
Poisson Learning: Graph Based Semi-Supervised Learning At Very Low Label Rates

Jeff Calder¹ Brendan Cook¹ Matthew Thorpe² Dejan Slepčev³

Abstract

We propose a new framework, called *Poisson learning*, for graph based semi-supervised learning at very low label rates. Poisson learning is motivated by the need to address the degeneracy of Laplacian semi-supervised learning in this regime. The method replaces the assignment of label values at training points with the placement of sources and sinks, and solves the resulting Poisson equation on the graph. The outcomes are provably more stable and informative than those of Laplacian learning. Poisson learning is efficient and simple to implement, and we present numerical experiments showing the method is superior to other recent approaches to semi-supervised learning at low label rates on MNIST, FashionMNIST, and Cifar-10. We also propose a graph-cut enhancement of Poisson learning, called *Poisson MBO*, that gives higher accuracy and can incorporate prior knowledge of relative class sizes.

1. Introduction

Semi-supervised learning uses both labeled and unlabeled data in learning tasks. For problems where very few labels are available, geometric or topological structure in unlabeled data can be used to greatly improve the performance of classification, regression, or clustering algorithms. One of the most widely used methods in graph-based semi-supervised learning is Laplace learning, originally proposed in (Zhu et al., 2003), which seeks a graph harmonic function that extends the labels. Laplace learning, and variants thereof, have been widely applied in semi-supervised learning (Zhou et al., 2005; 2004a;b; Ando & Zhang, 2007) and manifold

ranking (He et al., 2004; Yang et al., 2013; Xu et al., 2011), among many other problems.

This paper is concerned with graph-based semi-supervised learning at very low label rates. In this setting, it has been observed that Laplace learning can give very poor classification results (Nadler et al., 2009; El Alaoui et al., 2016). The poor results are often attributed to the fact that the solutions develop localized spikes near the labeled points and are almost constant far from the labeled points. In particular, label values are not propagated well by the Laplacian learning approach. To address this issue, recent work has suggested to consider p -Laplace learning (El Alaoui et al., 2016). Several works have rigorously studied p -Laplace regularization with few labels (Slepčev & Thorpe, 2019; Calder, 2018; 2019), and recent numerical results show that $p > 2$ is superior to Laplace learning at low label rates (Flores et al., 2019). The case of $p = \infty$ is called Lipschitz learning (Kyng et al., 2015), which seeks the absolutely minimal Lipschitz extension of the training data. Other methods to address low label rate problems include higher order Laplacian regularization (Zhou & Belkin, 2011) and spectral cutoffs (Belkin & Niyogi, 2002).

While p -Laplace learning improves upon Laplace learning, the method is more computationally burdensome than Laplace learning, since the optimality conditions are non-linear. Other recent approaches have aimed to re-weight the graph more heavily near labels, in order to give them wider influence when the labeling rate is very low. One way to re-weight the graph is the Weighted Nonlocal Laplacian (WNLL) (Shi et al., 2017), which amplifies the weights of edges directly connected to labeled nodes. The WNLL achieves better results at moderately low label rates, but still performs poorly at very low label rates (Flores et al., 2019). To address this, (Calder & Slepčev, 2019) proposed the Properly Weighted Laplacian, which re-weights the graph in a way that is well-posed at arbitrarily low label rates.

Much of the recent work on low label rate problems has focused on formulating and implementing new learning approaches that are well-posed with few labels. The exact nature of the degeneracy in Laplace learning, and the question of how the tails of the spikes propagate label information, has not been studied and is still poorly understood.

¹School of Mathematics, University of Minnesota, Minneapolis, USA. ²Department of Mathematics, University of Manchester, Manchester, UK. ³Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, USA.. Correspondence to: Jeff Calder <jwcalder@umn.edu>.

For some problems, the performance of Laplacian learning is good (Zhu et al., 2003), while for other problems it is catastrophic, yielding very poor classification results similar to random guessing (Shi et al., 2017; Flores et al., 2019).

In this paper, we carefully analyze Laplace learning at very low label rates, and we discover that *nearly all* of the degeneracy of Laplace learning is due to a large constant bias in the solution of the Laplace equation that is present only at low label rates. In order to overcome this problem we introduce a new algorithm, we call *Poisson learning*, that gives very good classification performance down to extremely low label rates. We give a random walk interpretation of Poisson learning that shows how the method uses information from the random walkers only *before* they reach the mixing time of the random walk and forget their initial condition. We also propose a graph-cut enhancement of Poisson learning, called *Poisson MBO*, that can incorporate knowledge about class sizes, and further improves the class-label accuracy.

The rest of the paper is organized as follows. In Section 2 we first briefly introduce Poisson learning, and then provide a detailed motivation for the algorithm and a theoretical analysis from both the variational and random walk perspectives. The Poisson MBO approach is presented in Section 2.4. In Section 3 we present the step-by-step algorithms and discuss implementation details for the Poisson and Poisson MBO algorithms. In Section 4 we present numerical experiments with semi-supervised classification on the MNIST, FashionMNIST, and Cifar-10 datasets. The proofs of the results are available in the supplementary materials.

2. Poisson learning

Let $X = \{x_1, x_2, \dots, x_n\}$ denote the vertices of a graph with edge weights $w_{ij} \geq 0$ between x_i and x_j . We assume the graph is symmetric, so $w_{ij} = w_{ji}$. We define the degree $d_i = \sum_{j=1}^n w_{ij}$. For a multi-class classification problem with k classes, we let the standard basis vector $\mathbf{e}_i \in \mathbb{R}^k$ represent the i^{th} class. We assume the first m vertices x_1, x_2, \dots, x_m are given labels $y_1, y_2, \dots, y_m \in \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_k\}$, where $m < n$. The task of graph-based semi-supervised learning is to extend the labels to the rest of the vertices $x_{m+1}, x_{m+2}, \dots, x_n$.

The well-known Laplace learning algorithm (Zhu et al., 2003) extends the labels by solving the problem

$$\left. \begin{aligned} \mathcal{L}u(x_i) &= 0, & \text{if } m+1 \leq i \leq n, \\ u(x_i) &= y_i, & \text{if } 1 \leq i \leq m, \end{aligned} \right\} \quad (2.1)$$

where \mathcal{L} is the unnormalized graph Laplacian given by

$$\mathcal{L}u(x_i) = \sum_{j=1}^n w_{ij}(u(x_i) - u(x_j)).$$

Here, $u : X \rightarrow \mathbb{R}^k$ and we write the components of u as $u(x_i) = (u_1(x_i), u_2(x_i), \dots, u_k(x_i))$. The label decision for vertex x_i is determined by the largest component of $u(x_i)$

$$\ell(x_i) = \arg \max_{j \in \{1, \dots, k\}} \{u_j(x_i)\}. \quad (2.2)$$

We note that Laplace learning is also called *label propagation (LP)* (Zhu et al., 2005), since the Laplace equation (2.1) can be solved by repeatedly replacing $u(x_i)$ with the weighted average of its neighbors, which can be viewed as dynamically propagating labels.

At very low label rates, we propose to replace the problem (2.1) by *Poisson learning*: Let $\bar{y} = \frac{1}{m} \sum_{j=1}^m y_j$ be the average label vector and let $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ if $i \neq j$. One computes the solution of the Poisson equation

$$\mathcal{L}u(x_i) = \sum_{j=1}^m (y_j - \bar{y})\delta_{ij} \quad \text{for } i = 1, \dots, n \quad (2.3)$$

satisfying $\sum_{i=1}^n d_i u(x_i) = 0$. While the label decision can be taken to be the same as (2.2), it is also simple to account for unbalanced classes or training data with the modified label decision

$$\ell(x_i) = \arg \max_{j \in \{1, \dots, k\}} \{s_j u_j(x_i)\}, \quad (2.4)$$

where $s_j = b_j / (\bar{y} \cdot \mathbf{e}_j)$ and b_j is the fraction of data belonging to class j . We explain this label decision in Remark 2.2. The Poisson equation (2.3) can be solved efficiently with a simple iteration given in Algorithm 1.

Technically speaking, in Laplace learning, the labels are imposed as boundary conditions in a Laplace equation, while in Poisson learning, the labels appears as a source term in a graph Poisson equation. In the sections below, we explain why Poisson learning is a good idea for problems with very few labels. In particular, we give random walk and variational interpretations of Poisson learning, and we illustrate how Poisson learning arises as the low label rate limit of Laplace learning.

2.1. Random walk interpretation

We present a random walk interpretation of Poisson learning and compare to the random walk interpretation of Laplace learning to explain its poor performance at low label rates. We note that Laplace learning works very well in practice for semi-supervised learning problems with a moderate amount of labeled data. For example, on the MNIST dataset we obtained around 95% accuracy at 16 labels per class (0.23% label rate). However, at very low label rates the performance is poor. At 1 label per class, we find the average performance is around 16% accuracy. This phenomenon has been observed in other works recently (Nadler et al.,

2009; El Alaoui et al., 2016). However, a clear understanding of the issues with Laplace learning at low label rates was lacking. The clearest understanding of this phenomenon comes from the random walk interpretation, and this leads directly to the foundations for Poisson learning.

Let $x \in X$ and let $X_0^x, X_1^x, X_2^x, \dots$ be a random walk on X starting at $X_0^x = x$ with transition probabilities

$$\mathbb{P}(X_k^x = x_j | X_{k-1}^x = x_i) = d_i^{-1} w_{ij}.$$

Let u be the solution of the *Laplace* learning problem (2.1). Define the stopping time to be the first time the walk hits a label, that is

$$\tau = \inf\{k \geq 0 : X_k^x \in \{x_1, x_2, \dots, x_m\}\}.$$

Let $i_\tau \leq m$ denote the index of the point X_τ^x , so $X_\tau^x = x_{i_\tau}$. Then, by Doob's optimal stopping theorem, we have

$$u(x) = \mathbb{E}[y_{i_\tau}]. \quad (2.5)$$

This gives the standard representation formula for the solution of Laplace learning (2.1). The interpretation is that we release a random walker from x and let it walk until it hits a labeled vertex and then record that label. We average over many random walkers to get the value of $u(x)$.

When there are insufficiently many labels, the stopping time τ is so large that we have passed the *mixing time* of the random walk, and the distribution of X_τ^x is very close to the invariant distribution of the random walk $\pi(x_i) = d_i / \sum_i d_i$. In this case, (2.5) gives that

$$u(x) \approx \frac{\sum_{i=1}^m d_i y_i}{\sum_{j=1}^m d_j} =: \bar{y}_w. \quad (2.6)$$

Of course, the function u is not exactly constant, and instead it is approximately constant with sharp spikes at the labeled vertices; see Figure 1. Previous work has proved rigorously that Laplace learning degenerates in this way at low label rates (Slepčev & Thorpe, 2019; Calder, 2018).

It is important to note that the constant vector \bar{y}_w depends only on the degrees of the *labeled* nodes in the graph, which is very sensitive to local graph structure. When Laplace learning returns a nearly constant label function, it can be catastrophic for classification, since most datapoints are assigned the same label. This explains the 16% accuracy in the MNIST experiment described above.

In contrast, the Poisson equation (2.3) has a random walk interpretation that involves the Green's function for a random walk, and is in some sense dual to Laplace learning. The source term on the right hand side of (2.3) represents random walkers being released from labeled points and exploring the graph, while carrying their label information




Figure 1. Demonstration of spikes in Laplace learning. The graph consists of $n = 10^4$ independent uniform random variables on $[0, 1]^2$ and two points are given labels of 0 and 1. Most values of the solution u of Laplace learning are very close to 0.5.

with them. For $T > 0$ let us define

$$w_T(x_i) = \mathbb{E} \left[\sum_{k=0}^T \sum_{j=1}^m y_j \mathbb{1}_{\{X_k^{x_j} = x_i\}} \right].$$

Essentially, each time the random walk starting from x_j , denoted $X_k^{x_j}$, visits x_i we record the label y_j and sum this over all visits from $0 \leq k \leq T$. For short time T , this quantity is meaningful, but since the random walk is recurrent (it is a finite graph), $w_T(x_i) \rightarrow \infty$ as $T \rightarrow \infty$. If we normalize by $1/T$, we still have the same issue as with Laplace learning, where we are only measuring the invariant distribution. Indeed, we note that

$$w_T(x_i) = \sum_{k=0}^T \sum_{j=1}^m y_j \mathbb{P}(X_k^{x_j} = x_i)$$

$$\text{and } \lim_{k \rightarrow \infty} \mathbb{P}(X_k^{x_j} = x_i) = \frac{d_i}{\sum_{i=1}^m d_i}.$$

Therefore, when k is large we have

$$\sum_{j=1}^m y_j \mathbb{P}(X_k^{x_j} = x_i) \approx \frac{d_i}{\sum_{i=1}^m d_i} \sum_{j=1}^m y_j,$$

and so the tail of the sum defining $w_T(x_i)$ is recording a blind average of labels.

