ESTIMATING PERIMETER USING GRAPH CUTS

NICOLÁS GARCÍA TRILLOS, DEJAN SLEPČEV, AND JAMES VON BRECHT

ABSTRACT. We investigate the estimation of the perimeter of a set by a graph cut of a random geometric graph. For $\Omega \subset D = (0,1)^d$, with $d \ge 2$, we are given n random i.i.d. points on D whose membership in Ω is known. We consider the sample as a random geometric graph with connection distance $\varepsilon > 0$. We estimate the perimeter of Ω (relative to D) by the, appropriately rescaled, graph cut between the vertices in Ω and the vertices in $D \setminus \Omega$. We obtain bias and variance estimates on the error, which are optimal in scaling with respect to n and ε . We consider two scaling regimes: the dense (when the average degree of the vertices goes to ∞) and the sparse one (when the degree goes to 0). In the dense regime there is a crossover in the nature of approximation at dimension d = 5: we show that in low dimensions d = 2, 3, 4 one can obtain confidence intervals for the approximation error, while in higher dimensions one can only obtain error estimates for testing the hypothesis that the perimeter is less than a given number.

1. INTRODUCTION

This paper investigates the use of random-graph cuts to obtain empirical estimates of the perimeter of a domain $\Omega \subset D := (0,1)^d$ for $d \ge 2$. Let $\mathbf{x}_1, \ldots, \mathbf{x}_n, \ldots$ denote a sequence of independent random points uniformly distributed on the unit cube D and let $V_n := {\mathbf{x}_1, \ldots, \mathbf{x}_n}$.

The problem of estimating the perimeter of Ω based on knowing which points of V_n belong to Ω is a classical question, see [4, 8, 9, 17, 23, 24] for recent contributions, and see Subsection 1.2 below which contains a discussion about related work. Here we consider an estimator of the perimeter that is based on a geometric graph constructed from the point cloud V_n . More precisely, we select $\varepsilon_n > 0$, and connect two points in the cloud if they are within distance ε_n of each other; then we consider an appropriately scaled 'cut' determined by the number of edges in the graph that connect points in Ω with points that belong to Ω^c . This type of estimator is natural to consider, since graph cuts arise as a discretization of the perimeter in many applications such as clustering [3, 7, 14, 15, 18, 25, 26, 27, 28].

While our choice of the estimator is based on its relevance to statistics and machine learning tasks, one of its important features is that it has a small bias. Indeed, the expectation of our estimator provides a second order approximation (in terms of the natural parameter ε_n) of the true perimeter of Ω under some regularity conditions on the boundary of the set (see (7) below); this turns out to be a sharp estimate for the bias. We also obtain precise estimates for the variance of the estimator; for these estimates to hold, we do not need any regularity assumptions on the boundary of the set Ω except that it has finite perimeter in the most general sense. Furthermore, we show that our estimator converges a.s. for remarkably sparse graphs (and indeed in settings which are sparser than for previously considered estimators). The a.s. convergence holds with no regularity assumptions on Ω (other than the fact that it has finite perimeter in the most general

Date: February 10, 2016.

¹⁹⁹¹ Mathematics Subject Classification. 60D05, 62G20, 68R10.

Key words and phrases. perimeter, nonparametric estimation, graph cut, point cloud, random geometric graph, concentration inequality.

sense). Finally, assuming some smoothness on the boundary of Ω , we are able to establish (in the dense graph regime $\frac{1}{n^{1/d}} \ll \varepsilon \ll 1$) the asymptotic distribution of the error. These estimates lead to asymptotic confidence intervals and bounds on the type I and type II errors for hypothesis tests associated to the perimeter of a set.

1.1. Set-up and main results. Let us now be more precise about the setting we consider in this paper. We consider random geometric graphs with vertex set V_n and radius $\varepsilon_n > 0$. That is, graphs where \mathbf{x}_i and \mathbf{x}_j are connected by an edge if $||\mathbf{x}_i - \mathbf{x}_j|| \le \varepsilon_n$. The graph cut between $A \subseteq V_n$ and A^c is given by

$$\operatorname{Cut}_{\varepsilon_n}(A, A^c) := \sum_{\mathbf{x}_i \in A} \sum_{\mathbf{x}_j \in V \setminus A} \mathbf{1}_{\{\|\mathbf{x}_i - \mathbf{x}_j\| \le \varepsilon_n\}}.$$

We define the graph perimeter as a rescaling of the graph cut: For any $\Omega \subseteq D$

(1)
$$\operatorname{GPer}_{n,\varepsilon_n}(\Omega) := \frac{2}{n(n-1)\varepsilon_n^{d+1}}\operatorname{Cut}(V_n \cap \Omega, V_n \cap \Omega^c).$$

See Figure 1 for an illustration of this construction. The scaling is chosen so that $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ becomes a consistent estimator for the true (continuum) perimeter.



