

# A Quest for Structure: Jointly Learning the Graph Structure and Semi-Supervised Classification

Xuan Wu\*  
School of Computer Science  
Carnegie Mellon University  
xuanw1@andrew.cmu.edu

Lingxiao Zhao\*  
H. John Heinz III College  
Carnegie Mellon University  
lingxiao@cmu.edu

Leman Akoglu  
H. John Heinz III College  
Carnegie Mellon University  
lakoglu@andrew.cmu.edu

## ABSTRACT

Semi-supervised learning (SSL) is effectively used for numerous classification problems, thanks to its ability to make use of abundant unlabeled data. The main assumption of various SSL algorithms is that the nearby points on the data manifold are likely to share a label. Graph-based SSL constructs a graph from point-cloud data as an approximation to the underlying manifold, followed by label inference. It is no surprise that the quality of the constructed graph in capturing the essential structure of the data is critical to the accuracy of the subsequent inference step [6].

How should one construct a graph from the input point-cloud data for graph-based SSL? In this work we introduce a new, parallel graph learning framework (called PG-LEARN) for the graph construction step of SSL. Our solution has two main ingredients: (1) a gradient-based optimization of the edge weights (more specifically, different kernel bandwidths in each dimension) based on a validation loss function, and (2) a parallel hyperparameter search algorithm with an adaptive resource allocation scheme. In essence, (1) allows us to search around a (random) initial hyperparameter configuration for a better one with lower validation loss. Since the search space of hyperparameters is huge for high-dimensional problems, (2) empowers our gradient-based search to go through as many different initial configurations as possible, where runs for relatively unpromising starting configurations are terminated early to allocate the time for others. As such, PG-LEARN is a carefully-designed hybrid of random and adaptive search. Through experiments on multi-class classification problems, we show that PG-LEARN significantly outperforms a variety of existing graph construction schemes in accuracy (per fixed time budget for hyperparameter tuning), and scales more effectively to high dimensional problems.

## ACM Reference Format:

Xuan Wu, Lingxiao Zhao, and Leman Akoglu. 2018. A Quest for Structure: Jointly Learning the Graph Structure and Semi-Supervised Classification. In *The 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*, October 22–26, 2018, Torino, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3269206.3271692>

\*X. Wu & L. Zhao equally contributed to the paper.

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](mailto:permissions@acm.org).

CIKM '18, October 22–26, 2018, Torino, Italy

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-6014-2/18/10...\$15.00

<https://doi.org/10.1145/3269206.3271692>

## 1 INTRODUCTION

Graph-based semi-supervised learning (SSL) algorithms, based on graph min-cuts [2], local and global consistency [26], and harmonic energy minimization [27], have been used widely for classification and regression problems. These employ the manifold assumption to take advantage of the unlabeled data, which dictates the label (or value) function to change smoothly on the data manifold.

The data manifold is modeled by a graph structure. In some cases, this graph is explicit; for example, explicit social network connections between individuals have been used in predicting their political orientation [4], age [19], income [9], etc. In others (e.g., image classification), the data is in (feature) vector form, where a graph is to be constructed from point-cloud data. In this graph, nodes correspond to labeled and unlabeled data points and edge weights encode pairwise similarities. Often, some graph sparsification scheme is also used to ensure that the SSL algorithm runs efficiently. Then, labeling is done in such a way that instances connected by large weights are assigned similar labels.

In essence, graph-based SSL for non-graph data consists of two steps: (1) graph construction, and (2) label inference. It is well-understood in areas such as clustering and outlier detection that the choice of the similarity measure has considerable effect on the outcomes. Specifically, Maier et al. demonstrate the critical influence of graph construction on graph-based clustering [18]. Graph-based SSL is no exception. A similar study by de Sousa et al. find that “SSL algorithms are strongly affected by the graph sparsification parameter value and the choice of the adjacency graph construction and weighted matrix generation methods” [6].

Interestingly, however, the (1)st step—graph construction for SSL—is notably under-emphasized in the literature as compared to the (2)nd step—label inference algorithms. Most practitioners default to using a similarity measure such as radial basis function (RBF), coupled with sparsification by  $\epsilon$ -neighborhood (where node pairs only within distance  $\epsilon$  are connected) or  $k$ NN (where each node is connected to its  $k$  nearest neighbors). Hyper-parameters, such as RBF bandwidth  $\sigma$  and  $\epsilon$  (or  $k$ ), are then selected by grid search based on cross-validation error.

There exist some work on graph construction for SSL beyond  $\epsilon$ - and  $k$ NN-graphs, which we review in §2. Roughly, related work can be split into unsupervised and supervised techniques. All of them suffer from one or more drawbacks in terms of efficient search, scalability, and graph quality for the given SSL task. More specifically, unsupervised methods do not leverage the available labeled data for learning the graph. On the supervised side, most methods are not task-driven, that is, they do not take into account the given SSL task to evaluate graph quality and guide the graph construction, or do not effectively scale to high dimensional data in terms of

both runtime and memory. Most importantly, the graph learning problem is typically non-convex and comes with a prohibitively large search space that should be explored strategically, which is not addressed by existing work.

In this work, we address the problem of graph (structure) learning for SSL, suitable and scalable to high dimensional problems. We set out to perform the graph construction and label inference steps of semi-supervised learning *simultaneously*. To this end, we learn different RBF bandwidths  $\sigma_{1:d}$  for each dimension, by adaptively minimizing a function of validation loss using (iterative) gradient descent. In essence, these different bandwidths become model hyperparameters that provide a more general edge weighting function, which in turn can more flexibly capture the underlying data manifold. Moreover, it is a form of feature selection/importance learning that becomes essential in high dimensions with noisy features. On the other hand, this introduces a scale problem for high-dimensional datasets as we discussed earlier, that is, a large search space with numerous hyperparameters to tune.

Our solution to the scale problem is a Parallel Graph Learning algorithm, called PG-LEARN, which is a hybrid of random search and adaptive search. It is motivated by the successive halving strategy [15], which has been recently proposed for efficient hyperparameter optimization for *iterative* machine learning algorithms.

Our main contributions are summarized as follows.

- **Graph learning for SSL:** We propose an efficient and effective gradient-based graph (structure) learning algorithm, called PG-LEARN (for Parallel Graph Learning), which *jointly* optimizes both steps of graph-based SSL: graph construction and label inference.
- **Parallel graph search with adaptive strategy:** In high dimensions, it becomes critical to effectively explore the (large) search space. To this end, we couple our (1) *iterative/sequential* gradient-based *local* search with (2) a *parallel*, resource-adaptive, *random* search scheme. In this hybrid, the gradient search runs in parallel with different random initializations, the relatively unpromising fraction of which is terminated early to allocate the time for other initializations in the search space. In effect, (2) empowers (1) to explore the search space more efficiently.
- **Efficiency and scalability:** We use tensor-form gradient (which is more compact and efficient), and make full use of the sparsity of  $k$ NN graph to reduce runtime and memory requirements. Overall, PG-LEARN scales linearly in dimensionality  $d$  and log-linearly in number of samples  $n$  computationally, while memory complexity is linear in both  $d$  and  $n$ .

Experiments on multi-class classification tasks show that the proposed PG-LEARN significantly outperforms a variety of existing graph construction schemes in terms of test accuracy per fixed time budget for hyperparameter search, and further tackles high dimensional, noisy problems more effectively. **Reproducibility:** The source code can be found at project page <https://pg-learn.github.io/>. All datasets used in experiments are publicly available (See §5.1).

## 2 RELATED WORK

Among existing work, a group of graph construction methods are unsupervised, which do not leverage any information from the labeled data. The most typical ones include similarity-based methods

such as  $\epsilon$ -neighborhood graphs,  $k$  nearest neighbor ( $k$ NN) graphs and mutual variants. Jebara et al. introduced the b-matching method [13] toward a balanced graph in which all nodes have the same degree. There are also self-representation based approaches, like locally linear embedding (LLE) [21], low-rank representation (LRR) [17], and variants [3, 5, 25], which model each instance to be a weighted linear combination of other instances where nodes with non-zero coefficients are connected. Karasuyama and Mamitsuka [14] extend the LLE idea by restricting the regression coefficients (i.e., edge weights) to be derived from Gaussian kernels that forces the weights to be positive and greatly reduces the number of free parameters. Zhu et al. [27] proposed to learn different  $\sigma_d$  hyperparameters per dimension for the Gaussian kernel by minimizing the entropy of the solution on unlabeled instances via gradient descent. Wang et al. [22] focused on the scalability of graph construction by improving Anchor Graph Regularization algorithms, which transform the similarity among samples into similarity between samples and anchor points.