The discussion above suggests that we should subtract off this average tail behavior from w_T , so that we only record the *short-time* behavior of the random walk, before the mixing time is reached. We also normalize by d_i , which leads us to define

$$u_T(x_i) = \mathbb{E} \left[\sum_{k=0}^T \frac{1}{d_i} \sum_{j=1}^m (y_j - \bar{y}) \mathbb{1}_{\{X_k^{x_j} = x_i\}} \right], \quad (2.7)$$

where $\bar{y} = \frac{1}{m} \sum_{j=1}^m y_j$. It turns out that as $T \rightarrow \infty$, the function $u_T(x_i)$ converges to the solution of (2.3).

Theorem 2.1. For every $T \geq 0$ we have

$$u_{T+1}(x_i) = u_T(x_i) + d_i^{-1} \left(\sum_{j=1}^m (y_j - \bar{y}) \delta_{ij} - \mathcal{L}u_T(x_i) \right).$$

If the graph G is connected and the Markov chain induced by the random walk is aperiodic, then $u_T \rightarrow u$ as $T \rightarrow \infty$, where u is the unique solution of the Poisson equation (2.3) satisfying $\sum_{i=1}^n d_i u(x_i) = 0$.

Theorem 2.1 gives the foundation for Poisson learning through the random walk perspective, and in fact, it also gives a numerical method for computing the solution (see Algorithm 1).

Remark 2.2. The representation formula (2.7) for the solution of Poisson learning (2.3) shows that the solution u is a linear function of the label vectors y_1, \dots, y_m . That is, for any $A \in \mathbb{R}^{k \times k}$, the solution $u_A : X \rightarrow \mathbb{R}^k$ of

$$\mathcal{L}u_A(x_i) = \sum_{j=1}^m (Ay_j - A\bar{y}) \delta_{ij} \quad \text{for } i = 1, \dots, n$$

satisfying $\sum_{i=1}^n d_i u_A(x_i) = 0$ is exactly $u_A = Au$, where u is the solution of (2.3). This shows that any reweighting of the point sources, by which we mean $y_j \mapsto Ay_j$, is equivalent to reweighting the solution by $u \mapsto Au$.

If we set $A = \text{diag}(s_1, \dots, s_k)$, then Au corresponds to multiplying the point sources for class i by the weight s_i . We can use this reweighting to account for unbalanced classes, or a discrepancy between the balancing of training and testing data, in the following way. Let n_j be the number of training examples from class j , and let b_j denote the true fraction of data in class j . We can choose s_j so that $n_j s_j = b_j$ to ensure that the mass of the point sources for each class, weighted by s_j , is proportional to the true fraction of data in that class. Since n_j is proportional to $\bar{y} \cdot \mathbf{e}_j$, this explains our modified label decision (2.4).

2.2. Variational interpretation

We can also interpret Poisson learning (2.3) as a gradient regularized variational problem. Before proceeding, we briefly review some facts about calculus on graphs. Let $\ell^2(X)$ denote the space of functions $u : X \rightarrow \mathbb{R}^k$ equipped with the inner product

$$(u, v)_{\ell^2(X)} = \sum_{i=1}^n u(x_i) \cdot v(x_i).$$

This induces a norm $\|u\|_{\ell^2(X)}^2 = (u, u)_{\ell^2(X)}$. We also define the space of mean-zero functions

$$\ell_0^2(X) = \left\{ u \in \ell^2(X) : \sum_{i=1}^n d_i u(x_i) = 0 \right\}$$

We define a *vector field* on the graph to be an *antisymmetric* function $V : X^2 \rightarrow \mathbb{R}^k$ (i.e., $V(x_i, x_j) = -V(x_j, x_i)$). The *gradient* of a function $u \in \ell^2(X)$, denoted for simplicity as ∇u , is defined to be the vector field

$$\nabla u(x_i, x_j) = u(x_j) - u(x_i).$$

The inner product between vector fields V and W is

$$(V, W)_{\ell^2(X^2)} = \frac{1}{2} \sum_{i,j=1}^n w_{ij} V(x_i, x_j) \cdot W(x_i, x_j).$$

and the norm of V is $\|V\|_{\ell^2(X^2)}^2 = (V, V)_{\ell^2(X^2)}$.

We now consider the variational problem

$$\min_{u \in \ell_0^2(X)} \left\{ \frac{1}{2} \|\nabla u\|_{\ell^2(X^2)}^2 - \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j) \right\}. \quad (2.8)$$

The following theorem makes the connection between Poisson learning and the variational problem (2.8).

Theorem 2.3. Assume G is connected. Then there exists a unique minimizer $u \in \ell_0^2(X)$ of (2.8), and u satisfies the Poisson learning equation (2.3).

Theorem 2.3 shows that the Poisson learning equation (2.3) arises as the necessary conditions for a gradient regularized variational problem with an affine loss function. We contrast this with the solution of Laplace learning (2.1), which is the minimizer of the variational problem

$$\min_{u \in \ell^2(X)} \left\{ \|\nabla u\|_{\ell^2(X^2)}^2 : u(x_i) = y_i, 1 \leq i \leq m \right\}. \quad (2.9)$$

Thus, while both Laplace and Poisson learning are gradient regularized variational problems, the key difference is how each algorithm handles the labeled data; Laplace learning enforces hard label constraints while Poisson learning adds an affine loss function to the energy. Of course, many variants of Laplace learning have been proposed with various types of soft label constraints in place of hard constraints. These variations perform similarly poorly to Laplace learning at low label rates, and the key feature of Poisson learning is the *affine* loss function that can be easily centered.

We note that it would be natural to add a weight $\mu > 0$ to one of the terms in (2.8) to trade-off the importance of the two terms in the variational problem. However, as we show in the supplementary material (see Lemma A.3), this weight would have no effect on the final label decision. We also mention that the variational problem (2.8) has a natural ℓ^p generalization that we also explore in the supplementary material (Section A.2).

2.3. Laplace learning at low label rates

Finally, to further motivate Poisson learning, we connect Poisson learning with the limit of Laplace learning at very

low label rates. At low label rates, the solution of Laplace learning (2.1) concentrates around a constant \bar{y}_w , for which we gave an interpretation of via random walks in Section 2.1. Near labeled nodes, u has sharp spikes (recall Figure 1) in order to attain the labels. From the variational perspective, the constant function has zero cost, and the cost of spikes is very small, so this configuration is less expensive than continuously attaining the labels $u(x_i) = y_i$.

A natural question concerns whether the spikes in Laplace learning contain useful information, or whether they are too localized and do not propagate information well. To test this, we changed the label decision (2.2) in the MNIST experiment described in Section 2.1 to subtract off the tail constant \bar{y}_w identified in (2.6)

$$\ell(x_i) = \arg \max_{j \in \{1, \dots, k\}} \{u_j(x) - \bar{y}_w \cdot \mathbf{e}_j\}.$$

where \bar{y}_w is defined in (2.6). Thus, we are centering the function u at zero. At 1 label per class (10 labeled images total), the accuracy improved from 16% to 85.9%! Hence, the difference $u - \bar{y}_w$ contains enough information to make informed classification decisions, and therefore the spikes contain useful information.

This indicates that much of the poor performance of Laplace learning can be explained by a large shift bias that occurs only at very low label rates. Fixing this seems as simple as applying the appropriate shift before making a decision on a label, but this does not lead to a well-grounded method, since the shifted function $u - \bar{y}_w$ is exactly the graph-harmonic extension of the shifted labels $y_j - \bar{y}_w$. Why should we have to use a harmonic extension of the *wrong labels* in order to achieve a better result? On the other hand, Poisson learning, which we introduced above, provides an intuitive and well-grounded way of fixing the shift bias in Laplace learning.

To see the connection to Poisson learning, let us assume the solution u of the Laplace learning equation (2.1) is nearly equal to the constant vector $\bar{y}_w \in \mathbb{R}^k$ from (2.6) at all unlabeled points x_{m+1}, \dots, x_n . For any *labeled* node x_i with $i = 1, \dots, m$ we can compute (assuming $w_{ij} = 0$ for all $j \in \{1, \dots, m\}$) that

$$\begin{aligned} \mathcal{L}u(x_i) &= \sum_{j=1}^n w_{ij}(u(x_i) - u(x_j)) \\ &\approx \sum_{j=m+1}^n w_{ij}(y_i - \bar{y}_w) = d_i(y_i - \bar{y}_w). \end{aligned}$$

Since $\mathcal{L}u(x_i) = 0$ for $i = m+1, \dots, n$, we find that u approximately satisfies the Poisson equation

$$\mathcal{L}u(x_i) = \sum_{j=1}^m d_j(y_j - \bar{y}_w)\delta_{ij} \quad \text{for } i = 1, \dots, n. \quad (2.10)$$

This gives a connection, at a heuristic level, been Laplace equations with hard constraints, and Poisson equations with point sources, for problems with very low label rates. We note that since constant functions are in the kernel of \mathcal{L} , $u - \bar{y}_w$ also satisfies (2.10). We also note that the labels, and the constant \bar{y}_w , are weighted by the degree d_j , which does not appear in our Poisson learning equation (2.3). We have found that both models give good results, but that (2.3) works slightly better, which is likely due to the rigorous foundation of (2.3) via random walks.

2.4. The Poisson MBO algorithm

Poisson learning provides a robust method for propagating label information that is stable at very low label rates. After applying Poisson learning to propagate labels, we propose a graph-cut method to incrementally adjust the decision boundary so as to improve the label accuracy and account for prior knowledge of class sizes. The graph-cut method we propose is to apply several steps of gradient descent on the graph-cut problem

$$\min_{\substack{u: X \rightarrow S_k \\ (u)_X = b}} \left\{ \frac{1}{2} \|\nabla u\|_{\ell^2(X^2)}^2 - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j) \right\}, \quad (2.11)$$

where $S_k = \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_k\}$, $b \in \mathbb{R}^k$ is given, and $(u)_X := \frac{1}{n} \sum_{i=1}^n u(x_i)$. Since we are restricting $u(x_i) \in S_k$, the term $\frac{1}{2} \|\nabla u\|_{\ell^2(X^2)}^2$ is exactly the graph-cut energy of the classification given by u . Likewise, the components of the average $(u)_X$ represent the fraction of points assigned to each class. The constraint $(u)_X = b$ therefore allows us to incorporate prior knowledge about relative class sizes through the vector $b \in \mathbb{R}^k$, which should have positive entries and sum to one. If there exists $u : X \rightarrow S_k$ with $(u)_X = b$, then (2.11) admits a solution, which in general may not be unique.

On its own, the graph-cut problem (2.11) can admit many local minimizers that would yield poor classification results. The phenomenon is similar to the degeneracy in Laplace learning at low label rates, since it is very inexpensive to violate any of the label constraints. Our overall plan is to first use Poisson learning to robustly propagate the labels, and then project onto the constraint set for (2.11) and perform several steps of gradient-descent on (2.11) to improve the classification accuracy. While Poisson learning propagates the labels in a robust way, the cut energy is more suitable for locating the exact decision boundary.

To relax the discrete problem (2.11), we approximate the graph-cut energy with the Ginzburg-Landau approximation

$$\min_{\substack{u \in \ell^2(X) \\ (u)_X = b}} \left\{ \text{GL}_\tau(u) - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j) \right\}, \quad (2.12)$$

where

$$\text{GL}_\tau(u) = \frac{1}{2} \|\nabla u\|_{\ell^2(X^2)}^2 + \frac{1}{\tau} \sum_{i=1}^n \prod_{j=1}^k |u(x_i) - \mathbf{e}_j|^2.$$

The Ginzburg-Landau approximation allows $u \in \ell^2(X)$ to take on any real values, instead of discrete values $u \in S_k$, making the approximation (2.12) easier to solve computationally. The graph Ginzburg-Landau approximation GL_τ has been used previously for graph-based semi-supervised learning in (Garcia-Cardona et al., 2014), and other works have rigorously studied how GL_τ approximates graph-cut energies in the scalar setting (Van Gennip & Bertozzi, 2012). Here, we extend the results to the vector multi-class setting.