FIGURE 1. Edges of the cut between Ω and Ω^c are represented by bold lines, while other edges are dashed lines. Total number of vertices is n = 200 and connectivity radius $\varepsilon = 0.13$.

One of our interests lies in determining how well does $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ estimate the relative perimeter of Ω in D. We first investigate for which scaling of ε_n on n does the convergence hold almost surely as $n \to \infty$. In other words, we want to understand the relation between edgesparsity of random geometric graphs and point-wise convergence of the graph perimeter to the continuum perimeter in the almost sure sense. We consider this question for a very broad family of sets $\Omega \subset D$, which are only assumed to have finite relative perimeter in D in the general sense of [2]. That is, we define the relative perimeter of Ω with respect to D to be

(2)
$$\operatorname{Per}(\Omega) = \sup\left\{\int_{\Omega} \operatorname{div}(v) \, \mathrm{d}x : (\forall x \in D) \ \|v(x)\| \le 1, \ v \in C_c^{\infty}(D, \mathbb{R}^d)\right\}.$$

If Ω has a smooth relative boundary then $Per(\Omega)$ is nothing but the surface area of $\partial \Omega \cap D$. We remark that the notion of the perimeter we use is more general than the notion of Hausdorff measure of the boundary, $\mathcal{H}^{d-1}(\partial \Omega \cap D)$, and than the Minkowski content, which are the ones more typically used in the statistics literature [4, 8, 9]. In particular, as we see below, we work with consistent nonparametric estimators in the most general setting available.

As we recall below in (17–19), it is known that when $\varepsilon_n \to 0$ as $n \to \infty$ then the bias of the estimator vanishes in the limit:

(3)
$$\mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega)) \to \sigma_d \operatorname{Per}(\Omega) \quad \text{as} \quad n \to \infty$$

The scaling factor σ_d satisfies

(4)
$$\sigma_d := \int_{||z|| \le 1} |z_1| \, \mathrm{d}z = \frac{2s_{d-2}}{(d+1)(d-1)}$$

where z_1 denotes the first component of the vector $z \in \mathbb{R}^d$ and s_{d-2} is the area of the (d-2)dimensional unit sphere (the boundary of the unit ball in \mathbb{R}^{d-1}). We refer to the normalizing quantity σ_d as the *surface tension*.

We obtain the following estimates on the deviation of the graph perimeter $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ from its mean. Let

(5)
$$f(n) := \begin{cases} \frac{1}{\sqrt{n\varepsilon_n}} & \text{if } \frac{1}{n^{1/d}} \le \varepsilon_n \\ \frac{1}{n\varepsilon_n^{(d+1)/2}} & \text{if } \frac{1}{n^{2/(d+1)}} \le \varepsilon_n \le \frac{1}{n^{1/d}}. \end{cases}$$

Theorem 1.1. Let $p \ge 1$ and and let $\Omega \subseteq D$ be a set with finite perimeter. Assume $\varepsilon_n \to 0$ as $n \to \infty$. Then,

(6)
$$\mathbb{E}(|\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega))|^p) \le C(f(n))^p$$

where C is a constant that depends only on p and the true perimeter $\operatorname{Per}(\Omega)$ of Ω . In particular, if $n^{-\frac{2}{(d+1)}} \ll \varepsilon_n \ll 1$, then

$$\operatorname{GPer}_{n,\varepsilon_n}(\Omega) \to \sigma_d \operatorname{Per}(\Omega), \text{ almost surely as } n \to \infty.$$

The last part of the previous theorem follows from (3), the moment estimates (6), Markov's inequality, and Borel-Cantelli Lemma which imply that

$$\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \mathbb{E}\left(\operatorname{GPer}_{n,\varepsilon_n}(\Omega)\right) \to 0 \quad \text{a.s.}$$

Remark 1.2. We note that the a.s. convergence holds for rather sparse graphs (see Figure 2). Namely the typical degree of a node is $\omega_d n \varepsilon^d$, where ω_d is the volume of the unit ball in d dimensions. When $n^{-\frac{2}{(d+1)}} \ll \varepsilon_n \ll n^{-\frac{1}{d}}$ the a.s. convergence holds, while the average degree of a vertex converges to zero. The convergence is still possible because the expected number of edges crossing $\partial\Omega$ is still a quantity converging to infinity.



