

End-to-End Learning and Optimization on Graphs

BRYAN WILDER, Harvard University
ERIC EWING, University of Southern California
BISTRA DILKINA, University of Southern California
MILIND TAMBE, Harvard University

Real-world applications often combine learning and optimization problems on graphs. For instance, our objective may be to cluster the graph in order to detect meaningful communities (or solve other common graph optimization problems such as facility location, maxcut, and so on). However, graphs or related attributes are often only partially observed, introducing learning problems such as link prediction which must be solved prior to optimization. We propose an approach to integrate a differentiable proxy for common graph optimization problems into training of machine learning models for tasks such as link prediction. This allows the model to focus specifically on the downstream task that its predictions will be used for. Experimental results show that our end-to-end system obtains better performance on example optimization tasks than can be obtained by combining state of the art link prediction methods with expert-designed graph optimization algorithms.

Additional Key Words and Phrases: graph optimization; clustering; link prediction

ACM Reference Format:

Bryan Wilder, Eric Ewing, Bistra Dilkina, and Milind Tambe. 2018. End-to-End Learning and Optimization on Graphs. 1, 1 (June 2018), 8 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

While deep learning has proven enormously successful at a range of tasks, an expanding area of interest concerns systems that can flexibly combine learning with optimization. Examples include recent attempts to solve combinatorial optimization problems using neural architectures [Bello et al. 2016; Khalil et al. 2017; Kool et al. 2019; Vinyals et al. 2015], as well as work which incorporates explicit optimization algorithms into larger differentiable systems [Amos and Kolter 2017; Donti et al. 2017; Wilder et al. 2019]. The ability to combine learning and optimization promises improved performance for real-world problems which require decisions to be made on the basis of machine learning predictions by enabling end-to-end training which focuses the learned model on the decision problem at hand.

Authors' addresses: Bryan Wilder, bwilder@g.harvard.edu, Harvard University; Eric Ewing, ericewin@usc.edu, University of Southern California; Bistra Dilkina, dilkina@usc.edu, University of Southern California; Milind Tambe, milind_tambe@harvard.edu, Harvard University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

XXXX-XXXX/2018/6-ART \$15.00

<https://doi.org/10.1145/1122445.1122456>

We focus on graph optimization problems, an expansive subclass of combinatorial optimization. While graph optimization is ubiquitous across domains, complete applications must also solve machine learning challenges. For instance, the input graph is usually incomplete; some edges may be unobserved or nodes may have attributes that are only partially known. Recent work has introduced sophisticated methods for tasks such as link prediction and semi-supervised classification [Hamilton et al. 2017; Kipf and Welling 2017; Perozzi et al. 2014; Schlichtkrull et al. 2018; Zhang and Chen 2018], but these methods are developed in isolation of downstream optimization tasks. Most current solutions use a two-stage approach which first trains a model using a standard loss and then plugs the model's predictions into an optimization algorithm ([Bahulkar et al. 2018; Berlusconi et al. 2016; Burgess et al. 2016; Tan et al. 2016; Yan and Gregory 2012]). However, predictions which minimize a standard loss function (e.g., cross-entropy) may be suboptimal for specific optimization tasks, especially in difficult settings where even the best model is imperfect.

A preferable approach is to directly incorporate the downstream optimization problem into the training of the machine learning model [Donti et al. 2017; Wilder et al. 2019], which requires a differentiable layer that produces a solution to the optimization problem. To date, there are two main approaches to differentiable optimization. First, training a generic neural network to directly output a solution to the optimization problem. This approach often requires a large amount of data and results in suboptimal optimization performance because the network needs to discover algorithmic structure entirely from scratch. Second, hand-crafting a differentiable solver for the particular optimization problem (using, e.g., the LP relaxation of an integral problem [Wilder et al. 2019]) and including this solver as a layer in the network. This approach requires manual effort to develop a differentiable solver for each particular problem and often results in cumbersome systems that must, e.g., call a LP solver in every forward pass.

We propose a new approach that gets the best of both worlds by incorporating a more generic form of algorithmic structure into the network, which can then be automatically fine-tuned to solve a range of optimization tasks. In particular, we use a differentiable version of k -means clustering. Clustering is motivated by the observation that graph neural networks embed nodes into a continuous space, allowing us to approximate optimization over the discrete graph with optimization in continuous embedding space. We then interpret the cluster assignments as a solution to the discrete problem. We instantiate this approach for two classes of optimization problems: those that require *partitioning* the graph (e.g., community detection or maxcut), and those that require *selecting a subset of K nodes* (facility location, influence maximization, immunization, etc).

We don't claim that clustering is the right algorithmic structure for all tasks, but it is sufficient for many problems as shown in this paper.

In short, we make three contributions. First, we introduce a general framework for integrating graph learning and optimization, with optimization in continuous space as a proxy for the discrete problem. Second, we show how to differentiate through the clustering layer, allowing it to be used in deep learning systems. Third, we show experimental improvements over both two-stage baselines as well as alternate end-to-end approaches on a range of example datasets and optimization problems.

