

PROPERLY-WEIGHTED GRAPH LAPLACIAN FOR SEMI-SUPERVISED LEARNING

JEFF CALDER

Department of Mathematics, University of Minnesota

DEJAN SLEPČEV

Department of Mathematical Sciences, Carnegie Mellon University

ABSTRACT. The performance of traditional graph Laplacian methods for semi-supervised learning degrades substantially as the ratio of labeled to unlabeled data decreases, due to a degeneracy in the graph Laplacian. Several approaches have been proposed recently to address this, however we show that some of them remain ill-posed in the large-data limit.

In this paper, we show a way to correctly set the weights in Laplacian regularization so that the estimator remains well posed and stable in the large-sample limit. We prove that our semi-supervised learning algorithm converges, in the infinite sample size limit, to the smooth solution of a continuum variational problem that attains the labeled values continuously. Our method is fast and easy to implement.

1. INTRODUCTION

For many applications of machine learning, such as medical image classification and speech recognition, labeling data requires human input and is expensive [13], while unlabeled data is relatively cheap. *Semi-supervised* learning aims to exploit this dichotomy by utilizing the geometric or topological properties of the unlabeled data, in conjunction with the labeled data, to obtain better learning algorithms. A significant portion of the semi-supervised literature is on *transductive* learning, whereby a function is learned only at the unlabeled points, and not as a parameterized function on an ambient space. In the transductive setting, graph based algorithms, such the graph Laplacian-based learning pioneered by [53], are widely used and have achieved great success [3, 26, 27, 44–48, 50, 52].

Using graph Laplacians to propagate information from labeled to unlabeled points is one of the earliest and most popular approaches [53]. The constrained version of the graph Laplacian learning problem is to minimize over all $u : \mathcal{X} \rightarrow \mathbb{R}$

$$(1) \quad \begin{aligned} GL(u) &= \sum_{x,y \in \mathcal{X}} w_{xy} (u(x) - u(y))^2 \\ \text{subject to constraint} \quad &u(x) = g(x) \text{ for all } x \in \Gamma \end{aligned}$$

E-mail addresses: jcalder@umn.edu, slepcev@math.cmu.edu.

1991 Mathematics Subject Classification. 49J55, 35J20, 35B65, 62G20, 65N12.

Key words and phrases. semi-supervised learning, label propagation, asymptotic consistency, PDEs on graphs, Gamma-convergence.




FIGURE 1. Example of the degeneracy of graph Laplacian learning with few labels. The graph is a sequence of $n = 10^5$ *i.i.d.* random variables drawn from the unit box $[0, 1]^2$ in \mathbb{R}^2 , and two labels are given $g(0, 0.5) = 0$ and $g(1, 0.5) = 1$.

where the data points \mathcal{X} form the vertices of a graph with edge weights w_{xy} and $\Gamma \subset \mathcal{X}$ are the labeled nodes with label function $g : \Gamma \rightarrow \mathbb{R}$. The minimizer u of (1) is the learned function, which extends the given labels g on Γ to the remainder of the graph. In classification contexts, the values of u are often rounded to the nearest label. The method amounts to minimizing a Dirichlet energy on the graph, subject to a Dirichlet condition $u = g$ on Γ . Minimizers u are harmonic functions on the graph, and thus the problem can be viewed as harmonic extension.

It has been observed [18, 33] that when the size of Γ (the labeled points) is small, the performance of graph Laplacian learning algorithms degrades substantially. In practice, the learned function u fails to attain the conditions $u = g$ on Γ continuously, and degenerates into a constant label function that provides little information about the machine learning problem. Figure 1 gives an example of this issue. There are several ways to explain this degeneracy. First, in the limit of infinite data, the variational problem (1) is consistent with the continuum Dirichlet problem

$$(2) \quad \min_u \int_{\Omega} |\nabla u|^2 dx,$$

subject to a boundary condition $u = g$ on $\Gamma \subset \Omega \subset \mathbb{R}^d$. If Γ is finite this problem is ill-posed since the trace of an $H^1(\Omega)$ function at a point is not well-defined. In particular, there are minimizing sequences for the constrained problem converging to a constant function outside of Γ for which the Dirichlet energy converges to zero. In particular the minimum is not attained. From another perspective, minimizers of the continuum Dirichlet problem (2) satisfy Laplace's equation $\Delta u = 0$ with Dirichlet condition $u = g$ on Γ , and Laplace's equation is not well-posed without some boundary regularity (an exterior sphere condition), which does not hold for isolated points. In both cases, we are simply observing that the capacity of a point is zero in dimensions $d \geq 2$.

Several methods have been proposed recently to address the degeneracy of Laplacian learning with few labels. In [18], a class of p -Laplacian learning algorithms was proposed, which replace the exponent 2 in (1) with $p > 2$. The p -Laplacian models were considered previously for other applications [2, 9, 18, 49], and the $p \rightarrow \infty$ case, which is called Lipschitz learning, was considered

in [29, 32]. The idea behind the p -Laplacian models is that the continuum variational problem is now the p -Dirichlet problem

$$(3) \quad \min_u \int_{\Omega} |\nabla u|^p dx,$$

and for $p > d$ the Sobolev embedding $W^{1,p}(\Omega) \hookrightarrow C^{0,\alpha}(\bar{\Omega})$ allows the assignment of boundary values at isolated points. The p -Laplacian models, including the $p = \infty$ version, were proven to be well-posed in the limit of infinite unlabeled data and finite labeled data precisely when $p > d$ in [10, 11, 39]. The disadvantage of p -Laplacian models is the nonlinearity renders them more computationally challenging to solve, compared with standard Laplacian regularization. Other approaches include higher order Laplacian regularization [6, 17, 51] and using a spectral cut-off [5].

The approach most closely related to our work is the *weighted nonlocal Laplacian* of Shi, Osher, and Zhu [36], which replaces the learning problem (1) with

$$(4) \quad \min_{u: \mathcal{X} \rightarrow \mathbb{R}} \sum_{x \in \mathcal{X} \setminus \Gamma} \sum_{y \in \mathcal{X}} w_{xy} (u(x) - u(y))^2 + \mu \sum_{x \in \Gamma} \sum_{y \in \mathcal{X}} w_{xy} (g(x) - u(y))^2,$$

where $\mu > 0$ is selected as the ratio of unlabeled to labeled data. The method increases the weights of edges adjacent to labels, which encourages the label function to be flat near labels. The authors show in [36] that the method produces superior results, compared to the standard graph Laplacian, for classification with very few labels. Furthermore, since the method is a standard graph Laplacian with a modified weight matrix, it has similar computational complexity to Laplacian learning, and is fast compared to the non-linear p -Laplace methods, for example. However, as we prove in this paper, the weighted nonlocal Laplacian of [36] becomes ill-posed (degenerate) in the limit of infinite unlabeled and finite labeled data. This is a direct consequence of Corollary 3.8. Numerical simulations in Section 5 illustrate the way in which the method becomes degenerate. The issue is the same as for Laplacian learning, since the weights are modified only locally near label points and the size of this neighborhood shrinks to zero in the large sample size limit.

1.0.1. *Properly weighted laplacian.* In this paper, we show how to properly weight the graph Laplacian so that it remains well-posed in the limit of infinite unlabeled and finite labeled data. Our method, roughly speaking, modifies the problem to one of the form:

$$(5) \quad \begin{aligned} & \text{Minimize} \quad \sum_{x,y \in \mathcal{X}} \gamma(x) w_{xy} (u(x) - u(y))^2 \quad \text{over } u : \mathcal{X} \rightarrow \mathbb{R}, \\ & \text{subject to constraint} \quad u(x) = g(x) \text{ for all } x \in \Gamma \end{aligned}$$

where $\gamma(x) = \text{dist}(x, \Gamma)^{-\alpha}$ and $\alpha > d - 2$ (see Section 1.1 for precise definitions). Here, we are modifying the weights not just of edges connecting to points of Γ , but also in a neighborhood of Γ . We show that this model is stable as the number unlabeled data points increases to infinity, under appropriate scaling of the graph construction. In particular we show that the minimizers of the graph problem above converge as the number of unlabeled data points increases to the minimizer of a ‘‘continuum learning problem’’. We give the precise assumptions on the discrete model below and describe the continuum problem in Section 2. Here we give a brief explanation as to why $\alpha > d - 2$ is the natural scaling for the weight.

To illustrate what is happening near a labeled point, consider $\Gamma = \{0\}$ and take the domain from which the points are sampled to be the unit ball $\Omega = B(0, 1)$ in \mathbb{R}^d . The continuum

variational problem corresponding to (5) involves minimizing

$$(6) \quad I[u] = \int_{B(0,1)} |x|^{-\alpha} |\nabla u|^2 dx.$$

The Euler-Lagrange equation satisfied by minimizers of I is

$$(7) \quad \operatorname{div}(|x|^{-\alpha} \nabla u) = 0.$$

This equation has a radial solution $u(x) = |x|^{\alpha+2-d}$, which is continuous at $x = 0$ when $\alpha > d-2$. This suggests the solutions will assume this radial profile near labels, and the model will be well-posed for $\alpha > d-2$. Furthermore when $\alpha \geq d-1$ one can expect the solution to be Lipschitz near labels, and for $\alpha \geq d$ it should be differentiable at the labels. It is important to point out that the proper weighting changes the degenerate limiting continuum problem to one that is well-posed with “boundary” data at isolated points.

We now provide a precise description of the properly-weighted graph Laplacian.

1.1. Model and definitions. Let $\Omega \subset \mathbb{R}^d$ be open and bounded with a Lipschitz boundary. Let $\Gamma \subset \Omega$ be a finite collection of points along with a given label function $g : \Gamma \rightarrow \mathbb{R}$. Let x_1, x_2, \dots, x_n be independent and identically distributed random variables drawn probability measure μ with continuous density ρ supported on Ω , and set

$$X_n := \{x_1, x_2, \dots, x_n\},$$

and $\mathcal{X}_n := X_n \cup \Gamma$. To define the edge weights we use a radial kernel η with profile $\boldsymbol{\eta} : [0, \infty) \rightarrow [0, \infty)$ which is nonincreasing, continuous at 0 and satisfies

$$(8) \quad \begin{cases} \boldsymbol{\eta}(t) \geq 1, & \text{if } 0 \leq t \leq 1 \\ \boldsymbol{\eta}(t) = 0, & \text{if } t > 2. \end{cases}$$

All of the results we state can be extended to kernels which decay sufficiently fast, in particular the Gaussian. For $\varepsilon > 0$ we define the rescaled kernel

$$(9) \quad \eta_\varepsilon(x-y) = \frac{1}{\varepsilon^d} \boldsymbol{\eta}\left(\frac{|x-y|}{\varepsilon}\right).$$

We now introduce the penalization of the gradient, which is heavier near labeled points. Let $R > 0$ be the minimum distance between pairs of points in Γ :

$$R = \min\{|x-y| : x, y \in \Gamma, x \neq y\}.$$

For $r_0 > 0$ and $\alpha \geq 0$ let $\gamma \in C^\infty(\overline{\Omega} \setminus \Gamma)$ be any function satisfying $\gamma \geq 1$ on Ω and

$$(10) \quad \gamma(x) = 1 + \left(\frac{r_0}{\operatorname{dist}(x, \Gamma)}\right)^\alpha \text{ whenever } \operatorname{dist}(x, \Gamma) \leq \frac{R}{4},$$

where $\operatorname{dist}(x, \Gamma)$ denotes the Euclidean distance from x to the closest point in Γ . For $\zeta > 1$ we set

$$(11) \quad \gamma_\zeta(x) = \min\{\gamma(x), \zeta\}.$$

For $u \in L^2(\mathcal{X}_n)$ we define the energy

$$(12) \quad \mathcal{G}\mathcal{E}_{n,\varepsilon,\zeta}(u) = \frac{1}{n^2 \varepsilon^2} \sum_{x,y \in \mathcal{X}_n} \gamma_\zeta(x) \eta_\varepsilon(x-y) |u(x) - u(y)|^2.$$

The Laplacian learning problem is to

$$(13) \quad \text{minimize } \mathcal{GE}_{n,\varepsilon,\zeta}(u) \text{ over } \{u \in L^2(\mathcal{X}_n) \text{ and } u = g \text{ on } \Gamma\}.$$

We note that the unique minimizer $u \in L^2(\mathcal{X}_n)$ of (13) satisfies the optimality condition

$$(14) \quad \begin{cases} \mathcal{GL}_{n,\varepsilon,\zeta}u(x) = 0, & \text{if } x \in X_n (= \mathcal{X}_n \setminus \Gamma) \\ u(x) = g(x), & \text{if } x \in \Gamma, \end{cases}$$

where $\mathcal{GL}_{n,\varepsilon,\zeta} : L^2(\mathcal{X}_n) \rightarrow L^2(\mathcal{X}_n)$ is the graph Laplacian, given by

$$(15) \quad \mathcal{GL}_{n,\varepsilon,\zeta}u(x) = \frac{1}{2n\varepsilon^2} \sum_{y \in \mathcal{X}_n} (\gamma_\zeta(x) + \gamma_\zeta(y)) \eta_\varepsilon(x - y) (u(y) - u(x)).$$

Some remarks about the model are in order.

Remark 1.1. When considering the discrete functional ζ depends on n and converges to infinity (sufficiently fast) as $n \rightarrow \infty$. The constant r_0 represents the length scale of the crossover from the strong local penalization near Γ to uniform far-field penalization. The introduction of ζ is needed since $\gamma(x) = \infty$ on Γ and so using γ directly would impose a hard constraint on neighbors of labeled points. While we can allow $\zeta = \infty$ in our model by interpreting products $\infty \cdot 0$ as 0 , we wanted to allow for a model with far less stringent constraints on agreement with the labeled points in the immediate vicinity of Γ . We note that the critical distance to Γ , when γ_ζ crosses over from γ to ζ equals

$$(16) \quad r_\zeta = r_0(\zeta - 1)^{-1/\alpha} \quad \text{provided that } r_0(\zeta - 1)^{-1/\alpha} < \frac{R}{4}.$$

Remark 1.2. In practice, one can take (10) to be the definition of the weights on the whole domain Ω . We only need γ to be smooth for a part of our analysis in Section 2.2. The issue is that since the distance function $d(x, \Gamma)$ is not differentiable (it is only Lipschitz on $\Omega \setminus \Gamma$ if Γ has more than one point), γ cannot be both smooth and globally given by (10). To elaborate, γ appears as part of the diffusion coefficient in the limiting elliptic problem (see Eq. (21)). The solutions have nicer regularity properties when we take γ to be smooth, away from the labels. For the other results we only need that γ is bounded from below by a positive number and has singularities, with a particular growth rate, near the points of Γ .

Remark 1.3. Instead of truncating γ at the radius r_ζ to construct the weights γ_ζ , we can take a possibly discontinuous model of the form

$$(17) \quad \gamma_{\zeta,r}(x) = \begin{cases} \gamma(x), & \text{if } \text{dist}(x, \Gamma) > r \\ \zeta, & \text{if } \text{dist}(x, \Gamma) \leq r. \end{cases}$$

This model is more general, since we can set $r = r_\zeta$ to recover (11). Choosing $\zeta \gg 1 + (r_0/r)^\alpha$ places a larger penalty on the gradient in the inner region where $\text{dist}(x, \Gamma) \leq r$, compared to Eq. (11). This model is useful in the analysis of the graph based problem, and gives a shaper result for continuity at the labels (see Remark 4.3). In the limit as $n \rightarrow \infty$ we would take $\zeta \rightarrow \infty$ and $r \rightarrow 0$ with $r \geq r_\zeta$.

Remark 1.4. We remark that the discrete functional (12) can be rewritten as

$$(18) \quad \mathcal{GE}_{n,\varepsilon,\zeta}(u) = \frac{1}{2n^2\varepsilon^2} \sum_{x,y \in \mathcal{X}_n} (\gamma_\zeta(x) + \gamma_\zeta(y)) \eta_\varepsilon(x - y) |u(x) - u(y)|^2,$$

and so the problem has a symmetric weight matrix.

1.2. Outline. The continuum properly-weighted Dirichlet energy, which describes the asymptotic behavior of the properly-weighted graph Laplacian (13) is presented in Section 2 (equations (19) and (20)). To show that the continuum problem is well posed and to establish its basic properties, in Section 2 we also study properties of singularly weighted Sobolev spaces. In particular the Trace Theorem 2.2 plays a key role in showing that the data can be imposed on a set of isolated points, which enables us to show the well-posedness in Theorem 2.7. The Euler-Lagrange equation of the variational problem is the elliptic problem we study in Section 2.2. In particular we show that solutions are C^2 away from the labels and Hölder continuous globally.

In Section 3 we turn to asymptotics of the graph-based problems. We prove in Theorem 3.1 that the solutions of the graph-based learning problem (13), for the properly-weighted Laplacian, converge in the large sample size limit to the solution of a continuum variational problem (19)-(20). We achieve this by showing the Γ -convergence of the discrete variational problems to the corresponding continuum problem. We also prove a negative result, showing that the nonlocal weighted Laplacian [36] is degenerate (ill-posed) in the large data limit (with fixed number of labeled points). In Section 4.1 we prove that solutions of the graph-based learning problem for the properly-weighted Laplacian attain their labeled values continuously with high probability (Theorem 4.1). In Section 5 we present the results of numerical simulations illustrating the estimators obtained by our method, and its performance in classification tasks on synthetic data and in classifying handwritten digits from the MNIST dataset [30]. The classification problems on synthetic data contrast the stability of the properly-weighted Laplacian with the instability of the standard graph Laplacian and related methods. The MNIST experiments show superior performance of our method compared to the standard graph Laplacian, and similar performance to the weighted Laplacian of [36]. In the Appendix A we recall some background results used and show an auxiliary technical result.

1.3. Acknowledgements. Calder was supported by NSF grant DMS:1713691. Slepčev acknowledges the NSF support (grants DMS-1516677 and DMS-1814991). He is grateful to University of Minnesota, where this project started, for hospitality. He is also grateful to the Center for Nonlinear Analysis of CMU for its support.

