

HYPERGRAPH p -LAPLACIAN REGULARIZATION ON POINT CLOUDS FOR DATA INTERPOLATION

KEHAN SHI ^{*}, MARTIN BURGER ^{†,‡}

ABSTRACT. As a generalization of graphs, hypergraphs are widely used to model higher-order relations in data. This paper explores the benefit of the hypergraph structure for the interpolation of point cloud data that contain no explicit structural information. We define the ε_n -ball hypergraph and the k_n -nearest neighbor hypergraph on a point cloud and study the p -Laplacian regularization on the hypergraphs. We prove the variational consistency between the hypergraph p -Laplacian regularization and the continuum p -Laplacian regularization in a semisupervised setting when the number of points n goes to infinity while the number of labeled points remains fixed. A key improvement compared to the graph case is that the results rely on weaker assumptions on the upper bound of ε_n and k_n . To solve the convex but non-differentiable large-scale optimization problem, we utilize the stochastic primal-dual hybrid gradient algorithm. Numerical experiments on data interpolation verify that the hypergraph p -Laplacian regularization outperforms the graph p -Laplacian regularization in preventing the development of spikes at the labeled points.

1. INTRODUCTION

A hypergraph H is a triplet with the form $H = (V, E, W)$, where $V = \{x_i\}_{i=1}^n$ is the vertex set, $E = \{e_k\}_{k=1}^m$ denotes the set of hyperedges, and $W = \{w_k\}_{k=1}^m$ assigns weight $w_k > 0$ for hyperedge e_k . It is a generalization of a graph in the sense that each hyperedge $e_k \subset V$ contains an arbitrary number of vertices. Due to the applicability of modeling higher-order relations in data, hypergraphs have been extensively used in image processing [1, 2], bioinformatics [3, 4], social networks [5, 6], etc.

This paper focuses on the p -Laplacian regularization on hypergraph H ,

$$\sum_{k=1}^m w_k \max_{x_i, x_j \in e_k} |u(x_i) - u(x_j)|^p, \quad p > 1, \quad (1)$$

which was first proposed in [7] for clustering and semi-supervised learning (SSL). The authors defined the hypergraph total variation (i.e., energy (1) with $p = 1$) in analogy to the graph case [8], i.e., it is the Lovász extension of the hypergraph cut. Then a power $p > 1$ was introduced to avoid spiky solutions in SSL. Despite the successful applications of energy (1) and its variants [9, 10, 11], some of its mathematical properties are not clear. One of the fundamental questions is whether

^{*}DEPARTMENT OF MATHEMATICS, CHINA JILIANG UNIVERSITY, HANGZHOU 310018, CHINA

[†]COMPUTATIONAL IMAGING GROUP AND HELMHOLTZ IMAGING, DEUTSCHES ELEKTRONEN-SYNCHROTRON DESY, 22607 HAMBURG, GERMANY

[‡]FACHBEREICH MATHEMATIK, UNIVERSITÄT HAMBURG, 20146 HAMBURG, GERMANY

Key words and phrases. Hypergraph, p -Laplacian, continuum limit, stochastic primal-dual algorithm, data interpolation.

it is a discrete approximation of a continuum energy related to the p -Laplacian regularization

$$\int_{\Omega} |\nabla u|^p dx, \quad p > 1.$$

Since the latter is well-studied, a positive answer shall provide us with more insight into the energy (1).

The difficulty comes from the complicated structure of the hypergraph. In hypergraph-based learning, the first step is always to construct a hypergraph from the given data. This is neither trivial nor unique [12]. Typically, we construct the hyperedge using structural information, such as the attribute [13] or the network [14], that is explicitly contained in the data. This means that the formulation of energy (1) depends on the data, thus preventing us from establishing general results for it.

In this paper, we define the ε_n -ball hypergraph H_{n,ε_n} from a given point cloud $\Omega_n = \{x_i\}_{i=1}^n$ and investigate the hypergraph p -Laplacian regularization on it in a semisupervised setting. It is assumed that the first N points $\mathcal{O} := \{x_i\}_{i=1}^N$ are labeled with labels $\{y_i\}_{i=1}^N \subset \mathbb{R}$ and the remaining $n - N$ points are unlabeled and drawn from a probability measure μ supported in a bounded set $\Omega \subset \mathbb{R}^d$. The given point cloud Ω_n forms the vertices of hypergraph H_{n,ε_n} and each hyperedge is a subset of Ω_n consisting of a vertex and its ε_n -ball neighbors, i.e.,

$$V = \Omega_n, \quad E = \{e_k := B(x_k, \varepsilon_n), x_k \in \Omega_n\}_{k=1}^n, \quad \varepsilon_n > 0.$$

The homogeneous weight $W = \{w_k := 1\}_{k=1}^n$ is assigned. Then semi-supervised learning with p -Laplacian regularization (1) on the hypergraph H_{n,ε_n} is formulated as the problem of finding an estimator $u : \Omega_n \rightarrow \mathbb{R}$ by minimizing

$$\mathcal{E}_{n,\varepsilon_n}(u) = \frac{1}{n\varepsilon_n^p} \sum_{k=1}^n \max_{x_i, x_j \in B(x_k, \varepsilon_n)} |u(x_i) - u(x_j)|^p, \quad p > 1, \quad (2)$$

under constraint $u(x_i) = y_i$ for $x_i \in \mathcal{O}$. Here a scaling parameter $\frac{1}{n\varepsilon_n^p}$ is introduced to ensure that the energy is well-defined when $n \rightarrow \infty$. The method makes use of the smoothness assumption. Namely, vertices tend to share a label if they are close to each other such that they are in the same hyperedge.

We are interested in the asymptotic behavior of $\mathcal{E}_{n,\varepsilon_n}$ when the number of data points n goes to infinity while the number of labeled points N remains fixed. This corresponds to the learning problem that assigns labels to a large number of unlabeled data through a given few labeled data. It has tremendous practical value in applications where acquiring labeled data is expensive and time-consuming, while acquiring unlabeled data is relatively easy.

The continuum limit of $\mathcal{E}_{n,\varepsilon_n}$ as $n \rightarrow \infty$ reads

$$\mathcal{E}(u) = \mathcal{E}(u; \rho) = \begin{cases} 2^p \int_{\Omega} |\nabla u|^p \rho dx, & \text{if } u \in W^{1,p}(\Omega), \\ +\infty, & \text{otherwise,} \end{cases} \quad (3)$$

where ρ is the density of the probability measure μ with respect to the Lebesgue measure. To establish the connection between the discrete energy $\mathcal{E}_{n,\varepsilon_n}$ and the continuum energy \mathcal{E} , we utilize the method developed in [15, 16], where continuum limits of variational models on graphs were studied by tools of optimal transportation and TLP topology. More precisely, we prove that $\mathcal{E}_{n,\varepsilon_n}$ Γ -converges to \mathcal{E} in the

TL^p topology as $n \rightarrow \infty$ if the connection radius ε_n satisfies

$$\delta_n \ll \varepsilon_n \ll 1, \quad (4)$$

where

$$\delta_n = \begin{cases} \sqrt{\frac{\ln \ln(n)}{n}}, & \text{if } d = 1, \\ \frac{(\ln n)^{3/4}}{\sqrt{n}}, & \text{if } d = 2, \\ \frac{(\ln n)^{1/d}}{n^{1/d}}, & \text{if } d \geq 3, \end{cases} \quad (5)$$

is a constant depending on n and d .

In the case of the constrained energy

$$\mathcal{E}_{n,\varepsilon_n}^{con}(u) = \begin{cases} \mathcal{E}_{n,\varepsilon_n}(u), & \text{if } u(x_i) = y_i \text{ for } x_i \in \mathcal{O}, \\ +\infty, & \text{otherwise,} \end{cases} \quad (6)$$

where the training set $\{(x_i, y_i)\}_{i=1}^N$ is taken into account, the corresponding continuum limit becomes

$$\mathcal{E}^{con}(u) = \mathcal{E}^{con}(u; \rho) = \begin{cases} \mathcal{E}(u), & \text{if } u \in W^{1,p}(\Omega) \text{ and } u(x_i) = y_i \text{ for } x_i \in \mathcal{O}, \\ +\infty, & \text{otherwise.} \end{cases} \quad (7)$$

If $p > d$, Sobolev's embedding theorem [17] implies that the minimizer of \mathcal{E}^{con} is Hölder continuous and the constraint is well-defined. Under assumptions (4) and $p > d$, we also have that $\mathcal{E}_{n,\varepsilon_n}^{con}$ Γ -converges to \mathcal{E}^{con} in the TL^p topology as $n \rightarrow \infty$. As a consequence, the minimizer of $\mathcal{E}_{n,\varepsilon_n}^{con}$ converges to the minimizer of \mathcal{E}^{con} in the TL^p topology as $n \rightarrow \infty$. In other words, $\mathcal{E}_{n,\varepsilon_n}^{con}$ is a discrete approximation of the continuum functional \mathcal{E}^{con} on the point cloud. If $p \leq d$, $\mathcal{E}^{con} = \mathcal{E}$ and the minimizer of $\mathcal{E}_{n,\varepsilon_n}^{con}$ converges to a minimizer of \mathcal{E} instead, which is a constant. This means that $\mathcal{E}_{n,\varepsilon_n}^{con}$ is degenerate and develops spiky solutions for large n .

A similar result holds for the constrained p -Laplacian regularization on the random geometric graph under an additional assumption on the upper bound of the connection radius ε_n , i.e.,

$$\delta_n \ll \varepsilon_n \ll \left(\frac{1}{n}\right)^{\frac{1}{p}}.$$

See [16] and Section 1.1 for more details. This reveals that the hypergraph structure is beneficial in a semisupervised setting even for point cloud data that contain no explicit structural information.

The cardinality for the hyperedge of the ε_n -ball hypergraph generally varies significantly. This causes high computational costs in applications. A more feasible alternative is to consider the p -Laplacian regularization on the k_n -nearest neighbor (k_n -NN) hypergraph H_{n,k_n} , which reads

$$\mathcal{F}_{n,k_n}(u) = \frac{1}{n\bar{\varepsilon}_n^p} \sum_{k=1}^n \max_{x_i, x_j \overset{k_n}{\sim} x_k} |u(x_i) - u(x_j)|^p, \quad p > 1. \quad (8)$$

Here

$$\bar{\varepsilon}_n = \left(\frac{1}{\alpha_d} \frac{k_n}{n}\right)^{1/d}, \quad (9)$$

α_d denotes the volume of the d -dimensional unit ball, and $x_i \overset{k_n}{\sim} x_k$ denotes that vertex x_i is among the k_n nearest points to vertex x_k .

The motivation for the definition of $\bar{\varepsilon}_n$ is as follows. Let $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ be the empirical measure of Ω_n . If $\rho \equiv 1$ and n is large, then for any $x_k \in \Omega_n$ and a constant connection radius $\varepsilon_n > 0$,

$$\mu_n(B(x_k, \varepsilon_n)) \approx \mu(B(x_k, \varepsilon_n)) = \alpha_d \varepsilon_n^d.$$

This is what we are doing on the ε_n -ball hypergraph, while on the k_n -NN hypergraph, the connection radius ε_n is now a function of x_k such that

$$\mu_n(B(x_k, \varepsilon_n(x_k))) = \frac{k_n}{n}.$$

For sufficiently large n ,

$$\begin{aligned} \mu(B(x_{k_1}, \varepsilon_n(x_{k_1}))) &\approx \mu_n(B(x_{k_1}, \varepsilon_n(x_{k_1}))) = \frac{k_n}{n} \\ &= \mu_n(B(x_{k_2}, \varepsilon_n(x_{k_2}))) \approx \mu(B(x_{k_2}, \varepsilon_n(x_{k_2}))), \end{aligned}$$

which means that $\varepsilon_n(x_k)$ is approximately a constant function (denoted by $\bar{\varepsilon}_n$). Then we have

$$\alpha_d \bar{\varepsilon}_n^d = \mu(B(x_k, \bar{\varepsilon}_n)) \approx \mu_n(B(x_k, \bar{\varepsilon}_n)) \approx \mu_n(B(x_k, \varepsilon_n(x_k))) = \frac{k_n}{n},$$

from which we obtain (9).

Under the assumption

$$\delta_n \ll \bar{\varepsilon}_n \ll 1,$$

all theoretical results regarding the ε_n -ball hypergraph p -Laplacian regularization $\mathcal{E}_{n, \varepsilon_n}$ and $\mathcal{E}_{n, \varepsilon_n}^{con}$ also hold for \mathcal{F}_{n, k_n} and $\mathcal{F}_{n, k_n}^{con}$, where

$$\mathcal{F}_{n, k_n}^{con}(u) = \begin{cases} \mathcal{F}_{n, k_n}(u), & \text{if } u(x_i) = y_i \text{ for } x_i \in \mathcal{O}, \\ +\infty, & \text{otherwise.} \end{cases} \quad (10)$$

The only difference is that the continuum limits possess a different weight. Namely, \mathcal{F}_{n, k_n} and $\mathcal{F}_{n, k_n}^{con}$ ($p > d$) Γ -converge to $\mathcal{E}(\cdot; \rho^{1-p/d})$ and $\mathcal{E}^{con}(\cdot; \rho^{1-p/d})$ as $n \rightarrow \infty$ respectively.

Due to the maximum function, energy (6) and energy (10) are convex but non-differentiable for any $p \geq 1$. The primal-dual hybrid gradient (PDHG) algorithm [18] was used in [7] when $p = 1, 2$. Considering that the number of hyperedges n is usually large, the recent stochastic PDHG algorithm [19] provides us with a more efficient scheme, in which it updates a random subset of the separable dual variable in each iteration. The resulting algorithm updates the primal variable by projecting onto the training set, and updates the dual variable by solving the proximal operator of $\|\cdot\|_1^{p/p-1}$ ($p > 1$). The latter has no closed-form solution but can be solved exactly in a few iterations. In the case $p = 1$, the subproblem involving the dual variable becomes a projection onto the L^1 ball.

The theoretical results are partially verified by numerical experiments. By comparing the hypergraph p -Laplacian regularization and the graph p -Laplacian regularization for data interpolation in 1D, we observe that the hypergraph structure can better suppress spiky solutions. Experiments on higher-dimensional data interpolation, including SSL and image inpainting, are also presented to obtain the same conclusion.

At last, we mention that energy (2) and energy (8) consist of n terms, each of which is nothing but a p -th power of the objective function of the Lipschitz learning [20, 21] on a hyperedge. The Lipschitz learning (see (12) for its definition) has been

proven to be effective in suppressing spikes [22, 23]. This informally explains the numerical results that the hypergraph p -Laplacian regularization suppresses spikes better than the graph one.

This paper is organized as follows. We complete the introduction with a brief review of the graph p -Laplacian regularization. In Section 2, we present definitions of the ε_n -ball hypergraph and the k_n -NN hypergraph, and present mathematical tools needed for the following study. The continuum limits of the p -Laplacian regularization on the ε_n -ball hypergraph and the k_n -NN hypergraph in a semisupervised setting are studied in Section 3 and Section 4, respectively. Section 5 is devoted to the numerical algorithm, which is based on the stochastic primal-dual hybrid gradient method. Numerical experiments are presented in Section 6. We conclude this paper in Section 7.