Theorem 2.4. *Assume G is connected. Let $b \in \mathbb{R}^k$ and assume there exists $u : X \rightarrow S_k$ with $(u)_X = b$. For each $\tau > 0$ let u_τ be any solution of (2.12). Then, the sequence $(u_\tau)_\tau$ is precompact in $\ell^2(X)$ and any convergent subsequence u_{τ_m} converges to a solution of the graph-cut problem (2.11) as $\tau_m \rightarrow 0$. Furthermore, if the solution $u_0 : X \rightarrow S_k$ of (2.11) is unique, then $u_\tau \rightarrow u_0$ as $\tau \rightarrow 0$.*

Theorem 2.4 indicates that we can replace the graph-cut energy (2.11) with the simpler Ginzburg-Landau approximation (2.12). To descend on the energy (2.12), we use a time-spitting scheme that alternates gradient descent on

$$E_1(u) := \frac{1}{2} \|\nabla u\|_{\ell^2(X^2)}^2 - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j),$$

$$\text{and } E_2(u) := \frac{1}{\tau} \sum_{i=1}^n \prod_{j=1}^k |u(x_i) - \mathbf{e}_j|^2.$$

The first term E_1 is exactly the energy for Poisson learning (2.8), and gradient descent amounts to the iteration

$$u^{t+1}(x_i) = u^t(x_i) - dt \left(\mathcal{L}u^t(x_i) - \mu \sum_{j=1}^m (y_j - \bar{y}) \delta_{ij} \right).$$

We note that $\mathcal{L}u$ and the source term above both have zero mean value. Hence, the gradient descent equation for E_1 is *volume preserving*, i.e., $(u^{t+1})_X = (u^t)_X$. This would not be true for other fidelity terms, such as an ℓ^2 fidelity, and this volume conservation property plays an important role in ensuring the class size constraint $(u)_X = b$ in (2.12).

Gradient descent on the second term E_2 , when $\tau > 0$ is small, amounts to projecting each $u(x_i) \in \mathbb{R}^k$ to the closest label vector $\mathbf{e}_j \in S_k$, while preserving the volume constraint $(u)_X = b$. We approximate this by the following procedure: Let $\text{Proj}_{S_k} : \mathbb{R}^k \rightarrow S_k$ be the closest point projection, let $s_1, \dots, s_k > 0$ be positive weights, and set

$$u^{t+1}(x_i) = \text{Proj}_{S_k}(\text{diag}(s_1, \dots, s_k)u^t(x_i)), \quad (2.13)$$

where $\text{diag}(s_1, \dots, s_k)$ is the diagonal matrix with diagonal entries s_1, \dots, s_k . We use a simple gradient descent

scheme to choose the weights $s_1, \dots, s_k > 0$ so that the volume constraint $(u^{t+1})_X = b$ holds (see Steps 9-14 in Algorithm 2). By Remark 2.2, this procedure can be viewed as reweighting the point sources in the Poisson equation (2.3) so that the volume constraint holds. In particular, increasing or decreasing s_i grows or shrinks the size of class i .

We note that the work of (Jacobs et al., 2018) provides an alternative way to enforce explicit class balance constraints with a volume constrained MBO method based on auction dynamics. Their method uses a graph-cut based approach with a Voronoi-cell based initialization.

3. Poisson learning algorithms

We now present our proposed Poisson learning algorithms. The Python source code and simulation environment for reproducing our results is available online.¹

We let $\mathbf{W} = (w_{ij})_{i,j=1}^n$ denote our symmetric weight matrix. We treat all vectors as column vectors, and we let $\mathbb{1}$ and $\mathbf{0}$ denote the all-ones and all-zeros column vectors, respectively, of the appropriate size based on context. We assume that the first m data points x_1, x_2, \dots, x_m are given labels $y_1, y_2, \dots, y_m \in \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_k\}$, where the standard basis vector $\mathbf{e}_i \in \mathbb{R}^k$ represents the i^{th} class. We encode the class labels in a $k \times m$ matrix \mathbf{F} , whose j^{th} column is exactly y_j . Let $\mathbf{b} \in \mathbb{R}^k$ be the vector whose i^{th} entry b_i is the fraction of data points belonging to class i . If this information is not available, we set $\mathbf{b} = \frac{1}{k} \mathbb{1}$.

Poisson learning is summarized in Algorithm 1. The label decision for node i is $\ell_i = \arg \max_{1 \leq j \leq k} \mathbf{U}_{ij}$, and the reweighting in Step 11 implements the label decision (2.4). In all our results, we always set $\mathbf{b} = \frac{1}{k} \mathbb{1}$, so Poisson learning does not use prior knowledge of class sizes (the true value for \mathbf{b} is used in PoissonMBO below). The complexity is $O(TE)$, where E is the number of edges in the graph. We note that before the reweighting in Step 11, the Poisson learning algorithm computes exactly the function u_T defined in (2.7). In view of this, there is little to be gained from running the iterations beyond the *mixing time* of the random walk. This can be recorded within the loop in Steps 8-10 by adding the iteration $\mathbf{p}_{t+1} = \mathbf{W}\mathbf{D}^{-1}\mathbf{p}_t$, where the initial value \mathbf{p}_0 is the vector with ones in the positions of all labeled vertices, and zeros elsewhere. Up to a constant, \mathbf{p}_t is the probability distribution of a random walker starting from a random labeled node after t steps. Then the *mixing time* stopping condition is to run the iterations until

$$\|\mathbf{p}_t - \mathbf{p}_\infty\|_\infty \leq \varepsilon,$$

where $\mathbf{p}_\infty = \mathbf{W}\mathbb{1}/(\mathbb{1}^T\mathbf{W}\mathbb{1})$ is the invariant distribution. We use this stopping condition with $\varepsilon = 1/n$ in all experiments, which usually takes between 100 and 500 iterations.

¹Source Code: <https://github.com/jwcalder/GraphLearning>

Algorithm 1 PoissonLearning

```

1: Input:  $\mathbf{W}, \mathbf{F}, \mathbf{b}, T$ 
2: Output:  $\mathbf{U} \in \mathbb{R}^{n \times k}$ 
3:  $\mathbf{D} \leftarrow \text{diag}(\mathbf{W}\mathbf{1})$ 
4:  $\mathbf{L} \leftarrow \mathbf{D} - \mathbf{W}$ 
5:  $\bar{\mathbf{y}} \leftarrow \frac{1}{m}\mathbf{F}\mathbf{1}$ 
6:  $\mathbf{B} \leftarrow [\mathbf{F} - \bar{\mathbf{y}}, \mathbf{zeros}(k, n - m)]$ 
7:  $\mathbf{U} \leftarrow \mathbf{zeros}(n, k)$ 
8: for  $i = 1$  to  $T$  do
9:    $\mathbf{U} \leftarrow \mathbf{U} + \mathbf{D}^{-1}(\mathbf{B}^T - \mathbf{L}\mathbf{U})$ 
10: end for
11:  $\mathbf{U} \leftarrow \mathbf{U} \cdot \text{diag}(\mathbf{b}/\bar{\mathbf{y}})$ 

```

The Poisson MBO algorithm is summarized in Algorithm 2. The matrices \mathbf{D} , \mathbf{L} and \mathbf{B} are the same as in Poisson learning, and Poisson MBO requires an additional fidelity parameter μ and two parameters N_{inner} and N_{outer} . In all experiments in this paper, we set $\mu = 1$, $N_{\text{inner}} = 40$ and $N_{\text{outer}} = 20$. Steps 9-14 implement the volume constrained projection described in Section 2.4. We set the time step as $d\tau = 10$ and set the clipping values in Step 12 to $s_{\text{min}} = 0.5$ and $s_{\text{max}} = 2$. We tested on datasets with balanced classes, and on datasets with very unbalanced classes, one may wish to enlarge the interval $[s_{\text{min}}, s_{\text{max}}]$.

The additional complexity of PoissonMBO on top of Poisson learning is $O(N_{\text{inner}}N_{\text{outer}}E)$. On large datasets like MNIST, FashionMNIST and Cifar-10, our Poisson learning implementation in Python takes about 8 seconds to run on a standard laptop computer, and about 1 second with GPU acceleration.² The additional 20 iterations of PoissonMBO takes about 2 minutes on a laptop and 30 seconds on a GPU. These computational times do not include the time taken to construct the weight matrix.

4. Experimental Results

We tested Poisson learning on three datasets: MNIST (LeCun et al., 1998), FashionMNIST (Xiao et al., 2017) and Cifar-10 (Krizhevsky et al., 2009). FashionMNIST is a drop-in replacement for MNIST consisting of 10 classes of clothing items. To build good quality graphs, we trained autoencoders to extract important features from the data. For MNIST and FashionMNIST, we used variational autoencoders with 3 fully connected layers of sizes (784,400,20) and (784,400,30), respectively, followed by a symmetrically defined decoder. The autoencoder was trained for 100 epochs on each dataset. The autoencoder architecture, loss, and training, are similar to (Kingma & Welling, 2014). For Cifar-10, we used the AutoEncodingTransformations architecture from (Zhang et al., 2019), with all the default

²We used an NVIDIA RTX-2070 GPU, and it took 3 seconds to load data to/from the GPU and 1 second to solve Poisson learning.

Algorithm 2 PoissonMBO

```

1: Input:  $\mathbf{W}, \mathbf{F}, \mathbf{b}, T, N_{\text{inner}}, N_{\text{outer}}, \mu > 0$ 
2: Output:  $\mathbf{U} \in \mathbb{R}^{n \times k}$ 
3:  $\mathbf{U} \leftarrow \mu \cdot \text{PoissonLearning}(\mathbf{W}, \mathbf{F}, \mathbf{b}, T)$ 
4:  $dt \leftarrow 1/\max_{1 \leq i \leq n} \mathbf{D}_{ii}$ 
5: for  $i = 1$  to  $N_{\text{outer}}$  do
6:   for  $j = 1$  to  $N_{\text{inner}}$  do
7:      $\mathbf{U} \leftarrow \mathbf{U} - dt(\mathbf{L}\mathbf{U} - \mu\mathbf{B}^T)$ 
8:   end for
9:    $\mathbf{s} \leftarrow \mathbf{ones}(1, k)$ 
10:  for  $j = 1$  to 100 do
11:     $\hat{\mathbf{b}} \leftarrow \frac{1}{n}\mathbf{1}^T \mathbf{Proj}_{S_k}(\mathbf{U} \cdot \text{diag}(\mathbf{s}))$ 
12:     $\mathbf{s} \leftarrow \mathbf{max}(\mathbf{min}(\mathbf{s} + d\tau(\mathbf{b} - \hat{\mathbf{b}}), s_{\text{max}}), s_{\text{min}})$ 
13:  end for
14:   $\mathbf{U} \leftarrow \mathbf{Proj}_{S_k}(\mathbf{U} \cdot \text{diag}(\mathbf{s}))$ 
15: end for

```

parameters from their paper, and we normalized the features to unit-vectors.

We then constructed a graph over the latent feature space by connecting each image to its K -nearest neighbors with Gaussian weights given by

$$w_{ij} = \exp(-4|x_i - x_j|^2/d_K(x_i)^2),$$

where x_i represents the latent variables for image i , and $d_K(x_i)$ is the distance in the latent space between x_i and its K^{th} nearest neighbor. We used $K = 10$ in all experiments. The weight matrix was then symmetrized by replacing W with $W + W^T$. For Poisson learning, we additionally set $w_{ii} = 0$ for all i . Placing zeros on the diagonal does not change the solution the Poisson learning equation (2.3), but it does accelerate convergence of the iteration in Algorithm 1 by allowing the random walk to propagate faster.

We compare against Laplace learning (2.9) (Zhu et al., 2003), lazy random walks (Zhou & Schölkopf, 2004; Zhou et al., 2004a), multiclass MBO (Garcia-Cardona et al., 2014; Bertozzi & Flenner, 2012), weighted nonlocal Laplacian (WNLL) (Shi et al., 2017), volume constrained MBO (Jacobs et al., 2018), Centered Kernel Method (Mai & Couillet, 2018), sparse label propagation (Jung et al., 2016), and p -Laplace learning (Flores et al., 2019). In the volume constrained MBO method we used exact volume constraints and temperature of $T = 0.1$. In the Centered Kernel Method, we chose α to be 5% larger than the spectral norm of the centered weight matrix. For a baseline reference, we also compared against a nearest neighbor classifier that chooses the label of the closest labeled vertex with respect to the graph geodesic distance. In all experiments, we ran 100 trials randomly choosing which data points are labeled, with the exception of the p -Laplace and sparse label propagation methods, which are slower and were run for 10 trials. The same random label permutations were used for all methods.


Table 1. MNIST: Average accuracy scores over 100 trials with standard deviation in brackets.