2. ANALYSIS OF THE CONTINUUM PROBLEM

The continuum variational problem corresponding to the graph-based problem (13) is

$$(19) \quad \text{minimize } \mathcal{E}(u) \text{ over } \{u \in H_\gamma^1(\Omega) \text{ and } u = g \text{ on } \Gamma\},$$

where \mathcal{E} is given by

$$(20) \quad \mathcal{E}(u) = \frac{1}{2} \int_{\Omega} \gamma |\nabla u|^2 \rho^2 dx$$

and the weighted Sobolev Space $H_\gamma^1(\Omega)$ is defined by (24). It follows from Lemma 2.1 that for γ which grow near points of Γ as fast or faster than $\text{dist}(x, \Gamma)^{-\alpha}$, the functions in $H_\gamma^1(\Omega)$ have a trace at Γ (defined by (32)), which enables one to assign the condition $u = g$ on Γ in (19).

The Euler-Lagrange equation satisfied by minimizers of (20) is the elliptic equation

$$(21) \quad \begin{cases} -\text{div}(\gamma \rho^2 \nabla u) = 0 & \text{in } \Omega \setminus \Gamma \\ u = g & \text{on } \Gamma \\ \frac{\partial u}{\partial \nu} = 0 & \text{on } \partial\Omega. \end{cases}$$

In this section we study the variational problem (19) and the elliptic problem (21) rigorously. The theory is nonstandard due to the boundary condition $u = g$ on Γ , since Γ is a collection of isolated points and does not satisfy an exterior sphere condition. As a consequence of this analysis, we prove in Section 4.1 that solutions of the graph-based problem are continuous at the labels.

Before studying this problem, we need to perform a careful analysis of a particular weighted Sobolev space.

2.1. Weighted Sobolev spaces. In this section we study the Sobolev space with norm weighted by γ . While there exists a rich literature on Weighted Sobolev Spaces, we did not find the precise results we need. Below we develop a self-contained, but brief, description of the spaces with particular weights of interest.

For $u \in H^1(\Omega)$ we define

$$(22) \quad [u]_{H_\gamma^1(\Omega)}^2 = \int_{\Omega} \gamma |\nabla u|^2 dx,$$

and

$$(23) \quad \|u\|_{H_\gamma^1(\Omega)}^2 = \|u\|_{L^2(\Omega)}^2 + [u]_{H_\gamma^1(\Omega)}^2.$$

We define

$$(24) \quad H_\gamma^1(\Omega) = \left\{ u \in H^1(\Omega) : \|u\|_{H_\gamma^1(\Omega)} < \infty \right\},$$

and endow $H_\gamma^1(\Omega)$ with the norm $\|u\|_{H_\gamma^1(\Omega)}$. We also denote by $H_{\gamma,0}^1(\Omega)$ the closure of $C_c^\infty(\overline{\Omega} \setminus \Gamma)$ in $H_\gamma^1(\Omega)$. The space $H_\gamma^1(\Omega)$ is the natural function space on which to pose the variational problem (19).

Throughout this section we let B_r denote the open ball of radius $r > 0$ centered at the origin in \mathbb{R}^d . Whenever we consider the space $H_\gamma^1(B_r)$, we will implicitly assume the choice of $\gamma(x) = |x|^{-\alpha}$. Hence

$$(25) \quad [u]_{H_\gamma^1(B_r)}^2 = \int_{B_r} |\nabla u|^2 |x|^{-\alpha} dx.$$

We also use the notation $(u)_{x,r} = \int_{B(x,r)} u dx$ for the average of u over the ball $B(x,r)$, and $(u)_r := (u)_{0,r}$. We also assume in this section that Ω has a Lipschitz boundary.

First, we study the trace of $H_\gamma^1(\Omega)$ functions on Γ . Before proving a general trace theorem, we require a preliminary lemma.

Lemma 2.1. *Let $\alpha > d - 2$ and $u \in H_\gamma^1(B_r)$. Then $x = 0$ is a Lebesgue point for u , i.e., $u(0) = \lim_{\varepsilon \rightarrow 0} (u)_\varepsilon$, and*

$$(26) \quad |u(0) - (u)_\varepsilon| \leq C \varepsilon^{(\alpha+2-d)/2} [u]_{H_\gamma^1(B_\varepsilon)}$$

for all $0 < \varepsilon \leq r$.

Proof. We compute

$$\int_{B_\varepsilon} |\nabla u|^2 dx \leq \int_{B_\varepsilon} |\nabla u|^2 |x|^{-\alpha} r^\alpha dx \leq [u]_{H_\gamma^1(B_\varepsilon)}^2 r^\alpha.$$

By the Poincaré inequality we have

$$(27) \quad \int_{B_\varepsilon} (u - (u)_\varepsilon)^2 dx \leq C r^2 \int_{B_\varepsilon} |\nabla u|^2 dx \leq C [u]_{H_\gamma^1(B_\varepsilon)}^2 r^{\alpha+2-d}.$$

For $0 < s < t \leq r$ we have

$$\begin{aligned}
(28) \quad s^d((u)_s - (u)_t)^2 &\leq C \int_{B_s} ((u)_s - (u)_t)^2 dx \\
&\leq C \int_{B_s} ((u)_s - u)^2 dx + C \int_{B_s} ((u)_t - u)^2 dx \\
&\leq C[u]_{H_\gamma^1(B_t)}^2 (s^{\alpha+2} + t^{\alpha+2}).
\end{aligned}$$

For $0 < q < \varepsilon \leq r$ with $\varepsilon \leq 4q$ we can set $s = q$ and $t = \varepsilon$ above to obtain

$$(29) \quad |(u)_q - (u)_\varepsilon|^2 \leq C[u]_{H_\gamma^1(B_\varepsilon)}^2 \varepsilon^{\alpha+2-d}.$$

For $0 < q < \varepsilon \leq r$ with $\varepsilon > 4q$, let $k \in \mathbb{N}$ be the greatest integer smaller than $\log_2(\varepsilon/q)$. Since $\varepsilon > 2q$, we have $k \geq 1$. Choose $b > 0$ so that $b^k = \varepsilon/q$. Then

$$\log(b) = \frac{\log(\varepsilon/q)}{k} \geq \frac{\log(\varepsilon/q)}{\log_2(\varepsilon/q)} \geq \log(2)$$

and

$$\log(b) \leq \frac{\log(\varepsilon/q)}{\log_2(\varepsilon/q) - 1} = \frac{\log(2) \log_2(\varepsilon/q)}{\log_2(\varepsilon/q) - 1} \leq 2 \log(2),$$

since $\log_2(\varepsilon/q) > 2$. Therefore $2 \leq b \leq 4$. Let us set $\varepsilon_j = \varepsilon b^{-j}$ and $a_j = (u)_{\varepsilon_j}$. Then $\varepsilon_0 = \varepsilon$ and $\varepsilon_k = \varepsilon b^{-k} = q$. Setting $t = \varepsilon_j$ and $s = \varepsilon_{j+1}$ in (28) yields

$$|a_j - a_{j+1}|^2 \leq C[u]_{H_\gamma^1(B_\varepsilon)}^2 (\varepsilon_{j+1}^{\alpha+2-d} + \varepsilon_j^{\alpha+2} \varepsilon_{j+1}^{-d}) \leq C[u]_{H_\gamma^1(B_\varepsilon)}^2 b^{-j(\alpha+2-d)} \varepsilon^{\alpha+2-d}.$$

Therefore

$$(30) \quad |(u)_q - (u)_\varepsilon| \leq \sum_{j=0}^{k-1} |a_{j+1} - a_j| \leq C[u]_{H_\gamma^1(B_\varepsilon)} \varepsilon^{(\alpha+2-d)/2}$$

holds for all $k \geq 1$, where C is independent of u , ε and k .

In either case, we have established that

$$(31) \quad |(u)_q - (u)_\varepsilon|^2 \leq C[u]_{H_\gamma^1(B_\varepsilon)}^2 \varepsilon^{\alpha+2-d}$$

holds for all $0 < q < \varepsilon \leq r$. Thus, the sequence $\varepsilon \mapsto (u)_\varepsilon$ is Cauchy and converges to a real number as $\varepsilon \rightarrow 0$, which we will call $u(0)$. Sending $q \rightarrow 0$ in (31) completes the proof. \square

By Lemma 2.1, we can define the trace operator $\text{Tr} : H_\gamma^1(\Omega) \rightarrow \mathbb{R}^\Gamma$ by

$$(32) \quad \text{Tr}[u](x) = \lim_{\varepsilon \rightarrow 0} \int_{B(x,\varepsilon)} u dx \quad (x \in \Gamma).$$

We endow \mathbb{R}^Γ with the Euclidean norm. We now prove our main trace theorem.

Theorem 2.2 (Trace Theorem). *Assume $\alpha > d-2$. Then the trace operator $\text{Tr} : H_\gamma^1(\Omega) \rightarrow \mathbb{R}^\Gamma$ is bounded, and satisfies $\text{Tr}[u](x) = u(x)$ whenever u is continuous at $x \in \Gamma$. Furthermore, for every $u, v \in H_\gamma^1(\Omega)$ with $\|u - v\|_{L^2(\Omega)}^{2/(\alpha+2)} \leq R/2$ we have*

$$(33) \quad |\text{Tr}[u] - \text{Tr}[v]| \leq C(1 + [u]_{H_\gamma^1(\Omega)} + [v]_{H_\gamma^1(\Omega)}) \|u - v\|_{L^2(\Omega)}^{1-d/(\alpha+2)}.$$

Proof. By Lemma 2.1 each $x \in \Gamma$ is a Lebesgue point of u , and we have for $r \leq R/2$

$$|\operatorname{Tr}[u](x)| \leq \int_{B(x,r)} u \, dx + Cr^{(\alpha+2-d)/2} [u]_{H_\gamma^1(\Omega)} \leq Cr^{d/2} \|u\|_{L^2(\Omega)} + Cr^{(\alpha+2-d)/2} \|u\|_{H_\gamma^1(\Omega)}.$$

Fixing $r = R/2$ we have $|\operatorname{Tr}[u](x)| \leq C \|u\|_{H_\gamma^1(\Omega)}$, hence $\operatorname{Tr} : H_\gamma^1(\Omega) \rightarrow \mathbb{R}$ is bounded.

To prove (33), let $u, v \in H_\gamma^1(\Omega)$ and $x \in \Gamma$. For simplicity, we write $u(x)$ and $v(x)$ for $\operatorname{Tr}[u](x)$ and $\operatorname{Tr}[v](x)$, respectively. By Lemma 2.1 we have for $0 < \varepsilon \leq R/2$

$$\begin{aligned} |u(x) - v(x)| &\leq |u(x) - (u)_{x,\varepsilon}| + |v(x) - (v)_{x,\varepsilon}| + |(u)_{x,\varepsilon} - (v)_{x,\varepsilon}| \\ &\leq C\varepsilon^{(\alpha+2-d)/2} ([u]_{H_\gamma^1(B(x,\varepsilon))} + [v]_{H_\gamma^1(B(x,\varepsilon))}) + C\varepsilon^{-d} \int_{B_\varepsilon} |u - v| \, dx \\ &\leq C\varepsilon^{(\alpha+2-d)/2} ([u]_{H_\gamma^1(\Omega)} + [v]_{H_\gamma^1(\Omega)}) + C\varepsilon^{-d/2} \|u - v\|_{L^2(\Omega)}. \end{aligned}$$

Choosing $\varepsilon = \|u - v\|_{L^2(\Omega)}^{2/(\alpha+2)}$, we obtain

$$|u(x) - v(x)| \leq C(1 + [u]_{H_\gamma^1(\Omega)} + [v]_{H_\gamma^1(\Omega)}) \|u - v\|_{L^2(\Omega)}^{1-d/(\alpha+2)},$$

provided $\varepsilon \leq R/2$. □

We now examine the decay of the L^2 norm of trace zero functions.

Lemma 2.3. *Let $\alpha > d - 2$ and $u \in H_\gamma^1(B_r)$ with $\operatorname{Tr}[u](0) = 0$. Then*

$$(34) \quad \int_{\partial B_\varepsilon} u^2 \, dS + \int_{B_\varepsilon} u^2 \, dx \leq C\varepsilon^{\alpha+2-d} [u]_{H_\gamma^1(B_\varepsilon)}^2$$

for all $0 < \varepsilon \leq r$.

Proof. Since $\operatorname{Tr}[u](0) = 0$, Lemma 2.1 yields

$$\|(u)_\varepsilon\|_{L^2(B_\varepsilon)} \leq Cr^{d/2} |(u)_\varepsilon| \leq C\varepsilon^{(\alpha+2)/2} [u]_{H_\gamma^1(B_\varepsilon)}.$$

Recalling (27) from the proof of Lemma 2.1 we deduce

$$\|u - (u)_\varepsilon\|_{L^2(B_\varepsilon)} \leq C\varepsilon^{(\alpha+2)/2} [u]_{H_\gamma^1(B_\varepsilon)}.$$

Therefore

$$\int_{B_\varepsilon} u^2 \, dx = C\varepsilon^{-d} \|u\|_{L^2(B_\varepsilon)}^2 \leq C\varepsilon^{\alpha+2-d} [u]_{H_\gamma^1(B_\varepsilon)}^2,$$

which establishes one part of (34).

For the other part, we use a standard trace estimate that we include for completeness. We have

$$\begin{aligned} \varepsilon \int_{\partial B_\varepsilon} u^2 \, dS &= \int_{B_\varepsilon} \operatorname{div}(xu^2) \, dx \\ &= \int_{B_\varepsilon} du^2 + 2u \nabla u \cdot x \, dx \\ &\leq C \int_{B_\varepsilon} u^2 \, dx + C\varepsilon^2 \int_{B_\varepsilon} |\nabla u|^2 \, dx \\ &\leq C \int_{B_\varepsilon} u^2 \, dx + C\varepsilon^{\alpha+2} [u]_{H_\gamma^1(B_\varepsilon)}^2. \end{aligned}$$

Dividing both sides by ε^d we obtain

$$\int_{\partial B_\varepsilon} u^2 dS \leq C \int_{B_\varepsilon} u^2 dx + C\varepsilon^{\alpha+2-d} [u]_{H_\gamma^1(B_\varepsilon)}^2,$$

which completes the proof. \square

We now show that trace zero functions can be approximated in $H_\gamma^1(\Omega)$ by smooth functions compactly supported away from Γ .

Theorem 2.4 (Trace zero functions). *If $\alpha > d - 2$ then $u \in H_{\gamma,0}^1(\Omega)$ if and only if $u \in H_\gamma^1(\Omega)$ and $\text{Tr}[u] = 0$.*

Proof. If $u \in H_{\gamma,0}^1(\Omega)$, then there exists $u_k \in C_c^\infty(\overline{\Omega} \setminus \Gamma)$ so that $u_k \rightarrow u$ in $H_\gamma^1(\Omega)$. In particular, u_k is uniformly bounded in $H_\gamma^1(\Omega)$. Thus, by Theorem 2.2, we have $\text{Tr}[u](x) = \lim_{k \rightarrow \infty} u_k(x) = 0$ for each $x \in \Gamma$.

Conversely, let $u \in H_\gamma^1(\Omega)$ such that $\text{Tr}[u] = 0$. Without loss of generality, we may assume $\Omega = B_r$, $\Gamma = \{0\}$, and $\text{Tr}[u](0) = 0$. Choose a smooth nonincreasing function $\xi : [0, \infty) \rightarrow [0, 1]$ such that $\xi(t) = 1$ for $0 \leq t \leq 1$ and $\xi(t) = 0$ for $t \geq 2$. For a positive integer $k \geq 1$ define $\xi_k(x) = \xi(k|x|)$ and $w_k = u(1 - \xi_k)$. We compute

$$\begin{aligned} (35) \quad \int_{B_r} |\nabla w_k - \nabla u|^2 |x|^{-\alpha} dx &= \int_{B_r} \left| \xi_k \nabla u + k u \xi'(k|x|) \frac{x}{|x|} \right|^2 |x|^{-\alpha} dx \\ &\leq C \int_{B_{2/k}} |\nabla u|^2 |x|^{-\alpha} dx + Ck^2 \int_{B_{2/k} \setminus B_{1/k}} u^2 |x|^{-\alpha} dx \\ &\leq C \int_{B_{2/k}} |\nabla u|^2 |x|^{-\alpha} dx + Ck^{\alpha+2} \int_{B_{2/k}} u^2 dx \\ &\leq C \int_{B_{2/k}} |\nabla u|^2 |x|^{-\alpha} dx, \end{aligned}$$

the last line following from Lemma 2.3. Therefore $w_k \rightarrow u$ in $H_\gamma^1(B_r)$ as $k \rightarrow \infty$. To produce a smooth approximating sequence u_k , we simply mollify the sequence w_k . \square

As a corollary, we can prove density of smooth functions that are locally constant near Γ .

Corollary 2.5. *For any $\alpha > 0$ the set*

$$\mathcal{S} = \{u \in C^\infty(\overline{\Omega}) : (\exists s > 0)(\forall x \in \Gamma)(\forall z \in B(0, 1)) \ u(x + sz) = u(x)\}$$

is a dense subset of $H_\gamma^1(\Omega)$.

Proof. We split the proof into two cases.

Case 1: $\alpha > d - 2$. Let $u \in H_\gamma^1(\Omega)$. There exists $\psi \in \mathcal{S}$ such that $\text{Tr}[\psi] = \text{Tr}[u]$. Since $w := u - \psi \in H_{\gamma,0}^1(\Omega)$, there exists by Theorem 2.4 a sequence $\varphi_k \in C_c^\infty(\overline{\Omega} \setminus \Gamma)$ such that $\varphi_k \rightarrow w$ as $k \rightarrow \infty$. We simply note that $\psi_k := \varphi_k + \psi \in \mathcal{S}$ and $\psi_k \rightarrow u$ in $H_\gamma^1(\Omega)$ as $k \rightarrow \infty$.