1.1. Related works on the graph p -Laplacian regularization. The graph-based method has been widely used for SSL. It constructs a weighted graph G_n with vertex set $\Omega_n = \{x_i\}_{i=1}^n$ to represent the geometric structure in Ω_n . We follow the previous notation and assume that the first N vertices $\mathcal{O} = \{x_i\}_{i=1}^N$ are labeled with labels $\{y_i\}_{i=1}^N$. It is commonly assumed in the theoretical study that G_n is a random geometric graph. Namely, the unlabeled vertices $\{x_i\}_{i=N+1}^n$ are drawn from a probability measure μ and an edge $e_{i,j}$ between two vertices x_i and x_j exists if and only if their distance is smaller than a given connection radius ε_n . The weight $w_{i,j} > 0$ for edge $e_{i,j}$ is a nonincreasing function of the distance. Then an estimator is obtained by minimizing an objective functional defined on G_n under the constraint on the labeled subset.

The graph Laplacian regularization

$$\sum_{i,j=1}^n w_{i,j} |u(x_i) - u(x_j)|^2, \quad \text{subject to } u(x_i) = y_i \text{ for } x_i \in \mathcal{O},$$

was first proposed in [24] for SSL and is one of the most well-known approaches. It has attracted a lot of attention [25, 26] and has applications in clustering [27], manifold ranking [28], image processing [29], etc. Still, one of its drawbacks can not be neglected. It is degenerate in the sense that the minimizer becomes noninformative and develops spikes at the labeled points $x_i \in \mathcal{O}$ when the labeling rate $\frac{|\mathcal{O}|}{|\Omega_n|} = \frac{N}{n}$ is low [30]. Variants have been proposed as a remedy: In [31, 32], the authors proposed to assign more weights for edges adjacent to the labeled points. Instead of imposing the labels as boundary conditions, the authors of [33] constructed a source term with the labels and proposed the Poisson learning.

In [34], the authors suggested to use the p -Laplacian regularization [35, 36] with $p > d$, which has the following form

$$E_{n,\varepsilon_n}^{con}(u) = \begin{cases} \frac{1}{n^2\varepsilon_n^p} \sum_{i,j=1}^n w_{i,j} |u(x_i) - u(x_j)|^p, & \text{if } u(x_i) = y_i \text{ for } x_i \in \mathcal{O}, \\ +\infty, & \text{otherwise.} \end{cases} \quad (11)$$

A rigorous study [16] indicates that if $\delta_n \ll \varepsilon_n \ll (\frac{1}{n})^{\frac{1}{p}}$, energy $E_{n,\varepsilon_n}^{con}$ is variational consistent in the TL^p topology with the continuum energy

$$E^{con}(u) = \begin{cases} \sigma_p \int_{\Omega} |\nabla u(x)|^p \rho^2(x) dx, & \text{if } u \in W^{1,p}(\Omega) \text{ and } u(x_i) = y_i \text{ for } x_i \in \mathcal{O}, \\ +\infty, & \text{otherwise,} \end{cases}$$

where $\sigma_p > 0$ is a constant. Furthermore, the minimizer of $E_{n,\varepsilon_n}^{con}$ converges to the minimizer of E^{con} in the TL^p topology as $n \rightarrow \infty$. The assumption $\delta_n \ll \varepsilon_n \ll (\frac{1}{n})^{\frac{1}{p}}$ and the definition of δ_n (see (5)) imply that $p > d$. Consequently, by Sobolev's embedding theorem, the minimizer of E^{con} is Hölder continuous and the constraint in E^{con} is well-defined. If $p \leq d$ or $n\varepsilon_n^p \rightarrow \infty$, the continuum limiting energy of $E_{n,\varepsilon_n}^{con}$ is similar to E^{con} but without the constraint on the labeled subset, i.e.,

$$E(u) = \begin{cases} \sigma_p \int_{\Omega} |\nabla u(x)|^p \rho^2(x) dx, & \text{if } u \in W^{1,p}(\Omega), \\ +\infty, & \text{otherwise.} \end{cases}$$

This indicates that $E_{n,\varepsilon_n}^{con}$ forgets the labels for large n and explains the occurrence of spikes in the graph-based SSL.

The minimizer of energy $E_{n,\varepsilon_n}^{con}$ gains more smoothness as p increases. Recently, some authors studied the case $p \rightarrow +\infty$, which leads to the Lipschitz learning for SSL [20, 21],

$$\frac{1}{\varepsilon_n} \max_{x_i, x_j \in \Omega_n} w_{i,j} |u(x_i) - u(x_j)|, \quad \text{subject to } u(x_i) = y_i \text{ for } x_i \in \mathcal{O}. \quad (12)$$

As expected, the minimizer of the continuum limiting energy

$$E_{\infty}(u) = \begin{cases} \sigma_{\infty} \text{ess sup}_{x \in \Omega} |\nabla u(x)|, & \text{if } u \in W^{1,\infty}(\Omega) \text{ and } u(x_i) = y_i \text{ for } x_i \in \mathcal{O}, \\ \infty, & \text{otherwise,} \end{cases}$$

attains the labels continuously [22, 23]. However, (12) is less attractive for SSL since it forgets the distribution of the unlabeled points as $n \rightarrow \infty$ [34]. For the related ∞ -Laplacian equation, the author of [22] proved that the self-turning weights allow to remember the distribution.

Finally, we mention that the k_n -nearest neighbor (k_n -NN) graph has a more practical value than the random geometric graph due to its sparsity. In [37], the author studied the continuum limit of the total variation on the k_n -NN graph for clustering. The method is also valid for the graph p -Laplacian regularization with $p > 1$ on the k_n -NN graph. The basic idea is to replace the connection radius ε_n in $E_{n,\varepsilon_n}^{con}$ by $\bar{\varepsilon}_n$ defined in (9).

2. PRELIMINARIES

2.1. Settings. Let

$$\Omega_n = \{x_1, \dots, x_N, x_{N+1}, \dots, x_n\}$$

be a point cloud in the bounded domain $\Omega \subset \mathbb{R}^d$ ($d \geq 1$). We are given a training set $\{(x_i, y_i), i = 1, \dots, N\}$ with N labeled points $\mathcal{O} = \{x_i\}_{i=1}^N$, where each $y_i \in \{1, \dots, L\}$ for an integer $L \geq 2$ is the label of point $x_i \in \mathcal{O}$. The task of semi-supervised learning (SSL) is to assign labels from the label set $\{1, \dots, L\}$ to the remaining points $\{x_{N+1}, \dots, x_n\}$.

In this paper, we let the number of the labeled points N be fixed and study the asymptotic behavior of variational models defined on Ω_n when n goes to infinity. It is assumed that the unlabeled points $\{x_i\}_{i=N+1}^n$ are independently and identically distributed (i.i.d.) random samples of a probability measure μ on Ω . The density

ρ of μ (with respect to the Lebesgue measure) is a continuous function with positive lower and upper bounds, i.e.,

$$0 < \inf_{x \in \Omega} \rho(x) \leq \sup_{x \in \Omega} \rho(x) < \infty.$$

Throughout this paper, Ω is a connected and bounded domain with Lipschitz boundary $\partial\Omega$.

2.2. Hypergraphs. We consider regularizations on hypergraphs in a semisupervised setting. The first step is to generate a hypergraph from the given data set Ω_n . Clearly, it is not unique and affects the performance for SSL. A brief review of the generation of hypergraphs can be found in [12]. Since the data set Ω_n contains no explicit structural information, we construct the hypergraph by the distance-based method. More precisely, each hyperedge is a subset consisting of a vertex and its neighbors. The following two hypergraphs called the ε_n -ball hypergraph and the k_n -NN hypergraph are considered in this paper. Such hypergraphs have been used for image classification, image segmentation, and recommender systems [38, 39, 40].

The ε_n -ball hypergraph. Let ε_n be a constant representing the connection radius. The ε_n -ball hypergraph is defined as

$$H_{n,\varepsilon_n} = (V_n, E_{n,\varepsilon_n}, W_n),$$

with vertices $V_n = \Omega_n = \{x_k\}_{k=1}^n$, hyperedges $E_{n,\varepsilon_n} = \{e_k^{(\varepsilon_n)}\}_{k=1}^n$, and weights $W_n = \{w_k\}_{k=1}^n$, where

$$e_k^{(\varepsilon_n)} = \{x_j \in \Omega_n : |x_k - x_j| \leq \varepsilon_n\}$$

is the hyperedge corresponding to the centroid x_k . In this definition, we have no preference for different hyperedges. Consequently, we let $w_k = 1$ for any $k = 1, 2, \dots, n$. By considering H_{n,ε_n} in the p -Laplacian regularization (1), we obtain $\mathcal{E}_{n,\varepsilon_n}$ (defined in (2)) with an additional scaling parameter.

In the following, we assume that $\varepsilon_n \gg \delta_n$. Then, with probability one, the ε_n -ball graph is connected [41, 15]. Namely, for any $x_i, x_j \in \Omega_n$, there exist vertices $x_{k_1} = x_i, x_{k_2}, \dots, x_{k_{s-1}}, x_{k_s} = x_j$, such that $|x_{k_l} - x_{k_{l-1}}| \leq \varepsilon_n$ for any $2 \leq l \leq s$. This indicates that hyperedge $e_k^{(\varepsilon_n)}$ has cardinality $|e_k^{(\varepsilon_n)}| \geq 2$ for any $k = 1, 2, \dots, n$ and hypergraph H_{n,ε_n} is also connected.

The k_n -NN hypergraph. The k_n -nearest neighbor (k_n -NN) hypergraph is more frequently used in practice. In contrast to the ε_n -ball hypergraph, the k_n -NN hypergraph selects a fixed $k_n \geq 2$ neighbors for each vertex to form a hyperedge. More precisely, it is defined as

$$H_{n,k_n} = (V_n, E_{n,k_n}, W_n),$$

with a different hyperedge set $E_{n,k_n} = \{e_k^{(k_n)}\}_{k=1}^n$, where

$$e_k^{(k_n)} = \{x_j \in \Omega_n : x_j \overset{k_n}{\sim} x_k\},$$

and $x_j \overset{k_n}{\sim} x_k$ denotes that x_j is among the k_n -nearest neighbors of x_k . Hence H_{n,k_n} is a k_n uniform hypergraph.

Again, if $\bar{\varepsilon}_n \gg \delta_n$, the k_n -NN graph is connected with probability one [37]. Consequently, $|e_k^{(k_n)}| \geq 2$ for any $k = 1, 2, \dots, n$ and hypergraph H_{n,k_n} is also connected.

2.3. Mathematical tools. In this subsection, we present mathematical tools that will be needed for our study. We follow the idea of [16] and consider the convergence of minimizers of discrete functionals defined on hypergraphs to minimizers of continuum functionals. It relies on properties of Γ -convergence [42, 43], the optimal transportation from the probability measure μ to the associated empirical measure μ_n [44], and the TL^p space [15].

Definition 2.1. *Let X be a metric space and $F_n, F : X \rightarrow [-\infty, \infty]$ be functionals. We say that F_n Γ -converges to F as $n \rightarrow \infty$, denoted by $F_n \xrightarrow{\Gamma} F$, if for every $x \in X$,*

- *Liminf inequality: For every sequence $\{x_n\}_{n \in \mathbb{N}} \subset X$ converging to x ,*

$$\liminf_{n \rightarrow \infty} F_n(x_n) \geq F(x).$$

- *Limsup inequality: There exists a sequence $\{x_n\}_{n \in \mathbb{N}} \subset X$ converging to x such that*

$$\limsup_{n \rightarrow \infty} F_n(x_n) \leq F(x).$$

The Γ -convergence implies the convergence of minimizers, see [43, Theorem 1.21] for the proof.

Proposition 2.2. *Let X be a metric space, $F_n, F : X \rightarrow [0, \infty]$ be functionals, and $F_n \xrightarrow{\Gamma} F \not\equiv \infty$ as $n \rightarrow \infty$. If there exists a precompact sequence $\{x_n\}_{n \in \mathbb{N}}$ such that*

$$\lim_{n \rightarrow \infty} \left(F_n(x_n) - \inf_{x \in X} F_n(x) \right) = 0,$$

then

$$\lim_{n \rightarrow \infty} \inf_{x \in X} F_n(x) = \inf_{x \in X} F(x),$$

and any cluster point of $\{x_n\}_{n \in \mathbb{N}}$ is a minimizer of F .

Let $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ be the empirical measure of Ω_n . We call T_n a transportation map from μ to μ_n if it satisfies the push-forward condition

$$\mu_n = T_n \# \mu = \mu \circ T_n^{-1}.$$

It implies the change of variables,

$$\frac{1}{n} \sum_{i=1}^n \varphi(x_i) = \int_{\Omega} \varphi(x) d\mu_n(x) = \int_{\Omega} \varphi(T_n(x)) d\mu(x) = \int_{\Omega} \varphi(T_n(x)) \rho(x) dx, \quad (13)$$

and

$$\mu_n\text{-ess sup}_{x \in \Omega_n} \varphi(x) = \mu\text{-ess sup}_{y \in \Omega} \varphi(T_n(y)) = \text{ess sup}_{y \in \Omega} \varphi(T_n(y)), \quad (14)$$

for any function $\varphi : \Omega_n \rightarrow \mathbb{R}$.

The following proposition states the existence of the transportation map from μ to μ_n and its L^∞ estimates [16, 44].

Proposition 2.3. *Let $\{x_i\}_{i=1}^\infty$ be a sequence of independent random variables with distribution μ on Ω . Then, almost surely there exist transportation maps $\{T_n\}_{n=1}^\infty$ from μ to μ_n , such that*

$$c \leq \liminf_{n \rightarrow \infty} \frac{\|Id - T_n\|_{L^\infty(\Omega)}}{\delta_n} \leq \limsup_{n \rightarrow \infty} \frac{\|Id - T_n\|_{L^\infty(\Omega)}}{\delta_n} \leq C, \quad (15)$$

where δ_n is defined in (5).

In the following, we do not use the almost sure statement but consider our problem in the deterministic setting. More precisely, we assume that the unlabeled points are the realization of the random variables $\{x_i\}_{i=1}^\infty$ mentioned in Proposition 2.3 such that conclusion (15) holds.

To conclude this subsection, we recall the TL^p space [15], which provides an appropriate topology for the convergence of discrete functions to continuum functions.

Definition 2.4. *The space $TL^p(\Omega)$ is defined as*

$$TL^p(\Omega) = \{(\nu, g) : \nu \in \mathcal{P}(\bar{\Omega}), g \in L^p(\nu)\},$$

where $\mathcal{P}(\bar{\Omega})$ is the space of probability measures. It is endowed with the TL^p -metric

$$d_{TL^p}^p((\nu_1, g_1), (\nu_2, g_2)) = \inf_{\pi \in \Pi(\nu_1, \nu_2)} \left\{ \iint_{\Omega \times \Omega} |x - y|^p + |g_1(x) - g_2(y)|^p d\pi(x, y) \right\},$$

where $\Pi(\nu_1, \nu_2)$ is the set of transportation plans.