2 RELATED WORK

We build on a recent work on decision-focused learning [Demirovic et al. 2019; Donti et al. 2017; Wilder et al. 2019], which includes a solver for an optimization problem into training in order to improve performance on a downstream decision problem. Some work in structured prediction also integrates differentiable solvers for discrete problems (e.g., image segmentation [Djolonga and Krause 2017] or time series alignment [Mensch and Blondel 2018]). Our work differs in two ways. First, we tackle more difficult optimization problems. Previous work mostly focuses on convex problems [Donti et al. 2017] or discrete problems with near-lossless convex relations [Djolonga and Krause 2017; Wilder et al. 2019]. We focus on highly combinatorial problems where the methods of choice are hand-designed discrete algorithms. Second, in response to this difficulty, we differ methodologically in that we do not attempt to include a solver for the exact optimization problem at hand (or a close relaxation of it). Instead, we include a more generic algorithmic skeleton that is automatically finetuned to the optimization problem at hand.

There is also recent interest in training neural networks to solve combinatorial optimization problems [Bello et al. 2016; Khalil et al. 2017; Kool et al. 2019; Vinyals et al. 2015]. While we focus mostly on combining graph learning with optimization, our model can also be trained just to solve an optimization problem given complete information about the input. The main methodological difference is that we include more structure via a differentiable k -means layer instead of using more generic tools (e.g., feed-forward or attention layers). Another difference is that prior work mostly trains via reinforcement learning. By contrast, we use a differentiable approximation to the objective which removes the need for a policy gradient estimator. This is a benefit of our architecture, in which the final decision is fully differentiable in terms of the model parameters instead of requiring non-differentiable selection steps (as in [Bello et al. 2016; Khalil et al. 2017; Kool et al. 2019]). We give our end-to-end baseline ("GCN-e2e") the same advantage by training it with the same differentiable decision loss as our own model instead of forcing it to use noisier policy gradient estimates.

Finally, some work uses deep architectures as a part of a clustering algorithm [Guo et al. 2017; Law et al. 2017; Shaham et al. 2018; Tian et al. 2014], or includes a clustering step as a component of a deep network [Greff et al. 2016, 2017]. While some techniques are similar, the overall task we address and framework we propose are entirely distinct. Our aim is not to cluster a Euclidean dataset (as in [Guo

et al. 2017; Law et al. 2017; Shaham et al. 2018; Tian et al. 2014]), or to solve perceptual grouping problems (as in [Greff et al. 2016, 2017]). Rather, we propose an approach for graph optimization problems. Perhaps the closest of this work is Neural EM [Greff et al. 2017], which uses an unrolled EM algorithm to learn representations of visual objects. Rather than using EM to infer representations for objects, we use k -means in graph embedding space to solve an optimization problem. There is also some work which uses deep networks for graph clustering [Xie et al. 2016; Yang et al. 2016]. However, none of this work includes an explicit clustering algorithm in the network, and none consider our goal of integrating graph learning and optimization.

3 SETTING

We consider settings that combine learning and optimization. The input is a graph $G = (V, E)$, which is in some way partially observed. We will formalize our problem in terms of link prediction as an example, but our framework applies to other common graph learning problems (e.g., semi-supervised classification). In link prediction, the graph is not entirely known; instead, we observe only training edges $E^{train} \subset E$. Let A denote the adjacency matrix of the graph and A^{train} denote the adjacency matrix with only the training edges. The learning task is to predict A from A^{train} . In domains we consider, the motivation for performing link prediction, is to solve a decision problem for which the objective depends on the full graph. Specifically, we have a decision variable x , objective function $f(x, A)$, and a feasible set \mathcal{X} . We aim to solve the optimization problem

$$\max_{x \in \mathcal{X}} f(x, A). \quad (1)$$

However, A is unobserved. We can also consider an inductive setting in which we observe graphs A_1, \dots, A_m as training examples and then seek to predict edges for a partially observed graph from the same distribution. The most common approach to either setting is to train a model to reconstruct A from A^{train} using a standard loss function (e.g., cross-entropy), producing an estimate \hat{A} . The *two-stage* approach plugs \hat{A} into an optimization algorithm for Problem 1, maximizing $f(x, \hat{A})$.

We propose end-to-end models which map from A^{train} directly to a feasible decision x . The model will be trained to maximize $f(x, A^{train})$, i.e., the quality of its decision evaluated on the training data (instead of a loss $\ell(\hat{A}, A^{train})$ that measures purely predictive accuracy). One approach is to "learn away" the problem by training a standard model (e.g., a GCN) to map directly from A^{train} to x . However, this forces the model to entirely rediscover algorithmic concepts, while two-stage methods are able to exploit highly sophisticated optimization methods. We propose an alternative that embeds algorithmic structure into the learned model, getting the best of both worlds.