Case 2: $\alpha \leq d - 2$. In this case, $C^\infty(\overline{\Omega})$ is dense in $H_\gamma^1(\Omega)$ by a standard mollification argument, since the weighting kernel $|x|^{-\alpha}$ is integrable. Hence, for $u \in H_\gamma^1(\Omega)$ with $\alpha \leq d - 2$ there exists $\varphi_k \in C^\infty(\overline{\Omega})$ such that $\varphi_k \rightarrow u$ in $H_\gamma^1(\Omega)$. Since φ_k is smooth, we automatically have

$$\int_{\Omega} \text{dist}(x, \Gamma)^{1-d} |\nabla \varphi_k|^2 dx < \infty.$$

Thus, by case 1, there exists a sequence $\psi_{k,j} \in \mathcal{S}$ such that for each k , $\psi_{k,j} \rightarrow \varphi_k$ in H_γ^1 as $j \rightarrow \infty$, since $\alpha \leq d - 2 \leq d - 1$. The proof is completed with a diagonal argument. \square

Finally, we prove a Hardy-type inequality for trace zero functions in $H_\gamma^1(B_R)$.

Theorem 2.6 (Hardy's inequality). *Let $\alpha > d - 2$ and $u \in H_\gamma^1(B_r)$ with $\text{Tr}[u](0) = 0$. Then $\frac{u}{|x|^{(\alpha+2)/2}} \in L^2(B_r)$ and*

$$(36) \quad \int_{B_r} \frac{u^2}{|x|^{\alpha+2}} dx \leq C[u]_{H_\gamma^1(B_r)}^2.$$

Proof. By a change of variables we can reduce to the case of $r = 1$. We first note that

$$\text{div}\left(\frac{x}{|x|^{\alpha+2}}\right) = -\frac{\alpha + 2 - d}{|x|^{\alpha+2}}$$

for $x \neq 0$. Thus, for $\varepsilon > 0$ we have

$$\begin{aligned} \int_{B_1 \setminus B_\varepsilon} \frac{u^2}{|x|^{\alpha+2}} dx &= -(\alpha + 2 - d) \int_{B_1 \setminus B_\varepsilon} u^2 \text{div}\left(\frac{x}{|x|^{\alpha+2}}\right) dx \\ &= (\alpha + 2 - d) \left[2 \int_{B_1 \setminus B_\varepsilon} u \nabla u \cdot \frac{x}{|x|^{\alpha+2}} dx - \int_{\partial B_1} u^2 dS + \frac{1}{\varepsilon^{\alpha+1}} \int_{\partial B_\varepsilon} u^2 dS \right] \\ &\leq C \int_{B_1 \setminus B_\varepsilon} |u| |\nabla u| |x|^{-\alpha-1} dx + C[u]_{H_\gamma^1(B_1)}^2 + C \int_{B_\varepsilon} |x|^{-\alpha} |\nabla u|^2 dx, \end{aligned}$$

where the last line follows from Lemma 2.3 and the assumption $\alpha > d - 2$. Applying Cauchy's inequality to the first term and rearranging yields

$$\int_{B_1 \setminus B_\varepsilon} \frac{u^2}{|x|^{\alpha+2}} dx \leq C[u]_{H_\gamma^1(B_1)}^2 + C \int_{B_\varepsilon} |x|^{-\alpha} |\nabla u|^2 dx.$$

Sending $\varepsilon \rightarrow 0$ completes the proof. \square

We now establish the well posedness of the continuum properly-weighted Laplacian learning problem.

Theorem 2.7. *Assume $\alpha > d - 2$. Then the problem (19) has a unique solution.*

Proof. The existence follows by the direct method of the calculus of variations. Namely let $u_k, k = 1, 2, \dots$ be a minimizing sequence. By the Sobolev Embedding Theorem, u_k has a subsequence which converges weakly in $H_\gamma^1(\Omega)$ and in $L^2(\Omega)$ towards $u \in H_\gamma^1(\Omega)$. Since \mathcal{E} is convex, it is weakly lower-semicontinuous and thus $\mathcal{E}(u) \leq \liminf_{k \rightarrow \infty} \mathcal{E}(u_k)$. Furthermore note that (33) implies that $\text{Tr}(u_k)(z) \rightarrow \text{Tr}(u)(z)$ for every $z \in \Gamma$. Thus $u = g$ on Γ . We conclude that u is the desired minimizer. The uniqueness follows from convexity of \mathcal{E} , by a standard argument, which is recalled in the proof of Lemma 2.11. \square

2.2. Elliptic problem. We now study the elliptic Euler-Lagrange equation (21). We additionally assume in this section that Ω has a $C^{2,\alpha}$ boundary and $\rho \in C^{1,\sigma}(\Omega)$ for some $\sigma > 0$.

Definition 2.8. We say that $u \in H_\gamma^1(\Gamma)$ is a weak solution of (21) if

$$(37) \quad \int_{\Omega} \gamma \rho^2 \nabla u \cdot \nabla \varphi dx = 0$$

for all $\varphi \in H_{\gamma,0}^1(\Omega)$ and $\text{Tr}[u](x) = g(x)$ for all $x \in \Gamma$.

We first need a preliminary proposition on barrier functions.

Proposition 2.9 (Barrier). *Let $\alpha > d - 2$ and fix any $0 < \beta < \alpha + 2 - d$. Then there exists $c > 0$ depending on α, β, ρ and d such that $\varphi(x) = |x|^\beta$ satisfies*

$$(38) \quad \operatorname{div}(\rho^2(1 + |x|^{-\alpha})\nabla w) \leq -\frac{\beta}{2}(\alpha + 2 - \beta - d)\rho^2|x|^{-(\alpha+2-\beta)}$$

for all $0 < |x| \leq c$.

Proof. Since $\nabla\varphi(x) = \beta|x|^{\beta-2}x$ we have

$$\begin{aligned} \operatorname{div}(\rho^2(1 + |x|^{-\alpha})\nabla w) &= \beta \operatorname{div}(\rho^2|x|^{\beta-2}x) + \beta \operatorname{div}(\rho^2|x|^{\beta-\alpha-2}x) \\ &= 2\beta|x|^{\beta-\alpha-2}(1 + |x|^\alpha)\rho\nabla\rho \cdot x + \beta\rho^2 \operatorname{div}(|x|^{\beta-2}x) + \beta\rho^2 \operatorname{div}(|x|^{\beta-\alpha-2}x) \\ &= \beta\rho^2|x|^{\beta-\alpha-2}[2(1 + |x|^\alpha)\nabla \log \rho \cdot x + (d + \beta - 2)|x|^\alpha + d + \beta - \alpha - 2] \\ &\leq \beta\rho^2|x|^{\beta-\alpha-2}[C(1 + |x|^\alpha) + d + \beta - \alpha - 2]. \end{aligned}$$

Choosing $c > 0$ so that $C(1 + c^\alpha) \leq \frac{1}{2}(\alpha + 2 - \beta - d)$ completes the proof. \square

Theorem 2.10. *Let $\alpha > d - 2$. The elliptic equation (21) has a unique weak solution $u \in H_\gamma^1(\Omega)$. Furthermore, $u \in C(\overline{\Omega}) \cap C_{loc}^{2,\sigma}(\overline{\Omega} \setminus \Gamma)$ and satisfies for every $0 < \beta < \alpha + 2 - d$*

$$(39) \quad |u(x) - u(y)| \leq C(\beta)|x - y|^\beta \quad (x \in \overline{\Omega}, y \in \Gamma).$$

Proof. For $\varepsilon > 0$ set

$$\Omega_\varepsilon := \Omega \setminus \bigcup_{y \in \Gamma} \overline{B(y, \varepsilon)}$$

and let $u_\varepsilon \in C^{2,\sigma}(\overline{\Omega_\varepsilon})$ be the unique solution of the approximating problem

$$(40) \quad \begin{cases} -\operatorname{div}(\gamma\rho^2\nabla u_\varepsilon) = 0 & \text{in } \Omega_\varepsilon \\ u_\varepsilon = g(y) & \text{on } \partial B(y, \varepsilon) \text{ for all } y \in \Gamma \\ \frac{\partial u_\varepsilon}{\partial \nu} = 0 & \text{on } \partial\Omega. \end{cases}$$

It is a classical result that u_ε is the unique solution of the variational problem

$$(41) \quad \min \left\{ \int_{\Omega_\varepsilon} \gamma\rho^2 |\nabla u|^2 dx : u \in H^1(\Omega_\varepsilon) \text{ and } \forall y \in \Gamma, u = g(y) \text{ on } \partial B(y, \varepsilon) \right\}.$$

In particular, it follows that

$$(42) \quad \sup_{\varepsilon > 0} \int_{\Omega_\varepsilon} \gamma\rho^2 |\nabla u_\varepsilon|^2 dx < \infty.$$

By the maximum principle

$$(43) \quad \min_{\Gamma} g \leq u_\varepsilon \leq \max_{\Gamma} g.$$

Let $y \in \Gamma$. By Proposition 2.9, $w(x) := |x - y|^\beta$ satisfies

$$-\operatorname{div}(\gamma\rho^2\nabla w) \geq \frac{\beta}{2}(\alpha + 2 - \beta - d)|x - y|^{-(\alpha+2-\beta)} > 0$$

for $0 < |x| \leq c$, where c depends on α, β, ρ , and d . Thus, another application of the maximum principle yields

$$u_\varepsilon(x) \leq g(y) + C|x - y|^\beta$$

for all $x \in \overline{\Omega_\varepsilon}$, where C is independent of $\varepsilon > 0$. The other direction is similar, yielding

$$(44) \quad |u_\varepsilon(x) - g(y)| \leq C|x - y|^\beta \quad \text{for all } x \in \overline{\Omega_\varepsilon}.$$

By the Schauder estimates [24], for each $\delta > 0$ there exists a constant $C > 0$, independent of ε , such that

$$\|u_\varepsilon\|_{C^{2,\sigma}(\overline{\Omega_\delta})} \leq C$$

for all $0 < \varepsilon < \delta$. Therefore, there exists a subsequence u_{ε_k} and $u \in C_{loc}^{2,\sigma}(\overline{\Omega} \setminus \Gamma)$ such that $u_{\varepsilon_k} \rightarrow u$ in $C_{loc}^2(\overline{\Omega} \setminus \Gamma)$. In particular, u solves (21) classically and satisfies (39), due to (44). Thus $u \in C(\overline{\Omega})$ and $u = g$ on Γ . Finally, it follows from (42) that $u \in H_\gamma^1(\Omega)$, and so u is a weak solution of (21), as per Definition 2.8. Uniqueness of weak solutions follows a standard energy method argument. \square

Lemma 2.11. *The solution u of the variational problem (19) is the unique weak solution of the Euler-Lagrange equation (21).*

Proof. Let $u \in H_\gamma^1(\Omega)$ be the unique weak solution of (21), and let $w \in H_\gamma^1(\Omega)$ with $\text{Tr}[w] = \text{Tr}[u]$. Then by the definition of weak solution

$$\int_{\Omega} \gamma \rho^2 \nabla u \cdot \nabla (u - w) \, dx = 0.$$

Therefore

$$\begin{aligned} \int_{\Omega} \gamma \rho^2 |\nabla u|^2 \, dx &= \int_{\Omega} \gamma \rho^2 \nabla u \cdot \nabla w \, dx \\ &= \frac{1}{2} \int_{\Omega} \gamma \rho^2 |\nabla u|^2 \, dx + \frac{1}{2} \int_{\Omega} \gamma \rho^2 |\nabla w|^2 \, dx - \frac{1}{2} \int_{\Omega} \gamma \rho^2 |\nabla u - \nabla w|^2 \, dx, \end{aligned}$$

and so

$$\int_{\Omega} \gamma \rho^2 |\nabla u|^2 \, dx = \int_{\Omega} \gamma \rho^2 |\nabla w|^2 \, dx - \int_{\Omega} \gamma \rho^2 |\nabla u - \nabla w|^2 \, dx.$$

It follows that u is the unique solution of the variational problem (19). \square

3. DISCRETE TO CONTINUUM CONVERGENCE

Consider Ω , \mathcal{X}_n , and η which satisfy the assumptions of Section 1.1. Let $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ be the empirical measure of the sample. Let $d_\infty(\mu, \mu_n)$ be the ∞ -transportation distance between μ and μ_n , discussed in Appendix A.2.

We now state our main result. In order to compare discrete and continuum minimizers we use the TL^p topology introduced in [21]. We review the topology and its basic properties in Appendix A.3.

Theorem 3.1. *Let ε_n be a sequence of positive numbers converging to zero as $n \rightarrow \infty$ and such that $\varepsilon_n \gg d_\infty(\mu, \mu_n)$. Let $\zeta_n \in (1, \infty]$ be such that $\zeta_n \gg n\varepsilon_n^2$. Consider $\alpha > d - 2$. Let u_n be a sequence of minimizers of the problem (13) for $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}$. Then (μ_n, u_n) converges in TL^2 to (μ, u) where u is the minimizer of (19).*

Our approach to proving the theorem is via establishing the Γ -convergence of the discrete constrained functionals to the continuum ones. The overall approach to consistency of learning algorithms follows the one developed in [21, 23]. Ensuring that the discrete problem induces enough regularity for one to be able to show that the label values are preserved in the limit at points of Γ follows the general strategy of [39]. However the problems and proofs are rather different. We remark that one can also use the PDE-based approach of [11], however this would require a slightly more restrictive range on ε_n . However the PDE-based approach gives superior regularity of solutions which we exploit in Section 4.

Proof. Since $\varepsilon_n \gg d_\infty(\mu, \mu_n) \gtrsim n^{-1/d}$ it follows that $\zeta_n \rightarrow \infty$ as $n \rightarrow \infty$. By Lemma 3.5, discrete energy $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}$ Γ -converges to \mathcal{E} and the compactness property holds. Therefore (μ_n, u_n) converges along a subsequence in TL^2 metric to $(\tilde{\mu}, u)$. Since μ_n converges to μ in Wasserstein metric, $\tilde{\mu} = \mu$. The fact that u is the minimizer of (19) now follows directly from Γ -convergence of Proposition 3.4 below. Consequently, the fact that the whole sequence u_n converges to u follows from the uniqueness of the minimizer of (19). \square

Remark 3.2. While above we address only algebraically growing weights γ (see (10)) it is straightforward to modify the proofs to show that if γ grows faster than algebraically at labeled points (say $\gamma(x) = \exp(1/\text{dist}(x, \Gamma))$) the conclusion of the theorem hold (in any dimension $d \geq 2$).

Remark 3.3. In this paper we assume that the data measure is supported on the set Ω of full dimension. There are no substantial obstacles in extending the results to the manifold setting where the data are sampled from a measure which is supported on a smooth submanifold of \mathbb{R}^d . One would only need to adjust the statements using manifold analogues of the weighted dirichlet energy and the laplacian. The convergence of graph laplacian in the manifold setting has already been established in the standard setting [19]. In the manifold setting the dimension d in the results above should be replaced by the dimension of the data manifold.

Proposition 3.4. *Let ε_n be a sequence of positive numbers converging to zero as $n \rightarrow \infty$ and such that $\varepsilon_n \gg d_\infty(\mu, \mu_n)$. Let $\zeta_n \in (1, \infty]$ be such that $\zeta_n \gg n\varepsilon_n^2$. Let $\alpha > d - 2$. Then the constrained properly-weighted graph Dirichlet energy, defined on $TL^2(\Omega)$ by*

$$\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}^{\text{con}}(\tilde{\mu}_n, u_n) = \begin{cases} \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n) & \text{if } \tilde{\mu}_n = \mu_n \text{ and } u_n = g \text{ on } \Gamma \\ \infty & \text{else} \end{cases}$$

Γ -converges in TL^2 to the constrained continuum properly-weighted Dirichlet energy

$$\theta_\eta \mathcal{E}^{\text{con}}(\tilde{\mu}, u) = \begin{cases} \theta_\eta \mathcal{E}(u) & \text{if } \tilde{\mu} = \mu, u \in H_\gamma^1(\Omega) \text{ and } u = g \text{ on } \Gamma \\ \infty & \text{else,} \end{cases}$$

where the value of u on Γ is considered in the sense of the trace and

$$\theta_\eta = \frac{1}{d} \int_{\mathbb{R}^d} \eta(z) |z|^2 dz.$$

Proof. To show the limsup inequality recall that \mathcal{S} , the set of smooth functions which are constant in some neighborhood of Γ , is dense in $H_\gamma^1(\Omega)$, by Corollary 2.5. The fact that for every $f \in \mathcal{S}$, $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(f) \rightarrow \theta_\eta \mathcal{E}(f)$ follows by a standard argument, which was for example presented for total variation in Section 5 of [21]. The existence of a recovery sequence for arbitrary $f \in H_\gamma^1(\Omega)$ follows by a density argument.

To show the liminf inequality consider a sequence (μ_n, u_n) converging in TL^2 to (μ, u) . We can assume without a loss of generality that $u_n|_\Gamma = g$ and that $\liminf_{n \rightarrow \infty} \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n)$ is finite. Since $\varepsilon_n \gg d_\infty(\mu, \mu_n) \gtrsim n^{-1/d}$ it follows that $\zeta_n \rightarrow \infty$ as $n \rightarrow \infty$. By Lemma 3.5, discrete energy $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}$ Γ -converges to \mathcal{E} and the compactness property holds. Thus $u \in H_\gamma^1(\Omega)$ and

$$\liminf_{n \rightarrow \infty} \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n) \geq \theta_\eta \mathcal{E}(u).$$

What remains to be shown is that $u|_\Gamma = g$. The fact that $u|_\Gamma$ is a well defined object follows from Lemma 2.1. Let us assume that $\liminf_{n \rightarrow \infty} \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n) = \lim_{n \rightarrow \infty} \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n)$. For the general case one needs to consider a subsequence, which we omit for notational simplicity. We have that $E_{\max} = \sup_n \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n) < \infty$.