A second group of graph construction methods are supervised and make use of the labeled data in their optimization. Dhillon et al. [8] proposed a distance metric learning approach within a self-learning scheme to learn the similarity function. However, metric learning uses expensive SDP solvers that do not scale to very large dimensions. Rohban and Rabiee [20] proposed a supervised graph construction approach, showing that under certain manifold sampling rates, the optimal neighborhood graph is a subgraph of the  $k$ NN graph family. Similar to [27], Zhang and Lee [23] also tune  $\sigma_d$ 's for different dimensions using a gradient based method, where they minimize the leave-one-out prediction error on labeled data points. Their loss function, however, is specific to the binary classification problems. Li et al. [16] proposed a semi-supervised SVM formulation to derive a robust and non-deteriorated SSL by combining multiple graphs together, and it can be used to judge the quality of graphs. Zhuang et al. [28] incorporated labeling information to graph construction period for self-representation based approach by explicitly enforcing sample can only be represented by samples from the same class.

The above approaches to graph construction have a variety of drawbacks; and typically lack one or more of efficiency, scalability, and graph quality for the given SSL task. Specifically, Zhu et al.'s MinEnt [27] only maximizes confidence over unlabeled samples without using any label information; b-matching method [13] only creates a balanced sparse graph which is not a graph *learning* algorithm; self-representation based methods [3, 14, 17, 21, 25] assume each instance to be a weighted linear combination of other data points and connect those with non-zero coefficients, however such a graph is not necessarily suitable nor optimized specifically for the given SSL task; Anchor Graph Regularization [22] only stresses on scalability without considering the graph learning aspect; and several other graph learning algorithms connected with the SSL task [7, 23] are not scalable in both runtime and memory.

Our work differs from all existing graph construction algorithms in the following aspects: (1) PG-LEARN is a gradient-based task-driven graph learning method, which aims to find an optimized graph (evaluated over validation set) for a specific graph-based SSL task; (2) PG-LEARN achieves scalability over both dimensionality  $d$  and sample size  $n$  in terms of runtime and memory. Specifically, it

has  $O(nd)$  memory complexity and  $O(nd + n \log n)$  computational complexity for each gradient update. (3) Graph learning problem typically has a very large search space with a non-convex optimization objective, where initialization becomes extremely important. To this end, we design an efficient adaptive search framework outside the core of graph learning. This is not explicitly addressed by those prior work, whereas it is one of the key issues we focus on through the ideas of relative performance and early-termination.

### 3 PRELIMINARIES AND BACKGROUND

#### 3.1 Notation

Consider  $\mathcal{D} := \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_l, y_l), \mathbf{x}_{l+1}, \dots, \mathbf{x}_{l+u}\}$ , a data sample in which the first  $l$  examples are labeled, i.e.,  $\mathbf{x}_i \in \mathbb{R}^d$  has label  $y_i \in \mathbb{N}_c$  where  $c$  is the number of classes and  $\mathbb{N}_c := \{p \in \mathbb{N}^* | 1 \leq p \leq c\}$ . Let  $u := n - l$  be the number of unlabeled examples and  $Y \in \mathbb{B}^{n \times c}$  be a binary label matrix in which  $Y_{ij} = 1$  if and only if  $\mathbf{x}_i$  has label  $y_i = j$ .

The semi-supervised learning task is to assign labels  $\{y_{l+1}, \dots, y_{l+u}\}$  to the unlabeled instances.

#### 3.2 Graph Construction

A preliminary step to graph-based semi-supervised learning is the construction of a graph from the point-cloud data. The graph construction process generates a graph  $\mathcal{G}$  from  $\mathcal{D}$  in which each  $\mathbf{x}_i$  is a node of  $\mathcal{G}$ . To generate a weighted matrix  $W \in \mathbb{R}^{n \times n}$  from  $\mathcal{G}$ , one uses a similarity function  $\mathcal{K} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$  to compute the weights  $W_{ij} = \mathcal{K}(\mathbf{x}_i, \mathbf{x}_j)$ .

A widely used similarity function is the RBF (or Gaussian) kernel,  $\mathcal{K}(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|/(2\sigma^2))$ , in which  $\sigma \in \mathbb{R}_+^*$  is the kernel bandwidth parameter.

To sparsify the graph, two techniques are used most often. In  $\epsilon$ -neighborhood ( $\epsilon$ N) graphs, there exists an undirected edge between  $\mathbf{x}_i$  and  $\mathbf{x}_j$  if and only if  $\mathcal{K}(\mathbf{x}_i, \mathbf{x}_j) \geq \epsilon$ , where  $\epsilon \in \mathbb{R}_+^*$  is a free parameter.  $\epsilon$  thresholding is prone to generating disconnected or almost-complete graphs for an improper value of  $\epsilon$ . On the other hand, in the  $k$  nearest neighbors ( $k$ NN) approach, there exists an undirected edge between  $\mathbf{x}_i$  and  $\mathbf{x}_j$  if either  $\mathbf{x}_i$  or  $\mathbf{x}_j$  is one of the  $k$  closest examples to the other.  $k$ NN approach has the advantage of being robust to choosing an inappropriate fixed threshold.

In this work, we use a general kernel function that enables a more flexible graph family, in particular

$$\mathcal{K}(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\sum_{m=1}^d \frac{(\mathbf{x}_{im} - \mathbf{x}_{jm})^2}{\sigma_m^2}\right), \quad (1)$$

where  $\mathbf{x}_{im}$  is the  $m^{\text{th}}$  component of  $\mathbf{x}_i$ . We denote  $W_{ij} = \exp\left(-(\mathbf{x}_i - \mathbf{x}_j)^T A (\mathbf{x}_i - \mathbf{x}_j)\right)$ , where  $A := \text{diag}(\mathbf{a})$  is a diagonal matrix with  $A_{mm} = a_m = 1/\sigma_m^2$ , that corresponds to a metric in which different dimensions/features are given different “weights”, which allows a form of feature selection.<sup>1</sup> In addition, we employ  $k$ NN graph construction for sparsity.

<sup>1</sup>Setting  $A$  equal to (i) the identity, (ii) the (diagonal) variance, or (iii) the covariance matrix would compute similarity based on Euclidean, normalized Euclidean, or Mahalanobis distance, respectively.

Our goal is to *learn* both  $k$  as well as all the  $a_m$ 's, by means of which we aim to construct a graph that is suitable for the semi-supervised learning task at hand.

#### 3.3 Graph based Semi-Supervised Learning

Given the constructed graph  $\mathcal{G}$ , a graph-based SSL algorithm uses  $W$  and the label matrix  $Y$  to generate output matrix  $F$  by label diffusion in the weighted graph. Note that this paper focuses on the multi-class classification problem, hence  $F \in \mathbb{R}^{n \times c}$ .

There exist a number of SSL algorithms with various objectives. Perhaps the most widely used ones include the Gaussian Random Fields algorithm by Zhu et al. [27], Laplacian Support Vector Machine algorithm by Belkin et al. [1], and Local and Global Consistency (LGC) algorithm by Zhou et al. [26].

The topic of this paper is how to effectively learn the hyperparameters of graph construction. Therefore, we focus on how the performance of a given recognized SSL algorithm can be improved by means of learning the graph, rather than comparing the performance of different semi-supervised or supervised learning algorithms. To this end, we use the LGC algorithm [26] which we briefly review here. It is easy to follow the same way to generalize the graph learning ideas introduced in this paper for other popular SSL algorithms, such as Zhu et al.'s [27] and Belkin et al.'s [1] that have similar objectives to LGC, which we do not pursue further.

The LGC algorithm solves the optimization problem

$$\arg \min_{F \in \mathbb{R}^{n \times c}} \text{tr}((F - Y)^T (F - Y) + \alpha F^T L F), \quad (2)$$

where  $\text{tr}()$  denotes matrix trace,  $L := I_n - P$  is the normalized graph Laplacian, such that  $I_n$  is the  $n$ -by- $n$  identity matrix,  $P = D^{-1/2} W D^{-1/2}$ ,  $D := \text{diag}(W \mathbf{1}_n)$  and  $\mathbf{1}_n$  is the  $n$ -dimensional all-1's vector. Taking the derivative w.r.t.  $F$  and reorganizing the terms, we would get the closed-form solution  $F = (I_n + \alpha L)^{-1} Y$ .

The solution can also be found without explicitly taking any matrix inverse and instead using the power method [11], as

$$\begin{aligned} (I + \alpha L)F = Y &\Rightarrow F + \alpha F = \alpha P F + Y \Rightarrow F = \frac{\alpha}{1 + \alpha} P F + \frac{1}{1 + \alpha} Y \\ &\Rightarrow F^{(t+1)} \leftarrow \mu P F^{(t)} + (1 - \mu) Y. \end{aligned} \quad (3)$$

#### 3.4 Problem Statement

We address the problem of graph (structure) learning for SSL. Our goal is to estimate, for a given task, suitable hyperparameters within a flexible graph family. In particular, we aim to infer

- $A$ , containing the bandwidths (or weights)  $a_m$ 's for different dimensions in Eq. (1), as well as
- $k$ , for sparse  $k$ NN graph construction;

so as to better align the graph structure with the underlying (hidden) data manifold and the given SSL task.