4 APPROACH: CLUSTERNET

Our proposed ClusterNet system (Figure 1) merges two differentiable components into a system that is trained end-to-end. First, a *graph embedding* layer which uses A^{train} and any node features to embed the nodes of the graph into \mathbb{R}^p . In our experiments, we

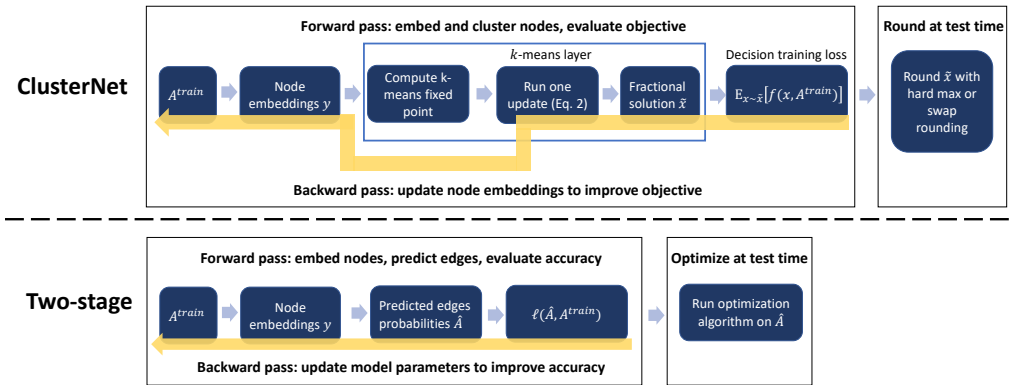


Fig. 1. Top: ClusterNet, our proposed system. Bottom: a typical two-stage approach.

use GCNs [Kipf and Welling 2017]. Second, a layer that performs *differentiable optimization*. This layer takes the continuous-space embeddings as input and uses them to produce a solution x to the graph optimization problem. Specifically, we propose to use a layer that implements a differentiable version of K -means clustering. This layer produces a soft assignment of the nodes to clusters, along with the cluster centers in embedding space.

The intuition is that cluster assignments can be interpreted as the solution to many common graph optimization problems. For instance, in community detection we can interpret the cluster assignments as assigning the nodes to communities. Or, in maxcut, we can use two clusters to assign nodes to either side of the cut. Another example is maximum coverage and related problems, where we attempt to select a set of K nodes which cover (are neighbors to) as many other nodes as possible. This problem can be approximated by clustering the nodes into K components and choosing nodes whose embedding is close to the center of each cluster. We do not claim that any of these problems is exactly reducible to K -means. Rather, the idea is that including K -means as a layer in the network provides a useful inductive bias. This algorithmic structure can be fine-tuned to specific problems by training the first component, which produces the embeddings, so that the learned representations induce clusterings with high objective value for the underlying downstream optimization task. We now explain the optimization layer of our system in greater detail. We start by detailing the forward and the backward pass for the clustering procedure, and then explain how the cluster assignments can be interpreted as solutions to the graph optimization problem.

4.1 Forward pass

Let x_j denote the embedding of node j and μ_k denote the center of cluster k . r_{jk} denotes the degree to which node j is assigned to cluster k . In traditional K -means, this is a binary quantity, but we will relax it to a fractional value such that $\sum_k r_{jk} = 1$ for all j . Specifically, we take $r_{jk} = \frac{\exp(-\beta \|x_j - \mu_k\|)}{\sum_\ell \exp(-\beta \|x_j - \mu_\ell\|)}$, which is a soft-min assignment of each point to the cluster centers based on distance. While our architecture can be used with any norm $\|\cdot\|$, we use the negative

cosine similarity due to its strong empirical performance. β is an inverse-temperature hyperparameter; taking $\beta \rightarrow \infty$ recovers the standard k -means assignment. We can optimize the cluster centers via an iterative process analogous to the typical k -means updates by alternately setting

$$\mu_k = \frac{\sum_j r_{jk} x_j}{\sum_j r_{jk}} \quad \forall k = 1 \dots K \quad (2)$$

$$r_{jk} = \frac{\exp(-\beta \|x_j - \mu_k\|)}{\sum_\ell \exp(-\beta \|x_j - \mu_\ell\|)} \quad \forall k = 1 \dots K, j = 1 \dots n. \quad (3)$$

These iterates converge to a fixed point where μ remains the same between successive updates [MacKay 2003].

4.2 Backward pass

We will use the implicit function theorem to analytically differentiate through the fixed point that the forward pass k -means iterates converge to, obtaining expressions for $\frac{\partial \mu}{\partial x}$ and $\frac{\partial r}{\partial x}$. Previous work [Donti et al. 2017; Wilder et al. 2019] has used the implicit function theorem to differentiate through the KKT conditions of optimization problems; here we take a more direct approach that characterizes the update process itself. Doing so allows us to backpropagate gradients from the decision loss to the component that produced the embeddings x . Define a function $f : \mathbb{R}^{Kp} \rightarrow \mathbb{R}$ as

$$f_{i,\ell}(\mu, x) = \mu_i^\ell - \frac{\sum_j r_{jk} x_j^\ell}{\sum_j r_{jk}} \quad (4)$$

Now, (μ, x) are a fixed point of the iterates if $f(\mu, x) = \mathbf{0}$. Applying the implicit function theorem yields that $\frac{\partial \mu}{\partial x} = - \left[\frac{\partial f(\mu, x)}{\partial \mu} \right]^{-1} \frac{\partial f(\mu, x)}{\partial x}$, from which $\frac{\partial r}{\partial x}$ can be easily obtained via the chain rule.