We first show that near points $z \in \Gamma$, the values of u_n remain, on average, close to $g(z)$. More precisely

$$E_{max} \geq \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u) \geq \frac{1}{2n^2\varepsilon_n^2} \sum_{x \in \mathcal{X}_n} \zeta_n \eta_{\varepsilon_n}(x-z) |u(x) - g(z)|^2.$$

and thus

$$\frac{1}{n} \sum_{x \in \mathcal{X}_n} \eta_{\varepsilon_n}(x-z) |u_n(x) - g(z)|^2 \leq 2E_{max} \frac{n\varepsilon_n^2}{\zeta_n}.$$

Since $\eta \geq 1$ on $B(0,1)$

$$\frac{1}{n} \sum_{x \in \mathcal{X}_n, |x-z| < \varepsilon_n} |u_n(x) - g(z)|^2 \leq 2E_{max} \frac{n\varepsilon_n^{d+2}}{\zeta_n}.$$

Let T_n be a sequence of transport maps satisfying the conclusions of Theorem A.3 and let $\tilde{\varepsilon}_n = \varepsilon_n - 2\|T_n - \text{Id}\|_{L^\infty(\Omega)}$. Then for a.e. $x \in B(z, \tilde{\varepsilon}_n)$, $T_n(x) \in B(z, \varepsilon_n)$ and thus

$$\int_{B(z, \tilde{\varepsilon}_n)} |u_n(T_n(x)) - g(z)|^2 \rho(x) dx \leq 2E_{max} \frac{n\varepsilon_n^{d+2}}{\zeta_n}.$$

Therefore

$$(45) \quad \int_{B(z, \tilde{\varepsilon}_n)} |u_n(T_n(x)) - g(z)|^2 \rho(x) dx \lesssim \frac{n\varepsilon_n^2}{\zeta_n} \ll 1$$

by the assumption on ζ_n .

By Lemma 3.9 we know that, for $\tilde{\zeta}_n = \min\left\{\zeta_n, \left(\frac{r_0}{2C}\right)^\alpha \left(\frac{n}{\ln n}\right)^{\alpha/d}\right\}$ where C is the constant from Theorem A.3

$$(46) \quad \mathcal{E}_{\tilde{\varepsilon}_n, \tilde{\zeta}_n}(u_n \circ T_n; \tilde{\eta}) \lesssim \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n; \eta)$$

Let J be a mollifier used in the proof of Lemma 3.10 and let $\tilde{u}_n = (u_n \circ T_n) * J_{\varepsilon_n}$. From (45) follows that

$$(47) \quad |\tilde{u}_n(z) - g(z)| \lesssim \frac{n\varepsilon_n^2}{\zeta_n}$$

for all $z \in \Gamma$. Combining the estimate of the lemma with (46) yields

$$\mathcal{E}(\tilde{u}_n; \gamma_{\tilde{\zeta}_n}, \Omega'_n) \lesssim \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n; \eta).$$

Finally by Lemma 3.11 there exist $v_n \in H_\gamma^1(\Omega'_n)$ such that $v_n|_\Gamma = \tilde{u}_n|_\Gamma$ and

$$(48) \quad \mathcal{E}(v_n, \Omega'_n) \lesssim \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n; \eta).$$

We note that by construction $d_{TL^2}((\mu_n, u_n), (\mu, \tilde{u}_n)) \rightarrow 0$ and thus $\tilde{u}_n \rightarrow u$ in $L^2(\Omega)$. By construction $\|v_n - \tilde{u}_n\|_{L^2(\Omega'_n)} \lesssim \|\tilde{u}_n\|_{L^2(\Gamma_{\tilde{\varepsilon}_n})}$, where

$$(49) \quad \Gamma_{\varepsilon_n} = \{x \in \Omega : \text{dist}(x, \Gamma) < \varepsilon_n\}.$$

Since $(\mu_n, u_n) \xrightarrow{TL^2} (\mu, u)$ it follows that

$$\int_{T_n^{-1}(\Gamma_{\varepsilon_n})} |u_n(T_n(x)) - u(x)|^2 d\mu(x) \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Since $u \in L^2(\mu)$, $\lim_{\delta \rightarrow 0} \sup\{\int_A u^2(x)dx : \mu(A) < \delta\} = 0$. Therefore

$$\|\tilde{u}_n\|_{L^2(\Gamma_{\varepsilon_n})} \leq 2 \int_{T_n^{-1}(\Gamma_{\varepsilon_n})} |u_n(T_n(x)) - u(x)|^2 + u^2(x) d\mu(x) \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Thus $\|v_n - \tilde{u}_n\|_{L^2(\Omega'_n)} \rightarrow 0$ and $n \rightarrow \infty$ and consequently $v_n \rightarrow u$ in $L^2(\mu)$. Since $v_n|_{\Gamma} = \tilde{u}_n|_{\Gamma}$, from (47) and the assumption on ζ_n follows that $v_n(z) \rightarrow g(z)$ as $n \rightarrow \infty$. Furthermore, (48) implies that v_n is a bounded sequence in $H^1_{\gamma}(K)$ for any compact subset $K \subset\subset \Omega$. Combining this with the fact that $v_n \rightarrow u$ in $L^2(\mu)$ implies via estimate (33) of the Trace Theorem that $v_n(z) \rightarrow \text{Tr } u(z)$ as $n \rightarrow \infty$ for all $z \in \Gamma$. Therefore $\text{Tr } u(z) = g(z)$ for all $z \in \Gamma$. \square

Lemma 3.5. *Assume $\alpha > 0$ and $\zeta_n \geq 1$. Under assumptions on \mathcal{X}_n and ε_n of Proposition 3.4, the discrete energy $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}$ Γ -converges with respect to TL^2 topology to the energy $\theta_{\eta}\mathcal{E}$, defined in (20) as $n \rightarrow \infty$ if $\zeta_n \rightarrow \infty$. Furthermore the compactness property holds for sequences bounded in L^{∞} .*

Proof. From the results in the literature [21, 22] follows that for any fixed $\zeta > 0$ the discrete energies $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta}$ Γ -converge to $\mathcal{E}(\cdot; \gamma_{\zeta})$ as $n \rightarrow \infty$, under standard assumptions on ε_n . To show the liminf inequality for general ζ_n consider a sequence (μ_n, u_n) TL^2 converging to (μ, u) . For any fixed k ,

$$\liminf_{n \rightarrow \infty} \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n) \geq \mathcal{E}(u; \zeta_k),$$

which implies the desired inequality by taking supremum over k . The limsup inequality follows by a simple diagonalization argument.

We recall from the literature (e.g [21] or Proposition 4.4 of [39]) that the compactness already holds for the weight $\gamma \equiv 1$. Thus the compactness for $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}$ follows by comparison. \square

3.0.1. The negative result.

Proposition 3.6. *Let ε_n be a sequence of positive numbers converging to zero as $n \rightarrow \infty$ and such that $\varepsilon_n \gg d_{\infty}(\mu, \mu_n)$. Let $\zeta_n \geq 1$ be a sequence converging to infinity. Consider $\alpha \leq d - 2$. Then the constrained energy $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}$, defined in Proposition 3.4, Γ -converges in TL^2 metric to the unconstrained continuum energy \mathcal{E} .*

Proof. The liminf part of the Γ -convergence claim follows from the liminf claim of Lemma 3.5.

To show the limsup inequality, we first observe that by localizing near the points of Γ , and given that limsup inequality holds for the unconstrained functional, the problem can be reduced to considering $\Gamma = \{0\}$, $u \equiv 0$, and the construction a sequence of functions $u_n \in L^2(\mu_n)$ such that $u_n(0) = 1$, $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n) \rightarrow 0$ as $n \rightarrow \infty$ and $u_n \rightarrow 0$ in TL^2 as $n \rightarrow \infty$.

We now make some observation about the continuum functional. Namely when $\alpha \leq d - 2$ then the function $\varphi(x) = \ln\left(\ln\left(\frac{1}{|x|}\right)\right)$ belongs to $H^1_{\gamma}(B(0, 1))$. Let $w_k = \max\{\min\{\frac{1}{k}\varphi(x), 1\}, 0\}$. Let $r_k > 0$ be such that $w_k = 1$ on $B(0, r_k)$. By mollifying we can obtain a smooth approximation v_k , $v_k = 1$ on $B(0, r_k/2)$ and $\|v_k\|_{H^1_{\gamma}(B(0,1))} \leq 2\|u_k\|_{H^1_{\gamma}(B(0,1))}$. Arguing as in Section 5 of [21], if one defines for each $k \in \mathbb{N}$, the sequence $u_n^k \in L^2(\mu_n)$ by $u_n^k(x_i) = v_n(x_i)$ for all $x_i \in V_n$ one has $u_n^k \rightarrow v_n$ in TL^2 and $\limsup_{n \rightarrow \infty} \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n^k) \geq \mathcal{E}(v_n)$. Since $v_k \rightarrow 0$ in $H^1_{\gamma}(B(0, 1))$ as $k \rightarrow \infty$, the conclusion follows by a diagonalization argument. \square

Corollary 3.7. *Let ε_n be a sequence of positive numbers converging to zero as $n \rightarrow \infty$ and such that $\varepsilon_n \gg d_{\infty}(\mu, \mu_n)$. Let $\zeta_n \geq 1$ be a sequence converging to ∞ as $n \rightarrow \infty$. Let $\alpha \leq d - 2$. Let u_n be a sequence of minimizers of the problem (13) for $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}$. Let c_n be the average of u_n (with respect to measure μ_n). Then $(\mu_n, u_n - c_n)$ converges in TL^2 to $(\mu, 0)$; in other words the information about the labels is forgotten in the limit.*

Proof. Assume the claim is false. Then there exists $\delta > 0$ and a subsequence u_{n_j} such that for all j , $d_{TL^2}((\mu_{n_j}, u_{n_j} - c_{n_j}), (\mu, 0)) > \delta$ for all j . By the maximum principle functions u_n are bounded by extremal values of g . Consequently, by Lemma 3.5, $u_{n_j} - c_{n_j}$ has a further convergent subsequence. Without a loss of generality we can assume that $u_{n_j} - c_{n_j}$ converges to some $v \in L^2(\mu)$. Then $\int v d\mu = \lim_{j \rightarrow \infty} \int u_{n_j} - c_{n_j} d\mu_{n_j} = 0$.

By the limsup part of Γ -convergence of Proposition 3.6 there exists a sequence $v_n \in L^2(\mu_n)$ such that $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}^{con}(v_n) \rightarrow 0$ as $n \rightarrow \infty$. Since u_{n_j} are minimizers $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_{n_j} - c_{n_j}) \rightarrow 0$ as $n \rightarrow \infty$. We conclude by the liminf part of Γ -convergence that $\mathcal{E}(v) = 0$. Since $\int v d\mu = 0$ this implies that $v \equiv 0$, which contradicts the assumption about the sequence. \square

We note that the analogue of the negative result in Corollary 3.7 for the standard graph Laplacian (corresponding to $\gamma \equiv 1$) was proved in [39][Theorem 2.1]. The following corollary then follows by the squeeze theorem for Γ -convergence.

Corollary 3.8. *Under the assumptions of Proposition 3.6 consider any sequence of graph based functionals \mathcal{F}_n such that for $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,1} \leq \mathcal{F}_n \leq \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}$ (where we note that $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,1}$ is just a convenient way to write the standard graph Laplacian). Let u_n be the minimizers of (13) for \mathcal{F}_n and let c_n be the average of u_n (with respect to measure μ_n). Then $(\mu_n, u_n - c_n)$ converges in TL^2 to $(\mu, 0)$.*

A particular consequence of this corollary is that the minimizers of the algorithm in [36] converge to a constant as $n \rightarrow \infty$.

3.1. Estimates for the discrete to continuum convergence. Here we establish several results needed in the proofs of the main results above. We follow a similar strategy as [39]. Let us define the nonlocal continuum energy as

$$(50) \quad \mathcal{E}_{\varepsilon,\zeta}(u) = \frac{1}{\varepsilon^2} \iint \gamma_\zeta(x)\eta_\varepsilon(x-y)|u(x) - u(y)|^2 d\mu(x)d\mu(y).$$

It serves as an intermediary between the discrete graph based functionals and the continuum derivative-based functionals.

Lemma 3.9 (discrete to nonlocal control). *Consider Ω , μ , η , ζ , and x_i as in Theorem 3.1. Let $\tilde{\eta}(|x|) = 1$ for $|x| \leq 1$ and $\tilde{\eta}(|x|) = 0$ otherwise, and so $\tilde{\eta} \leq \eta$. Let T_n be a sequence of transport maps satisfying the conclusions of Theorem A.3 and let $\tilde{\varepsilon}_n = \varepsilon_n - 2\|T_n - \text{Id}\|_{L^\infty(\Omega)}$. Define $\mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(\cdot; \eta)$ by (12), where we explicitly denote the dependence of η . Let $\tilde{\zeta}_n > 0$ be such that $\zeta_n \geq \tilde{\zeta}_n$ and $(\frac{n_0}{2C})^\alpha (\frac{n}{\ln n})^{\alpha/d} \geq \tilde{\zeta}_n$ where C is the constant from Theorem A.3. Then there exists $n_0 \in \mathbb{N}$ and a constant $\bar{C} > 0$ (independent of n and u_n) such that for all $n \geq n_0$*

$$\mathcal{E}_{\tilde{\varepsilon}_n, \tilde{\zeta}_n}(u_n \circ T_n; \tilde{\eta}) \leq \bar{C} \mathcal{G}\mathcal{E}_{n,\varepsilon_n,\zeta_n}(u_n; \eta)$$

Proof. If $\left| \frac{x-z}{\tilde{\varepsilon}_n} \right| < 1$ then

$$|T_n(x) - T_n(z)| \leq 2\|T_n - \text{Id}\|_{L^\infty(\Omega)} + |x - z| \leq 2\|T_n - \text{Id}\|_{L^\infty(\Omega)} + \tilde{\varepsilon}_n = \varepsilon_n.$$

So,

$$\left| \frac{x-z}{\tilde{\varepsilon}_n} \right| < 1 \quad \text{implies} \quad \left| \frac{T_n(x) - T_n(z)}{\varepsilon_n} \right| \leq 1$$

and therefore

$$\left| \frac{x-z}{\tilde{\varepsilon}_n} \right| < 1 \quad \text{implies} \quad \tilde{\eta}\left(\frac{|x-z|}{\tilde{\varepsilon}_n}\right) = 1 = \tilde{\eta}\left(\frac{|T_n(x) - T_n(z)|}{\varepsilon_n}\right).$$

Hence,

$$\tilde{\eta}\left(\frac{|x-z|}{\tilde{\varepsilon}_n}\right) \leq \tilde{\eta}\left(\frac{|T_n(x)-T_n(z)|}{\varepsilon_n}\right) \leq \eta\left(\frac{|T_n(x)-T_n(z)|}{\varepsilon_n}\right).$$

From the assumptions on $\tilde{\zeta}_n$ and T_n follows that

$$2\|T_n - \text{Id}\|_{L^\infty(\Omega)} \leq r_0(\tilde{\zeta}_n - 1)^{-\frac{1}{\alpha}} = r_{\tilde{\zeta}_n}$$

where $r_{\tilde{\zeta}_n}$ is the length scale such that $1 + \left(\frac{r_0}{\text{dist}(x, \Gamma)}\right)^\alpha > \tilde{\zeta}_n$ if $\text{dist}(x, \Gamma) < r_{\tilde{\zeta}_n}$. We claim that for a.e. $x \in \Omega$

$$(51) \quad \min\left\{1 + \left(\frac{r_0}{\text{dist}(x, \Gamma)}\right)^\alpha, \tilde{\zeta}_n\right\} \leq 2^\alpha \min\left\{1 + \left(\frac{r_0}{\text{dist}(T_n(x), \Gamma)}\right)^\alpha, \tilde{\zeta}_n\right\}$$

Namely if $d(x, \Gamma) \leq r_{\tilde{\zeta}_n}$ then $d(T_n, \Gamma) \leq |T_n(x) - x| + r_{\tilde{\zeta}_n} \leq 2r_{\tilde{\zeta}_n}$ for a.e. such x . Thus

$$1 + \left(\frac{r_0}{\text{dist}(T_n(x), \Gamma)}\right)^\alpha \geq 1 + \frac{1}{2^\alpha} \left(\frac{r_0}{r_{\tilde{\zeta}_n}}\right)^\alpha \geq \frac{1}{2^\alpha} \tilde{\zeta}_n$$

If $d(x, \Gamma) > r_{\tilde{\zeta}_n}$ then $d(x, \Gamma) \geq \frac{1}{2}d(T_n(x), \Gamma)$ for a.e. such x . Thus $\left(\frac{r_0}{\text{dist}(T_n(x), \Gamma)}\right)^\alpha \geq \frac{1}{2^\alpha} \left(\frac{r_0}{\text{dist}(x, \Gamma)}\right)^\alpha$.

Using (51) we conclude

$$\begin{aligned} & \mathcal{E}_{\tilde{\varepsilon}_n, \tilde{\zeta}_n}(u_n \circ T_n; \tilde{\eta}) \\ &= \frac{1}{\tilde{\varepsilon}_n^2} \iint \min\left\{1 + \left(\frac{r_0}{\text{dist}(x, \Gamma)}\right)^\alpha, \tilde{\zeta}_n\right\} \tilde{\eta}_{\tilde{\varepsilon}_n}(x-y) |u_n(T_n(x)) - u_n(T_n(y))|^2 d\mu(x) d\mu(y) \\ &\leq 2^\alpha \frac{\varepsilon_n^d}{\tilde{\varepsilon}_n^{d+2}} \iint \min\left\{1 + \left(\frac{r_0}{\text{dist}(T_n(x), \Gamma)}\right)^\alpha, \zeta_n\right\} \\ &\quad \eta_{\varepsilon_n}(|T_n(x) - T_n(y)|) |u_n(T_n(x)) - u_n(T_n(z))|^2 d\mu(x) d\mu(y) \\ &= 2^\alpha \frac{\varepsilon_n^{d+2}}{\tilde{\varepsilon}_n^{d+2}} \mathcal{G} \mathcal{E}_{n, \varepsilon_n, \zeta_n}(u_n; \eta). \end{aligned}$$

□

In the next lemma we show that boundedness of non-local energies implies regularity at scales greater than ε with weight $\gamma_{\tilde{\zeta}}$. This allows us to relate non-local bounds to local bounds after mollification using a mollifier $J \in C_c^\infty(\mathbb{R}^d, [0, \infty))$, with $\int_{\mathbb{R}^d} J(x) dx = 1$, and $J_\varepsilon(x) = \frac{1}{\varepsilon^d} J(x/\varepsilon)$.

