, A_t)$ (see Sec. 2.3 and Franceschi et al., 2017, for details). Finally, we take the single sample Monte Carlo estimator of (12) to update the parameters θ with projected gradient descent on the unit hypercube.

Computing the hypergradient by fully unrolling the dynamics may be too expensive both in time and memory³. We propose to truncate the computation and estimate the hypergradient every τ iterations, where τ is a parameter of the algorithm. This is essentially an adaptation of truncated back-propagation through time (Werbos, 1990; Williams & Peng, 1990) and can be seen as a short-horizon optimization

³Moreover, since we rely on biased estimations of the gradients, we do not expect to gain too much from a full computation.

Algorithm 1 LDS

```

1: Input data:  $X, Y, Y', A$ 
2: Input parameters:  $\eta, \tau, k$ 
3:  $[A \leftarrow \text{kNN}(X, k)]$  {Init.  $A$  to  $k$ NN graph if  $A = 0$ }
4:  $\theta \leftarrow A$  {Initialize  $P_\theta$  as a deterministic distribution}
5: while Stopping condition is not met do
6:    $t \leftarrow 0$ 
7:   while Inner objective decreases do
8:      $A_t \sim \text{Ber}(\theta)$  {Sample structure}
9:      $w_{\theta, t+1} \leftarrow \Phi_t(w_{\theta, t}, A_t)$  {Optimize inner objective}
10:     $t \leftarrow t + 1$ 
11:    if  $t = 0 \pmod{\tau}$  or  $\tau = 0$  then
12:       $G \leftarrow \text{computeHG}(F, Y, \theta, (w_{\theta, i})_{i=t-\tau}^t)$ 
13:       $\theta \leftarrow \text{Proj}_{\overline{\mathcal{H}}_N}[\theta - \eta G]$  {Optimize outer objective}
14:    end if
15:  end while
16: end while
17: return  $w, P_\theta$  {Best found weights and prob. distribution}
    
```

procedure with warm restart on w . A sketch of the method is presented in Algorithm 1, while a more complete version that includes details on the hypergradient computation can be found in the appendix. Inputs and operations in squared brackets are optional.

The algorithm contains stopping conditions at the outer and at the inner level. While it is natural to implement the latter with a decrease condition on the inner objective⁴, we find it useful to implement the first with a simple early stopping criterion. A fraction of the examples in the validation set is held-out to compute, in each outer iteration, the accuracy using the predictions of the empirically expected model (9). The optimization procedure terminates if there is no improvement for some consecutive outer loops. This helps avoiding overfitting the outer objective (6), which may be a concern in this context given the quantity of (hyper)parameters being optimized and the relative small size of the validation sets.

The hypergradients estimated with Algorithm 1 at each outer iteration are biased. The bias stems from both the straight-through estimator and from the truncation procedure introduced in lines 11-13 (Tallec & Ollivier, 2017). Nevertheless, we find empirically that the algorithm is able to make reasonable progress, finding configurations in the distribution space that are beneficial for the tasks at hand.

4. Experiments

We conducted a series of experiments with three main objectives. First, we evaluated LDS on node classification problems where a graph structure is available but where a certain fraction of edges is missing. Here, we compared

⁴We continue optimizing L until $L(w_{t-1}, A)(1 + \varepsilon) \geq L(w_{\theta, t}, A)$, for $\varepsilon > 0$ ($\varepsilon = 10^{-3}$ in the experiments). Since L is non-convex, we also use a patience window of p steps.