The following proposition provides a characterization of the convergence of sequences in $TL^p(\Omega)$ [15].

Proposition 2.5. *Let $(\nu, g) \in TL^p(\Omega)$ and $\{(\nu_n, g_n)\}_{n=1}^\infty \subset TL^p(\Omega)$, where ν is absolutely continuous with respect to the Lebesgue measure. Then $(\nu_n, g_n) \rightarrow (\nu, g)$ in $TL^p(\Omega)$ as $n \rightarrow \infty$ if and only if $\nu_n \rightarrow \nu$ weakly as $n \rightarrow \infty$ and there exists a sequence of transportation maps $\{T_n\}_{n=1}^\infty$ from ν to ν_n such that*

$$\lim_{n \rightarrow \infty} \int_{\Omega} |x - T_n(x)|^p d\nu(x) = 0,$$

and

$$\lim_{n \rightarrow \infty} \int_{\Omega} |g(x) - g_n(T_n(x))|^p d\nu(x) = 0.$$

In our case, the empirical measure μ_n converges weakly to μ as $n \rightarrow \infty$. Proposition 2.3 implies the existence of the transportation map T_n . To conclude the convergence of (μ_n, g_n) to (μ, g) in $TL^p(\Omega)$ we only need to verify the convergence of $g_n \circ T_n$ to g in $L^p(\Omega)$.

3. CONTINUUM LIMIT OF THE p -LAPLACIAN REGULARIZATION ON THE ε_n -BALL HYPERGRAPH

In this section, we study the continuum limit of the p -Laplacian regularization $\mathcal{E}_{n, \varepsilon_n}^{con}$ (see (6) for its definition) on the ε_n -ball hypergraph H_{n, ε_n} . The main result of this section is stated as follows.

Theorem 3.1. *Let $p > d$ and $\delta_n \ll \varepsilon_n \ll 1$. If u_n is a minimizer of $\mathcal{E}_{n, \varepsilon_n}^{con}$, then almost surely,*

$$(\mu_n, u_n) \rightarrow (\mu, u), \quad \text{in } TL^p(\Omega),$$

as $n \rightarrow \infty$ for a function $u \in L^p(\Omega)$. Furthermore, u is continuous,

$$\lim_{n \rightarrow \infty} \max_{x_k \in \Omega_n \cap \Omega'} |u_n(x_k) - u(x_k)| = 0, \quad (16)$$

for any Ω' compactly contained in Ω , and u is a minimizer of \mathcal{E}^{con} .

Proof. We first claim that $\mathcal{E}_{n,\varepsilon_n}^{con}$ admits at least one minimizer u_n for any fixed $n > 0$ and \mathcal{E}^{con} admits a unique minimizer.

Let us define the vector space

$$V = \{v : \Omega_n \rightarrow \mathbb{R}\},$$

and its subspace

$$V_{\mathcal{O}} = \{v \in V : v(x_i) = y_i \text{ for } x_i \in \mathcal{O}\}.$$

It follows from the Poincaré type inequality (see [45, Lemma 1] for $p = 2$, the proof works for any $p > 1$) that

$$\begin{aligned} \sum_{x_i \in \Omega_n \setminus \mathcal{O}} |u(x_i)|^p &\leq C \sum_{x_i \in \Omega_n} \sum_{x_j \in B(x_i, \varepsilon_n)} |u(x_i) - u(x_j)|^p + C \sum_{x_i \in \mathcal{O}} |u(x_i)|^p \\ &\leq C \mathcal{E}_{n,\varepsilon_n}(u) + C \sum_{x_i \in \mathcal{O}} |u(x_i)|^p, \end{aligned}$$

and $\mathcal{E}_{n,\varepsilon_n}$ is coercive on $V_{\mathcal{O}}$. Clearly, it is also lower semi-continuous on $V_{\mathcal{O}}$. Then $\mathcal{E}_{n,\varepsilon_n}$ admits a minimizer u_n on $V_{\mathcal{O}}$, which is also a minimizer of $\mathcal{E}_{n,\varepsilon_n}^{con}$. The existence of a unique minimizer for \mathcal{E}^{con} can be verified similarly, see also [16].

Let $M = \max_{1 \leq i \leq N} |y_i|$. For any function $f : \Omega \rightarrow \mathbb{R}$, we have

$$\mathcal{E}_{n,\varepsilon_n}^{con}(f_M) \leq \mathcal{E}_{n,\varepsilon_n}^{con}(f),$$

where $f_M := \max\{\min\{f, M\}, -M\}$ denotes the truncation of f . Consequently,

$$\|u_n\|_{L^\infty(\Omega)} \leq M, \tag{17}$$

for any $n > 0$. Notice that

$$\begin{aligned} \frac{1}{n^2 \varepsilon_n^p} \sum_{i=1}^n \sum_{j=1}^n \chi_{B(x_i, \varepsilon_n)}(x_j) |u_n(x_i) - u_n(x_j)|^p \\ \leq \frac{1}{n \varepsilon_n^p} \sum_{k=1}^n \max_{x_i, x_j \in B(x_k, \varepsilon_n)} |u_n(x_i) - u_n(x_j)|^p \\ = \mathcal{E}_{n,\varepsilon_n}(u_n) \leq \mathcal{E}_{n,\varepsilon_n}^{con}(u_n) < \infty. \end{aligned} \tag{18}$$

The left-hand side is just the graph p -Laplacian regularization with the indicator function as the kernel. By the compactness of the graph p -Laplacian regularization (see [16, Proposition 4.4]) and (17), there exists a subsequence of $\{u_n\}_{n \in \mathbb{N}}$, denoted by $\{u_{n_m}\}_{m \in \mathbb{N}}$, and a function $u \in L^p(\Omega)$, such that

$$(\mu_{n_m}, u_{n_m}) \rightarrow (\mu, u), \quad \text{in } TLP(\Omega), \tag{19}$$

as $m \rightarrow \infty$.

Let us assume that

$$\mathcal{E}_{n,\varepsilon_n}^{con} \xrightarrow{\Gamma} \mathcal{E}^{con}, \tag{20}$$

in the TLP metric on the set $\{(\nu, g) : \nu \in \mathcal{P}(\Omega), \|g\|_{L^\infty(\nu)} \leq M\}$, which will be proven in the following in Theorem 3.9. Then it follows from Proposition 2.2 that u is a minimizer of \mathcal{E}^{con} . Since the minimizer of \mathcal{E}^{con} is unique, the convergence in (19) holds along the whole sequence. The uniform convergence (16) will be proven in Lemma 3.8. \square

Remark 3.2. If $1 < p \leq d$, we still have (19). But now $\mathcal{E}^{con} = \mathcal{E}$ and (20) becomes

$$\mathcal{E}_{n,\varepsilon_n}^{con} \xrightarrow{\Gamma} \mathcal{E}.$$

By Proposition 2.2, u is a minimizer of \mathcal{E} . In this case we no longer have the uniform convergence (16).

The rest of this section is devoted to the proof of the Γ -convergence (20) and the uniform convergence (16).

3.1. Nonlocal to local convergence. In this section, we introduce the nonlocal functional

$$\mathcal{E}_\varepsilon(u) = \int_{\Omega} \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{y,z \in B(x,\varepsilon(x)) \cap \Omega} |u(z) - u(y)|^p \rho(x) dx$$

as a bridge between $\mathcal{E}_{n,\varepsilon_n}$ and \mathcal{E} , and prove the Γ -convergence of \mathcal{E}_ε to \mathcal{E} in $L^p(\Omega)$ as $\varepsilon \rightarrow 0$. Here we consider a general setting where the connection radius ε depends on the location $x \in \Omega$. This is helpful for Section 4 in which we study the continuum limit of the p -Laplacian regularization on the k_n -NN hypergraph.

The Γ -convergence of \mathcal{E}_ε to \mathcal{E} in $L^p(\Omega)$ as $\varepsilon \rightarrow 0$ is understood in the following sense. For any sequence of functions $\{\varepsilon_n(x)\}_{n \in \mathbb{N}}$ such that

$$1 \leq \frac{\sup_{x \in \Omega} \varepsilon_n(x)}{\inf_{x \in \Omega} \varepsilon_n(x)} \leq C < \infty, \quad \forall n > 0, \quad (21)$$

and

$$0 < \inf_{x \in \Omega} \varepsilon_n(x) \leq \sup_{x \in \Omega} \varepsilon_n(x) \rightarrow 0,$$

as $n \rightarrow \infty$, we have

$$\mathcal{E}_{\varepsilon_n} \xrightarrow{\Gamma} \mathcal{E},$$

in $L^p(\Omega)$ as $n \rightarrow \infty$. In this subsection, we simply use $\varepsilon \rightarrow 0$ to represent that $\sup_{x \in \Omega} \varepsilon_n(x) \rightarrow 0$ as $n \rightarrow \infty$ for any sequence of functions $\{\varepsilon_n(x)\}_{n \in \mathbb{N}}$.

The characteristic function $\eta : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by

$$\eta(x) = \begin{cases} 1, & \text{if } |x| \leq 1, \\ 0, & \text{if } |x| > 1, \end{cases}$$

will be used several times in the following.

Lemma 3.3. *Let $\{u_\varepsilon\}$ be a sequence of uniformly bounded C^2 functions. If $\nabla u_\varepsilon \rightarrow \nabla u$ in $L^p(\Omega)$ as $\varepsilon \rightarrow 0$, then*

$$\lim_{\varepsilon \rightarrow 0} \mathcal{E}_\varepsilon(u_\varepsilon) = \mathcal{E}(u).$$

Proof. By change of variables $y = x + \varepsilon(x)\hat{y}$, $z = x + \varepsilon(x)\hat{z}$ and Taylor's expansion,

$$\begin{aligned} \mathcal{E}_\varepsilon(u_\varepsilon) &= \int_{\Omega} \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{y,z \in \mathbb{R}^d} \eta\left(\frac{x-y}{\varepsilon(x)}\right) \eta\left(\frac{x-z}{\varepsilon(x)}\right) |u_\varepsilon(z) - u_\varepsilon(y)|^p \rho(x) dx - I_1 \\ &= \int_{\Omega} \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in \mathbb{R}^d} \eta(\hat{y}) \eta(\hat{z}) |u_\varepsilon(x + \varepsilon(x)\hat{z}) - u_\varepsilon(x + \varepsilon(x)\hat{y})|^p \rho(x) dx - I_1 \\ &= \int_{\Omega} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in \mathbb{R}^d} \eta(\hat{y}) \eta(\hat{z}) |\nabla u_\varepsilon(x) \cdot (\hat{z} - \hat{y}) + \mathcal{O}(\varepsilon(x))|^p \rho(x) dx - I_1 \\ &= I_2 - I_1, \end{aligned}$$

where

$$\begin{aligned} |I_1| &= \left| \int_{\Omega} \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{y,z \in \mathbb{R}^d} \eta\left(\frac{x-y}{\varepsilon(x)}\right) \eta\left(\frac{x-z}{\varepsilon(x)}\right) |u_{\varepsilon}(z) - u_{\varepsilon}(y)|^p \rho(x) dx \right. \\ &\quad \left. - \int_{\Omega} \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{y,z \in \Omega} \eta\left(\frac{x-y}{\varepsilon(x)}\right) \eta\left(\frac{x-z}{\varepsilon(x)}\right) |u_{\varepsilon}(z) - u_{\varepsilon}(y)|^p \rho(x) dx \right| \\ &\leq C \int_{\Omega} \operatorname{ess\,sup}_{y \in \mathbb{R}^d, z \in \mathbb{R}^d \setminus \Omega} \eta\left(\frac{x-y}{\varepsilon(x)}\right) \eta\left(\frac{x-z}{\varepsilon(x)}\right) dx, \end{aligned}$$

converges to zero as $\varepsilon \rightarrow 0$. Consequently,

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \mathcal{E}_{\varepsilon}(u_{\varepsilon}) &= \lim_{\varepsilon \rightarrow 0} I_2 = \int_{\Omega} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in \mathbb{R}^d} \eta(\hat{y}) \eta(\hat{z}) |\nabla u(x) \cdot (\hat{z} - \hat{y})|^p \rho(x) dx \\ &= 2^p \int_{\Omega} |\nabla u|^p \rho dx = \mathcal{E}(u). \end{aligned}$$

□

Lemma 3.3 implies the liminf inequality of the Γ -convergence, which is stated as follows.

Lemma 3.4. *If $u_{\varepsilon} \rightarrow u$ in $L^p(\Omega)$ as $\varepsilon \rightarrow 0$, then*

$$\liminf_{\varepsilon \rightarrow 0} \mathcal{E}_{\varepsilon}(u_{\varepsilon}) \geq \mathcal{E}(u).$$

Proof. Define $\Omega_{\delta} = \{x \in \Omega : \operatorname{dist}(x, \partial\Omega) > \delta\}$ for any $\delta > 0$, $\varepsilon_1 = \inf_{x \in \Omega} \varepsilon(x)$, and $\varepsilon_2 = \sup_{x \in \Omega} \varepsilon(x)$. Assume w.l.o.g. that $\mathcal{E}_{\varepsilon}(u_{\varepsilon}) < \infty$. For any $\alpha, \beta \in \mathbb{R}^d$ with $|\alpha| = |\beta| = r \in (0, 1)$, it follows from (21) that

$$\int_{\Omega_{\varepsilon_1}} \left| \frac{u_{\varepsilon}(x + \varepsilon_1 \alpha) - u_{\varepsilon}(x + \varepsilon_1 \beta)}{\varepsilon_1} \right|^p dx \leq \frac{1}{\inf \rho} \left(\frac{\varepsilon_2}{\varepsilon_1} \right)^p \mathcal{E}_{\varepsilon}(u_{\varepsilon}) < \infty.$$

Then there exists a function $g \in L^p(\Omega)$ such that (up to a subsequence),

$$\chi_{\Omega_{\varepsilon_1}}(\cdot) \frac{u_{\varepsilon}(\cdot + \varepsilon_1 \alpha) - u_{\varepsilon}(\cdot + \varepsilon_1 \beta)}{\varepsilon_1} \rightharpoonup g, \quad \text{in } L^p(\Omega),$$

as $\varepsilon \rightarrow 0$. Namely, for any $\varphi(x) \in C_0^{\infty}(\Omega)$, if ε_1 is sufficiently small,

$$\begin{aligned} \int_{\Omega_{\varepsilon_1}} \frac{u_{\varepsilon}(x + \varepsilon_1 \alpha) - u_{\varepsilon}(x + \varepsilon_1 \beta)}{\varepsilon_1} \varphi(x) dx &= \int_{\Omega} u_{\varepsilon}(x) \frac{\varphi(x - \varepsilon_1 \alpha) - \varphi(x - \varepsilon_1 \beta)}{\varepsilon_1} dx \\ &\rightarrow - \int_{\Omega} u(x) \nabla \varphi(x) \cdot (\alpha - \beta) dx = \int_{\Omega} g(x) \varphi(x) dx. \end{aligned}$$

By choosing $\alpha = r e_i$ and $\beta = -r e_i$, we have

$$-2r \int_{\Omega} u(x) \frac{\partial}{\partial x_i} \varphi(x) dx = \int_{\Omega} g(x) \varphi(x) dx,$$

which implies that $u \in W^{1,p}(\Omega)$.