Lemma 3.10 (nonlocal to weak local control). *There exists a constant $C \geq 1$ and a radially symmetric mollifier J with $\text{supp}(J) \subseteq \overline{B(0, 1)}$ such that for all $\varepsilon > 0$, $u \in L^2(\Omega)$, and any $\Omega' \subset\subset \Omega$ (i.e. for every Ω' that is compactly contained in Ω) with $\text{dist}(\Omega', \partial\Omega) > \varepsilon$ it holds that*

$$(52) \quad \mathcal{E}(u * J_\varepsilon; \gamma_{\tilde{\zeta}}, \Omega') \leq C \mathcal{E}_{\varepsilon, \tilde{\zeta}}(u; \Omega)$$

where for both functionals we explicitly denote the dependence of the domain.

Proof. Let J be a radially symmetric mollifier whose support is contained in $\overline{B(0,1)}$. There exists $\beta > 0$ such that $J \leq \beta\eta(|\cdot|)$ and $|\nabla J| \leq \beta\eta(|\cdot|)$. Let $u_\varepsilon = J_\varepsilon * u$. For arbitrary $x \in \Omega$ with $\text{dist}(x, \partial\Omega) > \varepsilon$ we have

$$\begin{aligned} |\nabla u_\varepsilon(x)| &= \left| \int_{\Omega} \nabla J_\varepsilon(x-z)u(z) dz \right| \\ &= \left| \int_{\Omega} \nabla J_\varepsilon(x-z)(u(z) - u(x)) dz - \int_{\mathbb{R}^d \setminus \Omega} \nabla J_\varepsilon(x-z)u(x) dz \right| \\ &\leq \frac{\beta}{\varepsilon^{d+1}} \int_{\Omega} \eta\left(\frac{|x-z|}{\varepsilon}\right) |u(z) - u(x)| dz + \frac{1}{\varepsilon^{d+1}} \int_{\mathbb{R}^d \setminus \Omega} \left| \nabla J\left(\frac{x-z}{\varepsilon}\right) \right| |u(x)| dz. \end{aligned}$$

where the second line follows from $\int_{\mathbb{R}^d} \nabla J(w) dw = 0$. For the second term we have

$$\frac{1}{\varepsilon^{d+1}} \int_{\mathbb{R}^d \setminus \Omega} \left| \nabla J\left(\frac{x-z}{\varepsilon}\right) \right| |u(x)| dz = 0$$

since for all $z \in \mathbb{R}^d \setminus \Omega$ and $x \in \Omega$ with $\text{dist}(x, \partial\Omega) > \varepsilon$ it follows that $|x-z| > \varepsilon$ and thus $\nabla J\left(\frac{x-z}{\varepsilon}\right) = 0$. Therefore, for $\theta_\eta = \int_{\mathbb{R}^d} \eta(|w|) dw$,

$$\begin{aligned} |\nabla u_\varepsilon(x)|^2 &\leq \beta^2 \left(\int_{\Omega} \frac{1}{\varepsilon} \eta_\varepsilon(|x-z|) |u(z) - u(x)| dz \right)^2 \\ &= \frac{\theta_\eta^2 \beta^2}{\varepsilon^2} \left(\int_{\Omega} \frac{\eta_\varepsilon(|x-z|)}{\theta_\eta} |u(z) - u(x)| dz \right)^2 \\ &\leq \theta_\eta \beta^2 \int_{\Omega} \eta_\varepsilon(|x-z|) \frac{|u(z) - u(x)|^2}{\varepsilon^2} dz \end{aligned}$$

by Jensen's inequality (since $\frac{1}{\theta_\eta} \int_{\mathbb{R}^d} \eta_\varepsilon(|x-z|) dz = 1$). Hence,

$$\begin{aligned} \int_{\Omega'} |\nabla u_\varepsilon(x)|^2 \gamma_\zeta(x) \rho^2(x) dx &\leq \theta_\eta \beta^2 \int_{\Omega} \int_{\Omega} \eta_\varepsilon(|x-z|) \left| \frac{u(z) - u(x)}{\varepsilon} \right|^2 \gamma_\zeta(x) \rho^2(x) dz dx \\ &\leq \frac{\theta_\eta \beta^2 \sup_{x \in \Omega} \rho(x)}{\inf_{x \in \Omega} \rho(x)} \mathcal{E}_{\varepsilon, \zeta}(u; \Omega) \end{aligned}$$

which completes the proof. \square

We now show that controlling the local energy with cut-off near the singularity is sufficient to be able to find a nearby (in H_γ^1) function which has a similarly bounded energy without a cut-off.

Lemma 3.11 (weak local to strong local control). *Assume $\tilde{\zeta} > 1$ and $r_\zeta \leq \frac{1}{2} \min\{|x-y| : x, y \in \Gamma, x \neq y\}$. Let $\bar{r} = r_{\tilde{\zeta}}$. Then there exists a constant $C > 0$ such that for every $u \in H^1(\Omega)$ there exists $v \in H_\gamma^1(\Omega)$ such that*

$$(53) \quad v|_{B(z, \bar{r}/2)} \equiv \frac{1}{|B(0, \bar{r}/2)|} \int_{B(x, \bar{r}/2)} u(x) dx \quad \text{for all } z \in \Gamma,$$

$$(54) \quad v = u \quad \text{on } \Omega \setminus \Gamma_{\bar{r}}$$

$$(55) \quad \mathcal{E}(v; \gamma) \leq C \mathcal{E}(u, \gamma_{\tilde{\zeta}}).$$

Proof. Using the finiteness of Γ , from Lemma A.4 via translations and a rescaling follows that there exists $c \geq 1$, and $v \in H^1(\Omega)$ satisfying (53) and (54) such that

$$\int_{\Gamma_{\bar{r}}} |\nabla v(x)|^2 dx \leq c \int_{\Gamma_{\bar{r}}} |\nabla u(x)|^2 dx.$$

Using that $\nabla v = 0$ on $\Gamma_{\frac{\bar{r}}{2}}$ and $1 + \left(\frac{r_0}{\bar{r}}\right)^\alpha = \tilde{\zeta}$ we obtain

$$\begin{aligned} \mathcal{E}(v; \gamma) &= \int_{\Gamma_{\bar{r}}} \gamma(x) |\nabla v(x)|^2 \rho(x) dx + \int_{\Omega \setminus \Gamma_{\bar{r}}} \gamma(x) |\nabla u(x)|^2 \rho(x) dx \\ &\leq \int_{\Gamma_{\bar{r}} \setminus \Gamma_{\frac{\bar{r}}{2}}} \left(1 + \left(\frac{2r_0}{\bar{r}}\right)^\alpha\right) |\nabla v(x)|^2 \rho(x) dx + \int_{\Omega \setminus \Gamma_{\bar{r}}} \gamma_{\tilde{\zeta}} |\nabla u(x)|^2 \rho(x) dx \\ &\leq 2^\alpha \int_{\Gamma_{\bar{r}} \setminus \Gamma_{\frac{\bar{r}}{2}}} \tilde{\zeta} |\nabla v(x)|^2 \rho(x) dx + \int_{\Omega \setminus \Gamma_{\bar{r}}} \gamma_{\tilde{\zeta}} |\nabla u(x)|^2 \rho(x) dx \\ &\leq 2^\alpha c \int_{\Gamma_{\bar{r}}} \tilde{\zeta} |\nabla u(x)|^2 \rho(x) dx + \int_{\Omega \setminus \Gamma_{\bar{r}}} \gamma_{\tilde{\zeta}} |\nabla u(x)|^2 \rho(x) dx \\ &\leq 2^\alpha c \int_{\Omega} \gamma_{\tilde{\zeta}} |\nabla u(x)|^2 \rho(x) dx \end{aligned}$$

□

4. REGULARITY OF MINIMIZERS OF THE GRAPH PROPERLY-WEIGHTED LAPLACIAN

4.1. Hölder estimate near labeled points. Our main result in this section is a type of Hölder estimate near the labeled points, which shows that solutions of the graph-based learning problem (14) attain their boundary values on Γ continuously, with high probability. The proof is a graph-based version of the barrier argument from Theorem 2.10 that established continuity at labels in the continuum PDE, given in Eq. (39). Barrier arguments for proving Hölder regularity of solutions of PDEs are standard techniques for first order equations, such as Hamilton-Jacobi equations [4]. Normally, barrier arguments do not work for second order elliptic equations (since fundamental solutions are unbounded), though there are a handful of exceptions, such as the p -Laplace equation for $p > d$ [11], level set equations for affine curvature motion [12], and our continuum equation (21).

Our proof uses the barrier $v(x) = C|x - y|^\beta$, which is a supersolution of the continuum PDE (21) for $0 < \beta < \alpha + 2 - d$, due to Proposition 2.9. For $\beta < 2$, the barrier has a singularity at $x = y$, which has to be treated carefully in the translation to the graph setting. We show in Lemma 4.6 that v is a supersolution on the graph with high probability away from a small ball $B(y, C\varepsilon)$. Due to the singularity in the barrier, we cannot prove the supersolution property within the ball $B(y, C\varepsilon)$. To fill in the gap within this ball, we require a local regularity result, given in Lemma 4.8, that relies on the variational structure of the problem. At a high level, the proof is similar to the proof of Hölder regularity of solutions to the graph-based game theoretic p -Laplace equation, given in [11], though many of the ingredients are different. In particular, in [11] there is no variational interpretation of the problem, and the local argument utilizes another barrier construction.

We now proceed to present the main results in this section. Throughout we always assume $n\varepsilon^d \geq 1$. Our main result is the following Hölder-type estimate.

Theorem 4.1 (Hölder estimate). *Assume $0 < \varepsilon \leq 1$, $\alpha > d - 2$ and $\zeta \geq 1 + \varepsilon^{-\alpha}$. Let $u \in L^2(\mathcal{X}_n)$ be the solution of (14) and fix any $0 < \beta < \alpha + 2 - d$. For each $z \in \Gamma$ the event*

that

$$(56) \quad |u(x_i) - u(z)| \leq C|x_i - z|^\beta + Cn^{1/2}\varepsilon^{1+\alpha/2}$$

holds for all $x_i \in \mathcal{X}_n$ occurs with probability at least $1 - C \exp(-cn\varepsilon^{d+4} + \log(n))$.

The proof of Theorem 4.1, given at the end of the section, relies on some preliminary results that we establish after a few remarks.

Remark 4.2. For the result in Theorem 4.1 to be useful, we must choose $\varepsilon_n \rightarrow 0$ so that $n\varepsilon_n^{d+4} \gg \log(n)$ and $n\varepsilon_n^{\alpha+2} \ll 1$. Therefore, we must have $\alpha > d + 2$ and

$$(57) \quad \left(\frac{\log(n)}{n}\right)^{1/(d+4)} \ll \varepsilon_n \ll \left(\frac{1}{n}\right)^{1/(\alpha+2)}.$$

Remark 4.3. If we replace γ_ζ with $\gamma_{\zeta, C\varepsilon}$, as defined in Remark 1.3, then we can improve Theorem 4.1 to read

$$(58) \quad |u(x_i) - u(z)| \leq C|x_i - z|^\beta + C\zeta^{-1/2}n^{1/2}\varepsilon,$$

under the same assumptions and with the same probability, except we also require $\zeta \geq 1 + C\varepsilon^{-\alpha}$. In this model, the restrictive upper bound in (57) is not required.

We now turn to the proof of Theorem 4.1. We first recall a useful lemma from [11].

Lemma 4.4 (Remark 7 from [11]). *Let $Y_1, Y_2, Y_3, \dots, Y_n$ be a sequence of i.i.d random variables on \mathbb{R}^d with Lebesgue density $\rho : \mathbb{R}^d \rightarrow \mathbb{R}$, let $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$ be bounded and Borel measurable with compact support in a ball $B(x, h)$ for some $h > 0$, and define*

$$Y = \sum_{i=1}^n \psi(Y_i).$$

Then for any $0 \leq \lambda \leq 1$

$$(59) \quad \mathbb{P}\left(|Y - \mathbb{E}(Y)| \geq C\|\psi\|_{L^\infty(B(x, h))}nh^d\lambda\right) \leq 2\exp(-cnh^d\lambda^2),$$

for all $0 < \lambda \leq 1$, where $C, c > 0$ are constants depending only on $\|\rho\|_{L^\infty(B(x, h))}$ and d .

We can use Lemma 4.4 to prove pointwise consistency for our properly-weighted graph Laplacian. It extends, in a refined form, the results of [11][Theorem 5]. It is related to well known results on the pointwise consistency of the graph Laplacian [38].

For simplicity we set

$$(60) \quad \Delta_\rho \varphi = \rho^{-1} \operatorname{div}(\gamma \rho^2 \nabla \varphi).$$

Theorem 4.5. *For $\delta > 0$, let $D_{n, \varepsilon, \delta}$ be the event that*

$$(61) \quad \left| \mathcal{GL}_{n, \varepsilon, \infty} \varphi(x_i) - \frac{1}{2} \sigma_\eta \Delta_\rho \varphi(x_i) \right| \leq C((\varepsilon \beta_1 + \varepsilon^2 \beta_2)M^{-(\alpha+2)} + \varepsilon \beta_2 M^{-(\alpha+1)} + \delta \beta_3 M^{-\alpha})$$

holds for all x_i with $2\varepsilon < \operatorname{dist}(x_i, \Gamma) \leq R/4$ and all $\varphi \in C^3(B(x_i, 2\varepsilon))$, where $\beta_k = \|\varphi\|_{C^k(B(x_i, 2\varepsilon))}$ and $M = \operatorname{dist}(x_i, \Gamma) - 2\varepsilon$. Then for $\varepsilon \leq \delta \leq 1$ we have

$$(62) \quad \mathbb{P}(D_{n, \varepsilon, \delta}) \geq 1 - C \exp(-c\delta^2 n \varepsilon^{d+2} + \log(n))$$

Proof. Let us write \mathcal{GL} in place of $\mathcal{GL}_{n, \varepsilon, \infty}$ for simplicity. We also define

$$w(x, y) = \frac{1}{2n\varepsilon^2} (\gamma(x) + \gamma(y)) \eta_\varepsilon(x - y).$$

Then

$$\mathcal{GL}u(x) = \sum_{y \in \mathcal{X}_n} w(x, y)(u(y) - u(x)).$$

By conditioning on the location of $x \in X_n$, we can assume without loss of generality that $x \in \Omega$ is a fixed (non-random) point, $B(x, 2\varepsilon) \subset \Omega$ and $\text{dist}(x, \Gamma) > 2\varepsilon$. Let $\varphi \in C^3(B(x, 2\varepsilon))$, $p = D\varphi(x)$ and $A = D^2\varphi(x)$. Note that

$$(63) \quad \begin{aligned} \mathcal{GL}\varphi(x) &= \sum_{i=1}^d p_i \sum_{y \in X_n} w(x, y)(y_i - x_i) + \frac{1}{2} \sum_{i,j=1}^d a_{ij} \sum_{y \in X_n} w(x, y)(y_i - x_i)(y_j - x_j) \\ &\quad + O(\varepsilon^3 \beta_3 \text{deg}(x)), \end{aligned}$$

where $\text{deg}(x)$ is the degree given by

$$\text{deg}(x) = \sum_{y \in \mathcal{X}_n} w(x, y).$$

Since $\text{dist}(y, \Gamma) \geq \text{dist}(x, \Gamma) - 2\varepsilon$ we have

$$\begin{aligned} w(x, y) &\leq \frac{C}{2n\varepsilon^{d+2}} \left(\frac{1}{\text{dist}(x, \Gamma)^\alpha} + \frac{1}{\text{dist}(y, \Gamma)^\alpha} \right) \\ &\leq \frac{C}{n\varepsilon^{d+2}(\text{dist}(x, \Gamma) - 2\varepsilon)^\alpha}. \end{aligned}$$

By Lemma 4.4

$$\left| \text{deg}(x) - n \int_{B(x, 2\varepsilon)} w(x, y) \rho(y) dy \right| \leq C\varepsilon^{-2}(\text{dist}(x, \Gamma) - 2\varepsilon)^{-\alpha},$$

holds with probability at least $1 - 2 \exp(-cn\varepsilon^d)$. This implies

$$\text{deg}(x) \leq C\varepsilon^{-2}(\text{dist}(x, \Gamma) - 2\varepsilon)^{-\alpha}.$$