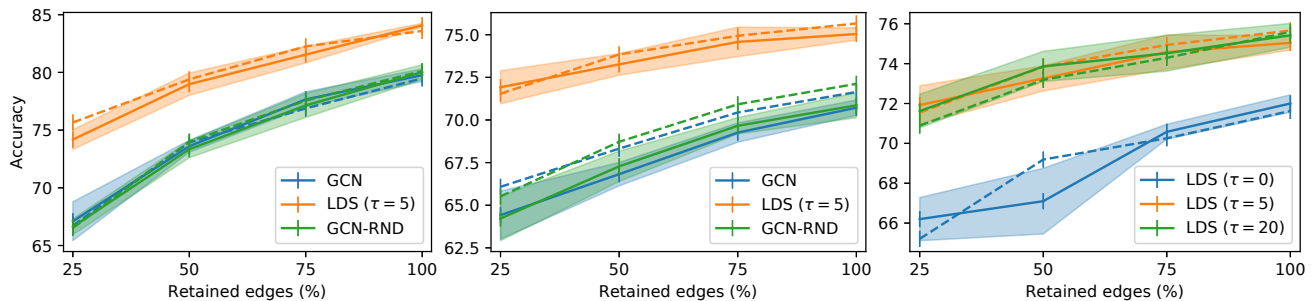


Figure 2. Mean accuracy \pm standard deviation on validation (early stopping; dashed lines) and test (solid lines) sets for edge deletion scenarios on Cora (left) and Citeseer (center). (Right) Validation of the number of steps τ used to compute the hypergradient (Citeseer); $\tau = 0$ corresponds to alternating minimization. All results are obtained from five runs with different random seeds.

LDS with graph-based learning algorithms including vanilla GCNs. Second, we wanted to validate our hypothesis that LDS can achieve competitive results on semi-supervised classification problems for which a graph is *not* available. To this end, we compared LDS with a number of existing semi-supervised classification approaches. We also compared LDS with algorithms that first create k -NN affinity graphs on the data set. Third, we analyzed the learned graph generative model to understand to what extent LDS is able to learn meaningful edge probability distributions even when a large fraction of edges is missing.

4.1. Datasets

Cora and Citeseer are two benchmark datasets that are commonly used to evaluate relational learners in general and GCNs in particular (Sen et al., 2008). The input features are bag of words and the task is node classification. We use the same dataset split and experimental setup of previous work (Yang et al., 2016; Kipf & Welling, 2017). To evaluate the robustness of LDS on incomplete graphs, we construct graphs with missing edges by randomly sampling 25%, 50%, and 75% of the edges. In addition to Cora and Citeseer where we removed all edges, we evaluate LDS on benchmark datasets that are available in scikit-learn (Pedregosa et al., 2011) such as Wine, Breast Cancer (Cancer), Digits, and 20 Newsgroup (20news). We take 10 classes from 20 Newsgroup and use words (TFIDF) with a frequency of more than 5% as features. We also use FMA, a dataset where 140 audio features are extracted from 7,994 music tracks and where the problem is genre classification (Defferrard et al., 2017). The statistics of the datasets are reported in the appendix.

4.2. Setup and Baselines

For the experiments on graphs with missing edges, we compare LDS to vanilla GCNs. In addition, we also conceived a method (GCN-RND) where we add randomly sampled edges at each optimization step of a vanilla GCN. With this method we intend to show that simply adding random edges to the standard training procedure of a GCN model (perhaps acting as a regularization technique) is not enough to

improve the generalization.

When a graph is completely missing, GCNs boil down to feed-forward neural networks. Therefore, we evaluate different strategies to induce a graph on both labeled and unlabeled samples by creating (1) a sparse Erdős-Rényi random graph (Erdos & Rényi, 1960) (Sparse-GCN); (2) a dense graph with equal edge probabilities (Dense-GCN); (3) a dense RBF kernel on the input features (RBF-GCN); and (4) a sparse k -nearest neighbor graph on the input features (k NN-GCN). For LDS we initialize the edge probabilities using the k -NN graph (k NN-LDS). We further include a dense version of LDS where we learn a dense similarity matrix (k NN-LDS (dense)). In this setting, we compare LDS to popular semi-supervised learning methods such as label propagation (LP) (Zhu et al., 2003), manifold regularization (ManiReg) (Belkin et al., 2006), and semi-supervised embedding (SemiEmb) (Weston et al., 2012). ManiReg and SemiEmb are given a k -NN graph as input for the Laplacian regularization. We also compare LDS to baselines that do not leverage a graph-structure such as logistic regression (LogReg), support vector machines (Linear and RBF SVM), random forests (RF), and feed-forward neural networks (FFNN). For comparison methods that need a k NN graph, $k \in \{2, 3, \dots, 20\}$ and the metric (Euclidean or Cosine) are tuned using validation accuracy. For k NN-LDS, k is tuned from 10 or 20.