### 4 PROPOSED METHOD: PG-LEARN

In this section, we present the formulation and efficient computation of our graph learning algorithm PG-LEARN, for Parallel Graph Learning for SSL.

In essence, the feature weights  $a_m$ 's and  $k$  are the model parameters that govern how the algorithm's performance generalizes to unlabeled data. Typical model selection approaches include random

search or grid search to find a configuration of the hyperparameters that yield the best cross-validation performance.

Unfortunately, the search space becomes prohibitively large for high-dimensional datasets that could render such methods futile. In such cases, one could instead carefully select the configurations in an adaptive manner. The general idea is to impose a smooth loss function  $g(\cdot)$  on the validation set over which  $A$  can be estimated using a gradient based method.

We present the main steps of our algorithm for adaptive hyperparameter search in Algorithm 1.

---

**Algorithm 1** GRADIENT (for Adaptive Hyperparameter Search)

---

- 1: Initialize  $k$  and  $\mathbf{a}$  (vector containing  $a_m$ 's);  $t := 0$
  - 2: **repeat**
  - 3:   Compute  $F^{(t)}$  using  $k$ NN graph on current  $a_m$ 's by (3)
  - 4:   Compute gradient  $\frac{\partial g}{\partial a_m}$  based on  $F^{(t)}$  by (5) for each  $a_m$
  - 5:   Update  $a_m$ 's by  $\mathbf{a}^{(t+1)} := \mathbf{a}^{(t)} - \gamma \frac{dg}{da}$ ;  $t := t + 1$
  - 6: **until**  $a_m$ 's have converged
- 

The initialization in step 1 can be done using some heuristics, although the most prevalent and easiest approach is a random guess. Given a fixed initial (random) configuration, we essentially perform an adaptive search that strives to find a better configuration around it, guided by the validation loss  $g(\cdot)$ . In Section 4.1, we introduce the specific function  $g(\cdot)$  that we use and how to compute its gradient.

While the gradient based optimization is likely to find a better configuration than where it started, the final performance of the SSL algorithm depends considerably on the initialization. Provided that the search space is quite large for high dimensional datasets, it is of paramount importance to try different random initializations in step 1, in other words, to run Algorithm 1 several times. As such, *the GRADIENT algorithm can be seen as an adaptive local search*, where we start at a random configuration and adaptively search in the vicinity for a better one.

As we discuss in Section 4.1, the gradient based updates are computationally demanding. This makes naively running Algorithm 1 several times expensive. There are however two properties that we can take considerable advantage of: (1) both the SSL algorithm (using the power method) as well as the gradient optimization are *iterative, any-time* algorithms (i.e., they can return an answer at any time that they are probed), and (2) different initializations can be run independently in *parallel*.

In particular, our search strategy is inspired by a general framework of parallel hyperparameter search designed for *iterative* machine learning algorithms that has been recently proposed by Jamieson and Talwalkar [12] and a follow-up by Li et al. [15]. This framework perfectly suits our SSL setting for the reasons (1) and (2) above. The idea is to start multiple (random) configurations in parallel threads, run them for a bounded amount of time, probe for their solutions, throw out the worst half (or some other pre-specified fraction), and repeat until one configurations remains. By this strategy of early termination, that is by quitting poor initializations early without running them to completion, the compute resources are effectively allocated to promising hyperparameter configurations. Beyond what has been proposed in [12], we start

new initializations on the idle threads whose jobs have been terminated in order to fully utilize the parallel threads. We describe the details of our parallel search in Section 4.2.

#### 4.1 Validation Loss $g(\cdot)$ & Gradient Updates

We base the learning of the hyperparameters of our kernel function ( $a_m$ 's in Eq. (1)) on minimizing some loss criterion on validation data. Let  $\mathcal{L} \subset \mathcal{D}$  denote the set of  $l$  labeled examples, and  $\mathcal{V} \subset \mathcal{L}$  a subset of the labeled examples designated as validation samples. A simple choice for the validation loss would be the labeling error, written as  $g_A(\mathcal{V}) = \sum_{v \in \mathcal{V}} (1 - F_{vc_v})$ , where  $c_v$  denotes the true class index for a validation instance  $v$ . Other possible choices for each  $v$  include  $-\log F_{vc_v}$ ,  $(1 - F_{vc_v})^x$ ,  $x^{-F_{vc_v}}$ , with  $x > 1$ .

In semi-supervised learning the labeled set is often small. This means the number of validation examples is also limited. To squeeze the most out of the validation set, we propose to use a *pairwise* learning-to-rank objective:

$$g_A(\mathcal{V}) = \sum_{c'=1}^c \sum_{\substack{(v,v'): v \in \mathcal{V}_{c'} \\ v' \in \mathcal{V} \setminus \mathcal{V}_{c'}}} -\log \sigma(F_{vc'} - F_{v'c'}) \quad (4)$$

where  $\mathcal{V}_{c'}$  denotes the validation nodes whose true class index is  $c'$  and  $\sigma(x) = \frac{\exp(x)}{1+\exp(x)}$  is the sigmoid function. The larger the difference  $(F_{vc'} - F_{v'c'})$ , or intuitively the more *confidently* the solution  $F$  ranks validation examples of class  $c'$  above other validation examples not in class  $c'$ , the better it is; since then  $\sigma(\cdot)$  would approach 1 and the loss to zero.

In short, we aim to find the hyperparameters  $A$  that minimize the total negative log likelihood of ordered validation pairs. The optimization is conducted by gradient descent. The gradient is computed as

$$\begin{aligned} \frac{\partial g}{\partial a_m} &= \frac{\partial \left( \sum_{c'=1}^c \sum_{(v,v'): v \in \mathcal{V}_{c'}, v' \in \mathcal{V} \setminus \mathcal{V}_{c'}} -F_{vv'} + \log(1 + \exp(F_{vv'})) \right)}{\partial a_m} \\ &= \sum_{c'=1}^c \sum_{(v,v'): v \in \mathcal{V}_{c'}, v' \in \mathcal{V} \setminus \mathcal{V}_{c'}} (o_{vv'} - 1) \left( \frac{\partial F_{vc'}}{\partial a_m} - \frac{\partial F_{v'c'}}{\partial a_m} \right) \end{aligned} \quad (5)$$

where we denote by  $F_{vv'} = (F_{vc'} - F_{v'c'})$  and  $o_{vv'} = \sigma(F_{vv'})$ .

The values  $\frac{\partial F_{vc'}}{\partial a_m}$  and  $\frac{\partial F_{v'c'}}{\partial a_m}$  for each class  $c'$  and  $v, v' \in \mathcal{V}$  can be read off of matrix  $\frac{\partial F}{\partial a_m}$ , which is given as

$$\frac{\partial F}{\partial a_m} = -(\mathbf{I}_n + \alpha L)^{-1} \frac{\partial (\mathbf{I}_n + \alpha L)}{\partial a_m} F = \alpha (\mathbf{I}_n + \alpha L)^{-1} \frac{\partial P}{\partial a_m} F, \quad (6)$$

using the equivalence  $dX^{-1} = -X^{-1}(dX)X^{-1}$ . Recall that  $P = D^{-1/2} W D^{-1/2}$  with  $P_{ij} = \frac{W_{ij}}{\sqrt{d_i d_j}}$ ;  $d_i$  being node  $i$ 's degree in  $\mathcal{G}$ .

We can then write the following:

$$\frac{\partial P_{ij}}{\partial a_m} = \frac{\partial W_{ij}}{\partial a_m} \frac{1}{\sqrt{d_i d_j}} - \frac{W_{ij}}{2} (d_i d_j)^{-3/2} \frac{\partial d_i d_j}{\partial a_m} \quad (7)$$

$$= \frac{\partial W_{ij}}{\partial a_m} \frac{P_{ij}}{W_{ij}} - \frac{W_{ij}}{2} \left( \frac{P_{ij}}{W_{ij}} \right)^3 \left( d_j \frac{\partial d_i}{\partial a_m} + d_i \frac{\partial d_j}{\partial a_m} \right) \quad (8)$$

$$= \frac{\partial W_{ij}}{\partial a_m} \frac{P_{ij}}{W_{ij}} - \frac{W_{ij}}{2} \left( \frac{P_{ij}}{W_{ij}} \right)^3 \left( \sum_n W_{in} \cdot \sum_n \frac{\partial W_{jn}}{\partial a_m} + \sum_n W_{jn} \cdot \sum_n \frac{\partial W_{in}}{\partial a_m} \right) \quad (9)$$

**4.1.1 Matrix-form gradient.** We can rewrite all element-wise gradients into a combined matrix-form gradient. The matrix-form is compact and can be computed more efficiently on platforms optimized for matrix operations (e.g., Matlab).