By another application of Lemma 4.4, both

$$\left| \sum_{y \in X_n} w(x, y)(y_i - x_i) - n \int_{B(x, 2\varepsilon)} w(x, y)(y_i - x_i) \rho(y) dy \right| \geq C\delta(\text{dist}(x, \Gamma) - 2\varepsilon)^{-\alpha},$$

and

$$\begin{aligned} \left| \sum_{y \in X_n} w(x, y)(y_i - x_i)(y_j - x_j) - n \int_{B(x, 2\varepsilon)} w(x, y)(y_i - x_i)(y_j - x_j) \rho(y) dy \right| \\ \geq C\delta\varepsilon(\text{dist}(x, \Gamma) - 2\varepsilon)^{-\alpha}, \end{aligned}$$

occur with probability at most $2 \exp(-c\delta^2 n\varepsilon^{d+2})$ provided $0 < \delta\varepsilon \leq 1$. Thus, if $\varepsilon \leq \delta \leq \varepsilon^{-1}$ we have

$$(64) \quad \begin{aligned} \mathcal{GL}\varphi(x) &= \frac{1}{2} \int_{B(0, 2)} (\gamma(x) + \gamma(x + z\varepsilon)) \rho(x + z\varepsilon) \boldsymbol{\eta}(|z|) \left(\frac{1}{\varepsilon} p \cdot z + \frac{1}{2} z \cdot Az \right) dz \\ &\quad + O(\delta\beta_3(\text{dist}(x, \Gamma) - 2\varepsilon)^{-\alpha}) \end{aligned}$$

holds for all $\varphi \in C^3(\mathbb{R}^d)$ with probability at least $1 - C \exp(-c\delta^2 n \varepsilon^{d+2})$. Note that

$$\begin{aligned} \frac{1}{2}(\gamma(x) + \gamma(x + z\varepsilon))\rho(x + z\varepsilon) &= \gamma(x)\rho(x) + \gamma(x)\nabla\rho(x) \cdot z\varepsilon + \frac{1}{2}\rho(x)\nabla\gamma(x) \cdot z\varepsilon \\ &\quad + O(\varepsilon^2\|\gamma\|_{C^2(B(x,2\varepsilon))} + \varepsilon^3\|\gamma\|_{C^1(B(x,2\varepsilon))}). \end{aligned}$$

We now have

$$\begin{aligned} \int_{B(0,2)} \gamma(x)\rho(x)\boldsymbol{\eta}(|z|) \left(\frac{1}{\varepsilon}p \cdot z + \frac{1}{2}z \cdot Az \right) dz &= \frac{1}{2}\gamma(x)\rho(x) \sum_{i,j=1}^d a_{ij} \int_{B(0,2)} \boldsymbol{\eta}(|z|)z_i z_j dz \\ &= \frac{1}{2}\gamma(x)\rho(x) \sum_{i=1}^d a_{ii} \int_{B(0,2)} \boldsymbol{\eta}(|z|)z_i^2 dz \\ &= \frac{\sigma_\eta}{2}\gamma(x)\rho(x)\text{Trace}(A), \end{aligned}$$

$$\begin{aligned} \int_{B(0,2)} \gamma(x)(\nabla\rho(x) \cdot z\varepsilon)\boldsymbol{\eta}(|z|) \left(\frac{1}{\varepsilon}p \cdot z + \frac{1}{2}z \cdot Az \right) dz \\ &= \gamma(x)\nabla\rho(x) \cdot \int_{B(0,2)} \boldsymbol{\eta}(|z|)(p \cdot z)z dz + O(\varepsilon\beta_2 \text{dist}(x, \Gamma)^{-\alpha}) \\ &= \gamma(x)\nabla\rho(x) \cdot \sum_{i=1}^d p_i \int_{B(0,2)} \boldsymbol{\eta}(|z|)z_i z dz + O(\varepsilon\beta_2 \text{dist}(x, \Gamma)^{-\alpha}) \\ &= \sigma_\eta\gamma(x)\nabla\rho(x) \cdot p + O(\varepsilon\beta_2 \text{dist}(x, \Gamma)^{-\alpha}), \end{aligned}$$

and

$$\begin{aligned} \frac{1}{2} \int_{B(0,2)} \rho(x)(\nabla\gamma(x) \cdot z\varepsilon)\boldsymbol{\eta}(|z|) \left(\frac{1}{\varepsilon}p \cdot z + \frac{1}{2}z \cdot Az \right) dz \\ &= \frac{\sigma_\eta}{2}\rho(x)\nabla\gamma(x) \cdot p + O(\varepsilon\beta_2\|\gamma\|_{C^1(B(x,2\varepsilon))}). \end{aligned}$$

Assembling these together with (64) we have that

$$\begin{aligned} \mathcal{GL}\varphi(x) &= \frac{\sigma_\eta}{2\rho(x)} \text{div}(\gamma\rho^2\nabla\varphi) + O\left(\varepsilon\beta_1\|\gamma\|_{C^2(B(x,2\varepsilon))} + \varepsilon\beta_2\|\gamma\|_{C^1(B(x,2\varepsilon))}\right) \\ &\quad + \varepsilon^2\beta_2\|\gamma\|_{C^2(B(x,2\varepsilon))} + \delta\beta_3(\text{dist}(x, \Gamma) - 2\varepsilon)^{-\alpha} \end{aligned}$$

holds for all $\varphi \in C^3(B(x, 2\varepsilon))$ with probability at least $1 - C \exp(-c\delta^2 n \varepsilon^{d+2})$. The proof is completed by computing

$$\|\gamma\|_{C^k(B(x,2\varepsilon))} \leq C(\text{dist}(x, \Gamma) - 2\varepsilon)^{-\alpha-k},$$

and applying a union bound over x_1, \dots, x_n . \square

We now establish that the function $|x|^\beta$ for $0 < \beta < \alpha + 2 - d$ serves as a barrier (e.g., is a supersolution) on the graph with high probability.

Lemma 4.6 (Barrier lemma). *Let $\alpha > d - 2$ and fix any $0 < \beta < \alpha + 2 - d$. For $y \in \Gamma$ define $\varphi(x) = |x - y|^\beta$. Then the event that*

$$(65) \quad \mathcal{GL}_{n,\varepsilon,\infty}\varphi(x_i) \leq -c|x_i - y|^{-(\alpha+2-\beta)}$$

for all x_i with $C\varepsilon < |x_i - y| \leq c$ occurs with probability at least $1 - C \exp(-cn\varepsilon^{d+4} + \log(n))$.

Proof. Let us write \mathcal{GL} in place of $\mathcal{GL}_{n,\varepsilon,\infty}$ for simplicity. We use Theorem 4.5 and Proposition 2.9. Note in Theorem 4.5 that if we restrict $3 \leq \varepsilon|x_i - y| \leq R/4$ then

$$M = \text{dist}(x_i, \Gamma) - 2\varepsilon = |x_i - y| - 2\varepsilon \geq \frac{1}{3}|x_i - y|.$$

Also, for $\beta_k = \|\varphi\|_{C^k(B(x_i, 2\varepsilon))}$ we compute

$$\beta_k \leq C\beta|x_i - y|^{\beta-k}.$$

Hence, setting $\delta = \varepsilon$ in Theorem 4.5 we obtain that

$$(66) \quad \left| \mathcal{GL}\varphi(x_i) - \frac{1}{2}\sigma_\eta\Delta_\rho\varphi(x_i) \right| \leq C\varepsilon|x_i - y|^{\beta-\alpha-3}(1 + \varepsilon|x_i - y|^{-1})$$

holds for all x_i with $3\varepsilon \leq |x_i - y| \leq r$ with probability at least

$$1 - C \exp\left(-cn\varepsilon^{d+4} + \log(n)\right).$$

For the rest of the proof we restrict to the event that (66) holds.

Note that since $\beta - \alpha - 3 < 0$ and $|x_i - y| \geq 3\varepsilon$, it follows from (66) that

$$(67) \quad \mathcal{GL}\varphi(x_i) \leq \frac{1}{2}\sigma_\eta\Delta_\rho\varphi(x_i) + C\varepsilon|x_i - y|^{\beta-\alpha-3}.$$

Combining this with Proposition 2.9 we have

$$\mathcal{GL}\varphi(x_i) \leq -c|x_i - y|^{\beta-\alpha-2} + C\varepsilon|x_i - y|^{\beta-\alpha-3} = -c|x_i - y|^{\beta-\alpha-2}(1 - C\varepsilon|x_i - y|^{-1}),$$

provided $3\varepsilon \leq |x_i - y| \leq c$. The proof is completed by restricting $|x_i - y| \geq 2C\varepsilon$. \square

The barrier lemma (Lemma 4.6) establishes the barrier property away from the local neighborhood $B(x, C\varepsilon)$. The singularity in the barrier (for $\alpha < d$) and the singularity in γ prevent us from pushing the barrier lemma inside this local neighborhood. Hence, the barrier can only be used to establish the following macroscopic continuity result.

Proposition 4.7 (Macroscopic Hölder estimate). *Let $u \in L^2(\mathcal{X}_n)$ be the solution of (14), let $\alpha > d - 2$, and fix any $0 < \beta < \alpha + 2 - d$. For each $y \in \Gamma$ the event that*

$$(68) \quad u(x_i) - u(y) \leq C|x_i - y|^\beta + \sup_{x \in \mathcal{X}_n \cap B(y, \delta_{\varepsilon, \zeta})} (u(x) - u(y))$$

holds for all $x_i \in \mathcal{X}_n$ occurs with probability at least $1 - C \exp(-cn\varepsilon^{d+4} + \log(n))$, where $\delta_{\varepsilon, \zeta} = \max\{C\varepsilon, r_\zeta + 2\varepsilon\}$.

Proof. We note the graph is connected with probability at least $1 - C \exp(-cn\varepsilon^d + \log(n))$. The proof uses the barrier function

$$(69) \quad \varphi(x) = \sup_{B(y, \delta_{\varepsilon, \zeta})} u + K|x - y|^\beta$$

constructed in Lemma 4.6 for a sufficiently large K , and the maximum principle on a connected graph. By Lemma 4.6 we have

$$(70) \quad \mathcal{GL}_{n,\varepsilon,\zeta}\varphi(x_i) \leq -cK|x_i - y|^{-(\alpha+2-\beta)}$$

for all x_i with $\delta_{\varepsilon, \zeta} \leq |x_i - y| \leq c$. By the maximum principle we have

$$\min_{\Gamma} g \leq u \leq \max_{\Gamma} g.$$

Therefore, we can choose K large enough so that $\varphi(x_i) > u(x_i)$ for $|x_i - y| \geq c$. We trivially have $u(x_i) < \varphi(x_i)$ for $|x_i - y| \leq \delta_{\varepsilon, \zeta}$. Since $\mathcal{GL}_{n,\varepsilon,\zeta}(u - \varphi) \leq 0$ for $\delta_{\varepsilon, \zeta} \leq |x_i - y| \leq c$, the maximum principle on a graph yields $u \leq \varphi$ on \mathcal{X}_n , which completes the proof. \square

We now establish a local regularity result that allows us to fill in the gap within the ball $B(x, C\varepsilon)$. The local result depends only on the variational structure of the problem, and does not use a barrier argument.

Proposition 4.8 (Local Hölder estimate). *Let $u \in L^2(\mathcal{X}_n)$. For each $z \in \Gamma$ the event that*

$$(71) \quad |u(x) - u(y)|^2 \leq \frac{Cn\varepsilon^2}{\min\{\gamma_\zeta(x), \gamma_\zeta(y)\}} \mathcal{GE}_{n,\varepsilon,\zeta}(u),$$

holds for all $x, y \in \mathcal{X}_n \cap B(z, r)$ with $|x - y| \leq \varepsilon$ occurs with probability at least $1 - C \exp(-cn\varepsilon^d + \log(n))$.

Proof. Let $z \in \Gamma$, and fix $r > 0$. Partition the cube $K := \prod_{i=1}^d [z_i - r, z_i + r]$ into hypercubes K_1, \dots, K_m of side length $h > 0$, where $m = (2r/h)^d$. Let Z_i denote the number of random variables falling in cube K_i . By Lemma 4.4 we have

$$(72) \quad \mathbb{P}(Z_i \leq \mathbb{E}[Z_i] - Cnh^d\lambda) \leq \exp(-cnh^d\lambda^2)$$

for any $0 < \lambda \leq 1$. Since $\mathbb{E}[Z_i] = nh^d$ we have

$$(73) \quad \mathbb{P}\left(\min_{1 \leq i \leq m} Z_i \leq \frac{1}{2}nh^d\right) \leq m \exp(-cnh^d).$$

Let $x, y \in \mathcal{X}_n \cap B(z, r)$ such that $|x - y| \leq \varepsilon$, and let $\bar{x} = (x + y)/2 \in B(z, r)$. Let K_i denote the cube to which \bar{x} belongs. Then for all $w \in K_i$ we have $|\bar{x} - w| \leq \sqrt{d}h$. Therefore, if $\sqrt{d}h \leq \varepsilon/2$ then

$$|x - w| \leq |x - \bar{x}| + |\bar{x} - w| \leq \frac{\varepsilon}{2} + \sqrt{d}h \leq \varepsilon,$$

and $|y - w| \leq \varepsilon$ for all $w \in K_i$. It follows that

$$K_i \subset B(x, \varepsilon) \cap B(y, \varepsilon).$$

For the remainder of the proof, we set $h = \varepsilon/(2\sqrt{d})$ and restrict ourselves to the event that

$$(74) \quad \min_{1 \leq i \leq m} Z_i \geq cn\varepsilon^d$$

Let

$$K = \mathcal{X}_n \cap B(x, \varepsilon) \cap B(y, \varepsilon).$$

Note that for any $z \in K$ we have

$$\min\{|u(x) - u(z)|, |u(y) - u(z)|\} \geq \frac{1}{2}|u(x) - u(y)|.$$

Now we have

$$\begin{aligned} \mathcal{GE}_{n,\varepsilon,\zeta}(u) &= \frac{1}{2n^2\varepsilon^2} \sum_{x,y \in \mathcal{X}_n} (\gamma_\zeta(x) + \gamma_\zeta(y)) \eta_\varepsilon(x - y) |u(x) - u(y)|^2 \\ &\geq \frac{1}{2n^2\varepsilon^2} \left[\sum_{z \in \mathcal{X}_n} \gamma_\zeta(x) \eta_\varepsilon(x - z) |u(x) - u(z)|^2 + \gamma_\zeta(y) \eta_\varepsilon(y - z) |u(y) - u(z)|^2 \right] \\ &\geq \frac{c \min\{\gamma_\zeta(x), \gamma_\zeta(y)\} |K|}{n^2\varepsilon^{d+2}} |u(x) - u(y)|^2. \end{aligned}$$

Since $K_i \cap \mathcal{X}_n \subset K$, we have $|K| \geq cn\varepsilon^d$, and hence

$$(75) \quad |u(x) - u(y)|^2 \leq \frac{Cn\varepsilon^2}{\min\{\gamma_\zeta(x), \gamma_\zeta(y)\}} \mathcal{GE}_{n,\varepsilon,\zeta}(u),$$

which completes the proof. \square

We are now equipped to give the proof of Theorem 4.1.

Proof of Theorem 4.1. The proof combines the macroscopic Hölder estimate (Proposition 4.7), and the local Hölder estimate (Proposition 4.8), and is split into two steps.

1. We note that

$$(76) \quad r_\zeta = \frac{r_0}{(\zeta - 1)^{1/\alpha}} \leq C\varepsilon,$$

as $\zeta \geq 1 + \varepsilon^{-\alpha}$. By (76) and Theorem 4.7, the event that

$$(77) \quad |u(x) - u(z)| \leq C|x - z|^\beta + \sup_{x \in \mathcal{X}_n \cap B(z, C\varepsilon)} |u(x) - u(y)|,$$

holds for all $x \in \mathcal{X}_n$ occurs with probability at least $1 - C \exp(-cn\varepsilon^{d+4} + \log(n))$.

2. We note that with probability at least $1 - C \exp(-cn\varepsilon^d + \log(n))$ we have $\mathcal{G}\mathcal{E}_{n\varepsilon, \zeta}(u) \leq C$ for a constant C . Therefore, by Proposition 4.8 we have that

$$(78) \quad |u(x) - u(y)|^2 \leq \frac{Cn\varepsilon^2}{\min_{B(z, C\varepsilon)} \gamma_\zeta}$$

holds for all $x, y \in \mathcal{X}_n \cap B(z, C\varepsilon)$ with $|x - y| \leq \varepsilon$ with probability at least $1 - C \exp(-cn\varepsilon^d + \log(n))$.

As in the proof of Proposition 4.8, we partition the cube $K := \prod_{i=1}^d [z_i - C\varepsilon, z_i + C\varepsilon]$ into hypercubes of side length $h < \varepsilon$, and find that all cubes have at least one point from \mathcal{X}_n with probability at least $1 - C \exp(-cn\varepsilon^d + \log(n))$. Thus, by traversing neighboring cubes, we can construct a path from $z \in \Gamma$ to any $x \in B(z, C\varepsilon)$ consisting of at most a constant number of points from $\mathcal{X}_n \cap B(z, C\varepsilon)$, with each step in the path smaller than ε . Applying (78) along the path yields

$$|u(x) - u(z)|^2 \leq \frac{Cn\varepsilon^2}{\min\{(C\varepsilon)^{-\alpha}, \zeta\}}$$

for all $x \in \mathcal{X}_n \cap B(z, C\varepsilon)$ with probability at least $1 - C \exp(-cn\varepsilon^d + \log(n))$. Since $\zeta \geq \varepsilon^{-\alpha}$ we deduce

$$\sup_{x \in \mathcal{X}_n \cap B(z, C\varepsilon)} |u(x) - u(z)|^2 \leq Cn\varepsilon^{2+\alpha},$$

which completes the proof. \square

5. NUMERICAL EXPERIMENTS

We present the results of several numerical experiments illustrating the properly-weighted Laplacian and comparing it with the nonlocal [36] and standard graph Laplacian on real and synthetic data. All experiments were performed in Matlab and use Matlab backslash to solve the graph Laplacian system. We mention there are indirect solvers that may be faster in certain applications, such as preconditioned conjugate gradient [25], algebraic multigrid [8, 25, 34], or more recent fast Laplacian solvers [29, 40]. Thus, the CPU times reported below have the potential to be improved substantially.