We use the two layers GCN given by Eq. (2) with 16 hidden neurons and ReLU activation. Given a set of labelled training instances V_{Train} (nodes or examples) we use the regularized cross-entropy loss

$$L(w, A) = - \sum_{v \in V_{\text{Train}}} y_v \circ \log [f_w(X, A)_v] + \rho \|w_1\|^2, \quad (13)$$

where y_v is the one-hot encoded target vector for the v -th instance, \circ denotes the element-wise multiplication and ρ is a non-negative coefficient. As additional regularization technique we apply dropout (Srivastava et al., 2014) with $\beta = 0.5$ as in previous work. We use Adam (Kingma & Ba, 2015) for optimizing L , tuning the learning rate γ from $\{0.005, 0.01, 0.02\}$. The same number of hidden neurons

Table 1. Test accuracy (\pm standard deviation) in percentage on various classification datasets. The best results and the statistical competitive ones (by paired t-test with $\alpha = 0.05$) are in bold. All experiments have been repeated with 5 different random seeds. We compare k NN-LDS to several supervised baselines and semi-supervised learning methods. No graph is provided as input. k NN-LDS achieves high accuracy results on most of the datasets and yields the highest gains on datasets with underlying graphs (Citeseer, Cora).

	Wine	Cancer	Digits	Citeseer	Cora	20news	FMA
LogReg	92.1 (1.3)	93.3 (0.5)	85.5 (1.5)	62.2 (0.0)	60.8 (0.0)	42.7 (1.7)	37.3 (0.7)
Linear SVM	93.9 (1.6)	90.6 (4.5)	87.1 (1.8)	58.3 (0.0)	58.9 (0.0)	40.3 (1.4)	35.7 (1.5)
RBF SVM	94.1 (2.9)	91.7 (3.1)	86.9 (3.2)	60.2 (0.0)	59.7 (0.0)	41.0 (1.1)	38.3 (1.0)
RF	93.7 (1.6)	92.1 (1.7)	83.1 (2.6)	60.7 (0.7)	58.7 (0.4)	40.0 (1.1)	37.9 (0.6)
FFNN	89.7 (1.9)	92.9 (1.2)	36.3 (10.3)	56.7 (1.7)	56.1 (1.6)	38.6 (1.4)	33.2 (1.3)
LP	89.8 (3.7)	76.6 (0.5)	91.9 (3.1)	23.2 (6.7)	37.8 (0.2)	35.3 (0.9)	14.1 (2.1)
ManiReg	90.5 (0.1)	81.8 (0.1)	83.9 (0.1)	67.7 (1.6)	62.3 (0.9)	46.6 (1.5)	34.2 (1.1)
SemiEmb	91.9 (0.1)	89.7 (0.1)	90.9 (0.1)	68.1 (0.1)	63.1 (0.1)	46.9 (0.1)	34.1 (1.9)
Sparse-GCN	63.5 (6.6)	72.5 (2.9)	13.4 (1.5)	33.1 (0.9)	30.6 (2.1)	24.7 (1.2)	23.4 (1.4)
Dense-GCN	90.6 (2.8)	90.5 (2.7)	35.6 (21.8)	58.4 (1.1)	59.1 (0.6)	40.1 (1.5)	34.5 (0.9)
RBF-GCN	90.6 (2.3)	92.6 (2.2)	70.8 (5.5)	58.1 (1.2)	57.1 (1.9)	39.3 (1.4)	33.7 (1.4)
k NN-GCN	93.2 (3.1)	93.8 (1.4)	91.3 (0.5)	68.3 (1.3)	66.5 (0.4)	41.3 (0.6)	37.8 (0.9)
k NN-LDS (dense)	97.5 (1.2)	94.9 (0.5)	92.1 (0.7)	70.9 (1.3)	70.9 (1.1)	45.6 (2.2)	38.6 (0.6)
k NN-LDS	97.3 (0.4)	94.4 (1.9)	92.5 (0.7)	71.5 (1.1)	71.5 (0.8)	46.4 (1.6)	39.7 (1.4)

and the same activation is used for SemiEmb and FFNN.