# LABELS PER CLASS	1	2	3	4	5
LAPLACE/LP (ZHU ET AL., 2003)	16.1 (6.2)	28.2 (10.3)	42.0 (12.4)	57.8 (12.3)	69.5 (12.2)
NEAREST NEIGHBOR	55.8 (5.1)	65.0 (3.2)	68.9 (3.2)	72.1 (2.8)	74.1 (2.4)
RANDOM WALK (ZHOU & SCHÖLKOPF, 2004)	66.4 (5.3)	76.2 (3.3)	80.0 (2.7)	82.8 (2.3)	84.5 (2.0)
MBO (GARCIA-CARDONA ET AL., 2014)	19.4 (6.2)	29.3 (6.9)	40.2 (7.4)	50.7 (6.0)	59.2 (6.0)
VOLUMEMBO (JACOBS ET AL., 2018)	89.9 (7.3)	95.6 (1.9)	96.2 (1.2)	96.6 (0.6)	96.7 (0.6)
WNLL (SHI ET AL., 2017)	55.8 (15.2)	82.8 (7.6)	90.5 (3.3)	93.6 (1.5)	94.6 (1.1)
CENTERED KERNEL (MAI & COUILLET, 2018)	19.1 (1.9)	24.2 (2.3)	28.8 (3.4)	32.6 (4.1)	35.6 (4.6)
SPARSE LP (JUNG ET AL., 2016)	14.0 (5.5)	14.0 (4.0)	14.5 (4.0)	18.0 (5.9)	16.2 (4.2)
P-LAPLACE (FLORES ET AL., 2019)	72.3 (9.1)	86.5 (3.9)	89.7 (1.6)	90.3 (1.6)	91.9 (1.0)
Poisson	90.2 (4.0)	93.6 (1.6)	94.5 (1.1)	94.9 (0.8)	95.3 (0.7)
PoissonMBO	96.5 (2.6)	97.2 (0.1)	97.2 (0.1)	97.2 (0.1)	97.2 (0.1)


Table 2. FashionMNIST: Average accuracy scores over 100 trials with standard deviation in brackets.

# LABELS PER CLASS	1	2	3	4	5
LAPLACE/LP (ZHU ET AL., 2003)	18.4 (7.3)	32.5 (8.2)	44.0 (8.6)	52.2 (6.2)	57.9 (6.7)
NEAREST NEIGHBOR	44.5 (4.2)	50.8 (3.5)	54.6 (3.0)	56.6 (2.5)	58.3 (2.4)
RANDOM WALK (ZHOU & SCHÖLKOPF, 2004)	49.0 (4.4)	55.6 (3.8)	59.4 (3.0)	61.6 (2.5)	63.4 (2.5)
MBO (GARCIA-CARDONA ET AL., 2014)	15.7 (4.1)	20.1 (4.6)	25.7 (4.9)	30.7 (4.9)	34.8 (4.3)
VOLUMEMBO (JACOBS ET AL., 2018)	54.7 (5.2)	61.7 (4.4)	66.1 (3.3)	68.5 (2.8)	70.1 (2.8)
WNLL (SHI ET AL., 2017)	44.6 (7.1)	59.1 (4.7)	64.7 (3.5)	67.4 (3.3)	70.0 (2.8)
CENTERED KERNEL (MAI & COUILLET, 2018)	11.8 (0.4)	13.1 (0.7)	14.3 (0.8)	15.2 (0.9)	16.3 (1.1)
SPARSE LP (JUNG ET AL., 2016)	14.1 (3.8)	16.5 (2.0)	13.7 (3.3)	13.8 (3.3)	16.1 (2.5)
P-LAPLACE (FLORES ET AL., 2019)	54.6 (4.0)	57.4 (3.8)	65.4 (2.8)	68.0 (2.9)	68.4 (0.5)
Poisson	60.8 (4.6)	66.1 (3.9)	69.6 (2.6)	71.2 (2.2)	72.4 (2.3)
PoissonMBO	62.0 (5.7)	67.2 (4.8)	70.4 (2.9)	72.1 (2.5)	73.1 (2.7)

# LABELS PER CLASS	10	20	40	80	160
LAPLACE/LP (ZHU ET AL., 2003)	70.6 (3.1)	76.5 (1.4)	79.2 (0.7)	80.9 (0.5)	82.3 (0.3)
NEAREST NEIGHBOR	62.9 (1.7)	66.9 (1.1)	70.0 (0.8)	72.5 (0.6)	74.7 (0.4)
RANDOM WALK (ZHOU & SCHÖLKOPF, 2004)	68.2 (1.6)	72.0 (1.0)	75.0 (0.7)	77.4 (0.5)	79.5 (0.3)
MBO (GARCIA-CARDONA ET AL., 2014)	52.7 (4.1)	67.3 (2.0)	75.7 (1.1)	79.6 (0.7)	81.6 (0.4)
VOLUMEMBO (JACOBS ET AL., 2018)	74.4 (1.5)	77.4 (1.0)	79.5 (0.7)	81.0 (0.5)	82.1 (0.3)
WNLL (SHI ET AL., 2017)	74.4 (1.6)	77.6 (1.1)	79.4 (0.6)	80.6 (0.4)	81.5 (0.3)
CENTERED KERNEL (MAI & COUILLET, 2018)	20.6 (1.5)	27.8 (2.3)	37.9 (2.6)	51.3 (3.3)	64.3 (2.6)
SPARSE LP (JUNG ET AL., 2016)	15.2 (2.5)	15.9 (2.0)	14.5 (1.5)	13.8 (1.4)	51.9 (2.1)
P-LAPLACE (FLORES ET AL., 2019)	73.0 (0.9)	76.2 (0.8)	78.0 (0.3)	79.7 (0.5)	80.9 (0.3)
Poisson	75.2 (1.5)	77.3 (1.1)	78.8 (0.7)	79.9 (0.6)	80.7 (0.5)
PoissonMBO	76.1 (1.4)	78.2 (1.1)	79.5 (0.7)	80.7 (0.6)	81.6 (0.5)



(a)



(b)

Figure 2. Accuracy of Poisson Learning for (a) different numbers of neighbors k used to construct the graph and (b) unbalanced training data. In (a) we used 5 labels per class and in (b) we used 1 label per class for the odd numbered classes, and $m = 1, 2, 3, 4, 5$ labels per class for the even numbered classes. Both figures show the difference in accuracy compared to $k = 10$ and balanced training data.

Table 3. Cifar-10: Average accuracy scores over 100 trials with standard deviation in brackets.

# LABELS PER CLASS	1	2	3	4	5
LAPLACE/LP (ZHU ET AL., 2003)	10.5 (1.3)	12.5 (4.4)	13.1 (3.8)	14.5 (4.7)	18.0 (6.9)
NEAREST NEIGHBOR	33.6 (4.4)	37.3 (3.3)	40.3 (3.0)	40.9 (2.7)	42.1 (2.4)
RANDOM WALK (ZHOU & SCHÖLKOPF, 2004)	37.1 (5.0)	42.1 (3.7)	45.8 (3.4)	47.0 (2.8)	48.8 (2.5)
MBO (GARCIA-CARDONA ET AL., 2014)	15.2 (4.1)	20.4 (4.8)	25.9 (4.1)	29.6 (4.3)	34.5 (4.2)
VOLUMEMBO (JACOBS ET AL., 2018)	40.3 (8.0)	47.2 (7.1)	52.2 (5.3)	53.3 (4.7)	55.9 (4.0)
WNLL (SHI ET AL., 2017)	20.8 (6.4)	34.5 (6.2)	42.1 (5.2)	46.1 (4.4)	50.2 (3.5)
CENTERED KERNEL (MAI & COUILLET, 2018)	13.8 (1.1)	15.5 (1.2)	17.3 (1.4)	18.8 (1.7)	20.4 (1.6)
SPARSE LP (JUNG ET AL., 2016)	10.4 (2.1)	11.1 (1.4)	11.8 (2.1)	12.8 (4.4)	13.6 (3.3)
P-LAPLACE (FLORES ET AL., 2019)	28.7 (6.6)	39.8 (6.4)	45.7 (2.6)	46.8 (1.7)	50.4 (2.9)
Poisson	41.6 (5.4)	46.9 (4.2)	51.1 (3.4)	52.5 (3.0)	54.5 (3.0)
PoissonMBO	42.1 (7.0)	49.1 (5.3)	53.8 (4.4)	55.6 (3.7)	57.4 (3.4)

# LABELS PER CLASS	10	20	40	80	160
LAPLACE/LP (ZHU ET AL., 2003)	31.5 (7.6)	49.7 (4.9)	59.6 (2.2)	63.9 (1.0)	66.7 (0.6)
NEAREST NEIGHBOR	45.5 (1.8)	48.7 (1.5)	51.2 (0.8)	53.7 (0.6)	56.1 (0.5)
RANDOM WALK (ZHOU & SCHÖLKOPF, 2004)	53.5 (1.9)	57.6 (1.1)	60.7 (0.7)	64.0 (0.4)	66.4 (0.3)
MBO (GARCIA-CARDONA ET AL., 2014)	47.3 (3.3)	58.6 (2.0)	63.8 (1.1)	66.4 (0.6)	68.5 (0.4)
VOLUMEMBO (JACOBS ET AL., 2018)	59.5 (2.0)	61.4 (0.9)	62.2 (0.7)	63.3 (0.5)	64.2 (0.4)
WNLL (SHI ET AL., 2017)	56.7 (2.3)	61.1 (1.1)	63.7 (0.7)	65.4 (0.4)	66.8 (0.3)
CENTERED KERNEL (MAI & COUILLET, 2018)	26.5 (2.0)	34.8 (2.1)	43.6 (2.2)	51.8 (2.3)	58.7 (1.5)
SPARSE LP (JUNG ET AL., 2016)	16.2 (2.6)	20.3 (1.4)	19.0 (1.1)	19.2 (1.1)	27.1 (1.7)
P-LAPLACE (FLORES ET AL., 2019)	54.6 (2.2)	60.7 (0.8)	63.9 (0.4)	66.0 (0.5)	67.9 (0.3)
Poisson	58.7 (1.8)	61.9 (1.0)	63.7 (0.7)	65.2 (0.6)	66.5 (0.5)
PoissonMBO	61.3 (1.8)	63.8 (1.0)	65.5 (0.6)	66.9 (0.5)	68.3 (0.4)

Tables 1, 2 and 3 show the average accuracy and standard deviation over all 100 trials for various low label rates. We also ran experiments at higher label rates on FashionMNIST and Cifar-10, which are reported in the lower half of their respective tables. We mention that in Tables 1, 2 and 3 the training data is balanced, so $\bar{y} = \frac{1}{10}\mathbf{1}$. Thus, the label decisions (2.4) and (2.2) are equivalent.

We see that in nearly all cases, PoissonMBO outperforms all other methods, with PoissonMBO typically outperforming Poisson learning by a few percentage points. The most drastic improvements are seen at the ultra low label rates, and at the moderate label rates shown in Tables 2 and 3, several other methods perform well. We note that VolumeMBO and PoissonMBO are the only methods that incorporate prior knowledge of class sizes, and are most suitable for direct comparison. Our results can be compared to the clustering results of 67.2% on FashionMNIST (McConville et al., 2019) and 41.2% on Cifar-10 (Ghasedi et al., 2019).

Figure 2(a) shows the accuracy of Poisson learning at 5 labels per class as a function of the number of neighbors K used in constructing the graph, showing that the algorithm is not particularly sensitive to this. Figure 2(b) shows the accuracy of Poisson learning for unbalanced training data. We take 1 label per class for half the classes and $m = 1, 2, 3, 4, 5$ labels per class for the other half. Since the training data is unbalanced, \bar{y} is not a constant vector and the label decision in Step 11 of Algorithm 1 (Poisson Learning) compensates for unbalanced training data. Note that in

Figure 2 we plot the difference in accuracy compared to (a) the baseline of $k = 10$ and (b) 1 label per class. In Figure 2 (b), we see an increase in accuracy when only half the classes get additional labels, though the increase is not as large as in Tables 1, 2 and 3 where all classes get additional labels.

5. Conclusion

We proposed a new framework for graph-based semi-supervised learning at very low label rates called *Poisson learning*. The method is efficient and simple to implement. We performed a detailed analysis of Poisson learning, giving random walk and variational interpretations. We also proposed a graph-cut enhancement of Poisson learning, called Poisson MBO, that can give further improvements. We presented numerical results showing that Poisson Learning outperforms all other methods for semi-supervised learning at low label rates on several common datasets.

Acknowledgements: Calder was supported by NSF-DMS grants 1713691,1944925, Cook was supported by a University of Minnesota Grant In Aid award, Thorpe was supported by the European Research Council under the European Union’s Horizon 2020 research and innovation programme grant No 777826 (NoMADS) and 647812, and Slepčev was supported by NSF-DMS grant 1814991. The authors thank Matt Jacobs for providing code for volume constrained MBO. The authors also thank the Center for Nonlinear Analysis (CNA) at CMU for its support.