FIGURE 2. Here we illustrate the "sparse" regime when $1/n^{1/d} \gg \varepsilon \gg (1/n)^{2/d+1}$ when the average degree goes to zero. Nevertheless since the number of edges in the cut still increases as n does the convergence of the cut to the perimeter still holds.

Remark 1.3. Given that we show the almost-sure consistency of our estimators for any arbitrary (but fixed) set for which the perimeter is finite, our construction provides a universal strongly convergent estimator which was the desired property listed as an open problem in [8] (for the estimator they considered).

We turn to estimating the bias of the empirical approximation: $|\mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega)) - \sigma_d \operatorname{Per}(\Omega)|$. We first characterize the mean of the graph perimeter, $\mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega))$, as the non-local perimeter $\operatorname{Per}_{\varepsilon_n}(\Omega)$ of Ω , which we introduce in (16); the non-local nature of the functional essentially has to do with the fact that it involves averages of finite differences as opposed to a *local* approach where one considers derivatives. We then proceed to estimate $|\operatorname{Per}_{\varepsilon_n}(\Omega) - \sigma_d \operatorname{Per}(\Omega)|$ explicitly. It proves straightforward to check that $|\operatorname{Per}_{\varepsilon_n}(\Omega) - \sigma_d \operatorname{Per}(\Omega)| = O(\varepsilon_n)$ for general subsets $\Omega \subseteq D$ with smooth relative boundary. However, we show that the error is actually quadratic in ε_n

(7)
$$|\mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega)) - \sigma_d \operatorname{Per}(\Omega)| = |\operatorname{Per}_{\varepsilon_n}(\Omega) - \sigma_d \operatorname{Per}(\Omega)| = O(\varepsilon_n^2)$$

under the extra condition that $dist(\Omega, \partial D) > 0$. This is the content of the next lemma, whose proof may be found in Appendix A.

Lemma 1.4. Let Ω be a set with smooth boundary, such that $\operatorname{dist}(\Omega, \partial D) > 0$. Let $\varepsilon > 0$ and let $\operatorname{Per}_{\varepsilon}(\Omega)$ be defined by (16). Then

(8)
$$\operatorname{Per}_{\varepsilon}(\Omega) = \sigma_d \operatorname{Per}(\Omega) + O(\varepsilon^2).$$

Remark 1.5. The assumption dist $(\Omega, \partial D) > 0$ in the above Lemma is needed in order to obtain bias of order ε^2 . If Ω touches the boundary ∂D the error of order ε^2 is not expected, as can be seen for example by considering the rectangle $\Omega := \{x = (x_1, \dots, x_d) \in D : x_1 \leq 1/2\}$, for which

the error is of order ε ; in this situation the error is completely due to the region where $\partial\Omega$ meets ∂D transversally. Thus, in general, for $\Omega \subseteq D$ with smooth relative boundary, the bias is of order ε . On the other hand, the smoothness of the boundary of Ω is only needed in the previous lemma to guarantee that curvature and its derivatives are well defined. Finally, the constant involved in the term $O(\varepsilon^2)$ depends on the reach of the set Ω and the intrinsic curvature of $\partial\Omega$ together with its derivatives; this can be seen from our computations in Appendix A.

Combining the bias and variance of estimates allows us to obtain the rates of convergence for the error $|\text{GPer}_{n,\varepsilon_n}(\Omega) - \sigma_d \operatorname{Per}(\Omega)|$. In particular we estimate the 'standard deviation'

$$\operatorname{std}(n) := \mathbb{E}\left((\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \sigma_d \operatorname{Per}(\Omega))^2 \right)^{1/2}$$

which we may quantify precisely by using the variance-bias decomposition

$$\operatorname{std}^2(n) = \operatorname{Var}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega)) + (\mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega)) - \sigma_d \operatorname{Per}(\Omega))^2.$$

Using the special case p = 2 of Theorem 1.1 to estimate the variance and using Lemma 1.4 to estimate the bias we obtain the following.