For LDS, we set the initial edge parameters $\theta_{i,j}$ to 0 except for the known edges (or those found by k NN) which we set to 1. We then let all the parameters (including those initially set to 1) to be optimized by the algorithm. We further split the validation set evenly to form the validation (A) and early stopping (B) sets. As outer objective we use the unregularized cross-entropy loss on (A) and optimize it with stochastic gradient descent, with exponentially decreasing learning rate. Initial experiments showed that accelerated optimization methods such as Adam or SGD with momentum underperform in this setting. We tune the step size η of the outer optimization loop and the number of updates τ used to compute the truncated hypergradient. Finally, we draw $S = 16$ samples to compute the output predictions (see Eq. (9)). For LDS and GCN, we apply early stopping with a window size of 20 steps.

LDS was implemented in TensorFlow (Abadi et al., 2015) and is available at <https://github.com/lucfra/LDS>. The implementations of the supervised baselines and LP are those from the scikit-learn python package (Pedregosa et al., 2011). GCN, ManiReg, and SemiEmb are implemented in Tensorflow. The hyperparameters for all the methods are selected through the validation accuracy.

4.3. Results

The results on the incomplete graphs are shown in Figure 2 for Cora (left) and Citeseer (center). For each percentage of retained edges the accuracy on validation (used for early stopping) and test sets are plotted. LDS achieves competitive results in all scenarios and accuracy gains of up to 7 percentage points. Notably, LDS improves the generaliza-

Table 2. Initial number of edges and expected number of sampled edges of learned graph by LDS.

% Edges	25%	50%	75%	100%
Cora Initial	1357	2714	4071	5429
Cora Learned	3635.6	4513.9	5476.9	6276.4
Citeseer Initial	1183	2366	3549	4732
Citeseer Learned	3457.4	4474.2	7842.5	6745.2

tion accuracy of GCN models also when the given graph is that of the respective dataset (100% of edges retained), by learning additional helpful edges. The accuracy of 84.1% and 75.0% for Cora and Citeseer, respectively, exceed all previous state-of-the-art results. Conversely, adding random edges does not help decreasing the generalization error. GCN and GCN-RND perform similarly which indicates that adding random edges to the graph is not helpful.

Figure 2 (right) depicts the impact of the number of iterations τ to compute the hypergradients. Taking multiple steps strongly outperforms alternating optimization⁵ (i.e. $\tau = 0$) in all settings. Increasing τ further to the value of 20, however, does not yield significant benefits, while increasing the computational cost.

In Table 2 we computed the expected number of edges in a sampled graph for Cora and Citeseer, to analyze the properties of the graphs sampled from the learned graph generator. The expected number of edges for LDS is higher than the original number which is to be expected since LDS has better accuracy results than the vanilla GCN in Figure

⁵For $\tau = 0$, one step of optimization of L w.r.t. w , fixing θ is interleaved with one step of minimization of F w.r.t. θ , fixing w . Even if computationally lighter, this approach disregards the nested structure of (6)-(7), not computing the first term of Eq. (4).