References

- Ando, R. K. and Zhang, T. Learning on graph with Laplacian regularization. In *Advances in Neural Information Processing Systems*, pp. 25–32, 2007.
- Belkin, M. and Niyogi, P. Using manifold structure for partially labelled classification. In *Advances in Neural Information Processing Systems*, 2002.
- Bertozzi, A. L. and Flenner, A. Diffuse interface models on graphs for classification of high dimensional data. *Multi-scale Modeling & Simulation*, 10(3):1090–1118, 2012.
- Calder, J. The game theoretic p-Laplacian and semi-supervised learning with few labels. *Nonlinearity*, 32(1), 2018.
- Calder, J. Consistency of Lipschitz learning with infinite unlabeled data and finite labeled data. *SIAM Journal on Mathematics of Data Science*, 1:780–812, 2019.
- Calder, J. and Slepčev, D. Properly-Weighted graph Laplacian for semi-supervised learning. *Applied Mathematics & Optimization*, Dec 2019.
- DiBenedetto, E. *Degenerate parabolic equations*. Springer Science & Business Media, 1993.
- El Alaoui, A., Cheng, X., Ramdas, A., Wainwright, M. J., and Jordan, M. I. Asymptotic behavior of ℓ_p -based Laplacian regularization in semi-supervised learning. In *Conference on Learning Theory*, pp. 879–906, 2016.
- Flores, M., Calder, J., and Lerman, G. Algorithms for L_p -based semi-supervised learning on graphs. *arXiv:1901.05031*, 2019.
- Garcia-Cardona, C., Merkurjev, E., Bertozzi, A. L., Flenner, A., and Percus, A. G. Multiclass data segmentation using diffuse interface methods on graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(8):1600–1613, 2014.
- Ghasedi, K., Wang, X., Deng, C., and Huang, H. Balanced self-paced learning for generative adversarial clustering network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4391–4400, 2019.
- He, J., Li, M., Zhang, H.-J., Tong, H., and Zhang, C. Manifold-ranking based image retrieval. In *Proceedings of the 12th annual ACM International Conference on Multimedia*, pp. 9–16. ACM, 2004.
- Hein, M., Audibert, J.-Y., and Luxburg, U. v. Graph Laplacians and their convergence on random neighborhood graphs. *Journal of Machine Learning Research*, 8(Jun): 1325–1368, 2007.
- Jacobs, M., Merkurjev, E., and Esedođlu, S. Auction dynamics: A volume constrained MBO scheme. *Journal of Computational Physics*, 354:288–310, 2018.
- Jung, A., Hero III, A. O., Mara, A., and Jahromi, S. Semi-supervised learning via sparse label propagation. *arXiv preprint arXiv:1612.01414*, 2016.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, 2014.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Kyng, R., Rao, A., Sachdeva, S., and Spielman, D. A. Algorithms for Lipschitz learning on graphs. In *Conference on Learning Theory*, pp. 1190–1223, 2015.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Mai, X. and Couillet, R. Random matrix-inspired improved semi-supervised learning on graphs. In *International Conference on Machine Learning*, 2018.
- McConville, R., Santos-Rodriguez, R., Piechocki, R. J., and Craddock, I. N2d:(not too) deep clustering via clustering the local manifold of an autoencoded embedding. *arXiv preprint arXiv:1908.05968*, 2019.
- Nadler, B., Srebro, N., and Zhou, X. Semi-supervised learning with the graph Laplacian: The limit of infinite unlabeled data. *Advances in Neural Information Processing Systems*, 22:1330–1338, 2009.
- Shi, Z., Osher, S., and Zhu, W. Weighted nonlocal Laplacian on interpolation from sparse data. *Journal of Scientific Computing*, 73(2-3):1164–1177, 2017.
- Slepčev, D. and Thorpe, M. Analysis of p-Laplacian regularization in semisupervised learning. *SIAM Journal on Mathematical Analysis*, 51(3):2085–2120, 2019.
- Van Gennip, Y. and Bertozzi, A. L. Gamma-convergence of graph Ginzburg-Landau functionals. *Advances in Differential Equations*, 17(11/12):1115–1180, 2012.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Xu, B., Bu, J., Chen, C., Cai, D., He, X., Liu, W., and Luo, J. Efficient manifold ranking for image retrieval. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 525–534. ACM, 2011.

- Yang, C., Zhang, L., Lu, H., Ruan, X., and Yang, M.-H. Saliency detection via graph-based manifold ranking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3166–3173, 2013.
- Zhang, L., Qi, G.-J., Wang, L., and Luo, J. Aet vs. aed: Unsupervised representation learning by auto-encoding transformations rather than data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2547–2555, 2019.
- Zhou, D. and Schölkopf, B. Learning from labeled and unlabeled data using random walks. In *Joint Pattern Recognition Symposium*, pp. 237–244. Springer, 2004.
- Zhou, D., Bousquet, O., Lal, T. N., Weston, J., and Schölkopf, B. Learning with local and global consistency. In *Advances in Neural Information Processing Systems*, pp. 321–328, 2004a.
- Zhou, D., Weston, J., Gretton, A., Bousquet, O., and Schölkopf, B. Ranking on data manifolds. In *Advances in Neural Information Processing Systems*, pp. 169–176, 2004b.
- Zhou, D., Huang, J., and Schölkopf, B. Learning from labeled and unlabeled data on a directed graph. In *Proceedings of the 22nd International Conference on Machine Learning*, pp. 1036–1043. ACM, 2005.
- Zhou, X. and Belkin, M. Semi-supervised learning by higher order regularization. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 892–900, 2011.
- Zhu, X., Ghahramani, Z., and Lafferty, J. D. Semi-supervised learning using Gaussian fields and harmonic functions. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 912–919, 2003.
- Zhu, X., Lafferty, J., and Rosenfeld, R. *Semi-supervised learning with graphs*. PhD thesis, Carnegie Mellon University, language technologies institute, school of ..., 2005.

Appendices

A. Proofs

We provide the proofs by section.

A.1. Proofs for Section 2.1

We recall $X_0^x, X_1^x, X_2^x, \dots$ is a random walk on X starting at $X_0^x = x$ with transition probabilities

$$\mathbb{P}(X_k^x = x_j \mid X_{k-1}^x = x_i) = \frac{w_{ij}}{d_i}.$$

Before giving the proof of Theorem 2.1, we recall some properties of random walks and Markov chains. The random walk described above induces a Markov chain with state space X . Since the graph is connected and X is finite, the Markov chain is *positive recurrent*. We also assume the Markov chain is *aperiodic*. This implies the distribution of the random walker converges to the invariant distribution of the Markov chain as $k \rightarrow \infty$. In particular, choose any initial distribution $p_0 \in \ell^2(X)$ such that $\sum_{i=1}^n p_0(x_i) = 1$ and $p_0 \geq 0$, and define

$$p_{k+1}(x_i) = \sum_{j=1}^n \frac{w_{ij}}{d_j} p_k(x_j). \quad (\text{A.1})$$

Then p_k is the distribution of the random walker after k steps. Since the Markov chain is positive recurrent and aperiodic we have that

$$\lim_{k \rightarrow \infty} p_k(x_i) = \pi(x_i)$$

for all i , where

$$\pi(x_i) = \frac{d_i}{\sum_{i=1}^n d_i}$$

is the invariant distribution of the Markov chain. It is simple to check that if $p_0 \in \ell^2(X)$ is any function (i.e., not necessarily a probability distribution), and we define p_k by the iteration (A.1), then

$$\lim_{k \rightarrow \infty} p_k(x_i) = \pi(x_i) \sum_{j=1}^n p_0(x_j). \quad (\text{A.2})$$

We now give the proof of Theorem 2.1.

Proof of Theorem 2.1. Define the normalized Green's function

$$G_T(x_i, x_j) = \frac{1}{d_i} \mathbb{E} \left[\sum_{k=0}^T \mathbb{1}_{\{X_k^{x_j} = x_i\}} \right] = \frac{1}{d_i} \sum_{k=0}^T \mathbb{P}(X_k^{x_j} = x_i).$$

Then we have

$$\begin{aligned} d_i G_T(x_i, x_j) &= \delta_{ij} + \sum_{k=1}^T \sum_{\ell=1}^n \frac{w_{\ell i}}{d_\ell} \mathbb{P}(X_{k-1}^{x_j} = x_\ell) \\ &= \delta_{ij} + \sum_{\ell=1}^n \frac{w_{\ell i}}{d_\ell} \sum_{k=1}^T \mathbb{P}(X_{k-1}^{x_j} = x_\ell) \\ &= \delta_{ij} + \sum_{\ell=1}^n \frac{w_{\ell i}}{d_\ell} \sum_{k=0}^{T-1} \mathbb{P}(X_k^{x_j} = x_\ell) \\ &= \delta_{ij} + \sum_{\ell=1}^n w_{\ell i} G_{T-1}(x_\ell, x_j). \end{aligned}$$

Therefore we have

$$d_i(G_T(x_i, x_j) - G_{T-1}(x_i, x_j)) + \mathcal{L}G_{T-1}(x_i, x_j) = \delta_{ij},$$

where the Laplacian \mathcal{L} is applied to the first variable of G_{T-1} while the second variable is fixed (i.e. $\mathcal{L}G_{T-1}(x_i, x_j) = [\mathcal{L}G_{T-1}(\cdot, x_j)]_{x_i}$). Since

$$u_T(x_i) = \sum_{j=1}^m (y_j - \bar{y}) G_T(x_i, x_j)$$

we have

$$d_i(u_T(x_i) - u_{T-1}(x_i)) + \mathcal{L}u_{T-1}(x_i) = \sum_{j=1}^m (y_j - \bar{y}) \delta_{ij}.$$

Summing both sides over $i = 1, \dots, n$ we find that

$$(u_T)_{d,X} = \sum_{i=1}^n d_i u_T(x_i) = \sum_{i=1}^n d_i u_{T-1}(x_i) = (u_{T-1})_{d,X},$$

where $d = (d_1, d_2, \dots, d_n)$ is the vector of degrees. Therefore $(u_T)_{d,X} = (u_{T-1})_{d,X} = \dots = (u_0)_{d,X}$. Noting that

$$d_i u_0(x_i) = \sum_{j=1}^m (y_j - \bar{y}) \delta_{ij},$$

we have $(u_0)_{d,X} = 0$, and so $(u_T)_{d,X} = 0$ for all $T \geq 0$. Let $u \in \ell^2(X)$ be the solution of

$$\mathcal{L}u(x_i) = \sum_{j=1}^m (y_j - \bar{y}) \delta_{ij}$$

satisfying $(u)_{d,X} = 0$. Define $v_T(x_i) = d_i(u_T(x_i) - u(x_i))$. We then check that v_T satisfies

$$v_T(x_i) = \sum_{j=1}^n \frac{w_{ij}}{d_j} v_{T-1}(x_j),$$

and $(v_T)_X = 0$. Since the random walk is aperiodic and the graph is connected, we have by (A.2) that $\lim_{T \rightarrow \infty} v_T(x_i) = \pi(x_i)(v_0)_X = 0$, which completes the proof. \square

A.2. Proofs for Section 2.2

We first review some additional calculus on graphs. The *graph divergence* of a vector field V is defined as

$$\operatorname{div} V(x_i) = \sum_{j=1}^n w_{ij} V(x_i, x_j).$$

The divergence is the negative adjoint of the gradient; that is, for every vector field $V \in \ell^2(X^2)$ and function $u \in \ell^2(X)$

$$(\nabla u, V)_{\ell^2(X^2)} = -(u, \operatorname{div} V)_{\ell^2(X)}. \quad (\text{A.3})$$

We also define $\|u\|_{\ell^p(X)}^p = \sum_{i=1}^n |u(x_i)|^p$ and

$$\|V\|_{\ell^p(X^2)}^p = \frac{1}{2} \sum_{i,j=1}^n w_{ij} |V(x_i, x_j)|^p,$$

where $|\cdot|$ is the Euclidean norm on \mathbb{R}^k .

The graph Laplacian $\mathcal{L}u$ of a function $u \in \ell^2(X)$ is defined as negative of the composition of gradient and divergence

$$\mathcal{L}u(x_i) = -\operatorname{div}(\nabla u)(x_i) = \sum_{j=1}^n w_{ij} (u(x_i) - u(x_j)).$$

The operator \mathcal{L} is the *unnormalized* graph Laplacian. Using (A.3) we have

$$(\mathcal{L}u, v)_{\ell^2(X)} = (-\operatorname{div} \nabla u, v)_{\ell^2(X)} = (\nabla u, \nabla v)_{\ell^2(X^2)}.$$

In particular $(\mathcal{L}u, v)_{\ell^2(X)} = (u, \mathcal{L}v)_{\ell^2(X)}$, and so the graph Laplacian \mathcal{L} is self-adjoint as an operator $\mathcal{L} : \ell^2(X) \rightarrow \ell^2(X)$. We also note that

$$(\mathcal{L}u, u)_{\ell^2(X)} = (\nabla u, \nabla u)_{\ell^2(X^2)} = \|\nabla u\|_{\ell^2(X^2)}^2,$$

that is, \mathcal{L} is positive semi-definite.