Exact backward pass: We now examine the process of calculating $\frac{\partial \mu}{\partial x}$. Both $\frac{\partial f(\mu, x)}{\partial x}$ and $\frac{\partial f(\mu, x)}{\partial \mu}$ can be easily calculated in closed form (see appendix¹). Computing the former requires time $O(nKp^2)$. Computing the latter requires $O(npK^2)$ time, after which it must be inverted (or else iterative methods must be used to compute the

¹<https://arxiv.org/abs/1905.13732>

product with its inverse). This requires time $O(K^3 p^3)$ since it is a matrix of size $(Kp) \times (Kp)$. While the exact backward pass may be feasible for some problems, it quickly becomes burdensome for large instances. We now propose a fast approximation.

Approximate backward pass: We start from the observation that $\frac{\partial f}{\partial \mu}$ will often be dominated by its diagonal terms (the identity matrix). The off-diagonal entries capture the extent to which updates to one entry of μ indirectly impact other entries via changes to the cluster assignments r . However, when the cluster assignments are relatively firm, r will not be highly sensitive to small changes to the cluster centers. We find to be typical empirically, especially since the optimal choice of the parameter β (which controls the hardness of the cluster assignments) is typically fairly high. Under these conditions, we can approximate $\frac{\partial f}{\partial \mu}$ by its diagonal, $\frac{\partial f}{\partial \mu} \approx I$. This in turn gives $\frac{\partial \mu}{\partial x} \approx -\frac{\partial f}{\partial x}$.

We can formally justify this approximation when the clusters are relatively balanced and well-separated. More precisely, define $c(j) = \arg \max_i r_{ji}$ to be the closest cluster to point j . Proposition 1 (proved in the appendix) shows that the quality of the diagonal approximation improves exponentially quickly in the product of two terms: β , the hardness of the cluster assignments, and δ , which measures how well separated the clusters are. α (defined below) measures the balance of the cluster sizes. We assume for convenience that the input is scaled so $\|x_j\|_1 \leq 1 \forall j$.

PROPOSITION 1. *Suppose that for all points j , $\|x_j - \mu_i\| - \|x_j - \mu_{c(j)}\| \geq \delta$ for all $i \neq c(j)$ and that for all clusters i , $\sum_{j=1}^n r_{ji} \geq \alpha n$. Moreover, suppose that $\beta \delta > \log \frac{2\beta K^2}{\alpha}$. Then,*

$$\left\| \frac{\partial f}{\partial \mu} - I \right\|_1 \leq \exp(-\delta \beta) \left(\frac{K^2 \beta}{\frac{1}{2}\alpha - K^2 \beta \exp(-\delta \beta)} \right)$$

where $\|\cdot\|_1$ is the operator 1-norm.

We now show that the approximate gradient obtained by taking $\frac{\partial f}{\partial \mu} = I$ can be calculated by unrolling a single iteration of the forward-pass updates from Equation 2 at convergence. Examining Equation 4, we see that the first term (μ_i^t) is constant with respect to x , since here μ is a fixed value. Hence,

$$-\frac{\partial f_k}{\partial x} = \frac{\partial}{\partial x} \frac{\sum_j r_{jk} x_j}{\sum_j r_{jk}}$$

which is just the update equation for μ_k . Since the forward-pass updates are written entirely in terms of differentiable functions, we can automatically compute the approximate backward pass with respect to x (i.e., compute products with our approximations to $\frac{\partial \mu}{\partial x}$ and $\frac{\partial r}{\partial x}$) by applying standard autodifferentiation tools to the final update of the forward pass. Compared to computing the exact analytical gradients, this avoids the need to explicitly reason about or invert $\frac{\partial f}{\partial \mu}$. The final iteration (the one which is differentiated through) requires time $O(npK)$, linear in the size of the data.

Compared to differentiating by unrolling the entire sequence of updates in the computational graph (as has been suggested for other problems [Andrychowicz et al. 2016; Domke 2012; Zheng et al. 2015]), our approach has two key advantages. First, it avoids storing

the entire history of updates and backpropagating through all of them. The runtime for our approximation is independent of the number of updates needed to reach convergence. Second, we can in fact use entirely non-differentiable operations to arrive at the fixed point, e.g., heuristics for the K -means problem, stochastic methods which only examine subsets of the data, etc. This allows the forward pass to scale to larger datasets since we can use the best algorithmic tools available, not just those that can be explicitly encoded in the autodifferentiation tool’s computational graph.

4.3 Obtaining solutions to the optimization problem

Having obtained the cluster assignments r , along with the centers μ , in a differentiable manner, we need a way to (1) differentially interpret the clustering as a soft solution to the optimization problem, (2) differentiate a relaxation of the objective value of the graph optimization problem in terms of that solution, and then (3) round to a discrete solution at test time. We give a generic means of accomplishing these three steps for two broad classes of problems: those that involve *partitioning the graph into K disjoint components*, and those that involve *selecting a subset of K nodes*.

Partitioning: (1) We can naturally interpret the cluster assignments r as a soft partitioning of the graph. (2) One generic continuous objective function (defined on soft partitions) follows from the random process of assigning each node j to a partition with probabilities given by r_j , repeating this process independently across all nodes. This gives the expected training decision loss $\ell = \mathbb{E}_{r^{hard} \sim r} [f(r^{hard}, A^{train})]$, where $r^{hard} \sim r$ denotes this random assignment. ℓ is now differentiable in terms of r , and can be computed in closed form via standard autodifferentiation tools for many problems of interest (see Section 5). (3) At test time, we simply apply a hard maximum to r to obtain each node’s assignment.