The matrix-form uses 3-d matrix (or tensor) representation. In the following, we use  $\odot$  to denote element-wise multiplication,  $\oslash$  element-wise division, and  $\otimes$  for element-wise power. In addition,  $\cdot$  denotes regular matrix dot product. For multiplication and division, a 3-d matrix should be viewed as a 2-d matrix with vector elements.

First we extend the derivative w.r.t.  $a_m$  in Eq. (9) into derivative w.r.t.  $\mathbf{a}$ :

$$\frac{\partial P_{ij}}{\partial \mathbf{a}} = \frac{\partial W_{ij}}{\partial \mathbf{a}} \frac{P_{ij}}{W_{ij}} - \frac{W_{ij}}{2} \left( \frac{P_{ij}}{W_{ij}} \right)^3 \left( \sum_n W_{in} \cdot \sum_n \frac{\partial W_{jn}}{\partial \mathbf{a}} + \sum_n W_{jn} \cdot \sum_n \frac{\partial W_{in}}{\partial \mathbf{a}} \right) \quad (10)$$

To write this equation concisely, let tensor  $\Omega$  be  $\frac{\partial W}{\partial \mathbf{a}}$ , a 2d-matrix with vector elements  $\Omega_{ij} = \frac{\partial W_{ij}}{\partial \mathbf{a}}$ , and let tensor  $\Delta X$  be the one with vector elements  $\Delta X_{ij} = (\mathbf{x}_i - \mathbf{x}_j)^2$ .

Then we can rewrite some equations using the above notation:

$$\sum_n W_{in} = (\mathbf{W} \cdot \mathbf{1}_n)_i \quad (11)$$

$$\sum_n \frac{\partial W_{jn}}{\partial \mathbf{a}} = (\Omega \cdot \mathbf{1}_n)_j \quad (12)$$

$$\sum_n W_{in} \cdot \sum_n \frac{\partial W_{jn}}{\partial \mathbf{a}} = (\mathbf{W} \cdot \mathbf{1}_n \cdot (\Omega \cdot \mathbf{1}_n)^T)_i \quad (13)$$

Now we can rewrite element-wise gradients in (10) into one matrix-form gradient:

$$\frac{dP}{d\mathbf{a}} = \Omega \odot (P \oslash W) - \frac{1}{2} P^{\otimes 3} \oslash W^{\otimes 2} \odot (\mathbf{W} \cdot \mathbf{1}_n \cdot (\Omega \cdot \mathbf{1}_n)^T + (\mathbf{W} \cdot \mathbf{1}_n \cdot (\Omega \cdot \mathbf{1}_n)^T)^T) \quad (14)$$

The only thing left is the computation of  $\Omega = \frac{\partial W}{\partial \mathbf{a}}$ . Notice that

$$\begin{aligned} \frac{\partial W_{ij}}{\partial a_m} &= \frac{\partial \exp(-\sum_{m=1}^d a_m (\mathbf{x}_{id} - \mathbf{x}_{jd})^2)}{\partial a_m} = -W_{ij} (\mathbf{x}_{id} - \mathbf{x}_{jd})^2 \\ &\Rightarrow \frac{\partial W_{ij}}{\partial \mathbf{a}} = -W_{ij} (\mathbf{x}_i - \mathbf{x}_j)^2 = -W_{ij} \Delta X_{ij} \\ &\Rightarrow \frac{dW}{d\mathbf{a}} = -W \odot \Delta X = \Omega \end{aligned} \quad (15)$$

All in all, we transform the element-wise gradients  $\frac{\partial P_{ij}}{\partial a_m}$  as given in Eq. (9) to compact tensor-form updates  $\frac{dP}{d\mathbf{a}}$  as in Eq. (14). The tensor-form gradient updates not only provide speed up, but also

can be expanded to make full use of the  $k$ NN graph sparsity. In particular,  $W$  is a  $k$ NN-sparse matrix with  $O(kn)$  non-zero elements. First, Eq. (15) for  $\Omega$  shows that we do not need to compute full  $\Delta X$  but only the elements in  $\Delta X$  corresponding to non-zero elements of  $W$ . Similarly, in Eq. (14), matrix  $P$  does not need to be fully computed, and the whole Eq. (14) can be computed sparsely.

**4.1.2 Complexity analysis.** We first analyze computational complexity in terms of two main components: constructing the  $k$ NN graph and computing  $F$  in line 3, and computing the gradient  $\frac{dF}{d\mathbf{a}}$  in line 4 of Algorithm 1 as outlined in this subsection.

Let us denote the number of non-zeros in  $W$ , i.e. the number of edges in the  $k$ NN graph, by  $e = nmz(W)$ . We assume  $kn \leq e \leq 2kn$  remains near-constant as  $\mathbf{a}$  changes over the GRADIENT iterations.

In line 4, we first construct tensor  $\Omega$  as in Eq. (15) in  $O(ed)$ . Computing  $\frac{dP}{d\mathbf{a}}$  as in Eq. (14) also takes  $O(ed)$ . Next, obtaining matrix  $\frac{\partial F}{\partial a_m}$  in Eq. (6) seemingly requires inverting  $(I_n + \alpha L)^{-1}$ . However, we  $\text{ifijcan}$  rewrite Eq. (6) as

$$(I_n + \alpha I_n - \alpha P) \frac{\partial F}{\partial a_m} = \alpha \frac{\partial P}{\partial a_m} F \Rightarrow \frac{\partial F}{\partial a_m} = \alpha (P - I_n) \frac{\partial F}{\partial a_m} + \alpha \frac{\partial P}{\partial a_m} F$$

which can be solved via the power method that takes  $t$  iterations in  $O(ect)$ . Computing  $\frac{\partial F}{\partial a_m}$  and plugging in Eq. (5) to get  $g(\cdot)$ 's gradient for all  $a_m$ 's then takes  $O(ectd)$ , or equivalently  $O(knctd)$ .

In line 3, updated  $a_m$ 's are used for weighted node similarities to compute  $k$ NNs for each instance. Nearest neighbor computation for all instances is inherently quadratic, which however can be sped up by approximation algorithms and data structures such as locality-sensitive hashing (LSH) [10]. To this end, we use a fast  $k$ NN graph construction algorithm that takes advantage of LSH and has  $O(n[dk^2 + \log n])$  complexity [24]; only quadratic in the (small)  $k$  but log-linear in  $n$ . Given the  $k$ NN graph,  $F$  can then be computed via (3) in  $O(ect')$  for  $t'$  iterations of the power method.

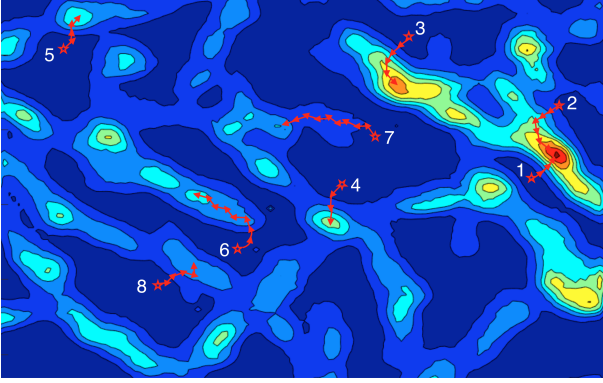
Overall, one iteration of Algo. 1 takes  $O(n[kctd + dk^2 + \log n])$ . Furthermore, if we consider  $k, c, t$  as constants, then the computational complexity can be written as  $O(n[d + \log n])$ .

In addition, memory requirement for each gradient update is  $O(knd)$ . The bottleneck is the construction of tensors  $\Omega$  and  $\Delta X$  with size- $d$  vector elements. As discussed earlier those are constructed sparsely, i.e., only the elements corresponding to non-zero entries of  $W$ , which is  $O(kn)$ , are stored.

## 4.2 Parallel Hyperparameter Search with Adaptive Resource Allocation

For high-dimensional datasets, the search space of hyperparameter configurations is huge. In essence, Algorithm 1 is an adaptive search around a single initial point in this space. As with many gradient-based optimization of high-dimensional non-convex functions with unknown smoothness, its performance depends on the initialization. Therefore, trying different initializations of Algorithm 1 is beneficial to improving performance.

An illustrative example over a 2-d search space is shown in Figure 1 (best in color). In this space most configurations yield poor validation accuracy, as would be the likely case in even higher dimensions. In the figure, eight random configurations are shown (with stars). The sequence of arrows from a configuration can be seen analogous to the iterations of a single run of Algorithm 1.



**Figure 1: (best in color)** The heatmap shows the validation error over an example 2-d search space with red corresponding to areas with lower error. Our approach is an inter-mix of random and adaptive search. We start at various random configurations (stars 1–8) and adaptively improve them (arrows depicting gradient updates), while strategically terminating unpromising ones (like 6, 7, and 8) early.