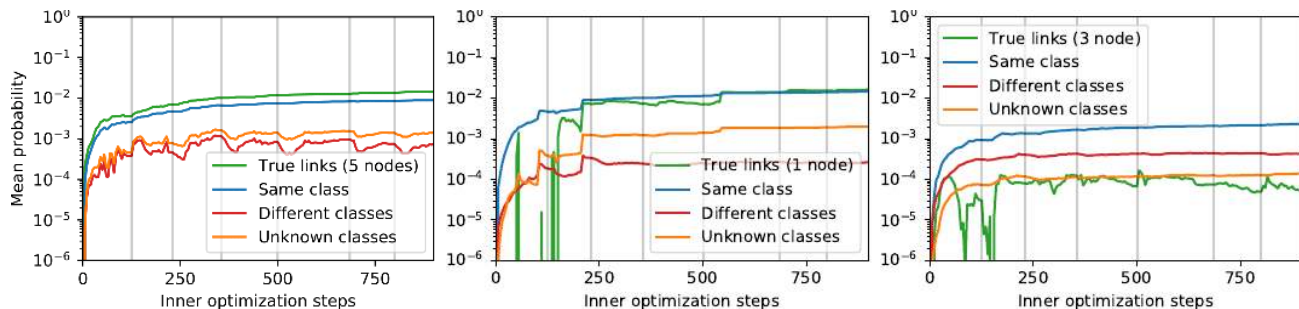


Figure 3. Mean edge probabilities to nodes aggregated w.r.t. four groups during LDS optimization, in \log_{10} scale for three example nodes. For each example node, all other nodes are grouped by the following criteria: (a) adjacent in the ground truth graph; (b) same class membership; (c) different class membership; and (d) unknown class membership. Probabilities are computed with LDS ($\tau = 5$) on Cora with 25% retained edges. From left to right, the example nodes belong to the training, validation, and test set, respectively. The vertical gray lines indicate when the inner optimization dynamics restarts, that is, when the weights of the GCN are reinitialized.

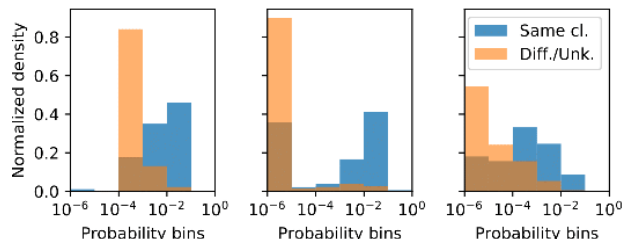


Figure 4. Normalized histograms of edges' probabilities for the same nodes of Figure 3.

2. Nevertheless, the learned graphs are still very sparse (e.g. for Cora, on average, less than 0.2% edges are present). This facilitates efficient learning of the GCN in the inner learning loop of LDS.

Table 1 lists the results for semi-supervised classification problems. The supervised learning baselines work well on some datasets such as Wine and Cancer but fail to provide competitive results on others such as Digits, Citeseer, Cora, and 20News. The semi-supervised learning baselines LP, ManiReg and SemiEmb can only improve the supervised learning baselines on 1, 3 and 4 datasets, respectively. The results for the GCN with different input graphs show that k NN-GCN works well and provides competitive results compared to the supervised baselines on all datasets. k NN-LDS significantly outperforms k NN-GCN on 4 out of the 7 datasets. In addition, k NN-LDS is among the most competitive methods on all datasets and yields the highest gains on datasets that have an underlying graph. Moreover, k NN-LDS performs slightly better than its dense counterpart where we learn a dense adjacency matrix. The added benefit of the sparse graph representation lies in the potential to scale to larger datasets.

In Figure 3, we show the evolution of mean edge probabilities during optimization on three types of nodes (train, validation, test) on the Cora dataset. LDS is able to learn a graph generative model that is, on average, attributing 10 to 100 times more probability to edges between samples sharing the same class label. LDS often attributes a higher

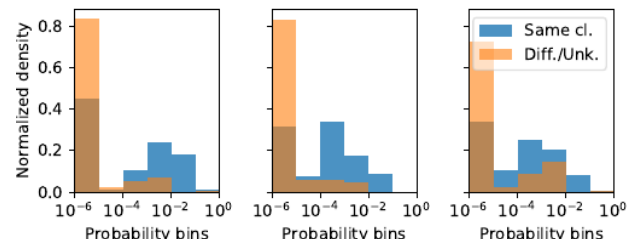


Figure 5. Histograms for three Citeseer test nodes, misclassified by k NN-GCN and rightly classified by k NN-LDS.

probability to edges that are present in the true held-out adjacency matrix (green lines in the plots). In Figure 4 we report the normalized histograms of the optimized edges probabilities for the same nodes of Figure 3, sorted into six bins in \log_{10} -scale. Edges are divided in two groups: edges between nodes of the same class (blue) and between nodes of unknown or different classes (orange). LDS is able to learn highly non-uniform edge probabilities that reflect the class membership of the nodes. Figure 5 shows similar qualitative results as Figure 4, this time for three Citeseer test nodes, misclassified by k NN-GCN and correctly classified by k NN-LDS. Again, the learned edge probabilities linking to nodes of the same classes is significantly different to those from different classes; but in this case the densities are more skewed toward the first bin. On the datasets we considered, what seems to matter is to capture a useful distribution (i.e. higher probability for links between same class) rather than pick exact links; of course for other datasets this may vary.