Subset selection: (1) Here, it is less obvious how to obtain a subset of K nodes from the cluster assignments. Our continuous solution will be a vector x , $0 \leq x \leq 1$, where $\|x\|_1 = K$. Intuitively, x_j is the probability of including x_j in the solution. Our approach obtains x_j by placing greater probability mass on nodes that are near the cluster centers. Specifically, each center μ_i is endowed with one unit of probability mass, which it allocates to the points x by as $a_{ij} = \text{softmax}(\eta \|x - \mu_i\|)_j$. The total probability allocated to node j is $b_j = \sum_{i=1}^K a_{ij}$. Since we may have $b_j > 1$, we pass b through a sigmoid function to cap the entries at 1; specifically, we take $x = 2 * \sigma(\gamma b) - 0.5$ where γ is a tunable parameter. If the resulting x exceeds the budget constraint ($\|x\|_1 > K$), we output $\frac{Kx}{\|x\|_1}$.

(2) We interpret this solution in terms of the objective similarly as above. Specifically, we consider the result of drawing a discrete solution $x^{hard} \sim x$ where every node j is included (i.e., set to 1) independently with probability x_j from the end of step (1). The training objective is then $\mathbb{E}_{x^{hard} \sim x} [f(x^{hard}, A^{train})]$. For many problems, this can again be computed and differentiated through in closed form (see Section 5).

(3) At test time, we need a feasible discrete vector x ; note that independently rounding the individual entries may produce a vector with more than K ones. Here, we apply a fairly generic approach

based on pipage rounding [Ageev and Sviridenko 2004], a randomized rounding scheme which has been applied to many problems (particularly those with submodular objectives). Pipage rounding can be implemented to produce a random feasible solution in time $O(n)$ [Karimi et al. 2017]; in practice we round several times and take the solution with the best decision loss on the observed edges. While pipage rounding has theoretical guarantees only for specific classes of functions, we find it to work well even in other domains (e.g., facility location). However, more domain-specific rounding methods can be applied if available.

5 EXPERIMENTAL RESULTS

We now show experiments on domains that combine link prediction with optimization.

Learning problem: In link prediction, we observe a partial graph and aim to infer which unobserved edges are present. In each of the experiments, we hold out 60% of the edges in the graph, with 40% observed during training. We used a graph dataset which is not included in our results to set our method’s hyperparameters, which were kept constant across datasets (see appendix for details). The learning task is to use the training edges to predict whether the remaining edges are present, after which we will solve an optimization problem on the predicted graph. The objective is to find a solution with high objective value measured on the *entire* graph, not just the training edges.

Optimization problems: We consider two optimization tasks, one from each of the broad classes introduced above. First, *community detection* aims to partition the nodes of the graph into K distinct subgroups which are dense internally, but with few edges across groups. Formally, the objective is to find a partition maximizing the modularity [Newman 2006b], defined as

$$\frac{1}{2m} \sum_{u,v \in V} \sum_{k=1}^K \left[A_{uv} - \frac{d_u d_v}{2m} \right] r_{uk} r_{vk}$$

. Here, d_v is the degree of node v , and r_{vk} is 1 if node v is assigned to community k and zero otherwise. This measures the number of edges within communities compared to the expected number if edges were placed randomly. Our clustering module has one cluster for each of the K communities. Defining B to be the modularity matrix with entries $B_{uv} = A_{uv} - \frac{d_u d_v}{2m}$, our training objective (the expected value of a partition sampled according to r) is $\frac{1}{2m} \text{Tr} [r^T B^{train} r]$.

Second, *facility location*, where the decision problem is to select a subset of K nodes from the graph, minimizing the maximum distance from any node to a facility (selected node). Letting $d(v, S)$ be the shortest path length from a vertex v to a set of vertices S , the objective is $f(S) = \min_{|S| \leq k} \max_{v \in V} d(v, S)$. To obtain the training loss, we take two steps. First, we replace $d(v, S)$ by $\mathbb{E}_{S \sim x} [d(v, S)]$, where $S \sim x$ denotes drawing a set from the product distribution with marginals x . This can easily be calculated in closed form [Karimi et al. 2017]. Second, we replace the min with a softmax.

Baseline learning methods: We instantiate ClusterNet using a 2-layer GCN for node embeddings, followed by a clustering layer. We compare to three families of baselines. *First*, GCN-2stage, the two stage approach which first trains a model for link prediction, and

then inputs the predicted graph into an optimization algorithm. For link prediction, we use the GCN-based system of [Schlichtkrull et al. 2018] (we also adopt their training procedure, including negative sampling and edge dropout). For the optimization algorithms, we use standard approaches for each domain, outlined below. *Second*, “train”, which runs each optimization algorithm only on the observed training subgraph (without attempting any link prediction). *Third*, GCN-e2e, an end-to-end approach which does not include explicit algorithm structure. We train a GCN-based network to directly predict the final decision variable (r or x) using the same training objectives as our own model. Empirically, we observed best performance with a 2-layer GCN. This baseline allows us to isolate the benefits of including algorithmic structure.