While it is beneficial to try as many random initializations as possible, especially in high dimensions, adaptive search is slow. A single gradient update by Algorithm 1 takes time in  $O(ectd)$  followed by reconstruction of the  $k$ NN graph. Therefore, it would be good to quit the algorithm early if the ongoing progress is not promising (e.g., poor initializations 6–8 in Figure 1) and simply try a new initialization. This would allow using the time efficiently for going through a larger number of configurations.

One way to realize such a scheme is called successive halving [12], which relies on an early-stopping strategy for iterative machine learning algorithms. The idea is quite simple and follows directly from its name: try out a set of hyperparameter configurations for some fixed amount of time (say in parallel threads), evaluate the performance of all configurations, keep the best half (terminate the worst half of the threads), and repeat until one configuration remains while allocating exponentially increasing amount of time after each round to not-yet-terminated, promising configurations (i.e., threads). Our proposed method is a parallel implementation of their general framework adapted to our problem, and further utilizes the idle threads that have been terminated.

Algorithm 2 gives the steps of our proposed method PG-LEARN, which calls the GRADIENT subroutine in Algorithm 1. Besides the input dataset  $\mathcal{D}$ , PG-LEARN requires three inputs: (1) budget  $B$ ; the maximum number of time units<sup>2</sup> that can be allocated to one thread (i.e., one initial hyperparameter configuration), (2) downsampling rate  $r$ ; an integer that controls the fraction of threads terminated (or equally, configurations discarded) in each round of PG-LEARN, and finally (3)  $T$ ; the number of parallel threads.

Concretely, PG-LEARN performs  $R = \lceil \log_r B \rceil$  rounds of elimination. At each round, the best  $1/r$  fraction of configurations are retained. Eliminations are done in exponentially increasing time intervals, that is, first round occurs at time  $B/r^R$ , second round at  $B/r^{R-1}$ , and so on.

<sup>2</sup>We assume time is represented in units, where a unit is the minimum amount of time one would run the models before comparing against each other.

---

### Algorithm 2 PG-LEARN (for Parallel Hyperparameter Search)

---

**Input:** Dataset  $\mathcal{D}$ , budget  $B$  time units, downsampling rate  $r$  ( $=$  by default), number of parallel threads  $T$

**Output:** Hyperparameter configuration  $(k, \mathbf{a})$

```

1:  $R = \lceil \log_r B \rceil$ ,  $d_1 = Br^{-R}$ 
2:  $C := \text{get\_hyperparameter\_configuration}(T)$ 
3:  $L := \{\text{run\_GRADIENT\_then\_return\_val\_loss}(c, d_1) : c \in C\}$ 
4: for  $i \in \{1, \dots, R\}$  do
5:    $C_{top} := \text{get\_top}(C, L, \lfloor T/r \rfloor)$ 
6:    $C_{new} := \text{get\_hyperparameter\_configuration}(T - \lfloor T/r \rfloor)$ 
7:    $d_i = B(r^{-(R-i)} - r^{-(R-i+1)})$ 
8:    $L_{top} := \{\text{resume\_GRADIENT\_then\_return\_val\_loss}(c, d_i)$ 
   for  $c \in C_{top}\}$ 
9:    $L_{new} := \{\text{run\_GRADIENT\_then\_return\_val\_loss}(c, d_i)$ 
   for  $c \in C_{new}\}$ 
10:   $C := C_{top} \cup C_{new}$ ,  $L := L_{top} \cup L_{new}$ 
11: end for
12: return  $c_{top} := \text{get\_top}(C, L, 1)$ 

```

---

After setting the number of elimination rounds  $R$  and the duration of the first round, denoted  $d_1$  (line 1), PG-LEARN starts by obtaining  $T$  initial hyperparameter configurations (line 2). Note that a configuration is a  $(k, \mathbf{a}_{1:d})$  pair. Our implementation<sup>3</sup> of PG-LEARN is parallel. As such, each thread draws their own configuration; uniformly at random. Then, each thread runs the GRADIENT (Algorithm 1) with their corresponding configuration for duration  $d_1$  and returns the validation loss (line 3).


At that point, PG-LEARN enters the rounds of elimination (line 4).  $L$  validation loss values across threads are gathered at the master node, which identifies the top  $\lfloor T/r \rfloor$  configurations  $C_{top}$  (or threads) with the lowest loss (line 5). The master then terminates the runs on the remaining threads and restarts them afresh with new configurations  $C_{new}$  (line 6). The second round is to run until  $B/r^{R-1}$ , or for  $B/r^{R-1} - B/r^R$  in duration. After the  $i$ th elimination, in general, we run the threads for duration  $d_i$  as given in line 7—notice that exponentially increasing amount of time is provided to “surviving” configurations over time. In particular, the threads with the promising configurations in  $C_{top}$  are *resumed* their runs from where they are left off with the GRADIENT iterations (line 8). The remaining threads start with the GRADIENT iterations using their new initialization (line 9). Together, this ensures full utilization of the threads at all times. Eliminations continue for  $R$  rounds, following the same procedure of resuming best threads and restarting from the rest of the threads (lines 4–11). After round  $R$ , all threads run until time budget  $B$  at which point the (single) configuration with the lowest validation error is returned (line 12).

The underlying principle of PG-LEARN exploits the insight that a hyperparameter configuration which is destined to yield good performance ultimately is more likely than not to perform in the top fraction of configurations even after a small number of iterations. In essence, even if the performance of a configuration after a small number of iterations of the GRADIENT (Algorithm 1) may not be representative of its ultimate performance after a large number of iterations in *absolute* terms, its *relative* performance in comparison to the alternatives is roughly maintained. We note that different

<sup>3</sup>We release all source code at <https://github.com/LingxiaoShawn/Pg-Learn>

configurations get to run different amounts of time before being tested against others for the first time (depending on the round they get introduced). This diversity offers some robustness against variable convergence rates of the  $g(\cdot)$  function at different random starting points.


**Example:** In Figure 2 we provide a simple example to illustrate PG-LEARN’s execution, using  $T = 8$  parallel threads, downsampling rate  $r = 2$  (equiv. to halving), and  $B = 16$  time units of processing budget. There are  $\lfloor \log_2 16 \rfloor = 4$  rounds of elimination at  $t = 1, 2, 4, 8$  respectively, with the final selection being made at  $t = B$ . It starts with 8 different initial configurations (depicted with circles) in parallel threads. At each round, bottom half (=4) of the threads with highest validation loss are terminated with their iterations of Algorithm 1 and restart running Algorithm 1 with a new initialization (depicted with a crossed-circle). Overall,  $T + (1 - 1/r)T \lfloor \log_r B \rfloor = 8 + 4 \lfloor \log_2 16 \rfloor = 24$  configurations are examined—a larger number as compared to the initial 8, thanks to the early-stopping and adaptive resource allocation strategy.



**Figure 2: Example execution of PG-LEARN with  $T = 8$  parallel threads, downsampling rate  $r = 2$ , and budget  $B = 16$  time units. At each “check point” in time (dashed vertical lines), (worst) half of the runs are discarded and corresponding threads restart Algorithm 1 with new random configurations of  $(k, \mathbf{a}_{1:d})$ . At the end, hyperparameters that yield the lowest  $g(\cdot)$  function value (i.e. validation loss) across all threads are returned (those by thread 4 in this example).**

Next in Figure 3 we show example runs on two different real-world datasets, depicting the progression of validation (blue) and test (red) accuracy over time, using  $T = 32$ ,  $r = 2$ ,  $B = 64$ ;  $\approx 15$  sec. unit-time. Thin curves depict those for individual threads. Notice the new initializations starting at different rounds, which progressively improve their validation accuracy over gradient updates (test acc. closely follows). Bold curves depict the overall-best validation accuracy (and corresponding test acc.) across all threads over time.

**Setting  $T$ ,  $B$ , and  $r$ :** Before we conclude the description of our proposed method, we briefly discuss the choices for its inputs. Number of threads  $T$  is a resource-driven input. Depending on the platform being utilized—single machine or a parallel architecture like Hadoop or Spark—PG-LEARN can be executed with as many parallel threads as physically available to the practitioner. Time units  $B$  should be chosen based on the upper bound of practically available time. For example, if one has time to run hyperparameter tuning for at most 3 hours and the minimum amount of time that is meaningful to execute gradient search of configurations before comparing them (i.e., unit time) is 5 minutes, then  $B$  becomes  $180/5 = 36$  units. Finally,  $r$  can be seen as a knob for greediness. A larger value of



**Figure 3: PG-LEARN’s val. (blue) and corresponding test (red) acc. vs. time on COIL (left) and MNIST (right) (see Table 1).**

$r$  corresponds to more aggressive elimination with fewer rounds; specifically, each round terminates  $T(r - 1)/r$  configurations for a total of  $\lfloor \log_r B \rfloor$  rounds. All in all,  $T$  and  $B$  are set based on practical *resource constraints*, physical and temporal, respectively. On the other hand,  $r$  can be set to a small integer, like 2 or 3, without results being very sensitive to the choice.