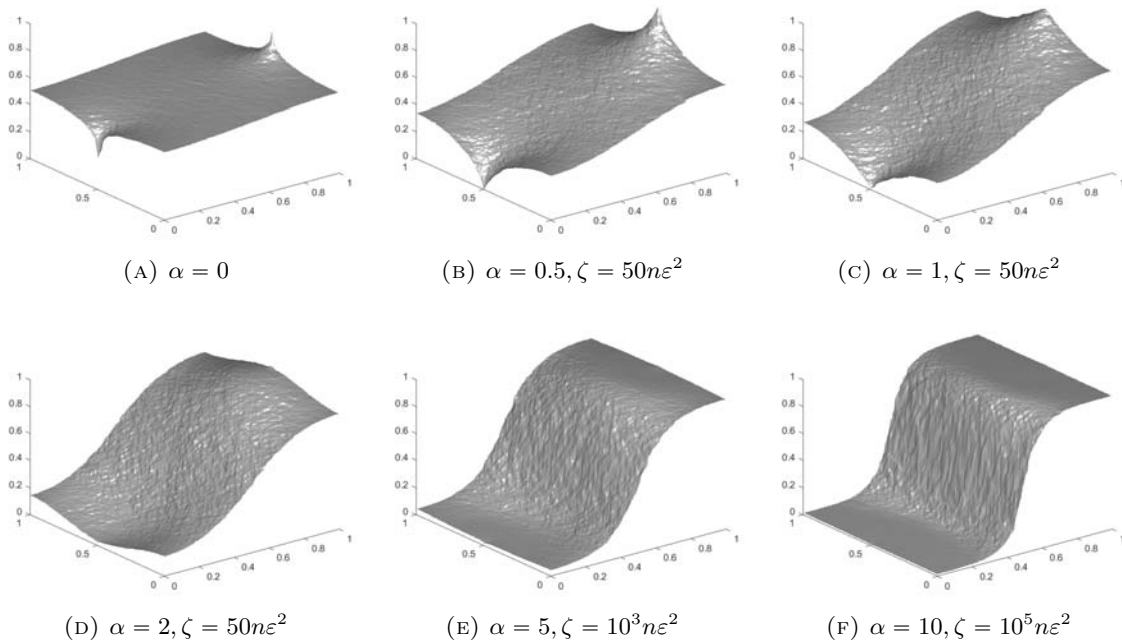


FIGURE 2. Comparison of (A) the standard graph Laplacian, and (B)–(F) our weighted graph Laplacian with various values of α and ζ . The graph is a random geometric graph (described in the text) on $[0, 1]^2$, and the labels are $g(0, 0.5) = 0$ and $g(1, 0.5) = 1$.




FIGURE 3. Comparison of (a) the standard graph Laplacian, (b) the nonlocal graph Laplacian [36], and (c) our properly-weighted Laplacian. The graph is a random geometric graph (described in the text) on $[0, 1]^3$, and the surfaces plotted are a slice of the learned function u near $x_3 = 0.5$. The labels are $g(0, 0.5, 0.5) = 0$ and $g(1, 0.5, 0.5) = 1$.

5.0.1. *Comparison of the profiles obtained.* First, we perform an experiment with two labels on the box $[0, 1]^d$ to illustrate our method and the differences with the nonlocal graph Laplacian [36]. The graph is a sequence of n *i.i.d.* random variables uniformly distributed on the unit box $[0, 1]^d$ in \mathbb{R}^d , and two labels are given $g(0, 0.5, \dots, 0.5) = 0$ and $g(1, 0.5, \dots, 0.5) = 1$.

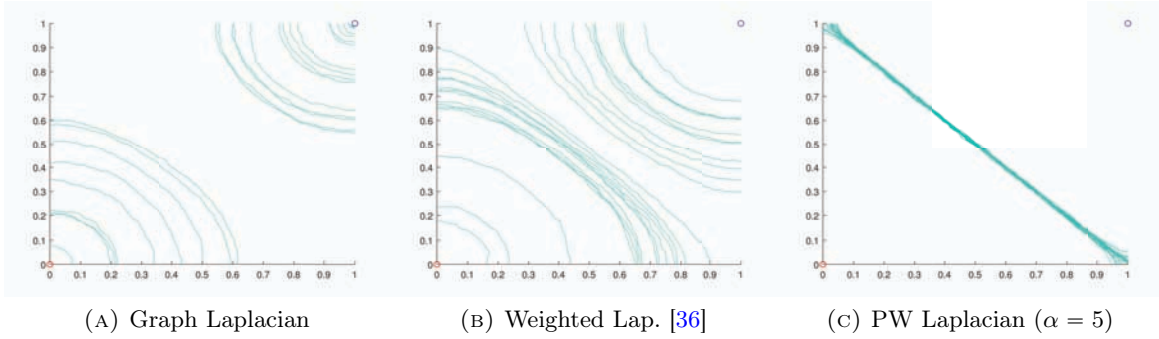


FIGURE 4. Comparison of decision boundaries plotted over 25 random trials for a synthetic classification problem for (A) the standard graph Laplacian, (B) the nonlocal graph Laplacian [36], and (C) our properly-weighted graph Laplacian. Two labels are given, $g(0,0) = 0$ (red point) and $g(1,1) = 1$ (blue point), and the graph is a uniform random geometric graph (described in the text).

We set $\varepsilon = 2/n^{1/d}$ and $r_0 = 1$. The weights follow a Gaussian distribution with $\sigma = \varepsilon/2$. In Figure 2 we show plots of the triangulated surface representing the learned function u on the graph for various values of α and ζ . Here, $n = 10^5$ and $d = 2$, and each simulation took approximately 1.5 seconds of CPU time. We notice that as α is increased, the learned functions are smoother in a vicinity of each label. The case of $\alpha = 0$ corresponds to the standard graph Laplacian, and returns an approximately constant label $u = 0.5$, which illustrates the degeneracy of the standard Laplacian with few labels. We note that as α is increased, we must increase ζ as well (recall (11)), otherwise the ball $B(x, r_\zeta)$ on which γ is truncated to value ζ becomes very large, and the method reduces to the standard graph Laplacian. This simply illustrates that the implicit rate in the condition $\zeta \gg n\varepsilon^2$ in Theorem 3.1 depends on α .

In Figure 3, we use the same model, but with $n = 2 \times 10^5$ points in dimension $d = 3$, and set $\zeta = 50n\varepsilon^2$. For visualization, we show the learned function restricted to a neighborhood of the slice $x_3 = 0.5$. Figure 3b illustrates the degeneracy of the nonlocal graph Laplacian [36], which returns a nearly constant label function. In contrast, our method, shown in Figure 3c, smoothly interpolates between the two labels. Each simulation in Figure 3 took approximately 25 seconds to compute.

5.0.2. *Comparison of decision boundaries in 2D.* We now give a synthetic classification example. The graph consists of $n = 10^5$ *i.i.d* uniform random variables on $[0, 1]^2$, and the weights are chosen to be Gaussian with $\sigma = \varepsilon/2$. We chose $\alpha = 5$, $\zeta = 10^6 n\varepsilon^2$, $r_0 = 1$, and $\varepsilon = 3/\sqrt{n}$. Two labels, $g(0,0) = 0$ and $g(1,1) = 1$ are provided. Figure 4 shows the decision boundaries (i.e., the level-set $\{u = 0.5\}$) over 25 trials for the standard graph Laplacian, the nonlocal Laplacian, and our method. Each trial took roughly 1.5 seconds to compute. We see that the nonlocal and standard Laplacian are highly sensitive to small variations in the graph, giving a wide variety of results over the 25 trials. This is a reflection of the degeneracy, or ill-posedness, in the small label regime, and suggests the methods are very sensitive to perturbations in the data. In contrast, our method very consistently divides the square along the diagonal.

5.0.3. *Comparison of classes obtained in 3D.* We consider samples of the measure supported on domain is $[0, 1]^3$ and with density 1 except for the strip $[0.45, 0.55] \times [0, 1] \times [0, 1]$ where the




FIGURE 5. Comparison of results for a synthetic classification problem for (A) the standard graph Laplacian, (B) the nonlocal graph Laplacian [36], and (C) our weighted graph Laplacian. The domain is $[0, 1]^3$ and the density is 1 except for the strip $[0.45, 0.55] \times [0, 1] \times [0, 1]$ where it is 0.6. The given labeled points are $g(0, 0.2, 0.2) = 0$ and $g(1, 0.2, 0.2) = 1$. There are $n = 50,000$ points in the domain. Connectivity distance for the graph construction is $3/n^{1/3}$ and for our method $\alpha = 5$.

density is 0.6. We considered 20 runs with $n = 50,000$ points in the domain. The given labeled points are $g(0, 0.2, 0.2) = 0$ and $g(1, 0.2, 0.2) = 1$. Due to the symmetry, the correct decision boundary is the plane $x_1 = 0.5$. We used a connectivity distance for the graph construction of $\varepsilon = 3/n^{1/3}$, which yielded a typical vertex degree of about 116. We consider Gaussian weights with $\sigma = \varepsilon/2$. We chose $\alpha = 5$, $\zeta = 10^6 n \varepsilon^2$, $r_0 = 1$ for our method. A typical result for one run is illustrated on Figure 5. The standard graph Laplacian produced very unstable results with the average of 49.8% misclassified points. The nonlocal Laplacian of [36] was also rather unstable with sometimes almost perfect decision boundary and sometimes large sections of misclassified points. On average it misclassified 11% of points. Our method was stable and in all experiments identified the correct boundary, with average classification error of 0.25%. We observed similar outcomes for a variety of sets of parameters.

5.0.4. *Comparison on the MNIST dataset.* Our last experiment considers classification of handwritten digits from the MNIST dataset, which consists of 70,000 grayscale 28x28 pixel images of handwritten digits 0–9 [30]. Figure 6 shows examples of some of the images in the MNIST dataset. MNIST is estimated to have intrinsic dimension between $d = 12$ and $d = 14$ [14, 28], which suggests a larger value for α is appropriate. We used all 70,000 MNIST images to construct the graph. Our construction is the same as in [36]; we connect each data point to its nearest 50 neighbors (in Euclidean distance), and assign Gaussian weights taking σ to be the distance to the 20th nearest neighbor. We symmetrize the graph by replacing the weight matrix W with $\frac{1}{2}(W^T + W)$, which is done automatically by the variational formulation (recall (18)). We then take L randomly chosen images from each class (digit) as labels, where $L = 1, 3, 5, 7, 10$, and provide the true labels for these digits. The semi-supervised learning algorithm performs 10 binary classifications, for each digit versus the rest, which generates functions $u^0, u^1, u^2, \dots, u^9$ on the graph. The label for each image x in the dataset is chosen as the index i for which $u^i(x)$ is maximal. The algorithm is standard in semi-supervised learning, and identical to the one used in [36].



FIGURE 6. Example of some of the handwritten digits from the MNIST dataset [30].

# Labels	10		30		50		70		100	
Method	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
Graph Laplacian	14.2%	6.3	24.3%	11.9	53%	10.9	68%	6.4	76.1%	7.6
Weighted Lap. [36]	67.9%	8.7	84.8%	2.7	88.8%	1.1	89.6%	1.3	90.9%	1.1
PW-Laplacian	68%	8.6	84.9%	2.7	88.8%	1.1	89.6%	1.3	90.9%	1.1

TABLE 1. Accuracy for classification of MNIST handwritten digits with 10, 30, 50, 70 and 100 labels via the standard graph Laplacian, the nonlocal weighted Laplacian [36], and our properly-weighted Laplacian. The results are averaged over 10 trials, and the mean and standard deviation of accuracy are reported.

For each value of $L \in \{1, 3, 5, 7, 10\}$, we ran the experiment described above 10 times, choosing the labels randomly (in the same way for each algorithm) every time. Each of the 500 trials took approximately 15 minutes to compute in Matlab. The mean and standard deviation of accuracy are shown in Table 1. Our method performs very similarly to the nonlocal Laplacian [36], and both significantly outperform the standard graph Laplacian. We ran our method for $\alpha = 2, 5, 10$, producing nearly identical results in all cases. We used $\zeta = 10^7$ and $r_0 = 0.1$. We found the results for our method were largely insensitive to many of the parameters in our algorithm; the accuracy begins to decrease when $\alpha < 1$ and when $r_0 > 1$. We note that the accuracy scores reported in Table 1 are much higher than those reported in [36]; we believe this is because the authors in [36] subsampled MNIST to 16,000 images. This observation speaks favorably to the semi-supervised paradigm that learning can be improved by access to additional unlabeled data. We note that the accuracy scores for our method and the nonlocal weighted Laplacian [36] are identical (to one significant digit) for 30, 50, 70, and 100 labels. Most data points in MNIST are relatively far from their nearest neighbors, and so our nonlocal weights have less effect, compared to the low dimensional examples presented above. For this reason, the weight matrix for our method is very similar to the nonlocal Laplacian [36]. We expect to see more of a difference in applications to larger

datasets. For example, it would be interesting (and challenging) to apply these techniques to a dataset like ImageNet [16], which consists of over 14 million natural images belonging to over 20,000 categories.

APPENDIX A. BACKGROUND MATERIAL

Here we recall some of the notions our work depends on and establish an auxiliary technical result.

A.1. Γ -Convergence. Γ -convergence was introduced by De Giorgi in 1970's to study limits of variational problems. We refer to [7, 15] for comprehensive introduction to Γ -convergence. We now recall the notion of Γ -convergence is in a random setting.

Definition A.1 (Γ -convergence). Let (Z, d) be a metric space, $L^0(Z; \mathbb{R} \cup \{\pm\infty\})$ be the set of measurable functions from Z to $\mathbb{R} \cup \{\pm\infty\}$, and $(\mathcal{X}, \mathbb{P})$ be a probability space. The function $\mathcal{X} \ni \omega \mapsto \mathcal{E}_n^{(\omega)} \in L^0(Z; \mathbb{R} \cup \{\pm\infty\})$ is a random variable. We say $\mathcal{E}_n^{(\omega)}$ Γ -converge almost surely on the domain Z to $\mathcal{E}_\infty : Z \rightarrow \mathbb{R} \cup \{\pm\infty\}$ with respect to d , and write $\mathcal{E}_\infty = \Gamma\text{-}\lim_{n \rightarrow \infty} \mathcal{E}_n^{(\omega)}$, if there exists a set $\mathcal{X}' \subset \mathcal{X}$ with $\mathbb{P}(\mathcal{X}') = 1$, such that for all $\omega \in \mathcal{X}'$ and all $f \in Z$:

- (i) (liminf inequality) for every sequence $\{u_n\}_{n=1, \dots}$ in Z converging to f

$$\mathcal{E}_\infty(f) \leq \liminf_{n \rightarrow \infty} \mathcal{E}_n^{(\omega)}(u_n), \text{ and}$$

- (ii) (recovery sequence) there exists a sequence $\{u_n\}_{n=1, 2, \dots}$ in Z converging to f such that

$$\mathcal{E}_\infty(f) \geq \limsup_{n \rightarrow \infty} \mathcal{E}_n^{(\omega)}(u_n).$$

For simplicity we suppress the dependence of ω in writing our functionals. The almost sure nature of the convergence in our claims is ensured by considering the set of realizations of $\{x_i\}_{i=1, \dots}$ such that the conclusions of Theorem A.3 hold (which they do almost surely).

An important result concerning Γ -convergence is that, together with accompanying compactness property it implies the convergence of minimizers. The proof can be found in [7, Theorem 1.21] or [15, Theorem 7.23].

Theorem A.2 (Convergence of Minimizers). *Let (Z, d) be a metric space and $\mathcal{E}_n : Z \rightarrow [0, \infty]$ be a sequence of functionals. Let u_n be a minimizing sequence for \mathcal{E}_n . If the set $\{u_n\}_{n=1, 2, \dots}$ is precompact and $\mathcal{E}_\infty = \Gamma\text{-}\lim_n \mathcal{E}_n$ where $\mathcal{E}_\infty : Z \rightarrow [0, \infty]$ is not identically ∞ then*

$$\min_Z \mathcal{E}_\infty = \lim_{n \rightarrow \infty} \inf_Z \mathcal{E}_n.$$

Furthermore any cluster point of $\{u_n\}_{n=1, 2, \dots}$ is a minimizer of \mathcal{E}_∞ .

The theorem is also true if we replace minimizers with approximate minimizers.

We note that Γ -convergence is defined for functionals on a common metric space. Section A.3 overviews the metric space we use to analyze the asymptotics of our semi-supervised learning models, in particular it allows us to go from discrete to continuum.

A.2. Optimal Transportation and Approximation of Measures. Here we recall the notion of optimal transportation between measures and the metric it introduces. Comprehensive treatment of the topic can be found in books of Villani [43] and Santambrogio [35].

Given $\Omega \subset \mathbb{R}^d$ is open and bounded, and probability measures μ and ν in $\mathcal{P}(\overline{\Omega})$ we define the set $\Pi(\mu, \nu)$ of transportation plans, or couplings, between μ and ν to be the set of probability

measures on the product space $\pi \in \mathcal{P}(\overline{\Omega} \times \overline{\Omega})$ whose first marginal is μ and second marginal is ν . We then define the p -optimal transportation distance (a.k.a. p -Wasserstein distance) by

$$d_p(\mu, \nu) = \begin{cases} \inf_{\pi \in \Pi(\mu, \nu)} \left(\int_{\Omega \times \Omega} |x - y|^p d\pi(x, y) \right)^{\frac{1}{p}} & \text{if } 1 \leq p < \infty \\ \inf_{\pi \in \Pi(\mu, \nu)} \pi\text{-esssup}_{(x, y)} |x - y| & \text{if } p = \infty. \end{cases}$$

If μ has a density with respect to Lebesgue measure on Ω , then the distance can be rewritten using transportation maps, $T : \Omega \rightarrow \Omega$, instead of transportation plans,

$$d_p(\mu, \nu) = \begin{cases} \inf_{T_{\#}\mu = \nu} \left(\int_{\Omega} |x - T(x)|^p d\mu(x) \right)^{\frac{1}{p}} & \text{if } 1 \leq p < \infty \\ \inf_{T_{\#}\mu = \nu} \mu\text{-esssup}_x |x - T(x)| & \text{if } p = \infty. \end{cases}$$

where $T_{\#}\mu = \nu$ means that the push forward of the measure μ by T is the measure ν , namely that T is Borel measurable and such that for all $U \subset \overline{\Omega}$, open, $\mu(T^{-1}(U)) = \nu(U)$.