Baseline optimization approaches: In each domain, we compare to expert-designed optimization algorithms found in the literature. In community detection, we compare to “CNM” [Clauset et al. 2004], an agglomerative approach, “Newman”, an approach that recursively partitions the graph [Newman 2006a], and “SC”, which performs spectral clustering [Von Luxburg 2007] on the modularity matrix. In facility location, we compare to “greedy”, the common heuristic of iteratively selecting the point with greatest marginal improvement in objective value, and “gonzalez” [Gonzalez 1985], an algorithm which iteratively selects the node furthest from the current set (which attains the optimal 2-approximation).

Datasets: We use several standard graph datasets: cora [Sen et al. 2008] (a citation network with 2,708 nodes), citeseer [Sen et al. 2008] (a citation network with 3,327 nodes), protein [Collection 2017c] (a protein interaction network with 3,133 nodes), adol [Collection 2017a] (an adolescent social network with 2,539 vertices), and fb [Collection 2017b; Leskovec and McAuley 2012] (an online social network with 2,888 nodes). For facility location, we use the largest connected component of the graph (since otherwise distances may be infinite). Cora and citeseer have node features (based on a bag-of-words representation of the document), which were given to all GCN-based methods. For the other datasets, we generated unsupervised node2vec features [Grover and Leskovec 2016] using the training edges.

5.1 Results on single graphs

We start out with results for the combined link prediction and optimization problem. Table 1 shows the objective value obtained by each approach on the full graph for community detection, with Table 2 showing facility location. We focus first on the “Learning + optimization” column which shows the combined link prediction/optimization task. We use $K = 5$; $K = 10$ is very similar and may be found in the appendix. ClusterNet outperforms the baselines in nearly all cases, often substantially. GCN-e2e learns to produce nontrivial solutions, often rivaling the other baseline methods. However, the explicit structure used by our approach ClusterNet results in much higher performance.

Interestingly, the two stage approach sometimes performs worse than the train-only baseline which optimizes just based on the training edges (without attempting to learn). This indicates that approaches which attempt to accurately reconstruct the graph can *sometimes* miss qualities which are important for optimization, and

Table 1. Performance on the community detection task

	Learning + optimization					Optimization				
	cora	cite.	prot.	adol	fb	cora	cite.	prot.	adol	fb
ClusterNet	0.54	0.55	0.29	0.49	0.30	0.72	0.73	0.52	0.58	0.76
GCN-e2e	0.16	0.02	0.13	0.12	0.13	0.19	0.03	0.16	0.20	0.23
Train-CNM	0.20	0.42	0.09	0.01	0.14	0.08	0.34	0.05	0.57	0.77
Train-Newman	0.09	0.15	0.15	0.15	0.08	0.20	0.23	0.29	0.30	0.55
Train-SC	0.03	0.02	0.03	0.23	0.19	0.09	0.05	0.06	0.49	0.61
GCN-2stage-CNM	0.17	0.21	0.18	0.28	0.13	-	-	-	-	-
GCN-2stage-Newman	0.00	0.00	0.00	0.14	0.02	-	-	-	-	-
GCN-2stage-SC	0.14	0.16	0.04	0.31	0.25	-	-	-	-	-

Table 2. Performance on the facility location task.

	Learning + optimization					Optimization				
	cora	cite.	prot.	adol	fb	cora	cite.	prot.	adol	fb
ClusterNet	10	14	6	6	4	9	14	6	5	3
GCN-e2e	12	15	8	6	5	11	14	7	6	5
Train-greedy	14	16	8	8	6	9	14	7	6	5
Train-gonzalez	12	17	8	6	6	10	15	7	7	3
GCN-2Stage-greedy	14	17	8	7	6	-	-	-	-	-
GCN-2Stage-gonzalez	13	17	8	6	6	-	-	-	-	-

in the worst case may simply add noise that overwhelms the signal in the training edges. In order to confirm that the two-stage method learned to make meaningful predictions, in the appendix we give AUC values for each dataset. The average AUC value is 0.7584, indicating that the two stage model does learn to make nontrivial predictions. However, the small amount of training data (only 40% of edges are observed) prevents it from perfectly reconstructing the true graph. This drives home the point that decision-focused learning methods can offer substantial benefits when highly accurate predictions are out of reach even for sophisticated methods.

We next examine an optimization-only task where the entire graph is available as input (the “Optimization” column of Tables 1 and Table 2). This tests ClusterNet’s ability to learn to solve combinatorial optimization problems compared to expert-designed algorithms, even when there is no partial information or learning problem in play. We find that ClusterNet is highly competitive, meeting and frequently exceeding the baselines. It is particularly effective for community detection, where we observe large (> 3x) improvements compared to the best baseline on some datasets. At facility location, our method always at least ties the baselines, and frequently improves on them. These experiments provide evidence that our approach, which is automatically specialized during training to optimize on a given graph, can rival and exceed hand-designed algorithms from the literature. The alternate learning approach, GCN-e2e, at best ties the baselines and typically underperforms. This underscores the benefit of including algorithmic structure as a part of the end-to-end architecture.

5.2 Generalizing across graphs

Next, we investigate whether our method can learn generalizable strategies for optimization: can we train the model on one set of graphs drawn from some distribution and then apply it to unseen graphs? We consider two graph distributions. First, a synthetic generator introduced by [Wilder et al. 2018], which is based on the spatial preferential attachment model [Barthélemy 2011] (details in the appendix). We use 20 training graphs, 10 validation, and 30 test. Second, a dataset obtained by splitting the pubmed graph into 20 components using metis [Karypis and Kumar 1998]. We fix 10 training graphs, 2 validation, and 8 test. At test time, only 40% of the edges in each graph are revealed, matching the “Learning + optimization” setup above.