## 5 EVALUATION

### 5.1 Datasets and Baselines

**Datasets:** We use the publicly available multi-class classification datasets listed in Table 1. COIL<sup>4</sup> (Columbia Object Image Library) contains images of different objects taken from various angles. Features are shuffled and downsampled pixel values from the red channel. USPS<sup>5</sup> is a standard dataset for handwritten digit recognition, with numeric pixel values scanned from the handwritten digits on envelopes from the U.S. Postal Service. MNIST<sup>6</sup> is another popular handwritten digit dataset, containing size-normalized and centered digit images. UMIST<sup>7</sup> face database is made up of images of 20 individuals with mixed race, gender, and appearance. Each individual takes a range of poses, from profile to frontal views. YALE<sup>8</sup> is a subset of the extended Yale Face Database B, which consists of frontal images under different illuminations from 5 individuals.

**Baselines:** We compare the accuracy of PG-LEARN against five baselines that use a variety of schemes, including the strawmen grid search and random guessing strategies, the seminal unsupervised gradient-based graph learning by Zhu et al., a self-representation based graph construction, and a metric learning based scheme. Specifically,

- (1) *Grid search* (GS):  $k$ -NN graph with RBF kernel where  $k$  and bandwidth  $\sigma$  are chosen via grid search,
- (2) *Rand<sub>d</sub> search* (RS):  $k$ -NN with RBF kernel where  $k$  and different bandwidths  $\mathbf{a}_{1:d}$  are randomly chosen,
- (3) *MinEnt*: Minimum Entropy based tuning of  $\mathbf{a}_{1:d}$ ’s as proposed by Zhu et al. [27] (generalized to multi-class),
- (4) *AEW*: Adaptive Edge Weighting by Karasuyama et al. [14] that estimates  $\mathbf{a}_{1:d}$ ’s through local linear reconstruction, and
- (5) *IDML*: Iterative self-learning scheme combined with distance metric learning by Dhillon et al. [8].

<sup>4</sup><http://olivier.chapelle.cc/ssl-book/index.html>, see ‘benchmark datasets’

<sup>5</sup><http://www.cs.huji.ac.il/~shais/datasets/ClassificationDatasets.html>

<sup>6</sup><http://yann.lecun.com/exdb/mnist/>

<sup>7</sup><https://www.sheffield.ac.uk/eee/research/iel/research/face>

<sup>8</sup><http://www.cad.zju.edu.cn/home/dengcai/Data/FaceData.html>

**Table 1: Summary of (multi-class) datasets used in this work.**

Name	#pts $n$	#dim $d$	#cls $c$	description
COIL	1500	241	6	objects with various shapes
USPS	1000	256	10	handwritten digits
MNIST	1000	784	10	handwritten digits
UMIST	575	644	20	faces (diff. race/gender/etc.)
YALE	320	1024	5	faces (diff. illuminations)

Note that *Grid* and *Rand<sub>d</sub>* are standard techniques employed by practitioners most typically. *MinEnt* is perhaps the first graph-learning strategy for SSL which was proposed as part of the Gaussian Random Fields SSL algorithm. It estimates hyperparameters by minimizing the entropy of the solution on unlabeled instances via gradient updates. *IDML* uses and iteratively enlarges the labeled data (via self-learning) to estimate the metric  $A$ ; which we restrict to a diagonal matrix, as our datasets are high dimensional and metric learning is prohibitively expensive for a full matrix. We generalized these baselines to multi-class and implemented them ourselves. We open-source (1)–(4) along with our PG-LEARN implementation.<sup>3</sup> Finally, *AEW* is one of the most recent techniques on graph learning, which extends the LLE [21] idea by restricting the regression coefficients (i.e., edge weights) to be derived from Gaussian kernels. We use their publicly-available implementation.<sup>9</sup>

## 5.2 Empirical Results

**5.2.1 Single-thread Experiments.** We first evaluate the proposed PG-LEARN against the baselines on a fair ground using a single thread, since the baselines do not leverage any parallelism. Single-thread PG-LEARN is simply the GRADIENT as given in Algo. 1.

**Setup:** For each dataset, we sample 10% of the points at random as the labeled set  $\mathcal{L}$ , under the constraint that all classes must be present in  $\mathcal{L}$  and treat the remaining unlabeled data as the test set. For each dataset, 10 versions with randomly drawn labeled sets are created and the average test accuracy across 10 runs is reported. Each run starts with a different random configuration of hyperparameters. For PG-LEARN, *Grid*, and *Rand<sub>d</sub>*, we choose (a small)  $k \in [5, 20]$ .  $\sigma$  for *Grid* and *MinEnt*<sup>10</sup>, and  $a_m$ 's for PG-LEARN, *Rand<sub>d</sub>*, and *AEW* are chosen from  $[0.1\bar{d}, 10\bar{d}]$ , where  $\bar{d}$  is the mean Euclidean distance across all pairs. Other hyperparameters of the baselines, like  $\epsilon$  for *MinEnt* and  $\gamma$  and  $\rho$  for *IDML*, are chosen as in their respective papers. Graph learning is performed for 15 minutes, around which all gradient-based methods have converged.

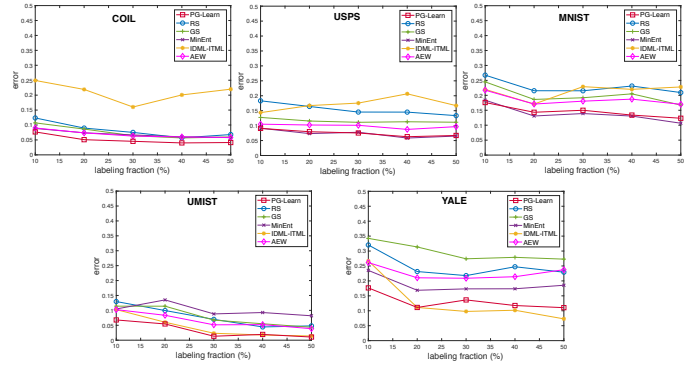
**Results:** Table 2 gives the average test accuracy of the methods on each dataset, avg'd over 10 runs with random labeled sets. PG-LEARN outperforms its competition significantly, according to the paired Wilcoxon signed rank test on a vast majority of the cases—only on the two handwritten digit recognition tasks there is no significant difference between PG-LEARN and *MinEnt*. Not only PG-LEARN is significantly superior to existing methods, its performance is desirably high in absolute terms. It achieves 93% prediction accuracy on the 20-class UMIST, and 82% on the  $2^{10}$ -dimensional YALE dataset.

<sup>9</sup><http://www.bic.kyoto-u.ac.jp/pathway/krsym/software/MSALP/MSALP.zip>

<sup>10</sup>*MinEnt* initializes  $\mathbf{a}$  uniformly, i.e., all  $a_m$ 's are set to the same  $\sigma$  initially [27].

**Table 2: Test accuracy with 10% labeled data, avg'd across 10 random samples; 15 mins of hyperparameter tuning on single thread. Symbols  $\blacktriangle$  ( $p < 0.005$ ) and  $\triangle$  ( $p < 0.01$ ) denote the cases where PG-LEARN is significantly better than the baseline w.r.t. the paired Wilcoxon signed rank test.**

Dataset	PG-LRN	MinEnt	IDML	AEW	Grid	Rand <sub>d</sub>
COIL	<b>0.9232</b>	0.9116 $\blacktriangle$	0.7508 $\blacktriangle$	0.9100 $\blacktriangle$	0.8929 $\blacktriangle$	0.8764 $\blacktriangle$
USPS	<b>0.9066</b>	<b>0.9088</b>	0.8565 $\blacktriangle$	0.8951 $\blacktriangle$	0.8732 $\blacktriangle$	0.8169 $\blacktriangle$
MNIST	<b>0.8241</b>	<b>0.8163</b>	0.7801 $\triangle$	0.7828 $\blacktriangle$	0.7550 $\blacktriangle$	0.7324 $\blacktriangle$
UMIST	<b>0.9321</b>	0.8954 $\blacktriangle$	0.8973 $\triangle$	0.8975 $\blacktriangle$	0.8859 $\blacktriangle$	0.8704 $\blacktriangle$
YALE	<b>0.8234</b>	0.7648 $\triangle$	0.7331 $\blacktriangle$	0.7386 $\blacktriangle$	0.6576 $\blacktriangle$	0.6797 $\blacktriangle$



**Figure 4: Test error (avg'd across 3 random samples) as labeled data percentage is increased up to 50%. PG-LEARN performs the best in many cases, and consistently ranks in top two among competitors on each dataset and each labeling %.**

Next we investigate how the prediction performance of the competing methods changes by varying labeling percentage. To this end, we repeat the experiments using up to 50% labeled data. As shown in Figure 4, test error tends to drop with increasing amount of labels as expected. PG-LEARN achieves the lowest error in many cases across datasets and labeling ratios. *MinEnt* is the closest competition on USPS and MNIST, which however ranks lower on UMIST and YALE. Similarly, *IDML* is close competition on UMIST and YALE, which however performs poorly on COIL and USPS. In contrast, PG-LEARN consistently performs near the top.