Assume that $J : \mathbb{R}^d \rightarrow [0, \infty)$ is a standard mollifier with $\operatorname{supp} J \subset \overline{B(0, 1)}$ and $\int_{\mathbb{R}^d} J(s) ds = 1$. Define $J_{\delta}(s) = \frac{1}{\delta^d} J(s/\delta)$. Then for a small $\delta > 0$ and a sufficiently

small $\varepsilon > 0$, it follows from a change of variables $\hat{y} = y + s$, $\hat{z} = z + s$, and Jensen's inequality that

$$\begin{aligned}
\mathcal{E}_\varepsilon(u_\varepsilon) &= \int_\Omega \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{y,z \in B(x,\varepsilon(x)) \cap \Omega} |u_\varepsilon(z) - u_\varepsilon(y)|^p \rho(x) dx \\
&= \int_{\mathbb{R}^d} J_\delta(s) \int_\Omega \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{y,z \in B(x,\varepsilon(x)) \cap \Omega} |u_\varepsilon(z) - u_\varepsilon(y)|^p \rho(x) dx ds \\
&\geq \int_{\mathbb{R}^d} \int_{\Omega_{2\delta}} J_\delta(s) \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in B(x+s,\varepsilon(x))} |u_\varepsilon(\hat{z} - s) - u_\varepsilon(\hat{y} - s)|^p \rho(x) dx ds \\
&\geq \int_{\Omega_{2\delta}} \frac{1}{(\varepsilon(x))^p} \inf_{|r| \leq \delta} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in B(x+r,\varepsilon(x))} \int_{\mathbb{R}^d} J_\delta(s) |u_\varepsilon(\hat{z} - s) - u_\varepsilon(\hat{y} - s)|^p ds \rho(x) dx \\
&\geq \int_{\Omega_{2\delta}} \frac{1}{(\varepsilon(x))^p} \inf_{|r| \leq \delta} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in B(x+r,\varepsilon(x))} \left| \int_{\mathbb{R}^d} J_\delta(s) (u_\varepsilon(\hat{z} - s) - u_\varepsilon(\hat{y} - s)) ds \right|^p \rho(x) dx \\
&= \int_{\Omega_{2\delta}} \frac{1}{(\varepsilon(x))^p} \inf_{|s| \leq \delta} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in B(x+s,\varepsilon(x))} |u_{\varepsilon,\delta}(\hat{z}) - u_{\varepsilon,\delta}(\hat{y})|^p \rho(x) dx
\end{aligned}$$

where $u_{\varepsilon,\delta} = (u_\varepsilon)_\delta = J_\delta * u_\varepsilon$. Clearly,

$$\nabla u_{\varepsilon,\delta} \rightarrow \nabla u_\delta, \quad \text{in } L^p(\Omega),$$

as $\varepsilon \rightarrow 0$ for a fixed $\delta > 0$,

$$\|\nabla u_{\varepsilon,\delta}\|_{L^\infty(\Omega)} + \|D^2 u_{\varepsilon,\delta}\|_{L^\infty(\Omega)} \leq C,$$

for a constant C depending not on ε , and

$$\nabla u_\delta \rightarrow \nabla u, \quad \text{in } L^p(\Omega),$$

as $\delta \rightarrow 0$. It is possible to consider a subsequence of $\{u_\delta\}$ such that the above convergence holds almost everywhere.

It follows from the proof of Lemma 3.3 and Lebesgue's dominated convergence theorem that

$$\begin{aligned}
&\liminf_{\varepsilon \rightarrow 0} \mathcal{E}_\varepsilon(u_\varepsilon) \\
&\geq \lim_{\delta \rightarrow 0} \liminf_{\varepsilon \rightarrow 0} \int_{\Omega_{2\delta}} \frac{1}{(\varepsilon(x))^p} \inf_{|s| \leq \delta} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in B(x+s,\varepsilon(x))} |u_{\varepsilon,\delta}(\hat{z}) - u_{\varepsilon,\delta}(\hat{y})|^p \rho(x) dx \\
&= \lim_{\delta \rightarrow 0} 2^p \int_{\Omega_{2\delta}} \inf_{|s| \leq \delta} |\nabla u_\delta(x+s)|^p \rho(x) dx = 2^p \int_\Omega |\nabla u|^p \rho dx = \mathcal{E}(u).
\end{aligned}$$

□

We now prove the limsup inequality of the Γ -convergence.

Lemma 3.5. *For any $u \in L^p(\Omega)$, there exists a sequence of functions $\{u_\varepsilon\} \subset L^p(\Omega)$ such that $u_\varepsilon \rightarrow u$ in $L^p(\Omega)$ as $\varepsilon \rightarrow 0$ and*

$$\limsup_{\varepsilon \rightarrow 0} \mathcal{E}_\varepsilon(u_\varepsilon) \leq \mathcal{E}(u).$$

Proof. Assume w.l.o.g. that $\mathcal{E}(u) < \infty$, i.e., $u \in W^{1,p}(\Omega)$. By [15, Remark 2.7], we only need to prove the limsup inequality for all u in a dense subset of $W^{1,p}(\Omega)$, e.g., $C^2(\bar{\Omega}) \cap W^{1,p}(\Omega)$.

Let $u \in C^2(\overline{\Omega}) \cap W^{1,p}(\Omega)$ and $u_\varepsilon = u$. Define $\Omega^\varepsilon = \{x \in \mathbb{R}^d : \text{dist}(x, \Omega) < \sup_{x \in \Omega} \varepsilon(x)\}$. By extension, u is well-defined on Ω^ε . By change of variables $y = x + \varepsilon(x)\hat{y}$, $z = x + \varepsilon(x)\hat{z}$, and Taylor's expansion,

$$\begin{aligned} \mathcal{E}_\varepsilon(u) &= \int_{\Omega} \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{y,z \in B(x,\varepsilon(x)) \cap \Omega} |u(z) - u(y)|^p \rho(x) dx \\ &= \int_{\Omega} \frac{1}{(\varepsilon(x))^p} \operatorname{ess\,sup}_{y,z \in B(x,\varepsilon(x)) \cap \Omega} \left| \int_0^1 \nabla u(y + t(z-y)) \cdot (z-y) dt \right|^p \rho(x) dx \\ &\leq \int_{\Omega} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in B(0,1)} \left| \int_0^1 \nabla u(x + \varepsilon(x)\hat{y} + \varepsilon(x)t(\hat{z} - \hat{y})) \cdot (\hat{z} - \hat{y}) dt \right|^p \rho(x) dx \\ &\leq \int_{\Omega} \operatorname{ess\,sup}_{\hat{y}, \hat{z} \in B(0,1)} |\nabla u(x) \cdot (\hat{z} - \hat{y})|^p \rho(x) dx + C \sup_{x \in \Omega} (\varepsilon(x))^p, \end{aligned}$$

where the constant C depends on $\|D^2u\|_{L^\infty(\Omega)}$. We conclude that

$$\limsup_{\varepsilon \rightarrow 0} \mathcal{E}_\varepsilon(u) \leq \mathcal{E}(u),$$

and completes the proof. \square

Combining Lemma 3.4 and Lemma 3.5 we arrive at the Γ -convergence of \mathcal{E}_ε to \mathcal{E} .

Theorem 3.6. *Let $1 < p < \infty$. Then*

$$\mathcal{E}_\varepsilon \xrightarrow{\Gamma} \mathcal{E},$$

in $L^p(\Omega)$ as $\varepsilon \rightarrow 0$.

3.2. Discrete to continuum convergence. The Γ -convergence of nonlocal functional \mathcal{E}_ε implies the Γ -convergence of discrete functional $\mathcal{E}_{n,\varepsilon_n}$. This is achieved by rewriting the discrete functional as a nonlocal functional via the transportation map T_n introduced in Proposition 2.3.

Theorem 3.7. *Let $1 < p < \infty$ and $\delta_n \ll \varepsilon_n \ll 1$. Then with probability one,*

$$\mathcal{E}_{n,\varepsilon_n} \xrightarrow{\Gamma} \mathcal{E},$$

in $TL^p(\Omega)$ as $n \rightarrow \infty$.

Proof. Let

$$\tilde{\varepsilon}_n := \varepsilon_n + 2\|T_n - Id\|_{L^\infty(\Omega)}, \quad \hat{\varepsilon}_n := \varepsilon_n - 2\|T_n - Id\|_{L^\infty(\Omega)}.$$

For any $x, y \in \Omega$,

$$|T_n(x) - T_n(y)| \leq \varepsilon_n \implies |x - y| \leq \tilde{\varepsilon}_n, \quad |x - y| \leq \hat{\varepsilon}_n \implies |T_n(x) - T_n(y)| \leq \varepsilon_n.$$

It follows that

$$\eta \left(\frac{x-y}{\hat{\varepsilon}_n} \right) \leq \eta \left(\frac{T_n(x) - T_n(y)}{\varepsilon_n} \right) \leq \eta \left(\frac{x-y}{\tilde{\varepsilon}_n} \right). \quad (22)$$

Moreover, (15) yields

$$\lim_{n \rightarrow \infty} \frac{\hat{\varepsilon}_n}{\varepsilon_n} = \lim_{n \rightarrow \infty} \frac{\tilde{\varepsilon}_n}{\varepsilon_n} = 1.$$

Notice from (13) and (14) that

$$\begin{aligned} & \mathcal{E}_{n,\varepsilon_n}(u_n) \\ &= \frac{1}{\varepsilon_n^p} \int_{\Omega} \operatorname{ess\,sup}_{y,z \in \Omega} \eta \left(\frac{T_n(x) - T_n(z)}{\varepsilon_n} \right) \eta \left(\frac{T_n(x) - T_n(y)}{\varepsilon_n} \right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p \rho(x) dx, \end{aligned}$$

where $\tilde{u}_n = u_n \circ T_n$.

Let us consider the liminf inequality. Assume that $(\mu_n, u_n) \rightarrow (\mu, u)$ in $TL^p(\Omega)$, i.e., $\tilde{u}_n \rightarrow u$ in $L^p(\Omega)$. It follows from (22) that

$$\mathcal{E}_{n,\varepsilon_n}(u_n) \geq \frac{1}{\varepsilon_n^p} \int_{\Omega} \operatorname{ess\,sup}_{y,z \in \Omega} \eta \left(\frac{x-z}{\hat{\varepsilon}_n} \right) \eta \left(\frac{x-y}{\hat{\varepsilon}_n} \right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p \rho(x) dx.$$

Then by Lemma 3.4,

$$\liminf_{n \rightarrow \infty} \mathcal{E}_{n,\varepsilon_n}(u_n) \geq \liminf_{n \rightarrow \infty} \left(\frac{\hat{\varepsilon}_n}{\varepsilon_n} \right)^p \mathcal{E}_{\hat{\varepsilon}_n}(\tilde{u}_n) \geq \mathcal{E}(u).$$

For the limsup inequality, we assume that $u \in W^{1,p}(\Omega)$ is Lipschitz. Let u_n be the restriction of u to the first n data points. It follows from (22) that

$$\begin{aligned} \mathcal{E}_{n,\varepsilon_n}(u_n) &\leq \frac{1}{\varepsilon_n^p} \int_{\Omega} \operatorname{ess\,sup}_{y,z \in \Omega} \eta \left(\frac{x-z}{\tilde{\varepsilon}_n} \right) \eta \left(\frac{x-y}{\tilde{\varepsilon}_n} \right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p \rho(x) dx \\ &= \frac{1}{\varepsilon_n^p} \int_{\Omega} \operatorname{ess\,sup}_{y,z \in \Omega} \eta \left(\frac{x-z}{\tilde{\varepsilon}_n} \right) \eta \left(\frac{x-y}{\tilde{\varepsilon}_n} \right) |u(z) - u(y)|^p \rho(x) dx + I, \end{aligned}$$

where, according to (15),

$$\begin{aligned} |I| &= \frac{1}{\varepsilon_n^p} \left| \int_{\Omega} \operatorname{ess\,sup}_{y,z \in \Omega} \eta \left(\frac{x-z}{\tilde{\varepsilon}_n} \right) \eta \left(\frac{x-y}{\tilde{\varepsilon}_n} \right) (|\tilde{u}_n(z) - \tilde{u}_n(y)|^p - |u(z) - u(y)|^p) \rho(x) dx \right| \\ &\leq \frac{C}{\varepsilon_n^p} \int_{\Omega} \operatorname{ess\,sup}_{z \in \Omega} \eta \left(\frac{x-z}{\tilde{\varepsilon}_n} \right) |\tilde{u}_n(z) - u(z)|^p dx \\ &\leq \frac{C}{\varepsilon_n^p} \int_{\Omega} \operatorname{ess\,sup}_{z \in \Omega} \eta \left(\frac{x-z}{\tilde{\varepsilon}_n} \right) |T_n(z) - z|^p dx \\ &\leq \frac{C}{\varepsilon_n^p} \|T_n - Id\|_{L^\infty(\Omega)}^p \end{aligned}$$

goes to zero as $n \rightarrow \infty$. The proof is completed by Lemma 3.5. Namely,

$$\limsup_{n \rightarrow \infty} \mathcal{E}_{n,\varepsilon_n}(u_n) \leq \limsup_{n \rightarrow \infty} \left(\frac{\tilde{\varepsilon}_n}{\varepsilon_n} \right)^p \mathcal{E}_{\tilde{\varepsilon}_n}(u) \leq \mathcal{E}(u).$$

□

We now go back to consider the constrained functionals. The following Lemma plays a key role.

Lemma 3.8. *Let $\{u_n\}_{n \in \mathbb{N}}$ be a sequence of functions on Ω_n such that*

$$\sup_{n \in \mathbb{N}} \mathcal{E}_{n,\varepsilon_n}(u_n) < \infty.$$

If $p > d$ and $(\mu_n, u_n) \rightarrow (\mu, u)$ in $TL^p(\Omega)$, then

$$\lim_{n \rightarrow \infty} \max_{x_k \in \Omega_n \cap \Omega'} |u_n(x_k) - u(x_k)| = 0,$$

for all Ω' compactly supported in Ω .