Table 3 shows the results. To start out, we do not conduct any fine-tuning to the test graphs, evaluating entirely the generalizability of the learned representations. ClusterNet outperforms all baseline methods on all tasks, except for facility location on pubmed where it places second. We conclude that the learned model successfully generalizes to completely unseen graphs. We next investigate (in the “finetune” section of Table 3) whether ClusterNet’s performance can be further improved by fine-tuning to the 40% of observed edges for each test graph (treating each test graph as an instance of the link prediction problem from Section 5.1, but with the model initialized to the parameters of the model learned over the training graphs). We see that ClusterNet’s performance typically improves, indicating that fine-tuning can allow us to extract additional gains if extra training time is available.

Interestingly, *only* fine-tuning (not using the training graphs at all) yields similar performance (the row “ClstrNet-ft-only”). While our earlier results show that ClusterNet can learn generalizable

Table 3. Inductive results. “%” is the fraction of test instances for which a method attains top performance (including ties). “Finetune” methods are excluded from this in the “No finetune” section.

	Community detection				Facility location				
	synthetic		pubmed		synthetic		pubmed		
	Avg.	%	Avg.	%	Avg.	%	Avg.	%	
No finetune					No finetune				
ClusterNet	0.57	26/30	0.30	7/8	ClusterNet	7.90	25/30	7.88	3/8
GCN-e2e	0.26	0/30	0.01	0/8	GCN-e2e	8.63	11/30	8.62	1/8
Train-CNM	0.14	0/30	0.16	1/8	Train-greedy	14.00	0/30	9.50	1/8
Train-Newman	0.24	0/30	0.17	0/8	Train-gonzalez	10.30	2/30	9.38	1/8
Train-SC	0.16	0/30	0.04	0/8	2Stage-greedy	9.60	3/30	10.00	0/8
2Stage-CNM	0.51	0/30	0.24	0/8	2Stage-gonz.	10.00	2/30	6.88	5/8
2Stage-Newman	0.01	0/30	0.01	0/8	ClstrNet-1train	7.93	12/30	7.88	2/8
2Stage-SC	0.52	4/30	0.15	0/8					
ClstrNet-1train	0.55	0/30	0.25	0/8					
Finetune					Finetune				
ClstrNet-ft	0.60	20/30	0.40	2/8	ClstrNet-ft	8.08	12/30	8.01	3/8
ClstrNet-ft-only	0.60	10/30	0.42	6/8	ClstrNet-ft-only	7.84	16/30	7.76	4/8

strategies, doing so may not be necessary when there is the opportunity to fine-tune. This allows a user to trade off between quality and runtime: without fine-tuning, applying our method at test time requires just a single forward pass through the network, which is extremely efficient. If additional computational cost at test time is acceptable, fine-tuning can be used to improve performance. Complete runtimes for all methods are shown in the appendix. ClusterNet’s forward pass (i.e., no fine-tuning) is extremely efficient, requiring at most 0.23 seconds on the largest network, and is *always faster than the baselines* (on identical hardware). Fine-tuning requires longer, on par with the slowest baseline.

We lastly investigate the reason why pretraining provides little to no improvement over only fine-tuning. Essentially, we find that ClusterNet is extremely sample-efficient: using only a single training graph results in nearly as good performance as the full training set (and still better than all of the baselines), as seen in the “ClstrNet-1train” row of Table 3. That is, ClusterNet is capable of learning optimization strategies that generalize with strong performance to completely unseen graphs after observing only a single training example. This underscores the benefits of including algorithmic structure as a part of the architecture, which guides the model towards learning meaningful strategies.

REFERENCES

Alexander A Ageev and Maxim I Sviridenko. 2004. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization* 8, 3 (2004), 307–328.

B. Amos and J. Kolter. 2017. Optnet: Differentiable optimization as a layer in neural networks. In *ICML*.

Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. 2016. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*. 3981–3989.

Ashwin Bahulkar, Boleslaw K Szymanski, N Orkun Baycik, and Thomas C Sharkey. 2018. Community Detection with Edge Augmentation in Criminal Networks. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*.

Marc Barthélemy. 2011. Spatial networks. *Physics Reports* 499, 1-3 (2011), 1–101.

Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. 2016. Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940* (2016).

Giulia Berlusconi, Francesco Calderoni, Nicola Parolini, Marco Verani, and Carlo Piccardi. 2016. Link prediction in criminal networks: A tool for criminal intelligence analysis. *PLoS one* 11, 4 (2016), e0154244.

Matthew Burgess, Eytan Adar, and Michael Cafarella. 2016. Link-prediction enhanced consensus clustering for complex networks. *PLoS one* 11, 5 (2016), e0153384.

Aaron Clauset, Mark EJ Newman, and Christopher Moore. 2004. Finding community structure in very large networks. *Physical review E* 70, 6 (2004), 066111.

Koblenz Network Collection. 2017a. Adolescent health. http://konect.uni-koblenz.de/networks/moreno_health.