The variational interpretation of Poisson learning can be directly extended to ℓ^p versions, so we proceed in generality here. For a function $u : X \rightarrow \mathbb{R}^k$ and a positive vector $a \in \mathbb{R}^n$ (meaning $a_i > 0$ for all $i = 1, \dots, n$) we define the weighted mean value

$$(u)_{a, X} := \frac{1}{\sum_{i=1}^n a_i} \sum_{i=1}^n a_i u(x_i).$$

We define the space of weighted mean-zero functions

$$\ell_{a,0}^p(X) = \{u \in \ell^p(X) : (u)_{a, X} = 0\}.$$

For $p \geq 1$ and $\mu > 0$ we consider the variational problem

$$\min_{u \in \ell_{a,0}^p(X)} \left\{ \frac{1}{p} \|\nabla u\|_{\ell^p(X^2)}^p - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j) \right\} \quad (\text{A.4})$$

where $\bar{y} = \frac{1}{m} \sum_{j=1}^m y_j$. This generalizes the variational problem (2.8) for Poisson learning, and the theorem below generalizes Theorem 2.3.

Theorem A.1. *Assume G is connected. For any $p > 1$, positive $a \in \mathbb{R}^n$, and $\mu \geq 0$, there exists a unique solution $u \in \ell_{a,0}^p(X)$ of (A.4). Furthermore, the minimizer u satisfies the graph p -Laplace equation*

$$-\operatorname{div} (|\nabla u|^{p-2} \nabla u)(x_i) = \mu \sum_{j=1}^m (y_j - \bar{y}) \delta_{ij}. \quad (\text{A.5})$$

We give the proof of Theorem A.1 below, after some remarks and other results.

Remark A.2. When $p = 1$, solutions of (A.4) may not exist for all $\mu \geq 0$, since the variational problem (A.4) may not be bounded from below. We can show that there exists $C > 0$ such that if $\mu < C$, the variational problem is bounded from below and our argument for existence in Theorem A.1 goes through.

It turns out that $\mu > 0$ is a redundant parameter when $p > 1$.

Lemma A.3. *Let $p > 1$ and for $\mu > 0$ let u_μ be the solution of (A.4). Then, $u_\mu = \mu^{1/(p-1)} u_1$.*

It follows from Lemma A.3 that when $p > 1$, the fidelity parameter $\mu > 0$ is *completely irrelevant* for classification problems, since the identity $u_\mu = \mu^{1/(p-1)} u_1$ implies that the label decision (2.2) gives the same labeling for *any* value of $\mu > 0$. Hence, in Poisson learning with $p > 1$ we always take $\mu = 1$. This remark is false for $p = 1$.

Before proving Theorem A.1 we first record a Poincaré inequality. The proof is standard but we include it for completeness. We can prove the Poincaré inequality for non-negative vectors $a \in \mathbb{R}^n$, meaning that $a_i \geq 0$ for every $i = 1, \dots, n$ as long as $\sum_{i=1}^n a_i > 0$.

Proposition A.4. *Assume G is connected, $a \in \mathbb{R}^d$ is non-negative with $\sum_{i=1}^n a_i > 0$, and $p \geq 1$. There exists $\lambda_p > 0$ such that*

$$\lambda_p \|u - (u)_{a, X}\|_{\ell^p(X)} \leq \|\nabla u\|_{\ell^p(X^2)}, \quad (\text{A.6})$$

for all $u \in \ell^p(X)$.

Proof. Define

$$\lambda_p = \min_{\substack{u \in \ell^p(X) \\ u \neq (u)_{a,X}}} \frac{\|\nabla u\|_{\ell^p(X^2)}}{\|u - (u)_{a,X}\|_{\ell^p(X)}}.$$

Then clearly (A.6) holds for this choice of λ_p , and $\lambda_p \geq 0$. If $\lambda_p = 0$, then there exists a sequence $u_k \in \ell^p(X)$ with $u_k \neq (u_k)_{a,X}$ such that

$$\frac{\|\nabla u_k\|_{\ell^p(X^2)}}{\|u_k - (u_k)_{a,X}\|_{\ell^p(X)}} \leq \frac{1}{k}.$$

We may assume that $(u_k)_{a,X} = 0$ and $\|u_k\|_{\ell^p(X)} = 1$, and so

$$\|\nabla u_k\|_{\ell^p(X^2)} \leq \frac{1}{k}. \quad (\text{A.7})$$

Since $|u_k(x)| \leq \|u_k\|_{\ell^p(X)} = 1$, the sequence u_k is uniformly bounded and by the Bolzano-Weierstrauss Theorem there exists a subsequence u_{k_j} such that $u_{k_j}(x_i)$ is a convergent sequence in \mathbb{R}^k for every i . We denote $u(x_i) = \lim_{j \rightarrow \infty} u_{k_j}(x_i)$. Since $\|u_{k_j}\|_{\ell^p(X)} = 1$ we have $\|u\|_{\ell^p(X)} = 1$, and thus $u \neq 0$. Similarly, since $(u_k)_{a,X} = 0$ we have $(u)_{a,X} = 0$ as well. On the other hand it follows from (A.7) that $\|\nabla u\|_{\ell^p(X^2)} = 0$, and so

$$w_{ij}(u(x_i) - u(x_j)) = 0 \quad \text{for all } i, j.$$

It follows that $u(x_i) = u(x_j)$ whenever $w_{ij} > 0$. Since the graph is connected, it follows that u is constant. Since $(u)_{a,X} = 0$ we must have $u \equiv 0$, which is a contradiction, since $\|u\|_{\ell^p(X)} = 1$. Therefore $\lambda_p > 0$, which completes the proof. \square

We can now prove Theorem A.1.

Proof of Theorem A.1. Let us write

$$I_p(u) = \frac{1}{p} \|\nabla u\|_{\ell^p(X^2)}^p - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j). \quad (\text{A.8})$$

By Proposition A.4 we have

$$I_p(u) \geq \frac{1}{p} \lambda_p^p \|u\|_{\ell^p(X)}^p - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j)$$

for $u \in \ell_{a,0}^p(X)$. By Hölder's inequality we have

$$\begin{aligned} \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j) &\leq \sum_{j=1}^m |y_j - \bar{y}| |u(x_j)| \\ &\leq \left(\sum_{j=1}^m |y_j - \bar{y}|^q \right)^{1/q} \left(\sum_{j=1}^m |u(x_j)|^p \right)^{1/p} \\ &\leq \left(\sum_{j=1}^m |y_j - \bar{y}|^q \right)^{1/q} \|u\|_{\ell^p(X)}, \end{aligned}$$

where $q = p/(p-1)$. Letting $C = \left(\sum_{j=1}^m |y_j - \bar{y}|^q \right)^{1/q}$ we have

$$I_p(u) \geq \frac{1}{p} \lambda_p^p \|u\|_{\ell^p(X)}^p - C \mu \|u\|_{\ell^p(X)}. \quad (\text{A.9})$$

Since $p > 1$, we see that I_p is bounded below.

Let $u_k \in \ell_{a,0}^p(X)$ be a minimizing sequence, that is, we take u_k so that

$$-\infty < \inf_{u \in \ell_{a,0}^p(X)} I_p(u) = \lim_{k \rightarrow \infty} I_p(u_k).$$

By (A.9) we have that

$$\frac{1}{p} \lambda_p^p \|u_k\|_{\ell^p(X)}^p - C\mu \|u_k\|_{\ell^p(X)} \leq \inf_{u \in \ell_{a,0}^p(X)} I_p(u) + 1,$$

for k sufficiently large. Since $p > 1$, it follows that there exists $M > 0$ such that $\|u_k\|_{\ell^p(X)} \leq M$ for all $k \geq 1$. Since $|u_k(x_i)| \leq \|u_k\|_{\ell^p(X)} \leq M$ for all $i = 1, \dots, n$, we can apply the Bolzano-Weierstrauss Theorem to extract a subsequence u_{k_j} such that $u_{k_j}(x_i)$ is a convergent sequence in \mathbb{R}^k for all $i = 1, \dots, n$. We denote by $u^*(x_i)$ the limit of $u_{k_j}(x_i)$ for all i . By continuity of I_p we have

$$\inf_{u \in \ell_{a,0}^p(X)} I_p(u) = \lim_{j \rightarrow \infty} I_p(u_{k_j}) = I_p(u^*),$$

and $(u^*)_{a,X} = 0$. This shows that there exists a solution of (A.4).

We now show that any solution of (A.4) satisfies $-\operatorname{div} (|\nabla u|^{p-2} \nabla u) = \mu f$. The proof follows from taking a variation. Let $v \in \ell_{a,0}^p(X)$ and consider $g(t) := I_p(u + tv)$, where I_p is defined in (A.8). Then g has a minimum at $t = 0$ and hence $g'(0) = 0$. We now compute

$$\begin{aligned} g'(0) &= \frac{d}{dt} \Big|_{t=0} \left\{ \frac{1}{p} \|\nabla u + t \nabla v\|_{\ell^p(X^2)}^p - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot (u(x_j) + tv(x_j)) \right\} \\ &= \frac{1}{2p} \sum_{i,j=1}^n w_{ij} \frac{d}{dt} \Big|_{t=0} |\nabla u(x_i, x_j) + t \nabla v(x_i, x_j)|^p - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot v(x_j) \\ &= \frac{1}{2} \sum_{i,j=1}^n w_{ij} |\nabla u(x_i, x_j)|^{p-2} \nabla u(x_i, x_j) \cdot \nabla v(x_i, x_j) - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot v(x_j) \\ &= (|\nabla u|^{p-2} \nabla u, \nabla v)_{\ell^2(X^2)} - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot v(x_j) \\ &= (-\operatorname{div} (|\nabla u|^{p-2} \nabla u), v)_{\ell^2(X)} - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot v(x_j) \\ &= (-\operatorname{div} (|\nabla u|^{p-2} \nabla u) - \mu f, v)_{\ell^2(X)}, \end{aligned}$$

where

$$f(x_i) = \sum_{j=1}^m (y_j - \bar{y}) \delta_{ij}.$$

We choose

$$v(x_i) = \frac{1}{a_i} (-\operatorname{div} (|\nabla u|^{p-2} \nabla u) (x_i) - \mu f(x_i))$$

then

$$(v)_{a,X} = \sum_{i=1}^n (-\operatorname{div} (|\nabla u|^{p-2} \nabla u) (x_i) - \mu f(x_i)) = 0$$

so $v \in \ell_{a,0}^p(X)$. Moreover, for this choice of v ,

$$0 = g'(0) = \sum_{i=1}^n \frac{1}{a_i} |\operatorname{div} (|\nabla u|^{p-2} \nabla u) (x_i) + \mu f(x_i)|^2 \geq \frac{1}{\max a_i} \|\operatorname{div} (|\nabla u|^{p-2} \nabla u) (x_i) + \mu f(x_i)\|_{\ell^2(X)}^2.$$

So, $-\operatorname{div} (|\nabla u|^{p-2} \nabla u) = \mu f$ as required.

To prove uniqueness, let $u, v \in \ell_{a,0}^p(X)$ be minimizers of (A.4). Then u and v satisfy (A.5) which we write as

$$-\operatorname{div}(|\nabla u|^{p-2}\nabla u) = \mu f.$$

Applying Lemma A.5 (below) we find that $\|u - v\|_{\ell^p(X)} = 0$ and so $u = v$. □

In the above proof we used a quantitative error estimate which is of interest in its own right. The estimate was on equations of the form

$$-\operatorname{div}(|\nabla u|^{p-2}\nabla u) = f$$

when $f \in \ell_0^p(X)$, where we use the notation: if $a \in \mathbb{R}^n$ is a constant vector (without loss of generality the vector of ones) then we write $(u)_X = (u)_{a,X} = \frac{1}{n} \sum_{i=1}^n u(x_i)$ and $\ell_0^p(X) = \{u \in \ell^p(X) : (u)_X = 0\}$.

Lemma A.5. *Let $p > 1$, $a \in \mathbb{R}^n$ be non-negative, and assume $u, v \in \ell_{a,0}^p(X)$ satisfy*

$$-\operatorname{div}(|\nabla u|^{p-2}\nabla u)(x_i) = f(x_i)$$

and

$$-\operatorname{div}(|\nabla v|^{p-2}\nabla v)(x_i) = g(x_i)$$

for all $i = 1, \dots, n$, where $f, g \in \ell_0^p(X)$. Then,

$$\|u - v\|_{\ell^p(X)} \leq \begin{cases} C\lambda_p^{-q}\|f - g\|_{\ell^q(X)}^{1/(p-1)} & \text{if } p \geq 2 \\ C\lambda_p^{-2}(\|\nabla u\|_{\ell^p(X)} + \|\nabla v\|_{\ell^p(X)})^{2-p}\|f - g\|_{\ell^2(X)} & \text{if } 1 < p < 2 \end{cases}$$

where C is a constant depending only on p and $q = \frac{p}{p-1}$.