Proof. For any fixed $x_k \in \Omega_n$, we have

$$\begin{aligned} \frac{1}{n\varepsilon_n^p} \sum_{j=1}^n \eta\left(\frac{x_k - x_j}{\varepsilon_n}\right) |u_n(x_k) - u_n(x_j)|^p &\leq \frac{1}{n\varepsilon_n^p} \sum_{j=1}^n \max_{x_k \in B(x_j, \varepsilon_n)} |u_n(x_k) - u_n(x_j)|^p \\ &\leq \mathcal{E}_{n, \varepsilon_n}(u_n) < \infty, \end{aligned}$$

which can be rewritten as

$$\frac{1}{\varepsilon_n^p} \int_{\Omega} \eta\left(\frac{x_k - T_n(y)}{\varepsilon_n}\right) |\tilde{u}_n(y) - u_n(x_k)|^p \rho(y) dy < \infty.$$

Let J be a standard mollifier and n be large such that

$$J_{\hat{\varepsilon}_n}(x_k - y) = \frac{1}{\hat{\varepsilon}_n^d} J\left(\frac{x_k - y}{\hat{\varepsilon}_n}\right) \leq \frac{C}{\hat{\varepsilon}_n^d} \eta\left(\frac{x_k - y}{\hat{\varepsilon}_n}\right) \leq \frac{C}{\hat{\varepsilon}_n^d} \eta\left(\frac{x_k - T_n(y)}{\varepsilon_n}\right), \quad \forall y \in \Omega.$$

We have

$$\frac{\hat{\varepsilon}_n^d}{\varepsilon_n^p} \int_{\Omega} J_{\hat{\varepsilon}_n}(x_k - y) |\tilde{u}_n(y) - u_n(x_k)|^p dy < \infty.$$

Consequently, by $p > d$ and Jensen's inequality,

$$\int_{\Omega} J_{\hat{\varepsilon}_n}(x_k - y) |\tilde{u}_n(y) - u_n(x_k)| dy \leq \left(\int_{\Omega} J_{\hat{\varepsilon}_n}(x_k - y) |\tilde{u}_n(y) - u_n(x_k)|^p dy \right)^{1/p} \rightarrow 0, \quad (23)$$

as $n \rightarrow \infty$.

The uniform boundedness of the graph p -Laplacian regularization (i.e., (18)) and the proof of Lemma 4.5 of [16] imply that $u \in W^{1,p}(\Omega') \hookrightarrow C^{1-\frac{n}{p}}(\overline{\Omega'})$ and

$$J_{\hat{\varepsilon}_n} * \tilde{u}_n \rightarrow u, \quad \text{uniformly on } \Omega', \quad (24)$$

as $n \rightarrow \infty$. Then the proof is completed by (23)–(24), i.e., for any $x_k \in \Omega_n \cap \Omega'$,

$$\begin{aligned} |u_n(x_k) - u(x_k)| &= \left| u_n(x_k) - \int_{\Omega} J_{\hat{\varepsilon}_n}(x_k - y) \tilde{u}_n(y) dy \right. \\ &\quad \left. + \int_{\Omega} J_{\hat{\varepsilon}_n}(x_k - y) \tilde{u}_n(y) dy - u(x_k) \right| \\ &\leq \int_{\Omega} J_{\hat{\varepsilon}_n}(x_k - y) |\tilde{u}_n(y) - u_n(x_k)| dy + \left| \int_{\Omega} J_{\hat{\varepsilon}_n}(x_k - y) \tilde{u}_n(y) dy - u(x_k) \right| \end{aligned}$$

converges to zero as $n \rightarrow \infty$. \square

The Γ -convergence of the constrained functional $\mathcal{E}_{n, \varepsilon_n}^{con}$ follows from Theorem 3.7 and Lemma 3.8 directly.

Theorem 3.9. *Let $p > d$ and $\delta_n \ll \varepsilon_n \ll 1$. Then with probability one,*

$$\mathcal{E}_{n, \varepsilon_n}^{con} \xrightarrow{\Gamma} \mathcal{E}^{con},$$

in the TL^p metric on the set $\{(\nu, g) : \nu \in \mathcal{P}(\Omega), \|g\|_{L^\infty(\nu)} \leq M\}$, where $M = \max_{1 \leq i \leq N} |y_i|$.

Proof. If $(\mu_n, u_n) \rightarrow (\mu, u)$ in $TL^p(\Omega)$ with $\|u_n\|_{L^\infty(\mu_n)} \leq M$ and $\|u\|_{L^\infty(\mu)} \leq M$, then

$$\liminf_{n \rightarrow \infty} \mathcal{E}_{n, \varepsilon_n}^{con}(u_n) \geq \liminf_{n \rightarrow \infty} \mathcal{E}_{n, \varepsilon_n}(u_n) \geq \mathcal{E}(u) = \mathcal{E}^{con}(u),$$

where the last equality follows from Lemma 3.8. This proves the liminf inequality.

Now let us consider the limsup inequality. Let $u \in W^{1,p}(\Omega)$ with $\|u\|_{L^\infty(\mu)} \leq M$. We further assume that $u(x_i) = y_i$ for $x_i \in \mathcal{O}$. Otherwise, there is nothing to prove. Define $u_n(x_i) = u(x_i)$ for $x_i \in \Omega_n$. Then

$$\limsup_{n \rightarrow \infty} \mathcal{E}_{n,\varepsilon_n}^{con}(u_n) = \limsup_{n \rightarrow \infty} \mathcal{E}_{n,\varepsilon_n}(u_n) \leq \mathcal{E}(u) = \mathcal{E}^{con}(u).$$

This completes the proof. \square

4. CONTINUUM LIMIT OF THE p -LAPLACIAN REGULARIZATION ON THE k_n -NN HYPERGRAPH

In this section, we generalize the previous results to the hypergraph p -Laplacian regularization on the k_n -NN hypergraph. Let us begin with the Γ -convergence of the unconstrained functional \mathcal{F}_{n,k_n} , which is defined in (8).

Theorem 4.1. *Let $1 < p < \infty$ and $\delta_n \ll \bar{\varepsilon}_n \ll 1$, where $\bar{\varepsilon}_n$ is defined in (9). Then with probability one,*

$$\mathcal{F}_{n,k_n} \xrightarrow{\Gamma} \mathcal{E}(\cdot; \rho^{1-p/d}),$$

in $TL^p(\Omega)$ as $n \rightarrow \infty$.

Proof. 1. Let $\varepsilon_n(x)$ be a function of $x \in \Omega$ such that

$$\mu(B(x, \varepsilon_n(x))) = \frac{k_n}{n}. \quad (25)$$

We first claim that (21) holds. Since Ω is bounded and has Lipschitz boundary, it satisfies the cone condition [17]. More precisely, there exists a finite cone V such that each $x \in \Omega$ is the vertex of a finite cone V_x contained in Ω and congruent to V . If n is large, it follows from (25) and (9) that

$$\alpha_d(\bar{\varepsilon}_n)^d = \mu(B(x, \varepsilon_n(x))) \geq |V_x \cap B(x, \varepsilon_n(x))| \geq C|B(x, \varepsilon_n(x))| = C\alpha_d(\varepsilon_n(x))^d,$$

for any $x \in \Omega$, where C depends only on V . Consequently,

$$\bar{\varepsilon}_n \leq \varepsilon_n(x) \leq C^{-\frac{1}{d}}\bar{\varepsilon}_n, \quad x \in \Omega,$$

which proves (21) and

$$1 \gg \varepsilon_n(x) \geq \bar{\varepsilon}_n \gg \delta_n.$$

Define

$$\hat{\varepsilon}_n(x) := \varepsilon_n(x) - 4\|T_n - Id\|_{L^\infty(\Omega)}, \quad \tilde{\varepsilon}_n(x) := \varepsilon_n(x) + 4\|T_n - Id\|_{L^\infty(\Omega)}.$$

The previous result and (15) yield

$$\lim_{n \rightarrow \infty} \frac{\hat{\varepsilon}_n(x)}{\varepsilon_n(x)} = \lim_{n \rightarrow \infty} \frac{\tilde{\varepsilon}_n(x)}{\varepsilon_n(x)} = 1, \quad (26)$$

uniformly for any $x \in \Omega$. If ρ is Lipschitz,

$$\left| \rho(x)(\varepsilon_n(x))^d - \bar{\varepsilon}_n^d \right| = \left| \rho(x)(\varepsilon_n(x))^d - \frac{1}{\alpha_d} \mu(B(x, \varepsilon_n(x))) \right| \leq C(\varepsilon_n(x))^{d+1}.$$

We further have

$$\lim_{n \rightarrow \infty} \frac{(\rho(x))^{1/d} \varepsilon_n(x)}{\bar{\varepsilon}_n} = 1, \quad (27)$$

uniformly for any $x \in \Omega$.

Let

$$\chi_{k_n}(x_i, x_j) = \begin{cases} 1, & \text{if } x_i \overset{k_n}{\sim} x_j, \\ 0, & \text{otherwise,} \end{cases} \quad \forall x_i, x_j \in \Omega_n.$$

We claim that

$$\eta\left(\frac{x-y}{\hat{\varepsilon}_n(x)}\right) \leq \chi_{k_n}(T_n(y), T_n(x)) \leq \eta\left(\frac{x-y}{\tilde{\varepsilon}_n(x)}\right), \quad \forall x, y \in \Omega. \quad (28)$$

In fact, for any $x, y \in \Omega$,

$$T_n(y) \overset{k_n}{\sim} T_n(x),$$

yields

$$\begin{aligned} \frac{k_n}{n} &\geq \mu_n(B(T_n(x), |T_n(x) - T_n(y)|)) \\ &\geq \mu(B(T_n(x), |T_n(x) - T_n(y)| - \|T_n - Id\|_{L^\infty(\Omega)})) \\ &\geq \mu(B(x, |T_n(x) - T_n(y)| - 2\|T_n - Id\|_{L^\infty(\Omega)})), \end{aligned}$$

where we utilize $T_n(B(x, r - \|T_n - Id\|)) \subset B(x, r)$ and

$$\mu_n(B(x, r)) \geq \mu_n(T_n(B(x, r - \|T_n - Id\|))) = \mu(B(x, r - \|T_n - Id\|))$$

for the second inequality. This and (25) imply that

$$|T_n(x) - T_n(y)| - 2\|T_n - Id\|_{L^\infty(\Omega)} \leq \varepsilon_n(x).$$

Consequently,

$|x-y| \leq |x-T_n(x)| + |y-T_n(y)| + |T_n(x)-T_n(y)| \leq \varepsilon_n(x) + 4\|T_n - Id\|_{L^\infty(\Omega)} = \tilde{\varepsilon}_n(x)$, from which we obtain the second inequality in (28). The first inequality can be verified similarly.

2. Let us now consider the liminf inequality. Assume that $(\mu_n, u_n) \rightarrow (\mu, u)$ in $TL^p(\Omega)$. Estimate (28) yields

$$\begin{aligned} &\mathcal{F}_{n, k_n}(u_n) \\ &= \frac{1}{n\bar{\varepsilon}_n^p} \sum_{k=1}^n \max_{x_i, x_j \overset{k}{\sim} x_k} |u_n(x_i) - u_n(x_j)|^p \\ &= \frac{1}{\bar{\varepsilon}_n^p} \int_{\Omega} \operatorname{ess\,sup}_{y, z \in \Omega} \chi_{k_n}(T_n(z), T_n(x)) \chi_{k_n}(T_n(y), T_n(x)) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p \rho(x) dx \\ &\geq \frac{1}{\bar{\varepsilon}_n^p} \int_{\Omega} \operatorname{ess\,sup}_{y, z \in \Omega} \eta\left(\frac{x-z}{\hat{\varepsilon}_n(x)}\right) \eta\left(\frac{x-y}{\hat{\varepsilon}_n(x)}\right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p \rho(x) dx \\ &= \int_{\Omega} \left| \frac{(\rho(x))^{1/d} \hat{\varepsilon}_n(x)}{\bar{\varepsilon}_n} \right|^p \\ &\quad \frac{1}{|\hat{\varepsilon}_n(x)|^p} \operatorname{ess\,sup}_{y, z \in \Omega} \eta\left(\frac{x-z}{\hat{\varepsilon}_n(x)}\right) \eta\left(\frac{x-y}{\hat{\varepsilon}_n(x)}\right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p (\rho(x))^{1-p/d} dx \\ &\geq \inf_{x \in \Omega} \left| \frac{(\rho(x))^{1/d} \hat{\varepsilon}_n(x)}{\bar{\varepsilon}_n} \right|^p \\ &\quad \int_{\Omega} \frac{1}{|\hat{\varepsilon}_n(x)|^p} \operatorname{ess\,sup}_{y, z \in \Omega} \eta\left(\frac{x-z}{\hat{\varepsilon}_n(x)}\right) \eta\left(\frac{x-y}{\hat{\varepsilon}_n(x)}\right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p (\rho(x))^{1-p/d} dx. \end{aligned} \quad (29)$$

It follows from (26)–(27) and Lemma 3.4 that

$$\begin{aligned} \liminf_{n \rightarrow \infty} \mathcal{F}_{n, k_n}(u_n) &\geq \liminf_{n \rightarrow \infty} \inf_{x \in \Omega} \left| \frac{(\rho(x))^{1/d} \hat{\varepsilon}_n(x)}{\bar{\varepsilon}_n} \right|^p \mathcal{E}_{\bar{\varepsilon}_n}(\tilde{u}_n; \rho^{1-p/d}) \\ &\geq \mathcal{E}(u; \rho^{1-p/d}). \end{aligned}$$

3. For the limsup inequality, we assume that $u \in W^{1,p}(\Omega)$ is Lipschitz. Let u_n be the restriction of u to the first n data points. Then (28) implies that

$$\begin{aligned} &\mathcal{F}_{n, k_n}(u_n) \\ &\leq \frac{1}{|\bar{\varepsilon}_n|^p} \int_{\Omega} \operatorname{ess\,sup}_{y, z \in \Omega} \eta \left(\frac{x-z}{\bar{\varepsilon}_n(x)} \right) \eta \left(\frac{x-y}{\bar{\varepsilon}_n(x)} \right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p \rho(x) dx \\ &\leq \sup_{x \in \Omega} \left| \frac{(\rho(x))^{1/d} \tilde{\varepsilon}_n(x)}{\bar{\varepsilon}_n} \right|^p \\ &\quad \int_{\Omega} \frac{1}{|\bar{\varepsilon}_n(x)|^p} \operatorname{ess\,sup}_{y, z \in \Omega} \eta \left(\frac{x-z}{\bar{\varepsilon}_n(x)} \right) \eta \left(\frac{x-y}{\bar{\varepsilon}_n(x)} \right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p (\rho(x))^{1-p/d} dx \\ &\leq \sup_{x \in \Omega} \left| \frac{(\rho(x))^{1/d} \tilde{\varepsilon}_n(x)}{\bar{\varepsilon}_n} \right|^p \\ &\quad \int_{\Omega} \frac{1}{|\bar{\varepsilon}_n(x)|^p} \operatorname{ess\,sup}_{y, z \in \Omega} \eta \left(\frac{x-z}{\bar{\varepsilon}_n(x)} \right) \eta \left(\frac{x-y}{\bar{\varepsilon}_n(x)} \right) |u(z) - u(y)|^p (\rho(x))^{1-p/d} dx + I. \end{aligned}$$

It is not difficult to see that

$$|I| \leq \frac{C}{\sup_{x \in \Omega} |\varepsilon_n(x)|^p} \|T_n - Id\|_{L^\infty(\Omega)}^p \rightarrow 0$$

as $n \rightarrow \infty$. Consequently, by Lemma 3.5,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathcal{F}_{n, k_n}(u_n) &\leq \limsup_{n \rightarrow \infty} \sup_{x \in \Omega} \left| \frac{(\rho(x))^{1/d} \tilde{\varepsilon}_n(x)}{\bar{\varepsilon}_n} \right|^p \mathcal{E}_{\bar{\varepsilon}_n}(u; \rho^{1-p/d}) \\ &\leq \mathcal{E}(u; \rho^{1-p/d}). \end{aligned}$$

4. To complete the proof for any continuous ρ , we approximate it by Lipschitz functions. In fact, for the liminf inequality, there exists a sequence of Lipschitz functions $\{\rho_k(x)\}$ such that $\rho_k(x) \nearrow \rho(x)$ as $k \rightarrow \infty$ for any $x \in \Omega$. Then

$$\begin{aligned} \liminf_{n \rightarrow \infty} \mathcal{F}_{n, k_n}(u_n; \rho) &\geq \lim_{k \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{F}_{n, k_n}(u_n; \rho_k) = \lim_{k \rightarrow \infty} \mathcal{E}(u; \rho_k^{1-p/d}) \\ &= \mathcal{E}(u; \rho^{1-p/d}), \end{aligned}$$

where the penultimate equality comes from step 2 and the last equality comes from the monotone convergence theorem.