We quantify the above more concretely, and provide the test accuracy for each labeling % in Table 3, averaged across random samples from all datasets, along with results of significance tests. We also give the average rank per method, as ranked by test error (hence, lower is better).

PG-LEARN significantly outperforms *all* competing methods in accuracy at *all* labeling ratios w.r.t. the paired Wilcoxon signed rank test at  $p = 0.01$ , as well as achieves the lowest rank w.r.t. test error. On average, *MinEnt* is the closest competition, followed by *AEW*. Despite being supervised, *IDML* does not perform on par. This may be due to labeled data not being sufficient to learn a proper metric in high dimensions, and/or the labels introduced during self-learning being noisy. We also find *Grid* and *Rand<sub>d</sub>* to rank at the bottom, suggesting that learning the graph structure provides advantage over these standard techniques.



**Table 3: Average test accuracy and rank (w.r.t. test error) of methods across datasets for varying labeling %.  $\blacktriangle$  ( $p < 0.005$ ) and  $\triangle$  ( $p < 0.01$ ) denote the cases where PG-LEARN is significantly better w.r.t. the paired Wilcoxon signed rank test.**

Labeled	PG-L	MinEnt	IDML	AEW	Grid	Rand <sub>d</sub>
10% acc.	<b>0.8819</b>	0.8594 $\blacktriangle$	0.8036 $\blacktriangle$	0.8448 $\blacktriangle$	0.8129 $\blacktriangle$	0.7952 $\blacktriangle$
rank	<b>1.20</b>	2.20	4.40	2.80	4.80	5.60
20% acc.	<b>0.8900</b>	0.8504 $\blacktriangle$	0.8118 $\blacktriangle$	0.8462 $\blacktriangle$	0.8099 $\blacktriangle$	0.8088 $\blacktriangle$
rank	<b>1.42</b>	2.83	4.17	2.92	4.83	4.83
30% acc.	<b>0.9085</b>	0.8636 $\blacktriangle$	0.8551 $\blacktriangle$	0.8613 $\blacktriangle$	0.8454 $\blacktriangle$	0.8386 $\blacktriangle$
rank	<b>1.33</b>	3.67	3.83	3.17	4.00	5.00
40% acc.	<b>0.9153</b>	0.8617 $\blacktriangle$	0.8323 $\blacktriangle$	0.8552 $\blacktriangle$	0.8381 $\blacktriangle$	0.8303 $\blacktriangle$
rank	<b>1.67</b>	3.67	3.50	3.67	4.00	4.50
50% acc.	<b>0.9251</b>	0.8700 $\triangle$	0.8647 $\blacktriangle$	0.8635 $\blacktriangle$	0.8556 $\blacktriangle$	0.8459 $\blacktriangle$
rank	<b>1.50</b>	3.17	3.83	3.67	4.00	4.83

**5.2.2 Parallel Experiments with Noisy Features.** Next we fully evaluate PG-LEARN in the parallel setting as proposed in Algo. 2. Graph learning is especially beneficial for SSL in noisy scenarios, where there exist irrelevant or noisy features that would cause simple graph construction methods like  $k$ NN and *Grid* go astray. To the effect of making the classification tasks more challenging, we *double* the feature space for each dataset, by adding 100% new noise features with values drawn randomly from standard  $Normal(0, 1)$ . Moreover, this provides a ground truth on the importance of features, based on which we are able to quantify how well our PG-LEARN recovers the necessary underlying relations by learning the appropriate feature weights.

**Setup:** We report results comparing PG-LEARN only with *MinEnt*, *Grid*, and *Rand<sub>d</sub>*—in this setup, *IDML* failed to learn a metric in several cases due to degeneracy and the authors’ implementation<sup>9</sup> of *AEW* gave out-of-memory errors in many cases. This however does not take away much, since *MinEnt* proved to be the second-best after PG-LEARN in the previous section (see Table 3) and *Grid* and *Rand<sub>d</sub>* are the typical methods used often in practice.

Given a budget  $B$  units of time and  $T$  parallel threads for our PG-LEARN, each competing method is executed for a total of  $BT$  units, i.e. all methods receive the same amount of processing time.<sup>11</sup> Specifically, *MinEnt* is started in  $T$  threads, each with a random initial configuration that runs until time is up (i.e., to completion, no early-terminations). *Grid* picks  $(k, \sigma)$  from the 2-d grid that we refine recursively, that is, split into finer resolution containing more cells as more allocated time remains, while *Rand<sub>d</sub>* continues picking random combinations of  $(k, \mathbf{a}_{1:d})$ . When the time is over, each method reports the hyperparameters that yield the highest validation accuracy, using which the test accuracy is computed.

**Results:** Table 4 presents the average test accuracy over 10 random samples from each dataset, using  $T = 32$ . We find that despite  $32\times$  more time, the baselines are crippled by the irrelevant features and increased dimensionality. In contrast, PG-LEARN maintains notably high accuracy that is significantly better than all the baselines on all datasets at  $p = 0.01$ .

<sup>11</sup>All experiments executed on a Linux server equipped with 96 Intel Xeon CPUs at 2.1 GHz and a total of 1 TB RAM, using Matlab R2015b Distributed Computing Server.

**Table 4: Test accuracy on datasets with 100% added noise features, avg’ed across 10 samples; 15 mins of hyperparameter tuning on  $T = 32$  threads. Symbols  $\blacktriangle$  ( $p < 0.005$ ) and  $\triangle$  ( $p < 0.01$ ) denote the cases where PG-LEARN is significantly better than the baseline w.r.t. the paired Wilcoxon signed rank test.**

Dataset	PG-LRN	MinEnt	Grid	Rand <sub>d</sub>
COIL	<b>0.9044</b>	0.8197 $\blacktriangle$	0.6311 $\blacktriangle$	0.6954 $\blacktriangle$
USPS	<b>0.9154</b>	0.8779 $\triangle$	0.8746 $\blacktriangle$	0.7619 $\blacktriangle$
MNIST	<b>0.8634</b>	0.8006 $\blacktriangle$	0.7932 $\blacktriangle$	0.6668 $\blacktriangle$
UMIST	<b>0.8789</b>	0.7756 $\blacktriangle$	0.7124 $\blacktriangle$	0.6405 $\blacktriangle$
YALE	<b>0.6859</b>	0.5671 $\blacktriangle$	0.5925 $\blacktriangle$	0.5298 $\blacktriangle$

Figure 5 (a) shows how the test error changes by time for all methods on average, and (b) depicts the validation and the corresponding test accuracies for PG-LEARN on an example run. We see that PG-LEARN gradually improves validation accuracy across threads over time, and test accuracy follows closely. As such, test error drops in time. *Grid* search has a near-flat curve as it uses the same kernel bandwidth on all dimensions, therefore, more time does not help in handling noise. *Rand<sub>d</sub>* error seems to drop slightly but stabilizes at a high value, demonstrating its limited ability to guess parameters in high dimensions with noise. Overall, PG-LEARN outperforms competition significantly in this high dimensional noisy setting as well. Its performance is particularly noteworthy on YALE, which has small  $n = 320$  but large  $2d > 2K$  half of which are noise.

Finally, Figure 6 shows PG-LEARN’s estimated hyperparameters,  $\mathbf{a}_{1:d}$  and  $\mathbf{a}_{(d+1):2d}$  (avg’ed over 10 samples), demonstrating that the noisy features  $(d + 1) : 2d$  receive notably lower weights.