Theorem 1.6. Let $\Omega \subset D$ be an open set with smooth boundary. Assume $n^{-\frac{2}{d+1}} \ll \varepsilon_n \ll 1$ and consider f(n) defined via (5). The error of approximating $\sigma_d \operatorname{Per}(\Omega)$ by $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ satisfies

$$\operatorname{std}(n) = O(f(n) + \varepsilon_n).$$

If we furthermore assume that Ω does not touch the boundary of D, that is $dist(\Omega, \partial D) > 0$, then a better estimate holds:

$$\operatorname{std}(n) = O(f(n) + \varepsilon_n^2).$$

A simple calculation using (5) allows one to choose a scaling of ε_n on n so that the error of the approximation is as small as possible. In particular, for a set Ω which touches the boundary $(\text{dist}(\Omega, \partial D) = 0)$ the optimal scaling of ε_n is

$$\varepsilon_n \sim \begin{cases} n^{-1/3} & \text{if } d \le 3\\ n^{-2/(d+3)} & \text{if } d \ge 3 \end{cases} \quad \text{giving} \quad \text{std}(n) = \begin{cases} n^{-1/3} & \text{if } d \le 3\\ n^{-2/(d+3)} & \text{if } d \ge 3 \end{cases}$$

We note that for d < 3 the optimal ε_n is achieved in the regime $n^{-1/d} \leq \varepsilon_n \ll 1$, while if d > 3 it is in the sparse regime.

If we consider sets Ω with smooth boundary but such that $dist(\Omega, \partial D) > 0$, then the optimal scaling of ε_n on n is as follows

$$\varepsilon_n \sim \begin{cases} n^{-2/5} & \text{if } d \le 5\\ n^{-4/(d+5)} & \text{if } d \ge 5 \end{cases} \quad \text{giving} \quad \operatorname{std}(n) = \begin{cases} n^{-1/5} & \text{if } d \le 5\\ n^{-2/(d+5)} & \text{if } d \ge 5. \end{cases}$$

Again we note that the optimal ε_n is in the sparse regime if d > 5.

This has implications to how well is the perimeter estimated by graph cuts in the graphs considered in most machine learning applications. Namely if $d \ge 5$ and the graph has average degree bounded from below, that is when $n^{-1/d} \le \varepsilon_n \ll 1$, then, since the bias bound is sharp, most of the error is due to the bias term.

We now consider obtaining confidence intervals for the value of the true perimeter $\operatorname{Per}(\Omega)$ based on the estimator $\frac{1}{\sigma_d} \operatorname{GPer}_{n,\varepsilon_n}(\Omega)$. We focus on the dense regime (See Figures 3 and 4), $\frac{1}{n^{1/d}} \ll \varepsilon_n \ll 1$, and first obtain the asymptotic distribution of $\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega))$.



FIGURE 3. Here we illustrate the "very dense" regime, $1 \gg \varepsilon \gg (\ln(n)/n)^{1/d}$, when the graphs are connected with high probability.

Since the term $\mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega)) = \operatorname{Per}_{\varepsilon_n}(\Omega)$ depends on the set Ω itself, which is unknown (we only assume we have an oracle access to it), we need the bias to be negligible compared to the standard deviation of our estimator in order to construct confidence intervals for $\operatorname{Per}(\Omega)$. From (6) the standard deviation of $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ scales as $\frac{1}{\sqrt{n\varepsilon_n}}$ while by (7) the bias scales as ε_n^2 ; these estimates lead to restrictions on the dimensions for which the bias is negligible with respect to the standard deviation. Namely this is possible for d = 2, 3, and 4.