While for the liminf inequality, there exists a sequence of Lipschitz functions $\{\rho_k(x)\}$ such that $\rho_k(x) \searrow \rho(x)$ as $k \rightarrow \infty$ for any $x \in \Omega$. Then

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathcal{F}_{n, k_n}(u_n; \rho) &\leq \lim_{k \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathcal{F}_{n, k_n}(u_n; \rho_k) = \lim_{k \rightarrow \infty} \mathcal{E}(u; \rho_k^{1-p/d}) \\ &= \mathcal{E}(u; \rho^{1-p/d}), \end{aligned}$$

where the penultimate equality comes from step 3 and the last equality comes from Lebesgue's dominated convergence theorem. \square

The Γ -convergence of the constrained functionals follows from Theorem 4.1 directly.

Theorem 4.2. *Let $p > d$ and $\delta_n \ll \bar{\varepsilon}_n \ll 1$. Then with probability one,*

$$\mathcal{F}_{n,k_n}^{con} \xrightarrow{\Gamma} \mathcal{E}^{con}(\cdot; \rho^{1-p/d}),$$

in the TL^p metric on the set $\{(\nu, g) : \nu \in \mathcal{P}(\Omega), \|g\|_{L^\infty(\nu)} \leq M\}$, where $M = \max_{1 \leq i \leq N} |y_i|$.

Proof. Assume that

$$\mathcal{E}_{n,\varepsilon_n}(u_n) \leq C\mathcal{F}_{n,k_n}(u_n), \quad (30)$$

for a constant ε_n that satisfies $\delta_n \ll \varepsilon_n \ll 1$. Then the proof of the theorem is exactly the same as the proof of Theorem 3.9, except that we use (30) to ensure that the conclusion of Lemma 3.8 holds.

To prove (30), we utilize (29). More precisely,

$$\begin{aligned} & \mathcal{F}_{n,k_n}(u_n) \\ & \geq \inf_{x \in \Omega} \left| \frac{(\rho(x))^{1/d} \hat{\varepsilon}_n(x)}{\bar{\varepsilon}_n} \right|^p \\ & \quad \int_{\Omega} \frac{1}{|\hat{\varepsilon}_n(x)|^p} \operatorname{ess\,sup}_{y,z \in \Omega} \eta \left(\frac{x-z}{\hat{\varepsilon}_n(x)} \right) \eta \left(\frac{x-y}{\hat{\varepsilon}_n(x)} \right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p (\rho(x))^{1-p/d} dx \\ & \geq \inf_{x \in \Omega} \left| \frac{(\rho(x))^{1/d} \hat{\varepsilon}_n(x)}{\bar{\varepsilon}_n} \right|^p \inf_{x \in \Omega} (\rho(x))^{-p/d} \\ & \quad \frac{\varepsilon_n^p}{\sup_{x \in \Omega} |\hat{\varepsilon}_n(x)|^p} \int_{\Omega} \frac{1}{\varepsilon_n^p} \operatorname{ess\,sup}_{y,z \in \Omega} \eta \left(\frac{x-z}{\varepsilon_n} \right) \eta \left(\frac{x-y}{\varepsilon_n} \right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p \rho(x) dx \\ & \geq \frac{C}{\varepsilon_n^p} \int_{\Omega} \operatorname{ess\,sup}_{y,z \in \Omega} \eta \left(\frac{T_n(x) - T_n(z)}{\varepsilon_n} \right) \eta \left(\frac{T_n(x) - T_n(y)}{\varepsilon_n} \right) |\tilde{u}_n(z) - \tilde{u}_n(y)|^p \rho(x) dx \\ & = C\mathcal{E}_{n,\varepsilon_n}(u_n), \end{aligned}$$

where $\tilde{\varepsilon}_n = \inf_{x \in \Omega} \hat{\varepsilon}_n(x)$ and $\varepsilon_n = \tilde{\varepsilon}_n - 2\|T_n - Id\|_{L^\infty(\Omega)}$. This completes the proof. \square

The main result of this section is a corollary of Theorem 4.2. The proof is identical to that of Theorem 3.1.

Theorem 4.3. *Let $p > d$ and $\delta_n \ll \bar{\varepsilon}_n \ll 1$. If u_n is a minimizer of $\mathcal{F}_{n,k_n}^{con}$, then almost surely,*

$$(\mu_n, u_n) \rightarrow (\mu, u), \quad \text{in } TL^p(\Omega),$$

as $n \rightarrow \infty$ for a function $u \in L^p(\Omega)$. Furthermore,

$$\lim_{n \rightarrow \infty} \max_{x_k \in \Omega_n \cap \Omega'} |u_n(x_k) - u(x_k)| = 0,$$

for any Ω' compactly contained in Ω and u is a minimizer of $\mathcal{E}^{con}(\cdot; \rho^{1-p/d})$.

We conclude this section with two remarks.

Remark 4.4. In this paper, we consider the hypergraph p -Laplacian regularization with $p > 1$ in a semisupervised setting. The results established in Section 3 and Section 4 can be generalized to $p = 1$ in an unsupervised setting. The continuum limits of $\mathcal{E}_{n,\varepsilon_n}$ and \mathcal{F}_{n,k_n} with $p = 1$ become weighted total variations.

Remark 4.5. Since the given data is a point cloud, we can compute the distance $|x_i - x_j|$ between any two vertices x_i and x_j and add this information to the objective function for a better performance. A straightforward way is to construct a weight $w_{i,j}$ as a nonincreasing function of the distance $|x_i - x_j|$ and replace $|u(x_i) - u(x_j)|^p$ in both $\mathcal{E}_{n,\varepsilon_n}$ and \mathcal{F}_{n,k_n} by $w_{i,j}|u(x_i) - u(x_j)|^p$. All results in Section 3 and Section 4 can be generalized to the new energies without any essential modification under the assumption that

$$w_{i,j} = \eta(|x_i - x_j|),$$

and $\eta : [0, \infty) \rightarrow [0, \infty)$ is positive and continuous at $x = 0$ and has compact support.

5. THE NUMERICAL ALGORITHM

In this section, we present an algorithm for solving the optimization problem

$$\min_{u \in \mathbb{R}^n} \frac{1}{p} \sum_{k=1}^n \max_{x_i, x_j \in e_k} w_{i,j} |u(x_i) - u(x_j)|^p \quad \text{s.t. } u(x_i) = y_i \text{ for } x_i \in \mathcal{O}, \quad (31)$$

where $e_k = \{x_j \in \Omega_n : |x_k - x_j| \leq \varepsilon_n\}$ for the ε_n -ball hypergraph or $e_k = \{x_j \in \Omega_n : x_j \stackrel{k_n}{\sim} x_k\}$ for the k_n -NN hypergraph.

Let

$$\mathbb{I}_{\mathcal{O}}(u) = \begin{cases} 0, & \text{if } u(x_i) = y_i \text{ for any } x_i \in \mathcal{O}, \\ \infty, & \text{otherwise,} \end{cases}$$

be an indicator function. Then problem (31) can be rewritten as an unconstrained problem

$$\min_{u \in \mathbb{R}^n} \sum_{k=1}^n g(A_k u) + \mathbb{I}_{\mathcal{O}}(u), \quad (32)$$

where

$$g(\beta) = \frac{1}{p} (\max |\beta|)^p, \quad \beta \in \mathbb{R}^m,$$

$m = \max \left\{ \frac{|e_k|(|e_k|-1)}{2} : k = 1, \dots, n \right\}$, and $A_k \in \mathbb{R}^{m \times n}$ such that

$$g(A_k u) = \frac{1}{p} \max_{x_i, x_j \in e_k} \left| w_{i,j}^{\frac{1}{p}} (u(x_i) - u(x_j)) \right|^p.$$

Clearly, both g and $\mathbb{I}_{\mathcal{O}}$ are proper, lower semicontinuous, and convex functions. A popular algorithm for solving (32) is the stochastic primal-dual hybrid gradient (SPDHG) algorithm [19] that rewrites the original optimal problem as a saddle point problem and updates the primal variable and the dual variable separately.

By introducing dual variables $\alpha_i \in \mathbb{R}^m$, $i = 1, \dots, n$, problem (32) is equivalent to the saddle point problem

$$\min_{u \in \mathbb{R}^n} \max_{\alpha_i \in \mathbb{R}^m, i=1, \dots, n} \sum_{i=1}^n \langle A_i u, \alpha_i \rangle - g^*(\alpha_i) + \mathbb{I}_{\mathcal{O}}(u), \quad (33)$$

where

$$g^*(\alpha_i) = \max_{\beta \in \mathbb{R}^m} \langle \beta, \alpha_i \rangle - g(\beta), \quad i = 1, 2, \dots, n,$$

denotes the Fenchel conjugate of g and we utilize the property $g^{**} = g$. Now SPDHG with serial sampling is summarized in Algorithm 1. Under the condition

$$\sigma\tau < \max_{i=1, \dots, n} \frac{p_i}{\|A_i\|^2},$$

it converges almost surely to a solution of problem (32) [46].

Algorithm 1 SPDHG for solving problem (33).

Require: Primal variable $u^{(0)}$, dual variables $\alpha_i^{(0)}$, $i = 1, 2, \dots, n$; step sizes τ, σ ; probabilities p_i , $i = 1, \dots, n$.

Initialization: $\bar{\alpha}_i^{(0)} = \alpha_i^{(0)}$, $i = 1, \dots, n$; $k = 0$.

while the stopping criterion is not satisfied **do**

 Update the primal variable:

$$\begin{aligned} u^{(k+1)} &= \text{prox}_{\tau\mathbb{I}_{\mathcal{O}}} \left(u^{(k)} - \tau \sum_{i=1}^n A_i^T \bar{\alpha}_i^{(k)} \right) \\ &:= \arg \min_{u \in \mathbb{R}^n} \tau \mathbb{I}_{\mathcal{O}}(u) + \frac{1}{2} \left\| u - \left(u^{(k)} - \tau \sum_{i=1}^n A_i^T \bar{\alpha}_i^{(k)} \right) \right\|_2^2. \end{aligned} \quad (34)$$

 Update dual variables: Randomly pick $i \in \{1, 2, \dots, n\}$ with probability p_i ,

$$\alpha_j^{(k+1)} = \begin{cases} \text{prox}_{\sigma g^*} \left(\alpha_i^{(k)} + \sigma A_i u^{(k+1)} \right), & \text{if } j = i, \\ \alpha_j^{(k)}, & \text{if } j \neq i. \end{cases} \quad (35)$$

 Extrapolation on dual variables:

$$\bar{\alpha}_j^{(k+1)} = \begin{cases} \alpha_i^{(k+1)} + \frac{1}{p_i} \left(\alpha_i^{(k+1)} - \alpha_i^{(k)} \right), & \text{if } j = i, \\ \alpha_j^{(k)}, & \text{if } j \neq i. \end{cases}$$

end while

return $u^{(k+1)}$.

Subproblem (34) has a closed form solution

$$u^{(k+1)}(x_j) = \begin{cases} y_j, & \text{if } x_j \in \mathcal{O}, \\ \left(u^{(k)} - \tau \sum_{i=1}^n A_i^T \bar{\alpha}_i^{(k)} \right)(x_j), & \text{otherwise.} \end{cases}$$

The rest is to find the Fenchel conjugate g^* and solve subproblem (35).

Lemma 5.1. *Let*

$$h(t) = \frac{1}{p} |t|^p, \quad t \in \mathbb{R}.$$

Then

$$h^*(s) = \begin{cases} 0, & \text{if } |s| \leq 1, \\ \infty, & \text{otherwise,} \end{cases} \quad \text{if } p = 1, \quad (36)$$

or

$$h^*(s) = \frac{1}{p'} |s|^{p'}, \quad \text{if } p > 1, \quad (37)$$

where $p' = \frac{p}{p-1}$. Moreover,

$$g^*(\alpha) = h^*(\|\alpha\|_1), \quad \alpha \in \mathbb{R}^m. \quad (38)$$

Proof. By the definition of the Fenchel conjugate, (36) and (37) can be verified directly. We prove (38) in the following.

For any $\alpha, \beta \in \mathbb{R}^m$,

$$\langle \alpha, \beta \rangle - g(\beta) \leq \|\alpha\|_1 \max |\beta| - \frac{1}{p} (\max |\beta|)^p \leq h^*(\|\alpha\|_1),$$

from which we obtain that $g^*(\alpha) \leq h^*(\|\alpha\|_1)$. On the other hand, for any $\alpha \in \mathbb{R}^m$, let $\gamma \in \mathbb{R}^m$ such that $\max |\gamma| = 1$ and $\langle \alpha, \gamma \rangle = \|\alpha\|_1$. Then

$$g^*(\alpha) \geq \max_{\substack{\beta=t\gamma \\ t \in \mathbb{R}}} \langle \alpha, \beta \rangle - g(\beta) = \max_{t \in \mathbb{R}} \|\alpha\|_1 t - \frac{1}{p} |t|^p = h^*(\|\alpha\|_1).$$

This completes the proof. \square

Now if $p = 1$, subproblem (35) with $j = i$ becomes a projection onto the L^1 ball of the following form

$$\alpha_i^{(k+1)} = \arg \min_{\substack{\alpha \in \mathbb{R}^m \\ \|\alpha\|_1 \leq 1}} \left\| \alpha - \left(\alpha_i^{(k)} + \sigma A_i u^{(k+1)} \right) \right\|_2^2,$$

which can be solved efficiently [47]. In the case of $p > 1$, the solution is given as follows.

Lemma 5.2. *Let $p > 1$. The solution of subproblem (35) with $j = i$ satisfies*

$$\alpha_i^{(k+1)} = \text{sign}(\alpha_i^{(k)} + \sigma A_i u^{(k+1)}) \max \left\{ \left| \alpha_i^{(k)} + \sigma A_i u^{(k+1)} \right| - \sigma \|\alpha_i^{(k+1)}\|_1^{p'-1}, 0 \right\}. \quad (39)$$

All operations involving vectors in (39) are understood elementwise.