When $p < \infty$ the metric d_p metrizes the weak* convergence of measures.

Optimal transportation plays an important role in comparing the discrete and continuum objects we study. In particular, we use sharp estimates on the ∞ -optimal transportation distance between a measure and the empirical measure of its sample. In the form below, for $d \geq 2$, they were established in [20], which extended the related results in [1, 31, 37, 41].

Theorem A.3. *For $d \geq 2$, let $\Omega \subset \mathbb{R}^d$ be open, connected and bounded with Lipschitz boundary. Let μ be a probability measure on Ω with density (with respect to Lebesgue) ρ which is bounded above and below by positive constants. Let x_1, x_2, \dots be a sequence of independent random variables with distribution μ and let μ_n be the empirical measure. Then, there exists constants $C \geq c > 0$ such that almost surely there exists a sequence of transportation maps $\{T_n\}_{n=1}^{\infty}$ from μ to μ_n with the property*

$$c \leq \liminf_{n \rightarrow \infty} \frac{\|T_n - \text{Id}\|_{L^\infty(\Omega)}}{\ell_n} \leq \limsup_{n \rightarrow \infty} \frac{\|T_n - \text{Id}\|_{L^\infty(\Omega)}}{\ell_n} \leq C$$

where

$$(79) \quad \ell_n = \begin{cases} \frac{(\ln n)^{\frac{3}{4}}}{\sqrt{n}} & \text{if } d = 2 \\ \frac{(\ln n)^{\frac{1}{d}}}{n^{\frac{1}{d}}} & \text{if } d \geq 3. \end{cases}$$

A.3. The TL^p Space. The discrete functionals we consider (e.g. $\mathcal{E}_{n, \varepsilon_n, \zeta_n}$) are defined for functions $u_n : \mathcal{X}_n \rightarrow \mathbb{R}$, while the limit functional \mathcal{E} acts on functions $f : \Omega \rightarrow \mathbb{R}$, where Ω is an open set. We can view u_n as elements of $L^p(\mu_n)$ where μ_n is the empirical measure of the sample $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$. Likewise $f \in L^p(\mu)$ where μ is the measure with density ρ from which the data points are sampled. In order to compare f and u_n in a way that is consistent with the L^p topology we use the TL^p space that was introduced in [21], where it was used to study the continuum limit of the graph total variation. Subsequent development of the TL^p space has been carried out in [22, 42].

To compare the functions u_n and f above we need to take into account their domains, or more precisely to account for μ and μ_n . For that purpose the space of configurations is defined to be

$$TL^p(\Omega) = \{(\mu, f) : \mu \in \mathcal{P}(\overline{\Omega}), f \in L^p(\mu)\}.$$

The metric on the space is

$$d_{TL^p}^p((\mu, f), (\nu, g)) = \inf \left\{ \int_{\Omega \times \Omega} |x - y|^p + |f(x) - g(y)|^p d\pi(x, y) : \pi \in \Pi(\mu, \nu) \right\}$$

where $\Pi(\mu, \nu)$ the set of transportation plans defined in Section A.2. We note that the minimizing π exists and that TL^p space is a metric space, [21].

As shown in [21], when μ has a density with respect to Lebesgue measure on Ω , then the distance can be rewritten using transportation maps T , instead of transportation plans,

$$d_{TL^p}^p((\mu, f), (\nu, g)) = \inf \left\{ \int_{\Omega} |x - T(x)|^p + |f(x) - g(T(x))|^p d\mu(x) : T_{\#}\mu = \nu \right\}$$

where the push forward of the measure $T_{\#}\mu$ is defined in Section A.2. This formula provides an interpretation of the distance in our setting. Namely, to compare functions $u_n : \mathcal{X}_n \rightarrow \mathbb{R}$ we define a mapping $T_n : \Omega \rightarrow \mathcal{X}_n$ and compare the functions $\tilde{f}_n = u_n \circ T_n$ and f in $L^p(\mu)$, while also accounting for the transport, namely the $|x - T_n(x)|^p$ term.

We remark that the $TL^p(\bar{\Omega})$ space is not complete, and that its completion was discussed in [21]. In the setting of this paper, since the corresponding measure is clear from context, we often say that u_n converges in TL^p to f as a short way to say that (μ_n, u_n) converges in TL^p to (μ, f) .

A.4. Local estimates for weighted Laplacian.

Lemma A.4. *There exists $C > 0$ such that for each $\tilde{u} \in H^1(B(0, 1))$ there exists $\tilde{v} \in H^1(B(0, 1))$ such that*

$$\begin{aligned} \tilde{v}|_{B(0, \frac{1}{2})} &\equiv \frac{1}{|B(0, \frac{1}{2})|} \int_{B(0, \frac{1}{2})} \tilde{u}(x) dx \\ \tilde{v}|_{\partial B(0, 1)} &= \tilde{u}|_{\partial B(0, 1)} \\ \|\nabla \tilde{v}\|_{L^2(B(0, 1))} &\leq C \|\nabla \tilde{u}\|_{L^2(B(0, 1))} \end{aligned}$$

where the value on the boundary is considered in sense of the $L^2(\partial B(0, 1))$ trace.

Proof. In the proof we use u and v instead of \tilde{u} and \tilde{v} . We first consider such modification of u when u is smooth. The conclusion for general u follows by density. Let J be a mollifier supported in $B(0, 1)$. For $x \in B(0, 1)$ let

$$v(x) = \begin{cases} u * J_{\frac{1}{4}}(0) & \text{if } |x| \leq \frac{1}{2} \\ u * J_{\frac{1-|x|}{2}}\left(2x - \frac{x}{|x|}\right) & \text{if } \frac{1}{2} < |x| < 1. \end{cases}$$

Note that v is a continuous function and that for $\frac{1}{2} < |x| < 1$

$$\begin{aligned} v(x) &= \left(\frac{2}{1-|x|}\right)^d \int_{B(0, 1)} J\left(\frac{2(2x - \frac{x}{|x|}) - z}{1-|x|}\right) u(z) dz \\ &= \int_{B(0, 1)} J(y) u\left(\frac{1-|x|}{2} y - 2x + \frac{x}{|x|}\right) dy. \end{aligned}$$

Consequently

$$\nabla v(x) = - \int_{B(0, 1)} J(y) \left(\frac{x}{2|x|} \otimes y + 2I - \frac{1}{|x|} I + \frac{x \otimes x}{|x|^3}\right) \nabla u\left(\frac{1-|x|}{2} y - 2x + \frac{x}{|x|}\right) dy$$

and thus

$$|\nabla v(x)| \leq 3 \int_{B(0,1)} J(y) \left| \nabla u \left(\frac{1-|x|}{2} y - 2x + \frac{x}{|x|} \right) \right| dy.$$

By integrating and using Jensen's inequality we obtain that

$$\int_{B(0,1)} |\nabla v(x)|^2 dx \leq 9 \int_{B(0,1) \setminus B(0, \frac{1}{2})} \int_{B(0,1)} J(y) \left| \nabla u \left(\frac{1-|x|}{2} y - 2x + \frac{x}{|x|} \right) \right|^2 dy dx$$

We split the integral into two subregions:

$$\int_{B(0, \frac{3}{4}) \setminus B(0, \frac{1}{2})} \int_{B(0,1)} J(y) \left| \nabla u \left(\frac{1-|x|}{2} y - 2x + \frac{x}{|x|} \right) \right|^2 dy dx \leq c_1 \|\nabla u\|_{L^2(B(0,1))}^2$$

by substituting $s(y) = \frac{1-|x|}{2} y - 2x + \frac{x}{|x|}$ and using that the Jacobian $\left| \det \left(\frac{\partial s}{\partial y} \right) \right| = \left(\frac{1-|x|}{2} \right)^d \geq \frac{1}{8^d}$.

On the other region we interchange the integrals and use the substitution $s(x) = \frac{1-|x|}{2} y - 2x + \frac{x}{|x|}$:

$$\begin{aligned} \int_{B(0,1)} J(y) \int_{B(0,1) \setminus B(0, \frac{3}{4})} \left| \nabla u \left(\frac{1-|x|}{2} y - 2x + \frac{x}{|x|} \right) \right|^2 dx dy \\ \leq \int_{B(0,1)} J(y) \int_{B(0,1)} |\nabla u(s)|^2 \left| \det \left(\frac{\partial x}{\partial s} \right) \right| ds dy \leq c_2 \|\nabla u\|_{L^2(B(0,1))}^2 \end{aligned}$$

where we estimate the Jacobian as follows. Let $M := -\frac{\partial s}{\partial x} = \frac{y \otimes x}{2|x|} + 2 \left(I - \frac{1}{|x|} \right) + \frac{x \otimes x}{|x|^3}$. Thus

$$MM^T = \left(2 - \frac{1}{|x|} \right)^2 + 2 \left(\left(I - \frac{x \otimes x}{|x|^3} \right) + \left(\frac{x}{|x|^2} + \frac{y}{2} \right) \otimes \left(\frac{x}{|x|^2} + \frac{y}{2} \right) \right) \geq \left(2 - \frac{1}{|x|} \right)^2 I \geq \frac{1}{4} I.$$

Hence $\left| \det \left(\frac{\partial x}{\partial s} \right) \right| \leq \frac{1}{\sqrt{\det(MM^T)}} \leq 2^d$.

Combining the estimates gives $\int_{B(0,1)} |\nabla v(x)|^2 dx \leq C \int_{B(0,1)} |\nabla u(x)|^2 dx$. \square

REFERENCES

- [1] M. Ajtai, J. Komlós, and G. Tusnády. On optimal matchings. *Combinatorica*, 4(4):259–264, 1984.
- [2] M. Alamgir and U. V. Luxburg. Phase transition in the family of p-resistances. In *Advances in Neural Information Processing Systems*, pages 379–387, 2011.
- [3] R. K. Ando and T. Zhang. Learning on graph with laplacian regularization. *Advances in Neural Information Processing Systems*, 19:25, 2007.
- [4] M. Bardi and I. Capuzzo-Dolcetta. *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*. Springer Science & Business Media, 2008.
- [5] M. Belkin and P. Niyogi. Using manifold structure for partially labeled classification. In *Advances in Neural Information Processing Systems (NIPS)*, pages 953–960, 2003.
- [6] A. Bertozzi, X. Luo, A. Stuart, and K. Zygalakis. Uncertainty quantification in graph-based classification of high dimensional data. *SIAM/ASA Journal on Uncertainty Quantification*, 6(2):568–595, 2018.
- [7] A. Braides. Γ -convergence for beginners, volume 22 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, Oxford, 2002.
- [8] A. Brandt and O. E. Livne. *Multigrid techniques: 1984 guide with applications to fluid dynamics*, volume 67. SIAM, 2011.
- [9] N. Bridge and X. Zhu. p-voltages: Laplacian regularization for semi-supervised learning on high-dimensional data. In *Eleventh Workshop on Mining and Learning with Graphs (MLG2013)*, 2013.
- [10] J. Calder. Consistency of Lipschitz learning with infinite unlabeled and finite labeled data. *arXiv:1710.10364*, 2017.

- [11] J. Calder. The game theoretic p-Laplacian and semi-supervised learning with few labels. *arXiv:1711.10144*, 2017.
- [12] J. Calder and C. K. Smart. The limit shape of convex hull peeling. *arXiv preprint arXiv:1805.08278*, 2018.
- [13] O. Chapelle, B. Scholkopf, and A. Zien. *Semi-supervised learning*. MIT, 2006.
- [14] J. A. Costa and A. O. Hero. Determining intrinsic dimension and entropy of high-dimensional shape spaces. In *Statistics and Analysis of Shapes*, pages 231–252. Springer, 2006.
- [15] G. Dal Maso. *An introduction to Γ -convergence*, volume 8 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser Boston, Inc., Boston, MA, 1993.
- [16] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009.
- [17] M. M. Dunlop, D. Slepčev, A. M. Stuart, and M. Thorpe. Large data and zero noise limits of graph-based semi-supervised learning algorithms. *arXiv preprint arXiv:1805.09450*, 2018.
- [18] A. El Alaoui, X. Cheng, A. Ramdas, M. J. Wainwright, and M. I. Jordan. Asymptotic behavior of lp-based Laplacian regularization in semi-supervised learning. In *29th Annual Conference on Learning Theory*, pages 879–906, 2016.
- [19] N. García Trillos, M. Gerlach, M. Hein, and D. Slepčev. Error estimates for spectral convergence of the graph Laplacian on random geometric graphs towards the Laplace-Beltrami operator. *arXiv preprint arXiv:1801.10108*, 2018.
- [20] N. García Trillos and D. Slepčev. On the rate of convergence of empirical measures in ∞ -transportation distance. *Canad. J. Math.*, 67(6):1358–1383, 2015.
- [21] N. García Trillos and D. Slepčev. Continuum limit of total variation on point clouds. *Archive for Rational Mechanics and Analysis*, 220(1):193–241, 2016.
- [22] N. García Trillos and D. Slepčev. A variational approach to the consistency of spectral clustering. *Appl. Comput. Harmon. Anal.*, 45(2):239–281, 2018.
- [23] N. García Trillos, D. Slepčev, J. von Brecht, T. Laurent, and X. Bresson. Consistency of Cheeger and ratio graph cuts. *J. Mach. Learn. Res.*, 17:Paper No. 181, 46, 2016.
- [24] D. Gilbarg and N. S. Trudinger. *Elliptic partial differential equations of second order*. springer, 2015.
- [25] A. Greenbaum. *Iterative methods for solving linear systems*, volume 17. Siam, 1997.
- [26] J. He, M. Li, H.-J. Zhang, H. Tong, and C. Zhang. Manifold-ranking based image retrieval. In *Proceedings of the 12th Annual ACM International Conference on Multimedia*, pages 9–16. ACM, 2004.
- [27] J. He, M. Li, H.-J. Zhang, H. Tong, and C. Zhang. Generalized manifold-ranking-based image retrieval. *IEEE Transactions on image processing*, 15(10):3170–3177, 2006.
- [28] M. Hein and J.-Y. Audibert. Intrinsic dimensionality estimation of submanifolds in Rd. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 289–296. ACM, 2005.
- [29] R. Kyng, A. Rao, S. Sachdeva, and D. A. Spielman. Algorithms for lipschitz learning on graphs. In *Proceedings of The 28th Conference on Learning Theory*, pages 1190–1223, 2015.
- [30] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [31] T. Leighton and P. Shor. Tight bounds for minimax grid matching with applications to the average case analysis of algorithms. *Combinatorica*, 9(2):161–187, 1989.
- [32] U. v. Luxburg and O. Bousquet. Distance-based classification with lipschitz functions. *Journal of Machine Learning Research*, 5(Jun):669–695, 2004.
- [33] B. Nadler, N. Srebro, and X. Zhou. Semi-supervised learning with the graph Laplacian: The limit of infinite unlabelled data. In *Neural Information Processing Systems (NIPS)*, 2009.
- [34] J. W. Ruge and K. Stüben. Algebraic multigrid. In *Multigrid methods*, pages 73–130. SIAM, 1987.
- [35] F. Santambrogio. *Optimal transport for applied mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser/Springer, Cham, 2015. Calculus of variations, PDEs, and modeling.
- [36] Z. Shi, S. Osher, and W. Zhu. Weighted nonlocal laplacian on interpolation from sparse data. *Journal of Scientific Computing*, 73(2-3):1164–1177, 2017.
- [37] P. W. Shor and J. E. Yukich. Minimax grid matching and empirical measures. *Ann. Probab.*, 19(3):1338–1348, 1991.
- [38] A. Singer. From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, 2006.

- [39] D. Slepčev and M. Thorpe. Analysis of p-Laplacian regularization in semi-supervised learning. *arXiv:1707.06213*, 2017.
- [40] D. A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90. ACM, 2004.
- [41] M. Talagrand. *Upper and lower bounds of stochastic processes*, volume 60 of *Modern Surveys in Mathematics*. Springer-Verlag, Berlin Heidelberg, 2014.
- [42] M. Thorpe, S. Park, S. Kolouri, G. K. Rohde, and D. Slepčev. A transportation L^p distance for signal analysis. *J. Math. Imaging Vision*, 59(2):187–210, 2017.
- [43] C. Villani. *Optimal transport*, volume 338 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2009. Old and new.
- [44] Y. Wang, M. A. Cheema, X. Lin, and Q. Zhang. Multi-manifold ranking: Using multiple features for better image retrieval. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 449–460. Springer, 2013.
- [45] B. Xu, J. Bu, C. Chen, D. Cai, X. He, W. Liu, and J. Luo. Efficient manifold ranking for image retrieval. In *Proceedings of the 34th international ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 525–534. ACM, 2011.
- [46] C. Yang, L. Zhang, H. Lu, X. Ruan, and M.-H. Yang. Saliency detection via graph-based manifold ranking. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3166–3173, 2013.
- [47] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. *Advances in Neural Information Processing Systems*, 16(16):321–328, 2004.
- [48] D. Zhou, J. Huang, and B. Schölkopf. Learning from labeled and unlabeled data on a directed graph. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 1036–1043. ACM, 2005.
- [49] D. Zhou and B. Schölkopf. Regularization on discrete spaces. In *Proceedings of the 27th DAGM Conference on Pattern Recognition*, PR’05, pages 361–368, Berlin, Heidelberg, 2005. Springer-Verlag.
- [50] D. Zhou, J. Weston, A. Gretton, O. Bousquet, and B. Schölkopf. Ranking on data manifolds. *Advances in Neural Information Processing Systems*, 16:169–176, 2004.
- [51] X. Zhou and M. Belkin. Semi-supervised learning by higher order regularization. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 892–900, 2011.
- [52] X. Zhou, M. Belkin, and N. Srebro. An iterated graph laplacian approach for ranking on manifolds. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 877–885. ACM, 2011.
- [53] X. Zhu, Z. Ghahramani, J. Lafferty, et al. Semi-supervised learning using Gaussian fields and harmonic functions. In *International Conference on Machine Learning*, volume 3, pages 912–919, 2003.