Remark A.6. If $-\operatorname{div}(|\nabla u|^{p-2}\nabla u) = f$ then we can write $(|\nabla u|^{p-2}\nabla u, \nabla \varphi)_{\ell^2(X^2)} = (f, \varphi)_{\ell^2(X)}$ for any $\varphi \in \ell^2(X)$. Choosing $\varphi = u$ implies $\|\nabla u\|_{\ell^p(X^2)}^p = (f, u)_{\ell^2(X)} \leq \|f\|_{\ell^q(X)}\|u\|_{\ell^p(X)}$ so we could write the bound for $p \in (1, 2)$ in the above lemma without $\|\nabla u\|_{\ell^p(X)}$ and $\|\nabla v\|_{\ell^p(X)}$ on the right hand side.

Proof. For $p \geq 2$ we use the identity

$$|a - b|^p \leq C(|a|^{p-2}a - |b|^{p-2}b) \cdot (a - b)$$

for vectors $a, b \in \mathbb{R}^k$ for some constant C depending only on p (which can be found in Lemma 4.4 Chapter I (DiBenedetto, 1993)) to obtain

$$\begin{aligned} \|\nabla u - \nabla v\|_{\ell^p(X^2)}^p &= \frac{1}{2} \sum_{i,j=1}^n w_{ij} |\nabla u(x_i, x_j) - \nabla v(x_i, x_j)|^p \\ &\leq C \sum_{i,j=1}^n w_{ij} (|\nabla u(x_i, x_j)|^{p-2}\nabla u(x_i, x_j) - |\nabla v(x_i, x_j)|^{p-2}\nabla v(x_i, x_j)) \cdot (\nabla u(x_i, x_j) - \nabla v(x_i, x_j)) \\ &= C(|\nabla u|^{p-2}\nabla u - |\nabla v|^{p-2}\nabla v, \nabla(u - v))_{\ell^2(X^2)} \\ &= C(-\operatorname{div}(|\nabla u|^{p-2}\nabla u) + \operatorname{div}(|\nabla v|^{p-2}\nabla v), u - v)_{\ell^2(X)} \\ &= C(f - g, u - v)_{\ell^2(X)} \\ &\leq C\|f - g\|_{\ell^q(X)}\|u - v\|_{\ell^p(X)}, \end{aligned}$$

where in the last line we used Hölder's inequality, $\frac{1}{p} + \frac{1}{q} = 1$, and the value of C may change from line-to-line. By Proposition A.4 we have

$$\lambda_p^p \|u - v\|_{\ell^p(X)}^p \leq \|\nabla u - \nabla v\|_{\ell^p(X^2)}^p \leq C\|f - g\|_{\ell^q(X)}\|u - v\|_{\ell^p(X)}.$$

Therefore we deduce

$$\|u - v\|_{\ell^p(X)} \leq C\lambda_p^{-q}\|f - g\|_{\ell^q(X)}^{1/(p-1)}.$$

Now for $1 < p < 2$ we follow the proof of Lemma 4.4 in Chapter I (DiBenedetto, 1993) to infer

$$\begin{aligned} (|a|^{p-2}a - |b|^{p-2}b) \cdot (a - b) &= \int_0^1 |sa + (1-s)b|^{p-2} |a - b|^2 ds \\ &\quad + (p-2) \int_0^1 |sa + (1-s)b|^{p-4} |(sa + (1-s)b) \cdot (a - b)|^2 ds \end{aligned}$$

for any $a, b \in \mathbb{R}^k$. Hence, by the Cauchy Schwarz inequality,

$$\begin{aligned} (|a|^{p-2}a - |b|^{p-2}b) \cdot (a - b) &\geq (p-1) \int_0^1 |sa + (1-s)b|^{p-2} |a - b|^2 ds \\ &\geq (p-1) |a - b|^2 \int_0^1 \frac{1}{(s|a| + (1-s)|b|)^{2-p}} ds \\ &\geq \frac{(p-1) |a - b|^2}{(|a| + |b|)^{2-p}}. \end{aligned}$$

In the sequel we make use of the inequality

$$(|a|^{p-2}a - |b|^{p-2}b) \cdot (a - b) \geq \frac{C |a - b|^2}{(|a| + |b|)^{2-p}}.$$

By Hölder's inequality and the above inequality we have (where again the constant C may change from line-to-line)

$$\begin{aligned} \|\nabla u - \nabla v\|_{\ell^p(X^2)}^p &= \frac{1}{2} \sum_{i,j=1}^n w_{ij} |\nabla u(x_i, x_j) - \nabla v(x_i, x_j)|^p \\ &\leq \left(\frac{1}{2} \sum_{i,j=1}^n \frac{w_{ij} |\nabla u(x_i, x_j) - \nabla v(x_i, x_j)|^2}{(|\nabla u(x_i, x_j)| + |\nabla v(x_i, x_j)|)^{2-p}} \right)^{\frac{p}{2}} \left(\frac{1}{2} \sum_{i,j=1}^n w_{ij} (|\nabla u(x_i, x_j)| + |\nabla v(x_i, x_j)|)^p \right)^{\frac{2-p}{2}} \\ &\leq C \left(\sum_{i,j=1}^n w_{ij} (|\nabla u(x_i, x_j)|^{p-2} \nabla u(x_i, x_j) - |\nabla v(x_i, x_j)|^{p-2} \nabla v(x_i, x_j)) \cdot (\nabla u(x_i, x_j) - \nabla v(x_i, x_j)) \right)^{\frac{p}{2}} \\ &\quad \times (\|\nabla u\|_{\ell^p(X^2)} + \|\nabla v\|_{\ell^p(X^2)})^{\frac{(2-p)p}{2}} \\ &= C (|\nabla u|^{p-2} \nabla u - |\nabla v|^{p-2} \nabla v, \nabla(u - v))_{\ell^2(X^2)}^{\frac{p}{2}} (\|\nabla u\|_{\ell^p(X^2)} + \|\nabla v\|_{\ell^p(X^2)})^{\frac{(2-p)p}{2}} \\ &= C (-\operatorname{div}(|\nabla u|^{p-2} \nabla u) + \operatorname{div}(|\nabla v|^{p-2} \nabla v), u - v)_{\ell^2(X)}^{\frac{p}{2}} (\|\nabla u\|_{\ell^p(X^2)} + \|\nabla v\|_{\ell^p(X^2)})^{\frac{(2-p)p}{2}} \\ &= C (f - g, u - v)_{\ell^2(X)}^{\frac{p}{2}} (\|\nabla u\|_{\ell^p(X^2)} + \|\nabla v\|_{\ell^p(X^2)})^{\frac{(2-p)p}{2}} \\ &\leq C \|f - g\|_{\ell^2(X)}^{\frac{p}{2}} \|u - v\|_{\ell^2(X)}^{\frac{p}{2}} (\|\nabla u\|_{\ell^p(X^2)} + \|\nabla v\|_{\ell^p(X^2)})^{\frac{(2-p)p}{2}}. \end{aligned}$$

Combining the above with Proposition A.4 we have

$$\lambda_p^p \|u - v\|_{\ell^p(X)}^{\frac{p}{2}} \leq C \|f - g\|_{\ell^2(X)}^{\frac{p}{2}} (\|\nabla u\|_{\ell^p(X^2)} + \|\nabla v\|_{\ell^p(X^2)})^{\frac{(2-p)p}{2}}$$

which implies the result. □

The final proof from Section 2.2 is Lemma A.3.

Proof of Lemma A.3. Let us write

$$I_{p,\mu}(u) = \frac{1}{p} \|\nabla u\|_{\ell^p(X^2)}^p - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j).$$

We note that

$$I_{p,\mu}(\mu^{1/(p-1)}u) = \mu^{p/(p-1)}I_{p,1}(u).$$

Therefore

$$I_{p,\mu}(u_\mu) = \mu^{p/(p-1)}I_{p,1}(u_\mu\mu^{-1/(p-1)}) \geq \mu^{p/(p-1)}I_{p,1}(u_1).$$

On the other hand

$$\mu^{p/(p-1)}I_{p,1}(u_1) = I_{p,\mu}(\mu^{1/(p-1)}u_1) \geq I_{p,\mu}(u_\mu)$$

Therefore

$$I_{p,\mu}(\mu^{1/(p-1)}u_1) = I_{p,\mu}(u_\mu).$$

By uniqueness in Theorem A.1 we have $u_\mu = \mu^{1/(p-1)}u_1$, which completes the proof. \square

A.3. Proofs for Section 2.4

We now turn our attention to the Ginzburg–Landau approximation of the graph cut problem (2.11).

Proof of Theorem 2.4. Let us redefine GL_τ in a more general form,

$$\text{GL}_\tau(u) = \frac{1}{2}\|\nabla u\|_{\ell^2(X^2)}^2 + \frac{1}{\tau}\sum_{i=1}^n V(u(x_i))$$

where $V : \mathbb{R}^k \rightarrow [0, +\infty)$ is continuous and $V(t) = 0$ if and only if $t \in S_k$. Of course, the choice of $V(t) = \prod_{j=1}^k |t - \mathbf{e}_j|^2$ satisfies these assumptions. We let

$$\mathcal{E}_\tau(u) = \begin{cases} \text{GL}_\tau(u) - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j) & \text{if } (u)_X = b \\ +\infty & \text{else,} \end{cases}$$

$$\mathcal{E}_0(u) = \begin{cases} \frac{1}{2}\|\nabla u\|_{\ell^2(X^2)}^2 - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j) & \text{if } (u)_X = b \text{ and } u : X \rightarrow S_k \\ +\infty & \text{else.} \end{cases}$$

The theorem can be restated as showing that minimisers u_τ of \mathcal{E}_τ contain convergent subsequences, and any convergent subsequence converges to a minimiser of \mathcal{E}_0 . We divide the proof into two steps, in the first step we show that the sequence of minimisers $\{u_\tau\}_{\tau>0}$ is precompact, in the second step we show that any convergent subsequence is converging to a minimiser of \mathcal{E}_0 .

1. Compactness. We first show that any sequence $\{u'_\tau\}_{\tau>0}$ and $M \in \mathbb{R}$ satisfying $\sup_{\tau>0} \mathcal{E}_\tau(u'_\tau) \leq M$ is precompact. By Proposition A.4 and the Cauchy–Schwarz inequality

$$\begin{aligned} M &\geq \frac{\lambda_2^2}{2}\|u'_\tau - b\|_{\ell^2(X)}^2 + \underbrace{\frac{1}{\tau}\sum_{i=1}^n V(u'_\tau(x_i))}_{\geq 0} - \underbrace{\mu\sqrt{\sum_{j=1}^m (y_j - \bar{y})^2}}_{=:C} \|u'_\tau\|_{\ell^2(X)} \\ &\geq \frac{\lambda_2^2}{2}\|u'_\tau - b\|_{\ell^2(X)}^2 - C\mu\|u'_\tau - b\|_{\ell^2(X)} - C\mu\|b\|_{\ell^2(X)}. \end{aligned}$$

Hence,

$$\|u'_\tau - b\|_{\ell^2(X)} \leq \frac{C\mu}{\lambda_2^2} \left(1 + \sqrt{1 + \frac{2\lambda_2^2(M + C\mu\|b\|_{\ell^2(X)})}{C^2\mu^2}} \right) =: \tilde{C}$$

so $\{u'_\tau\}_{\tau>0}$ is bounded in $\ell^2(X)$ and therefore, by the Bolzano–Weierstrass Theorem, precompact.

To show that minimisers $\{u_\tau\}_{\tau>0}$ are precompact it is enough to show that there exists $M \in \mathbb{R}$ such that $\sup_{\tau>0} \mathcal{E}_\tau(u_\tau) \leq M$. This follows easily as we take $u \in \ell^2(X)$ satisfying $\sum_{i=1}^n u(x_i) = b$ and $u(x_i) \in S_k$ for all $i = 1, 2, \dots, n$ as a candidate. We have

$$\mathcal{E}_\tau(u_\tau) \leq \mathcal{E}_\tau(u) = \frac{1}{2}\|\nabla u\|_{\ell^2(X^2)}^2 - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot u(x_j) =: M.$$

Now we have shown that there exists convergent subsequences we show that any limit must be a minimiser of \mathcal{E}_0 .

2. Converging Subsequences. Let u_0 be a cluster point of $\{u_\tau\}_{\tau>0}$, i.e. there exists a subsequence such that $u_{\tau_m} \rightarrow u_0$ as $m \rightarrow \infty$. Pick any $v \in \ell^2(X)$ with $\mathcal{E}_0(v) < +\infty$. We will show

- (a) $\mathcal{E}_\tau(v) = \mathcal{E}_0(v)$,
- (b) $\liminf_{\tau \rightarrow 0} \mathcal{E}_\tau(u_\tau) \geq \mathcal{E}_0(u_0)$.