Proof. For simplicity of notation, let $\beta = \alpha_i^{(k)} + \sigma A_i u^{(k+1)}$ and $\alpha^* = \alpha_i^{(k+1)}$. Then subproblem (35) becomes

$$\alpha^* = \arg \min_{\alpha \in \mathbb{R}^m} \frac{1}{2} \|\alpha - \beta\|_2^2 + \frac{\sigma}{p'} (\|\alpha\|_1)^{p'}, \quad p > 1.$$

Since the right-hand side is non-separable, we introduce a new variable $t = \|\alpha\|_1$ and consider the Lagrangian

$$\frac{1}{2} \|\alpha - \beta\|_2^2 + \frac{\sigma}{p'} t^{p'} + \lambda (\|\alpha\|_1 - t).$$

The saddle point $(\alpha^*, \lambda^*, t^*)$ satisfies

$$\begin{cases} \alpha^* \in \arg \min_{\alpha \in \mathbb{R}^m} \frac{1}{2} \|\alpha - \beta\|_2^2 + \lambda^* \|\alpha\|_1, \\ \lambda^* = \sigma (t^*)^{p'-1}, \\ t^* = \|\alpha^*\|_1. \end{cases}$$

The L^1 regularized problem has a closed-form solution $\alpha^* = \text{sign}(\beta) \max\{|\beta| - \lambda^*, 0\}$. By substituting $\lambda^* = \sigma (t^*)^{p'-1} = \sigma \|\alpha^*\|_1^{p'-1}$ into it, we arrive at

$$\alpha^* = \text{sign}(\beta) \max\{|\beta| - \sigma \|\alpha^*\|_1^{p'-1}, 0\}, \quad (40)$$

which is exactly (39) and is the solution of subproblem (35). \square

The last step is to solve $\alpha_i^{(k+1)}$ from (39), which can be done in a few iterations. For the rest of this section, the subscript j of a vector denotes the j -th element of it. To explain the basic idea, we continue to use the notation of (40) for simplicity.

It follows from (40) that $\alpha^* = \text{sign}(\beta)|\alpha^*|$ and $|\alpha^*|$ depends only on $|\beta|$. W.l.o.g., we assume that $\beta \geq 0$ and β is ordered in a nondecreasing order. As a result, $\alpha^* \geq 0$ and α^* is also in a nondecreasing order.

Lemma 5.3. *Let $\beta \in \mathbb{R}^m$ be a nonnegative vector in a nondecreasing order and $\beta \neq 0$. If $\sigma > 0$ and $p > 1$, there exists a unique vector $\alpha \in \mathbb{R}^m$ such that*

$$\alpha = \beta - \sigma \left(\sum_{j=1}^m \alpha_j \right)^{p'-1}, \quad (41)$$

and

$$\lambda := \sigma \left(\sum_{j=1}^m \alpha_j \right)^{p'-1} > 0.$$

If $\alpha \geq 0$, then $\alpha^* = \alpha$, where α^* is the solution of equation (40). Otherwise, we have

$$\lambda \leq \lambda^* := \sigma \|\alpha^*\|_1^{p'-1}. \quad (42)$$

Proof. It follows from (41) that

$$s + m\sigma s^{p'-1} - \sum_{j=1}^m \beta_j = 0, \quad (43)$$

where $s = \sum_{j=1}^m \alpha_j$. Since $s + m\sigma s^{p'-1}$ is an increasing function on \mathbb{R}^+ and $\sum_{j=1}^m \beta_j > 0$, we obtain that equation (43) admits a unique solution $s > 0$. By substituting it into (41), we obtain a solution α with $\lambda > 0$.

If $\alpha \geq 0$, then α also satisfies equation (40) and $\alpha^* = \alpha$. Otherwise, there exists at least one zero element in α^* . Assume that

$$0 = \alpha_1^* = \dots = \alpha_k^* < \alpha_{k+1}^* \leq \dots \leq \alpha_m^*.$$

Then (40) yields

$$\beta_k \leq \lambda^* \leq \beta_{k+1}, \quad (44)$$

and

$$s^* + (m-k)\sigma s^{*p'-1} - \sum_{j=k+1}^m \beta_j = 0, \quad (45)$$

where $s^* = \|\alpha^*\|_1$. Now (43)–(45) imply that

$$\begin{aligned}
m\lambda &= m\sigma s^{p'-1} = \sum_{j=1}^m \beta_j - s = \sum_{j=k+1}^m \beta_j + \sum_{j=1}^k \beta_j - s \\
&= s^* + (m-k)\sigma s^{*p'-1} + \sum_{j=1}^k \beta_j - s \\
&= (m-k)\lambda^* + \sum_{j=1}^k \beta_j + s^* - s \\
&\leq m\lambda^* + \left(\frac{\lambda^*}{\sigma}\right)^{\frac{1}{p'-1}} - \left(\frac{\lambda}{\sigma}\right)^{\frac{1}{p'-1}},
\end{aligned}$$

or equivalently,

$$m\lambda + \left(\frac{\lambda}{\sigma}\right)^{p-1} \leq m\lambda^* + \left(\frac{\lambda^*}{\sigma}\right)^{p-1}. \quad (46)$$

From the monotonicity of $m\lambda + \left(\frac{\lambda}{\sigma}\right)^{p-1}$, we conclude that $\lambda \leq \lambda^*$. \square

Lemma 5.3 provides an algorithm for solving (39). The key idea is that (42) implies that $\alpha_j^* = 0$ for any $j \in \{j : \alpha_j < 0\}$. Consequently, we can use Lemma 5.3 again to solve the remaining elements of α^* . It should be noted that the proof of Lemma 5.3 does not rely on the monotonicity of β . The details are summarized in Algorithm 2.

Algorithm 2 Solving $\text{prox}_{\sigma g^*}(\beta)$ for subproblem (35) with $p > 1$.

Require: $\beta \in \mathbb{R}^m$; $\sigma > 0$, $p > 1$.

Initialization: $\beta^+ = |\beta|$, $\alpha = \beta^+$; stopping criterion $c = -1$.

while $c < 0$ **do**

Find positive elements: $\gamma = \{\beta_j^+ : j \in I\} \in \mathbb{R}^b$, where $I = \{j : \alpha_j > 0\}$.

Solve $s + b\sigma s^{p'-1} - \|\gamma\|_1 = 0$ on \mathbb{R}^+ .

Obtain the threshold: $\lambda = \sigma s^{p'-1}$.

Update the solution: $\alpha = \max\{\beta^+ - \lambda, 0\}$.

$c = \min \gamma - \lambda$.

end while

Fix the negative part: $\alpha_j = -\alpha_j$ for any $j \in \{j : \beta_j < 0\}$.

return α .

6. EXPERIMENTAL RESULTS

In this section, we present numerical experiments of the hypergraph p -Laplacian regularization (HpL) for data interpolation. The purpose is to verify the theoretical results established in previous sections and make comparisons with the graph p -Laplacian regularization (GpL). See (11) and (31) for the definition of GpL and HpL.

We start with experiments in 1D and then discuss performance of GpL and HpL on some higher-dimensional data interpolation problems, including semi-supervised learning and image inpainting.

6.1. Experiments in 1D. Let $d = 1$ and $n = 1280$. We select n random numbers under the standard uniform distribution on the interval $(0, 1)$ and let Ω_n be the set of random numbers. The set of labeled points \mathcal{O} consists of 6 points in Ω_n , which are marked by red circles in Figure 1–3. We shall discuss the difference between GpL and HpL for the interpolation of signals defined on Ω_n with given labeled points \mathcal{O} . The weight function $w_{i,j} = 1$ is used for both GpL and HpL.

Let us first consider the case of the ε_n -ball graph/hypergraph and $p = 2$. According to [16], the minimizer of GpL with the constraint on \mathcal{O} is an approximation of the minimizer of the continuum p -Laplacian regularization with the constraint on \mathcal{O} (the minimizer is a piecewise linear and continuous function that passes through the labeled points) for moderate connection radius ε_n (i.e., $\delta_n \ll \varepsilon_n \ll (\frac{1}{n})^{1/p}$). This is verified by the results shown in the first row of Figure 1. We consider the case of $\varepsilon_n \geq 0.006$ since the graph becomes disconnected when $\varepsilon_n \leq 0.005$. When $\varepsilon_n = 0.006$, GpL gives an approximation of the continuous solution. With an increase in ε_n , the error of GpL increases. It develops spikes at the labeled points for large ε_n .

The second row of Figure 1, which shows the results of HpL with different ε_n , is quite different from the first row. As the connection radius ε_n increases, the error of HpL also increases. However, it alleviates the development of spikes even for large ε_n (see the 3rd and the 4th labeled points). This partially verifies the conclusion of Theorem 3.1, showing that the convergence of the hypergraph p -Laplacian regularization to the continuum p -Laplacian regularization in a semisupervised setting requires a weaker assumption on the upper bound of ε_n (i.e., $\delta_n \ll \varepsilon_n \ll 1$). Two spikes (the 2nd and the 5th labeled points) appear in the results because the assumption $\varepsilon_n \ll 1$ is no longer fulfilled.

A similar conclusion can be obtained for the k_n -NN graph/hypergraph. In Figure 2, we present the results of GpL and HpL with different k_n . Again we observe that GpL develops spikes at the labeled points for large k_n while HpL performs better. This coincides with the result we established in Theorem 4.3.

Theoretically, the minimizer of HpL gains more smoothness as the power p increases. Figure 3 shows the results of HpL with different p . It is observed that a higher p also prevents the occurrence of spikes. A similar conclusion has been obtained for GpL [34]. In the case of $p = 1$ where the assumption $p > d$ is not fulfilled, HpL gives noninformative solutions.

At the end of this subsection, we mention that the k_n -NN hypergraph is more desirable in applications. Let us go back to the first column of Figure 1 and Figure 2, where the smallest ε_n and k_n are chosen such that the hypergraph is connected. The error of HpL on the k_n -NN hypergraph is much smaller than the error on the ε_n -ball hypergraph. Besides, minimizing HpL on the k_n -NN hypergraph requires less computational cost. We shall utilize the k_n -NN hypergraph in the following experiments of SSL and image inpainting and let $p = 2$ for both GpL and HpL.

6.2. Semi-supervised learning. We test GpL and HpL on the well-known MNIST dataset of handwritten digits [48]. It contains 70000 gray-scale images, each of which is 28×28 in size, representing a digit between 0 – 9. Figure 4 shows some images in the MNIST dataset. Following [24], we regard each image as a point in $\mathbb{R}^{28 \times 28}$. The details of SSL with GpL/HpL for the MNIST dataset are summarized in Algorithm 3.




FIGURE 1. Results of GpL and HpL with $p = 2$ and different ε_n .
First row: GpL; second row: HpL.




FIGURE 2. Results of GpL and HpL with $p = 2$ and different k_n .
First row: GpL; second row: HpL.




FIGURE 3. Results of HpL with different p . First row: ε_n -ball hypergraph with $\varepsilon_n = 0.048$; second row: k_n -NN hypergraph with $k_n = 72$.



FIGURE 4. Some images from the MNIST dataset.

Algorithm 3 SSL with GpL/HpL for the MNIST dataset.

Require: The MNIST dataset Ω_n , the labeled set \mathcal{O} with labels $u(x_i) = y_i, x_i \in \mathcal{O}$, $y_i \in \{0, 1, \dots, 9\}$.

for $k = 0 : 9$ **do**

Find the constraint: For any $x_i \in \mathcal{O}$,

$$u_k(x_i) = \begin{cases} 1, & \text{if } y_i = k, \\ 0, & \text{otherwise.} \end{cases} \quad (47)$$

Solve u_k on Ω_n by GpL/HpL with constraint (47).

end for

Label points in $\Omega_n \setminus \mathcal{O}$: For any $x_i \in \Omega_n \setminus \mathcal{O}$,

$$l = \arg \max_{0 \leq k \leq 9} u_k(x_i), \quad u(x_i) = l.$$

return u .

TABLE 1. Mean accuracy (%) and standard deviation of GpL and HpL for the classification of the MNIST dataset with different labeling rates.

	0.05%	0.1%	0.2%	0.5%	1%	2%	5%
GpL	16.2±7.5	32.1±13.3	56.6±13.4	86.1±3.7	92.4±0.9	93.9±0.4	95.0±0.2
HpL	50.2±8.7	69.8±6.2	80.7±5.0	87.9±2.6	91.6±0.7	92.9±0.4	94.4±0.2

For both GpL and HpL, we choose $k_n = 21$ and the weight function

$$w_{i,j} = \exp\left(-\frac{\|x_i - x_j\|^2}{\sigma(x_i)^2}\right),$$

where $\sigma(x_i)$ is the distance between x_i and its the 21th nearest neighbor. This is calculated by the `vfeat` toolbox [49].

Table 1 presents the classification accuracy of GpL and HpL with 0.05% – 5% randomly chosen labeled points. We repeat the experiment 20 times and show the

Algorithm 4 Image inpainting with GpL/HpL.

Require: Partially observed pixels of an image f : $\{f_{i,j}, (i,j) \in X\}$; an initial guess $u^{(0)}$.

Initialization: $k = 0$.

for $k = 0 : K - 1$ **do**

Construct point cloud $\Omega_n(u^{(k)})$.

Construct the training set $\mathcal{O}(u^{(k)}) = \{P_{i,j} \in \Omega_n(u^{(k)}), (i,j) \in X\}$ and $u^{(k+1)}(P_{i,j}) = f_{i,j}$ for any $P_{i,j} \in \mathcal{O}(u^{(k)})$.

Solve $u^{(k+1)}$ on $\Omega_n(u^{(k)})$ by GpL/HpL with the given training set.

end for

return $u^{(K-1)}$.




FIGURE 5. Test images with size 256×256 for image inpainting.

mean accuracy and the standard deviation. When the labeling rate is extremely low ($\leq 0.1\%$), GpL fails to give meaningful results. HpL has a significantly better performance, although it is not satisfactorily accurate. This is mainly attributed to HpL's ability of better inhibiting spikes. However, as the labeling rate increases, the gap between the two algorithms becomes smaller. GpL even has a higher mean accuracy of 1% than HpL in the case of 2% labeling rate.

6.3. Image inpainting. In this subsection, we consider image inpainting from sparse data. We follow [31] to construct a point cloud from an image. More precisely, let $f \in \mathbb{R}^{N_1 \times N_2}$ be a gray-scale image. We define image patch $P_{i,j}(f)$ to be a rectangle centered at pixel (i,j) of size $s_1 \times s_2$. A mirror extension of f is needed for pixels that are near the boundary. Now the point cloud generated by f is defined as

$$\Omega_n(f) = \{P_{i,j}(f) : 1 \leq i \leq N_1, 1 \leq j \leq N_2\} \subset \mathbb{R}^{s_1 \times s_2}.$$

The image patch $P_{i,j}(f)$ and the intensity $f_{i,j}$ of f at pixel (i,j) can be connected via a function u defined on $\Omega_n(f)$, i.e.,

$$u(P_{i,j}(f)) = f_{i,j}.$$

Then if $\Omega_n(f)$ is known and $\mathcal{O}(f) \subset \Omega_n(f)$, image inpainting with given u on $\mathcal{O}(f)$ becomes a data interpolation problem in $\mathbb{R}^{s_1 \times s_2}$. In practice, the point cloud is unknown. For this reason, we start with an initial guess of f and update the point

TABLE 2. The PSNR and SSIM values of results of GpL and HpL for different test images and different sampling rates.