5. Related work

Semi-supervised Learning. Early works on graph-based semi-supervised learning use graph Laplacian regularization and include label propagation (LP) (Zhu et al., 2003), manifold regularization (ManiReg) (Belkin et al., 2006), and semi-supervised embedding (SemiEmb) (Weston et al., 2012). These methods assume a given graph whose edges represent some similarity between nodes. Later, (Yang et al., 2016) proposed a method that uses graphs not for regular-

ization but rather for embedding learning by jointly classification and graph context prediction. Kipf & Welling (2017) presented the first GCN for semi-supervised learning. There are now numerous GCN variants all of which assume a given graph structure. Contrary to all existing graph-based semi-supervised learning approaches, LDS is able to work even when the graph is incomplete or missing.

Graph synthesis and generation. LDS learns a probabilistic generative model for graphs. The earliest probabilistic generative model for graphs was the Erdős-Rényi random graph model (Erdős & Rényi, 1960), where edge probabilities are modelled as identically distributed and mutually independent Bernoullis. Several network models have been proposed to model well particular graph properties such as degree distribution (Leskovec et al., 2005) or network diameter (Watts & Strogatz, 1998). Leskovec et al. (2010) proposed a generative model based on the Kronecker product that takes a real graph as input and generates graphs that have similar properties. Recently, deep learning based approaches have been proposed for graph generation (You et al., 2018; Li et al., 2018; Grover et al., 2018; De Cao & Kipf, 2018). The goal of these methods, however, is to learn a sophisticated generative model that reflects the properties of the training graphs. LDS, on the other hand, learns graph generative models as a means to perform well on classification problems and its input is not a collection of graphs. More recent work proposed an unsupervised model that learns to infer interactions between entities while simultaneously learning the dynamics of physical systems such as spring systems (Kipf et al., 2018). Contrary to LDS, the method is specific to dynamical interacting systems, is unsupervised, and uses a variational encoder-decoder. Finally, we note that Johnson (2017) proposed a fully differentiable neural model able to process and produce graph structures at both input, representation and output levels; training the model requires, however, supervision in terms of ground truth graphs.

Link prediction. Link prediction is a decades-old problem (Liben-Nowell & Kleinberg, 2007). Several survey papers cover the large body of work ranging from link prediction in social networks to knowledge base completion (Lü & Zhou, 2011; Nickel et al., 2016). While a majority of the methods are based on some similarity measure between node pairs, there has been a number of neural network based methods (Zhang & Chen, 2017; 2018). The problem we study in this paper is related to link prediction as we also want to learn or extend a graph. However, existing link prediction methods do not simultaneously learn a GNN node classifier. Statistical relational learning (SRL) (Getoor & Taskar, 2007) models often perform both link prediction and node classification through the existence of binary and unary predicates. However, SRL models are inherently intractable and the structure and parameter learning steps are

independent.

Gradient estimation for discrete random variables. Due to the intractable nature of the two bilevel objectives, LDS needs to estimate the hypergradients through a stochastic computational graph (Schulman et al., 2015). Using the score function estimator, also known as REINFORCE (Williams, 1992), would treat the outer objective as a black-box function and would not exploit F being differentiable w.r.t. the sampled adjacency matrices and inner optimization dynamics. Conversely, the path-wise estimator is not readily applicable, since the random variables are discrete. LDS borrows from a solution proposed before (Bengio et al., 2013), at the cost of having biased estimates. Recently, Jang et al. (2017); Maddison et al. (2017) presented an approach based on continuous relaxations to reduce variance, which Tucker et al. (2017) combined with REINFORCE to obtain an unbiased estimator. Grathwohl et al. (2018) further introduced surrogate models to construct control variates for black-box functions. Unfortunately, these latter methods require to compute the function in the interior of the hypercube, possibly in multiple points (Tucker et al., 2017). This would introduce additional computational overhead⁶.