Assuming (a) and (b) hold then, by (a),

$$\mathcal{E}_0(v) = \mathcal{E}_{\tau_m}(v) \geq \mathcal{E}_{\tau_m}(u_{\tau_m}).$$

Taking the limit as $m \rightarrow \infty$, and applying (b) we have

$$\mathcal{E}_0(v) \geq \liminf_{m \rightarrow \infty} \mathcal{E}_{\tau_m}(u_{\tau_m}) \geq \mathcal{E}_0(u_0).$$

It follows that for all $v \in \ell^2(X)$ we have $\mathcal{E}_0(u_0) \leq \mathcal{E}_0(v)$, hence u_0 is a minimiser of \mathcal{E}_0 .

To show (a), we easily notice that

$$\mathcal{E}_\tau(v) = \frac{1}{2} \|\nabla v\|_{\ell^2(X^2)}^2 + \frac{1}{\tau} \sum_{i=1}^n \underbrace{V(v(x_i))}_{=0} - \mu \sum_{j=1}^m (y_j - \bar{y}) \cdot v(x_j) = \mathcal{E}_0(v).$$

For (b) we without loss of generality assume that $u_\tau \rightarrow u_0$ and

$$\liminf_{\tau \rightarrow 0} \mathcal{E}_\tau(u_\tau) = \lim_{\tau \rightarrow 0} \mathcal{E}_\tau(u_\tau) < +\infty.$$

As $\sum_{i=1}^n u_\tau(x_i) = b$ for all $\tau > 0$ and $u_\tau(x_i) \rightarrow u_0(x_i)$ for every $i \in \{1, \dots, n\}$ then $(u_0)_X = \sum_{i=1}^n u_0(x_i) = b$. And since $V(u_\tau(x_i)) \leq \tau \mathcal{E}_\tau(u_\tau) \rightarrow 0$ then we have $V(u_0(x_i)) = 0$, hence $u_0(x_i) \in S_k$. Now,

$$\mathcal{E}_\tau(u_\tau) = \underbrace{\frac{1}{2} \|\nabla u_\tau\|_{\ell^2(X^2)}^2}_{\rightarrow \frac{1}{2} \|\nabla u_0\|_{\ell^2(X^2)}^2} + \frac{1}{\tau} \sum_{i=1}^n \underbrace{V(u_\tau(x_i))}_{\geq 0} - \mu \underbrace{\sum_{j=1}^m (y_j - \bar{y}) \cdot u_\tau(x_j)}_{\rightarrow \sum_{j=1}^m (y_j - \bar{y}) \cdot u_0(x_j)}.$$

So $\liminf_{\tau \rightarrow 0} \mathcal{E}_\tau(u_\tau) \geq \mathcal{E}_0(u_0)$ as required. □

Remark A.7. If (a) and (b) in the proof of Theorem 2.4 are strengthened to

- (a') for all $v \in \ell^2(X)$ there exists $v_\tau \rightarrow v$ such that $\lim_{\tau \rightarrow 0} \mathcal{E}_\tau(v_\tau) = \mathcal{E}_0(v)$,
- (b') for all $v \in \ell^2(X)$ and for all $v_\tau \rightarrow v$ then $\liminf_{\tau \rightarrow 0} \mathcal{E}_\tau(v_\tau) \geq \mathcal{E}_0(v)$

then one says that \mathcal{E}_τ Γ -converges to \mathcal{E}_0 (and one can show that (a') and (b') hold in our case with a small modification of the above proof). The notion of Γ -convergence is fundamental in the calculus of variations and is considered the variational form of convergence as it implies (when combined with a compactness result) the convergence of minimisers.

B. Continuum limits

We briefly discuss continuum limits for the Poisson learning problem (2.3). We take $p = 2$ for simplicity, but the arguments extend similarly to other values of $p \geq 1$. In order to analyze continuum limits of graph-based learning algorithms, we make the *manifold assumption*, and assume G is a random geometric graph sampled from an underlying manifold. To be precise, we assume the vertices of the graph corresponding to unlabeled points x_1, \dots, x_n are a sequence of *i.i.d.* random variables drawn from a d -dimensional compact, closed, and connected manifold \mathcal{M} embedded in \mathbb{R}^D , where $d \ll D$. We assume the probability distribution of the random variables has the form $d\mu = \rho d\text{Vol}_{\mathcal{M}}$, where $\text{Vol}_{\mathcal{M}}$ is the volume form on the manifold, and ρ is a smooth density. For the labeled vertices in the graph, we take a fixed finite set of points $\Gamma \subset \mathcal{M}$. The vertices of the random geometric graph are

$$X_n := \{x_1, \dots, x_n\} \cup \Gamma.$$

We define the edge weights in the graph by

$$w_{xy} = \eta_\varepsilon(|x - y|),$$

where $\varepsilon > 0$ is the length scale on which we connect neighbors, $|x - y|$ is Euclidean distance in \mathbb{R}^D , and $\eta : [0, \infty) \rightarrow [0, \infty)$ is smooth with compact support, and $\eta_\varepsilon(t) = \frac{1}{\varepsilon^d} \eta\left(\frac{t}{\varepsilon}\right)$. We denote the solution of the Poisson learning problem (2.3) for this random geometric graph by $u_{n,\varepsilon}(x)$.

The normalized graph Laplacian is given by

$$\mathcal{L}_{n,\varepsilon} u(x) = \frac{2}{\sigma_\eta n \varepsilon^2} \sum_{y \in X_n} \eta_\varepsilon(|x - y|) (u(x) - u(y)),$$

where $\sigma_\eta = \int_{\mathbb{R}^d} |z_1|^2 \eta(|z|) dz$. It is well-known (see, e.g., (Hein et al., 2007)), that $\mathcal{L}_{n,\varepsilon}$ is consistent with the (negative of) the weighted Laplace-Beltrami operator

$$\Delta_\rho := -\rho^{-1} \operatorname{div}_{\mathcal{M}}(\rho^2 \nabla_{\mathcal{M}} u),$$

where $\operatorname{div}_{\mathcal{M}}$ is the manifold divergence and $\nabla_{\mathcal{M}}$ is the manifold gradient. We write $\operatorname{div} = \operatorname{div}_{\mathcal{M}}$ and $\nabla = \nabla_{\mathcal{M}}$ now for convenience. In particular, for any $u \in C^3(\mathcal{M})$ we have

$$|\mathcal{L}_{n,\varepsilon} u(x) - \Delta_\rho u(x)| \leq C(\|u\|_{C^3(\mathcal{M})} + 1)(\lambda + \varepsilon)$$

holds for all $x \in X_n$ with probability at least $1 - Cn \exp(-cn\varepsilon^{d+2}\lambda^2)$ for any $0 < \lambda \leq 1$, where $C, c > 0$ are constants.

Using the normalised graph Laplacian in the Poisson learning problem (2.3) we write

$$\mathcal{L}_{n,\varepsilon} u_{n,\varepsilon}(x) = n \sum_{y \in \Gamma}^m (g(y) - \bar{y}) \delta_{x=y} \quad \text{for } x \in X_n, \quad (\text{B.1})$$

where $g(y) \in \mathbb{R}$ denotes the label associated to $y \in \Gamma$ and $\bar{y} = \frac{1}{|\Gamma|} \sum_{x \in \Gamma} g(x)$. We restrict to the scalar case (binary classification) for now. Note that the normalisation plays no role in the classification problem (2.2). To see what should happen in the continuum, as $n \rightarrow \infty$ and $\varepsilon \rightarrow 0$, we multiply both sides of (B.1) by a smooth test function $\varphi \in C^\infty(\mathcal{M})$, sum over $x \in X$, and divide by n to obtain

$$\frac{1}{n} (\mathcal{L}_{n,\varepsilon} u_{n,\varepsilon}, \varphi)_{\ell^2(X)} = \sum_{y \in \Gamma} (g(y) - \bar{y}) \varphi(y). \quad (\text{B.2})$$

Since $\mathcal{L}_{n,\varepsilon}$ is self-adjoint (symmetric), we have

$$(\mathcal{L}_{n,\varepsilon} u_{n,\varepsilon}, \varphi)_{\ell^2(X)} = (u_{n,\varepsilon}, \mathcal{L}_{n,\varepsilon} \varphi)_{\ell^2(X)} = (u_{n,\varepsilon}, \Delta_\rho \varphi)_{\ell^2(X)} + O((\lambda + \varepsilon) \|u_{n,\varepsilon}\|_{\ell^1(X)}).$$

We also note that

$$\sum_{y \in \Gamma} (g(y) - \bar{y}) \varphi(y) = \int_{\mathcal{M}} \sum_{y \in \Gamma} (g(y) - \bar{y}) \delta_y(x) \varphi(x) d\operatorname{Vol}_{\mathcal{M}}(x),$$

where δ_y is Dirac-Delta distribution centered at $y \in \mathcal{M}$, which has the property that

$$\int_{\mathcal{M}} \delta_y(x) \varphi(x) d\operatorname{Vol}_{\mathcal{M}}(x) = \varphi(y)$$

for every smooth $\varphi \in C^\infty(\mathcal{M})$. Combining these observations with (B.2) we see that

$$\frac{1}{n} (u_{n,\varepsilon}, \Delta_\rho \varphi)_{\ell^2(X)} + O\left(\frac{(\lambda + \varepsilon)}{n} \|u_{n,\varepsilon}\|_{\ell^1(X)}\right) = \int_{\mathcal{M}} \sum_{y \in \Gamma} (g(y) - \bar{y}) \delta_y(x) \varphi(x) d\operatorname{Vol}_{\mathcal{M}}(x).$$

If we can extend $u_{n,\varepsilon}$ to a function on \mathcal{M} in a suitable way, then the law of large numbers would yield

$$\frac{1}{n} (u_{n,\varepsilon}, \Delta_\rho \varphi)_{\ell^2(X)} \approx \int_{\mathcal{M}} u_{n,\varepsilon}(x) \rho(x) \Delta_\rho \varphi(x) d\operatorname{Vol}_{\mathcal{M}}(x).$$

Hence, if $u_{n,\varepsilon} \rightarrow u$ as $n \rightarrow \infty$ and $\varepsilon \rightarrow 0$ in a sufficiently strong sense, then the function $u : \mathcal{M} \rightarrow \mathbb{R}$ would satisfy

$$-\int_{\mathcal{M}} u \operatorname{div}(\rho^2 \nabla \varphi) \, d\operatorname{Vol}_{\mathcal{M}} = \int_{\mathcal{M}} \sum_{y \in \Gamma} (g(y) - \bar{y}) \delta_y(x) \varphi(x) \, d\operatorname{Vol}_{\mathcal{M}}(x)$$

for every smooth $\varphi \in C^\infty(\mathcal{M})$. If $u \in C^2(\mathcal{M})$, then we can integrate by parts on the left hand side to find that

$$-\int_{\mathcal{M}} \varphi \operatorname{div}(\rho^2 \nabla u) \, d\operatorname{Vol}_{\mathcal{M}} = \int_{\mathcal{M}} \sum_{y \in \Gamma} (g(y) - \bar{y}) \delta_y(x) \varphi(x) \, d\operatorname{Vol}_{\mathcal{M}}(x)$$

Since φ is arbitrary, this would show that u is the solution of the Poisson problem

$$-\operatorname{div}(\rho^2 \nabla u) = \sum_{y \in \Gamma} (g(y) - \bar{y}) \delta_y \quad \text{on } \mathcal{M}. \tag{B.3}$$

We conjecture that the solutions $u_{n,\varepsilon}$ converge to the solution of (B.3) as $n \rightarrow \infty$ and $\varepsilon \rightarrow 0$ with probability one.

Conjecture B.1. Assume ρ is smooth. Assume that $n \rightarrow \infty$ and $\varepsilon = \varepsilon_n \rightarrow 0$ so that

$$\lim_{n \rightarrow \infty} \frac{n\varepsilon^{d+2}}{\log n} = \infty.$$

Let $u \in C^\infty(\mathcal{M} \setminus \Gamma)$ be the solution of the Poisson equation (B.3) and $u_{n,\varepsilon}$ solve the graph Poisson problem (B.1). Then with probability one

$$\lim_{n \rightarrow \infty} \max_{\substack{x \in X_n \\ \operatorname{dist}(x, \Gamma) > \delta}} |u_{n,\varepsilon}(x) - u(x)| = 0$$

for all $\delta > 0$.

The conjecture states that $u_{n,\varepsilon}$ converges to u uniformly as long as one stays a positive distance away from the source points Γ , where the solution u is singular. We expect the conjecture to be true, since similar results are known to hold when the source term on the right hand side is a smooth function f . The fact that the right hand side in (B.3) is highly singular, involving delta-mass concentration, raises difficult technical problems that will require new insights that are far beyond the scope of this paper.

Remark B.2. If Conjecture B.1 is true, it shows that Poisson learning is consistent with a well-posed continuum PDE for arbitrarily low label rates. This is in stark contrast to Laplace learning, which does not have a well-posed continuum limit unless the number of labels grows to ∞ as $n \rightarrow \infty$ sufficiently fast. This partially explains the superior performance of Poisson learning for low label rate problems.