		PSNR				SSIM			
		5%	10%	15%	20%	5%	10%	15%	20%
Airplane	GpL	21.03	22.65	23.93	24.58	0.3246	0.4516	0.5499	0.6120
	HpL	21.33	23.09	24.55	25.24	0.3825	0.5125	0.6075	0.6618
Barbara	GpL	22.55	24.60	25.99	27.03	0.4436	0.5767	0.6571	0.7142
	HpL	23.15	25.51	27.00	27.95	0.5043	0.6354	0.7062	0.7553
Beans	GpL	22.66	25.12	26.35	28.26	0.4486	0.5435	0.6097	0.6620
	HpL	23.20	26.18	27.52	29.64	0.4911	0.5877	0.6464	0.6914
Boat	GpL	21.96	23.14	24.18	25.21	0.3199	0.4407	0.5408	0.6115
	HpL	22.19	23.51	24.60	25.81	0.3704	0.5000	0.5915	0.6591
Bridge	GpL	20.31	21.69	22.50	23.24	0.2790	0.4040	0.4874	0.5534
	HpL	20.41	21.92	22.74	23.53	0.3203	0.4496	0.5268	0.5906
C.man	GpL	20.88	22.01	23.43	24.02	0.2257	0.3434	0.4331	0.4943
	HpL	21.07	22.31	23.85	24.53	0.2621	0.3776	0.4688	0.5246
Clock	GpL	22.62	24.20	25.98	27.08	0.3061	0.4256	0.5285	0.5880
	HpL	22.92	24.72	26.65	27.70	0.3374	0.4667	0.5692	0.6227
Couple	GpL	21.55	23.35	24.32	25.38	0.2936	0.4653	0.5486	0.6351
	HpL	21.76	23.80	24.88	26.00	0.3532	0.5302	0.6074	0.6871
Hill	GpL	23.43	25.28	26.33	27.06	0.3234	0.4818	0.5671	0.6264
	HpL	23.81	25.98	26.95	27.76	0.3804	0.5425	0.6196	0.6730
House	GpL	24.55	26.99	28.76	30.19	0.2636	0.3895	0.4616	0.5276
	HpL	25.01	27.70	29.80	31.18	0.2981	0.4217	0.4902	0.5505
Man	GpL	22.32	24.07	25.01	25.93	0.3156	0.4506	0.5385	0.6034
	HpL	22.75	24.66	25.71	26.61	0.3772	0.5156	0.5971	0.6501
Peppers	GpL	22.33	24.46	25.70	27.05	0.4696	0.5956	0.6686	0.7263
	HpL	23.05	25.50	26.75	28.14	0.5357	0.6614	0.7242	0.7701
Average	GpL	22.18	23.96	25.21	26.25	0.3344	0.4640	0.5493	0.6129
	HpL	22.56	24.57	25.92	27.01	0.3844	0.5167	0.5962	0.6530

cloud with the restored image to reduce the error. The details of image inpainting with GpL and HpL are summarized in Algorithm 4.

To calculate the weight function for GpL and HpL, we define the semilocal patch

$$\bar{P}_{i,j}(f) = [P_{i,j}(f), \lambda \bar{x}],$$

with

$$\bar{x} = \left(\frac{i}{N_1}, \frac{j}{N_2} \right),$$

that includes the local coordinate. This restricts the search of the k_n nearest neighbors to a local area and accelerates the algorithm. Then the weight function is calculated as in the case of Section 6.2. We utilize $s_1 = s_2 = 11$ and $\lambda = 10$ for the experiments. The iteration number K in Algorithm 4 is chosen to be 15 for GpL. While for HpL, we let the result of GpL be the initial guess $u^{(0)}$ and let $K = 3$.




FIGURE 6. The restored results of GpL and HpL for four test images. From left to right: The test image, The observed pixels, GpL, and HpL. From top to bottom (image/sampling rate): House/5%, Peppers/10%, Cameraman/15%, Barbara/20%.

The test images are shown in Figure 5. In Table 2, we list the PSNR and SSIM values [50] of the results of GpL and HpL with sampling rates 5%, 10%, 15%, and 20%. HpL outperforms GpL for all test images and all sampling rates. From Figure 6, it can be seen that HpL causes fewer spikes than GpL. This is consistent with previous experimental results.

A main drawback of the new hypergraph p -Laplacian regularization is the high computational cost. In our definition, the hyperedges have the same number as the vertices, resulting in the objective function being a large-scale non-smooth function. There exists no simple algorithm for solving it, even when $p = 2$. In the experiment of image inpainting, we use the result of GpL as the initial guess for HpL, which saves a lot of computational time. Nevertheless, working with MATLAB 2020b on a desktop equipped with an Intel Core i7 3.20 GHz CPU, the average running time of GpL/HpL is about 165/262 seconds.

7. CONCLUSION AND FUTURE WORK

In this paper, we defined the ε_n -ball hypergraph and the k_n -nearest neighbor hypergraph from a point cloud and studied the p -Laplacian regularization on both hypergraphs in a semisupervised setting. It was shown that the hypergraph p -Laplacian regularization is variational consistent with the continuum p -Laplacian regularization. Compared to the graph p -Laplacian regularization, the new hypergraph regularization was proven both theoretically and numerically to be more effective in suppressing spiky solutions.

Some parts of the numerical results are still not well understood. For example, the staircasing behavior has been observed in experiments in 1D. The hypergraph model loses its competitiveness in semi-supervised learning with large labeling rates. Both phenomena are worth further exploring. Besides, developing new algorithms that require less computational cost is also one of future directions.

ACKNOWLEDGMENTS

The authors would like to thank the referee for the valuable comments and suggestions. KS is supported by China Scholarship Council. The authors acknowledge support from DESY (Hamburg, Germany), a member of the Helmholtz Association HGF.

REFERENCES

- [1] Songyang Zhang, Shuguang Cui, and Zhi Ding. Hypergraph-based image processing. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 216–220. IEEE, 2020.
- [2] Yawen Zeng, Qin Jin, Tengfei Bao, and Wenfeng Li. Multi-modal knowledge hypergraph for diverse image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 3376–3383, 2023.
- [3] Steffen Klamt, Utz-Uwe Haus, and Fabian Theis. Hypergraphs and cellular networks. *PLoS computational biology*, 5(5):e1000385, 2009.
- [4] Rob Patro and Carl Kingsford. Predicting protein interactions via parsimonious network history inference. *Bioinformatics*, 29(13):i237–i246, 2013.
- [5] Alessia Antelmi, Gennaro Cordasco, Carmine Spagnuolo, and Przemysław Szufel. Social influence maximization in hypergraphs. *Entropy*, 23(7):796, 2021.
- [6] Ariane Fazeney, Daniel Tenbrinck, and Martin Burger. Hypergraph p -Laplacians, scale spaces, and information flow in networks. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 677–690. Springer, 2023.
- [7] Matthias Hein, Simon Setzer, Leonardo Jost, and Syama Sundar Rangapuram. The total variation on hypergraphs-learning on hypergraphs revisited. *Advances in Neural Information Processing Systems*, 26, 2013.
- [8] Francis Bach et al. Learning with submodular functions: A convex optimization perspective. *Foundations and Trends® in Machine Learning*, 6(2-3):145–373, 2013.
- [9] Chenzi Zhang, Shuguang Hu, Zhihao Gavin Tang, and TH Hubert Chan. Re-visiting learning on hypergraphs: confidence interval and subgradient method. In *International Conference on Machine Learning*, pages 4026–4034. PMLR, 2017.
- [10] Pan Li and Olgica Milenkovic. Submodular hypergraphs: p -Laplacians, cheeger inequalities and spectral clustering. In *International Conference on Machine Learning*, pages 3014–3023. PMLR, 2018.
- [11] Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha Talukdar. Hypergcn: A new method for training graph convolutional networks on hypergraphs. *Advances in Neural Information Processing Systems*, 32, 2019.
- [12] Yue Gao, Zizhao Zhang, Haojie Lin, Xibin Zhao, Shaoyi Du, and Changqing Zou. Hypergraph learning: Methods and practices. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(5):2548–2566, 2020.

- [13] Sheng Huang, Mohamed Elhoseiny, Ahmed Elgammal, and Dan Yang. Learning hypergraph-regularized attribute predictors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 409–417, 2015.
- [14] Quan Fang, Jitao Sang, Changsheng Xu, and Yong Rui. Topic-sensitive influencer mining in interest-based social media networks via hypergraph learning. *IEEE Transactions on Multimedia*, 16(3):796–812, 2014.
- [15] Nicolás García Trillos and Dejan Slepčev. Continuum limit of total variation on point clouds. *Archive for Rational Mechanics and Analysis*, 220:193–241, 2016.
- [16] Dejan Slepcev and Matthew Thorpe. Analysis of p -Laplacian regularization in semisupervised learning. *SIAM Journal on Mathematical Analysis*, 51(3):2085–2120, 2019.
- [17] Robert A Adams and John JF Fournier. *Sobolev spaces*. Elsevier, 2003.
- [18] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40:120–145, 2011.
- [19] Antonin Chambolle, Matthias J Ehrhardt, Peter Richtárik, and Carola-Bibiane Schonlieb. Stochastic primal-dual hybrid gradient algorithm with arbitrary sampling and imaging applications. *SIAM Journal on Optimization*, 28(4):2783–2808, 2018.
- [20] Ulrike von Luxburg and Olivier Bousquet. Distance-based classification with Lipschitz functions. *J. Mach. Learn. Res.*, 5(Jun):669–695, 2004.
- [21] Rasmus Kyng, Anup Rao, Sushant Sachdeva, and Daniel A Spielman. Algorithms for Lipschitz learning on graphs. In *Conference on Learning Theory*, pages 1190–1223. PMLR, 2015.
- [22] Jeff Calder. Consistency of Lipschitz learning with infinite unlabeled data and finite labeled data. *SIAM Journal on Mathematics of Data Science*, 1(4):780–812, 2019.
- [23] Tim Roith and Leon Bungert. Continuum limit of Lipschitz learning on graphs. *Foundations of Computational Mathematics*, 23(2):393–431, 2023.
- [24] Xiaojin Zhu, Zoubin Ghahramani, and John D Lafferty. Semi-supervised learning using Gaussian fields and harmonic functions. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 912–919, 2003.
- [25] Dengyong Zhou, Jiayuan Huang, and Bernhard Schölkopf. Learning from labeled and unlabeled data on a directed graph. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 1036–1043, 2005.
- [26] Rie Ando and Tong Zhang. Learning on graph with Laplacian regularization. *Advances in Neural Information Processing Systems*, 19, 2006.
- [27] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17:395–416, 2007.
- [28] Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. Saliency detection via graph-based manifold ranking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3166–3173, 2013.
- [29] Guy Gilboa and Stanley Osher. Nonlocal linear image regularization and supervised segmentation. *Multiscale Modeling & Simulation*, 6(2):595–630, 2007.
- [30] Boaz Nadler, Nathan Srebro, and Xueyuan Zhou. Statistical analysis of semi-supervised learning: The limit of infinite unlabelled data. *Advances in Neural Information Processing Systems*, 22, 2009.
- [31] Zuoqiang Shi, Stanley Osher, and Wei Zhu. Weighted nonlocal Laplacian on interpolation from sparse data. *Journal of Scientific Computing*, 73:1164–1177, 2017.
- [32] Jeff Calder and Dejan Slepčev. Properly-weighted graph Laplacian for semi-supervised learning. *Applied mathematics & optimization*, 82:1111–1159, 2020.
- [33] Jeff Calder, Brendan Cook, Matthew Thorpe, and Dejan Slepcev. Poisson learning: Graph based semi-supervised learning at very low label rates. In *International Conference on Machine Learning*, pages 1306–1316. PMLR, 2020.
- [34] Ahmed El Alaoui, Xiang Cheng, Aaditya Ramdas, Martin J Wainwright, and Michael I Jordan. Asymptotic behavior of ℓ_p -based Laplacian regularization in semi-supervised learning. In *Conference on Learning Theory*, pages 879–906. PMLR, 2016.
- [35] Dengyong Zhou and Bernhard Schölkopf. Regularization on discrete spaces. In *Joint Pattern Recognition Symposium*, pages 361–368. Springer, 2005.
- [36] Abderrahim Elmoataz, Matthieu Toutain, and Daniel Tenbrinck. On the p -laplacian and ∞ -laplacian on graphs with applications in image and data processing. *SIAM Journal on Imaging Sciences*, 8(4):2412–2451, 2015.

- [37] Nicolas Garcia Trillos. Variational limits of k-nn graph-based functionals on data clouds. *SIAM Journal on Mathematics of Data Science*, 1(1):93–120, 2019.
- [38] Taisong Jin, Liujuan Cao, Baochang Zhang, Xiaoshuai Sun, Cheng Deng, and Rongrong Ji. Hypergraph induced convolutional manifold networks. In *IJCAI*, pages 2670–2676, 2019.
- [39] Lin Bie, Gang Xiao, Yipeng Li, and Yue Gao. Hyperg-ps: Voxel correlation modeling via hypergraph for lidar panoptic segmentation. *Fundamental Research*, 2025.
- [40] Shanshan Wan and Jiaqi Ding. A recommendation method based on multi-source heterogeneous hypergraphs and contrastive learning. *IEEE Access*, 2024.
- [41] Mathew Penrose. *Random geometric graphs*, volume 5. OUP Oxford, 2003.
- [42] Gianni Dal Maso. *An introduction to Γ -convergence*, volume 8. Springer Science & Business Media, 2012.
- [43] Andrea Braides. *Gamma-convergence for Beginners*, volume 22. Clarendon Press, 2002.
- [44] Nicolás Garcia Trillos and Dejan Slepčev. On the rate of convergence of empirical measures in ∞ -transportation distance. *Canadian Journal of Mathematics*, 67(6):1358–1383, 2015.
- [45] Imad El Bouchairi, Abderrahim El Moataz, and Jalal Fadili. Discrete p -bilaplacian operators on graphs. In *Image and Signal Processing: 9th International Conference, ICISP 2020, Marrakesh, Morocco, June 4–6, 2020, Proceedings 9*, pages 339–347. Springer, 2020.
- [46] Ahmet Alacaoglu, Olivier Fercoq, and Volkan Cevher. On the convergence of stochastic primal-dual hybrid gradient. *SIAM Journal on Optimization*, 32(2):1288–1318, 2022.
- [47] Laurent Condat. Fast projection onto the simplex and the l1 ball. *Mathematical Programming*, 158(1-2):575–585, 2016.
- [48] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [49] Andrea Vedaldi and Brian Fulkerson. Vlfeat: An open and portable library of computer vision algorithms. In *Proceedings of the 18th ACM International Conference on Multimedia*, pages 1469–1472, 2010.
- [50] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.