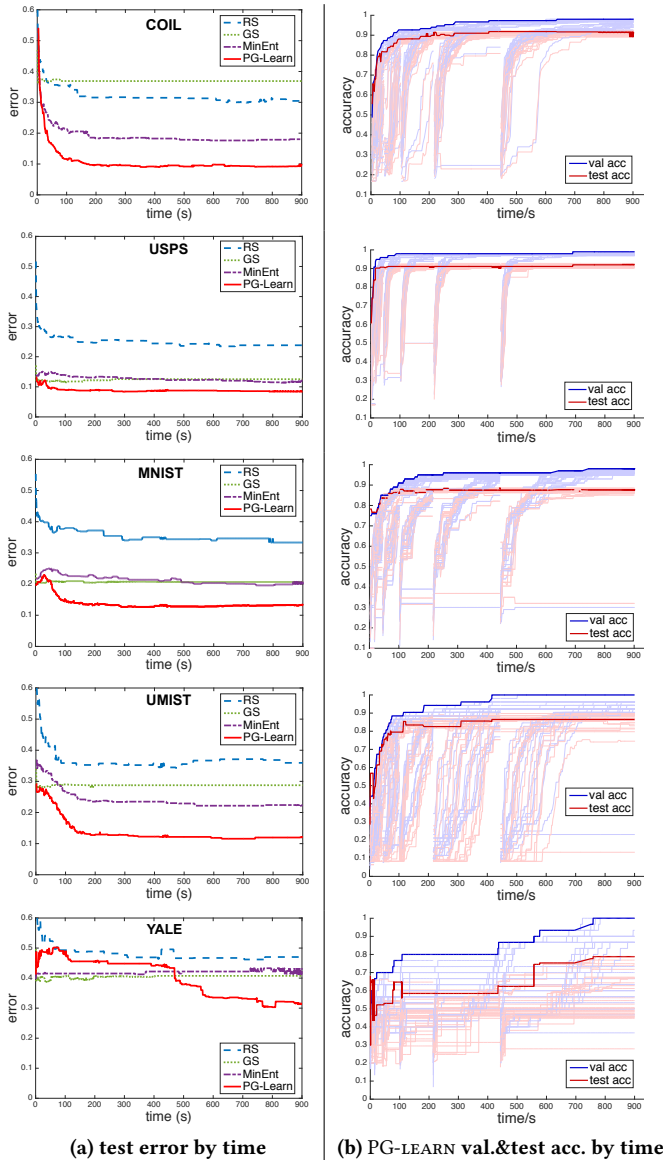
## 6 CONCLUSION

In this work we addressed the graph structure estimation problem as part of relational semi-supervised inference. It is now well-understood that graph construction from point-cloud data has critical impact on learning algorithms [6, 18]. To this end, we first proposed a learning-to-rank based objective parameterized by different weights per dimension and derived its gradient-based learning (§4.1). We then showed how to integrate this type of adaptive local search within a parallel framework that early-terminates searches based on relative performance, in order to dynamically allocate resources (time and processors) to those with promising configurations (§4.2). Put together, our solution PG-LEARN is a hybrid that strategically navigates the hyperparameter search space. What is more, PG-LEARN is scalable in dimensionality and number of samples both in terms of runtime and memory requirements. Finally, all source code and data can be publicly accessed at <https://pg-learn.github.io/>.

As future work we plan to deploy PG-LEARN on a distributed platform like Apache Spark, and generalize the ideas to other graph-based learning problems such as graph-regularized regression.

## ACKNOWLEDGMENTS


This research is sponsored by National Science Foundation CAREER 1452425 and IIS 1408287. Any conclusions expressed in this material are of the authors and do not necessarily reflect, either expressed or implied, the views of the funding parties.



**Figure 5: (a) Test error vs. time (avg'd across 10 runs w/ random samples) comparing PG-LEARN with baselines on noisy datasets; (b) PG-LEARN's validation and corresponding test accuracy over time as it executes Algo. 2 on 32 threads (1 run).**

## REFERENCES

- [1] Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani. 2006. Manifold Regularization: A Geometric Framework for Learning from Labeled and Unlabeled Examples. *Journal of Machine Learning Research* 7 (2006), 2399–2434.
- [2] Avrim Blum and Shuchi Chawla. 2001. Learning from Labeled and Unlabeled Data using Graph Mincuts. In *ICML*. 19–26.
- [3] Bin Cheng, Jianchao Yang, Shuicheng Yan, Yun Fu, and Thomas S. Huang. 2010. Learning With L1-Graph for Image Analysis. *IEEE Tran. Img. Proc.* 19, 4 (2010), 858–866.
- [4] Michael Conover, Bruno Goncalves, Jacob Ratkiewicz, Alessandro Flammini, and Filippo Menczer. 2011. Predicting the Political Alignment of Twitter Users. In *SocialCom/PASSAT*. IEEE, 192–199.



**Figure 6: Distribution of weights estimated by PG-LEARN, shown separately for the original and injected noisy features on each dataset. Notice that the latter is much lower, showing its competitiveness in hyperparameter estimation.**

- [5] Samuel I. Daitch, Jonathan A. Kelner, and Daniel A. Spielman. 2009. Fitting a Graph to Vector Data (*ICML*). ACM, New York, NY, USA, 201–208.
- [6] Celso André R. de Sousa, Solange O. Rezende, and Gustavo E. A. P. A. Batista. 2013. Influence of Graph Construction on Semi-supervised Learning. In *ECML/PKDD*. 160–175.
- [7] Paramveer S. Dhillon, Partha Pratim Talukdar, and Koby Crammer. 2010. Inference Driven Metric Learning for Graph Construction. *NESCAI (North East Student Symposium on Artificial Intelligence)*.
- [8] Paramveer S. Dhillon, Partha Pratim Talukdar, and Koby Crammer. 2010. Learning Better Data Representation using Inference-Driven Metric Learning (IDML). In *ACL*.
- [9] Lucie Flekova, Daniel Preotiu-Pietro, and Lyle H. Ungar. 2016. Exploring Stylistic Variation with Age and Income on Twitter. In *ACL*.
- [10] Aristides Gionis, Piotr Indyk, and Rajeev Motwani. 1999. Similarity Search in High Dimensions via Hashing. In *Vldb*. San Francisco, CA, USA, 518–529.
- [11] G.H. Golub and C.F. Van Loan. 1989. *Matrix Computations*. Johns Hopkins University Press.
- [12] Kevin G. Jamieson and Amee Talwalkar. 2016. Non-stochastic Best Arm Identification and Hyperparameter Optimization. In *AISTATS*. 240–248.
- [13] Tony Jebara, Jun Wang, and Shih-Fu Chang. 2009. Graph Construction and B-matching for Semi-supervised Learning (*ICML*). ACM, 441–448.
- [14] Masayuki Karasuyama and Hiroshi Mamitsuka. 2017. Adaptive edge weighting for graph-based learning algorithms. *Machine Learning* 106, 2 (2017), 307–335.
- [15] Lisha Li, Kevin G. Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Amee Talwalkar. 2016. Efficient Hyperparameter Optimization and Infinitely Many Armed Bandits. *CoRR* abs/1603.06560 (2016).
- [16] Yu-Feng Li, Shao-Bo Wang, and Zhi-Hua Zhou. 2016. Graph quality judgement: a large margin expedition. In *IJCAI*. 1725–1731.
- [17] Guangcan Liu, Zhouchen Lin, and Yong Yu. 2010. Robust Subspace Segmentation by Low-rank Representation. In *ICML*. 663–670.
- [18] Markus Maier, Ulrike von Luxburg, and Matthias Hein. 2008. Influence of graph construction on graph-based clustering measures. In *NIPS*. 1025–1032.
- [19] Bryan Perozzi and Steven Skiena. 2015. Exact Age Prediction in Social Networks. In *WWW (Companion Volume)*. ACM, 91–92.
- [20] Mohammad H. Rohban and Hamid R. Rabiee. 2012. Supervised neighborhood graph construction for semi-supervised classification. *Pattern Recognition* 45, 4 (2012), 1363–1372.
- [21] S.T. Roweis and L.K. Saul. 2000. Nonlinear dimensionality reduction by locally linear embedding. *Science* 290, 5500 (2000), 2323–2326.
- [22] Meng Wang, Weijie Fu, Shijie Hao, Dacheng Tao, and Xindong Wu. 2016. Scalable semi-supervised learning by efficient anchor graph regularization. *IEEE TKDE* 28, 7 (2016), 1864–1877.
- [23] Xinhua Zhang and Wee Sun Lee. 2006. Hyperparameter Learning for Graph Based Semi-supervised Learning Algorithms. In *NIPS*. MIT Press, 1585–1592.
- [24] Yan-Ming Zhang, Kaizhu Huang, Guanggang Geng, and Cheng-Lin Liu. 2013. Fast kNN Graph Construction with Locality Sensitive Hashing. In *ECML/PKDD*. 660–674.
- [25] Yan-Ming Zhang, Kaizhu Huang, Xinwen Hou, and Cheng-Lin Liu. 2014. Learning Locality Preserving Graph from Data. *IEEE Trans. Cybernetics* 44, 11 (2014), 2088–2098.
- [26] Dengyong Zhou, Olivier Bousquet, Thomas Navin Lal, Jason Weston, and Bernhard Schölkopf. 2003. Learning with Local and Global Consistency. In *NIPS*. MIT Press, 321–328.
- [27] Xiaojin Zhu, Zoubin Ghahramani, and John Lafferty. 2003. Semi-supervised learning using Gaussian fields and harmonic functions. In *ICML*. 912–919.
- [28] Liansheng Zhuang, Zihan Zhou, Shenghua Gao, Jingwen Yin, Zhouchen Lin, and Yi Ma. 2017. Label information guided graph construction for semi-supervised learning. *IEEE Transactions on Image Processing* 26, 9 (2017), 4182–4192.