Koblenz Network Collection. 2017b. Facebook (NIPS). <http://konect.uni-koblenz.de/networks/ego-facebook>.

Koblenz Network Collection. 2017c. Human protein (Vidal). <http://konect.uni-koblenz.de/networks/maayan-vidal>.

Emir Demirovic, Peter J Stuckey, James Bailey, Jeffrey Chan, Chris Leckie, Kotagiri Ramamohanarao, and Tias Guns. 2019. Prediction + optimisation for the knapsack problem. In *CPAIOR*.

Josip Djolonga and Andreas Krause. 2017. Differentiable learning of submodular models. In *NeurIPS*.

Justin Domke. 2012. Generic methods for optimization-based modeling. In *Artificial Intelligence and Statistics*. 318–326.

Priya Donti, Brandon Amos, and J Zico Kolter. 2017. Task-based end-to-end model learning in stochastic optimization. In *Advances in Neural Information Processing Systems*. 5484–5494.

Teofilo F Gonzalez. 1985. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science* 38 (1985), 293–306.

Klaus Greff, Antti Rasmus, Mathias Berglund, Tele Hao, Harri Valpola, and Jürgen Schmidhuber. 2016. Tagger: Deep unsupervised perceptual grouping. In *NeurIPS*.

Klaus Greff, Sjoerd van Steenkiste, and Jürgen Schmidhuber. 2017. Neural expectation maximization. In *NeurIPS*.

Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 855–864.

Xifeng Guo, Long Gao, Xinwang Liu, and Jianping Yin. 2017. Improved Deep Embedded Clustering with Local Structure Preservation. In *IJCAI*.

W. Hamilton, Z. Ying, and J. Leskovec. 2017. Inductive representation learning on large graphs. In *NIPS*.

Mohammad Karimi, Mario Lucic, Hamed Hassani, and Andreas Krause. 2017. Stochastic submodular maximization: The case of coverage functions. In *Advances in Neural Information Processing Systems*.

George Karypis and Vipin Kumar. 1998. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing* 20, 1 (1998), 359–392.

E. Khalil, H. Dai, Y. Zhang, B. Dilikina, and L. Song. 2017. Learning combinatorial optimization algorithms over graphs. In *NIPS*.

T. Kipf and M. Welling. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.

Wouter Kool, Herke van Hoof, and Max Welling. 2019. Attention, Learn to Solve Routing Problems!. In *ICLR*.

Marc T Law, Raquel Urtasun, and Richard S Zemel. 2017. Deep spectral clustering learning. In *ICML*.

- Jure Leskovec and Julian J Mcauley. 2012. Learning to discover social circles in ego networks. In *Advances in neural information processing systems*. 539–547.
- David JC MacKay. 2003. *Information theory, inference and learning algorithms*. Cambridge university press.
- Arthur Mensch and Mathieu Blondel. 2018. Differentiable Dynamic Programming for Structured Prediction and Attention. In *ICML*.
- Mark EJ Newman. 2006a. Finding community structure in networks using the eigenvectors of matrices. *Physical review E* 74, 3 (2006), 036104.
- Mark EJ Newman. 2006b. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences* 103, 23 (2006), 8577–8582.
- Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 701–710.
- M. Schlichtkrull, T. Kipf, P. Bloem, R. Van Den Berg, I. Titov, and M. Welling. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*.
- Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.
- Uri Shaham, Kelly Stanton, Henry Li, Boaz Nadler, Ronen Basri, and Yuval Kluger. 2018. Spectralnet: Spectral clustering using deep neural networks. In *ICLR*.
- Suo-Yi Tan, Jun Wu, Linyuan Lü, Meng-Jun Li, and Xin Lu. 2016. Efficient network disintegration under incomplete information: the comic effect of link prediction. *Scientific reports* 6 (2016), 22916.
- Fei Tian, Bin Gao, Qing Cui, Enhong Chen, and Tie-Yan Liu. 2014. Learning deep representations for graph clustering. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- O. Vinyals, M. Fortunato, and N. Jaitly. 2015. Pointer networks. In *NIPS*.
- Ulrike Von Luxburg. 2007. A tutorial on spectral clustering. *Statistics and computing* 17, 4 (2007), 395–416.
- B. Wilder, B. Dilkina, and M. Tambe. 2019. Melding the Data-Decisions Pipeline: Decision-Focused Learning for Combinatorial Optimization. In *AAAI*.
- Bryan Wilder, Han Ching Ou, Kayla de la Haye, and Milind Tambe. 2018. Optimizing network structure for preventative health. In *AAMAS*.
- Junyuan Xie, Ross Girshick, and Ali Farhadi. 2016. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*. 478–487.
- Bowen Yan and Steve Gregory. 2012. Detecting community structure in networks using edge prediction methods. *Journal of Statistical Mechanics: Theory and Experiment* 2012, 09 (2012), P09008.
- Liang Yang, Xiaochun Cao, Dongxiao He, Chuan Wang, Xiao Wang, and Weixiong Zhang. 2016. Modularity Based Community Detection with Deep Learning.. In *IJCAI*, Vol. 16. 2252–2258.
- M. Zhang and Y. Chen. 2018. Link prediction based on graph neural networks. In *NIPS*.
- Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. 2015. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE international conference on computer vision*. 1529–1537.