6. Conclusion

We propose LDS, a framework that simultaneously learns the graph structure and the parameters of a GNN. While we have used a specific GCN variant (Kipf & Welling, 2017) in the experiments, the method is more generally applicable to other GNNs. The strengths of LDS are its high accuracy gains on typical semi-supervised classification datasets at a reasonable computational cost. Moreover, due to the graph generative model LDS learns, the edge parameters have a probabilistic interpretation.

The method has its limitations. While relatively efficient, it cannot currently scale to large datasets: this would require an implementation that works with mini-batches of nodes. We evaluated LDS only in the transductive setting, when all data points (nodes) are available during training. Adding additional nodes after training (the inductive setting) would currently require retraining the entire model from scratch. When sampling graphs, we do not currently enforce the graphs to be connected. This is something we anticipate to improve the results, but this would require a more sophisticated sampling strategy. All of these shortcomings motivate future work. In addition, we hope that suitable variants of LDS algorithm will also be applied to other problems such as neural architecture search or to tune other discrete hyperparameters.

⁶Recall that F can be computed only after (approximately) solving the inner optimization problem.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <http://tensorflow.org/>. Software available from tensorflow.org.
- Belkin, M., Niyogi, P., and Sindhvani, V. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7:2399–2434, 2006.
- Bengio, Y., Léonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Bennett, K. P., Hu, J., Ji, X., Kunapuli, G., and Pang, J.-S. Model selection via bilevel optimization. In *Neural Networks, 2006. IJCNN'06. International Joint Conference on*, pp. 1922–1929. IEEE, 2006.
- Bottou, L. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pp. 177–186. Springer, 2010.
- Colson, B., Marcotte, P., and Savard, G. An overview of bilevel optimization. *Annals of operations research*, 153(1):235–256, 2007.
- De Cao, N. and Kipf, T. Molgan: An implicit generative model for small molecular graphs. *arXiv preprint arXiv:1805.11973*, 2018.
- Defferrard, M., Benzi, K., Vandergheynst, P., and Bresson, X. Fma: A dataset for music analysis. In *18th International Society for Music Information Retrieval Conference*, 2017. URL <https://arxiv.org/abs/1612.01840>.
- Domke, J. Generic methods for optimization-based modeling. In *Artificial Intelligence and Statistics*, pp. 318–326, 2012.
- Duran, A. G. and Niepert, M. Learning graph representations with embedding propagation. In *Advances in Neural Information Processing Systems*, pp. 5119–5130, 2017.
- Erdos, P. and Rényi, A. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.
- Feurer, M. and Hutter, F. Hyperparameter optimization. In Hutter, F., Kotthoff, L., and Vanschoren, J. (eds.), *Automatic Machine Learning: Methods, Systems, Challenges*, pp. 3–38. Springer, 2018. In press, available at <http://automl.org/book>.
- Flamary, R., Rakotomamonjy, A., and Gasso, G. Learning constrained task similarities in graph-regularized multi-task learning. *Regularization, Optimization, Kernels, and Support Vector Machines*, pp. 103, 2014.
- Franceschi, L., Donini, M., Frasconi, P., and Pontil, M. Forward and reverse gradient-based hyperparameter optimization. *ICML*, 2017.
- Franceschi, L., Frasconi, P., Salzo, S., Grazi, R., and Pontil, M. Bilevel programming for hyperparameter optimization and meta-learning. *ICML*, 2018.
- Frecon, J., Salzo, S., and Pontil, M. Bilevel learning of the group lasso structure. In *Advances in Neural Information Processing Systems 31*, pp. 8311–8321, 2018.
- Getoor, L. and Taskar, B. *Introduction to statistical relational learning*, volume 1. MIT press Cambridge, 2007.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017.
- Grathwohl, W., Choi, D., Wu, Y., Roeder, G., and Duvenaud, D. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. *ICLR*, 2018.
- Griewank, A. and Walther, A. *Evaluating derivatives: principles and techniques of algorithmic differentiation*, volume 105. Siam, 2008.
- Grover, A., Zweig, A., and Ermon, S. Graphite: Iterative generative modeling of graphs. *arXiv preprint arXiv:1803.10459*, 2018.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pp. 1024–1034, 2017.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. *ICLR*, 2017.
- Johnson, D. D. Learning graphical state transitions. *ICLR*, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *ICLR*, 2015.
- Kipf, T., Fetaya, E., Wang, K.-C., Welling, M., and Zemel, R. Neural relational inference for interacting systems. *arXiv preprint arXiv:1802.04687*, 2018.

- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *ICLR*, 2017.
- Leskovec, J., Kleinberg, J., and Faloutsos, C. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pp. 177–187. ACM, 2005.
- Leskovec, J., Chakrabarti, D., Kleinberg, J., Faloutsos, C., and Ghahramani, Z. Kronecker graphs: An approach to modeling networks. *Journal of Machine Learning Research*, 11(Feb):985–1042, 2010.
- Li, Y., Vinyals, O., Dyer, C., Pascanu, R., and Battaglia, P. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324*, 2018.
- Liben-Nowell, D. and Kleinberg, J. The link-prediction problem for social networks. *Journal of the American society for information science and technology*, 58(7):1019–1031, 2007.
- Lü, L. and Zhou, T. Link prediction in complex networks: A survey. *Physica A: statistical mechanics and its applications*, 390(6):1150–1170, 2011.
- Maclaurin, D., Duvenaud, D., and Adams, R. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning*, pp. 2113–2122, 2015.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *ICLR*, 2017.
- 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 *Proc. CVPR*, pp. 3, 2017.
- Muñoz-González, L., Biggio, B., Demontis, A., Paudice, A., Wongrassamee, V., Lupu, E. C., and Roli, F. Towards poisoning of deep learning algorithms with back-gradient optimization. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pp. 27–38. ACM, 2017.
- Nickel, M., Murphy, K., Tresp, V., and Gabrilovich, E. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1):11–33, 2016.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Roweis, S. T. and Saul, L. K. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- Schulman, J., Heess, N., Weber, T., and Abbeel, P. Gradient estimation using stochastic computation graphs. In *Advances in Neural Information Processing Systems*, pp. 3528–3536, 2015.
- Sen, P., Namata, G., Bilgic, M., Getoor, L., Galligher, B., and Eliassi-Rad, T. Collective classification in network data. *AI magazine*, 29(3):93, 2008.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Talleg, C. and Ollivier, Y. Unbiasing truncated backpropagation through time. *arXiv preprint arXiv:1705.08209*, 2017.
- Tenenbaum, J. B., De Silva, V., and Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- Tucker, G., Mnih, A., Maddison, C. J., Lawson, J., and Sohl-Dickstein, J. Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, pp. 2627–2636, 2017.
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. *ICLR*, 2018.
- Watts, D. J. and Strogatz, S. H. Collective dynamics of small-world networks. *nature*, 393(6684):440, 1998.
- Werbos, P. J. Backpropagation through time: What it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.
- Weston, J., Ratle, F., Mobahi, H., and Collobert, R. Deep learning via semi-supervised embedding. In *Neural Networks: Tricks of the Trade*, pp. 639–655. Springer, 2012.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Williams, R. J. and Peng, J. An efficient gradient-based algorithm for on-line training of recurrent network trajectories. *Neural computation*, 2(4):490–501, 1990.

- Yang, Z., Cohen, W. W., and Salakhutdinov, R. Revisiting semi-supervised learning with graph embeddings. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pp. 40–48, 2016.
- You, J., Ying, R., Ren, X., Hamilton, W. L., and Leskovec, J. Graphrnn: A deep generative model for graphs. *ICML*, 2018.
- Zhang, M. and Chen, Y. Weisfeiler-lehman neural machine for link prediction. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 575–583. ACM, 2017.
- Zhang, M. and Chen, Y. Link prediction based on graph neural networks. *arXiv preprint arXiv:1802.09691*, 2018.
- Zhu, X., Ghahramani, Z., and Lafferty, J. D. Semi-supervised learning using gaussian fields and harmonic functions. In *Proceedings of the 20th International conference on Machine learning (ICML-03)*, pp. 912–919, 2003.