Theorem 1.7. Let $\Omega \subseteq D$ be an open set with smooth boundary such that $dist(\Omega, \partial D) > 0$. Let ε_n be such that

$$\frac{1}{n^{1/d}} \ll \varepsilon_n \ll 1$$

Then,

$$\sqrt{\frac{n\varepsilon_n}{4C_d \operatorname{Per}(\Omega)}} \left(\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \operatorname{Per}_{\varepsilon_n}(\Omega) \right) \xrightarrow{w} N(0,1),$$

where C_d is given by

(9)
$$C_d := 2 \int_0^1 \left| B_d(0,1) \cap \left\{ x = (x_1, \dots, x_d) \in \mathbb{R}^d : x_d \ge t \right\} \right|^2 \mathrm{d}t.$$

If in addition the dimension d is either 2, 3 or 4 and if

(10)
$$\frac{1}{n^{1/d}} \ll \varepsilon_n \ll \frac{1}{n^{1/5}}$$

then,

(11)
$$\sqrt{\frac{n\varepsilon_n}{4C_d \operatorname{Per}(\Omega)}} \left(\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \sigma_d \operatorname{Per}(\Omega) \right) \xrightarrow{w} N(0,1)$$



FIGURE 4. Here we illustrate an intermediate, "dense", regime $(\ln(n)/n)^{1/d} \gg \varepsilon \gg (1/n)^{1/d}$. The average degree still increases, but the graphs are disconnected with high probability.

Naturally, the previous theorem implies that one can obtain confidence intervals for the value of $Per(\Omega)$ when the dimension d is 2, 3 or 4. Let us fix $\alpha \in (0, 0.5)$ and let Z_{α} be the $1 - \alpha$ quantile of the standard normal distribution. That is, Z_{α} is such that

$$\mathbb{P}\left(N(0,1) \le Z_{\alpha}\right) = 1 - \alpha.$$

Then provided that

$$\frac{1}{n^{1/d}} \ll \varepsilon_n \ll \frac{1}{n^{1/5}}$$

it follows that with probability converging to $1 - \alpha$, $Per(\Omega) \in (a_n^-, a_n^+)$, where

$$a_n^{\pm} := \frac{1}{\sigma_d} (\operatorname{GPer}_{n,\varepsilon_n}(\Omega) \pm \sqrt{\frac{4C_d \operatorname{Per}(\Omega)}{n\varepsilon_n}} Z_{\alpha/2}).$$

For $d \ge 5$, the confidence intervals can not be obtained unless one has some extra quantitative information on the smoothness of $\partial\Omega$ (like curvature bounds). Nevertheless since it is known that the nonlocal perimeter $\operatorname{Per}_{\varepsilon_n}(\Omega) = \mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega))$ is less than $\sigma_d \operatorname{Per}(\Omega)$ (see (30)), one can construct reliable test for the hypothesis that $\operatorname{Per}(\Omega)$ is *less* than a certain number ρ without using any quantitative estimates on the smoothness of $\partial\Omega$. Namely consider the hypothesis testing of

(12)
$$H_0: \operatorname{Per}(\Omega) \le \rho, \quad \text{vs} \quad H_A: \operatorname{Per}(\Omega) > \rho,$$

based on our estimator $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$. We consider the statistic:

(13)
$$l_n := \sqrt{\frac{n\varepsilon_n}{4C_d\rho}} \left(\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \sigma_d \rho \right)$$

The test consists on

(14) Accept
$$H_0$$
 if $l_n \leq Z_{\alpha}$, reject otherwise.

Proposition 1.8. Assume $d \ge 2$ and $\frac{1}{n^{1/d}} \ll \varepsilon_n \ll 1$. Then, the type I error satisfies

$$\mathbb{P}_{H_0}\left(l_n > Z_\alpha\right) \lesssim \alpha,$$

where \leq simply means that the left hand side of the above expression is asymptotically below α . The type II error satisfies

$$\mathbb{P}_{H_A}(l_n \le Z_\alpha) = O\left(\frac{1}{\sqrt{n\varepsilon_n}}\right).$$

Remark 1.9 (Extensions). The estimates obtained in Theorem 1.1 and Theorem 1.6 are not exclusive to the case where the points are uniformly distributed in the unit cube and to geometric graphs. In fact, with slight modifications to the proofs of Theorem 1.1 and Theorem 1.6, we can extend these results to more general situations. For example, if the data points are distributed according to some smooth density p that is supported on a regular, bounded domain $D \subset \mathbb{R}^d$ with p bounded below and above by positive constants, then the results still hold. In this case, the limiting value of $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ is a weighted perimeter $\operatorname{Per}(\Omega, p^2)$ (see for example [10]). Convergence is guaranteed for the same scaling for ε_n as in the uniform case. The fact that the weight is p^2 and not p (as may be a priori expected) ultimately comes from the fact that a graph cut is a double sum.

Furthermore if instead of weights $\mathbf{1}_{\{\|\mathbf{x}_i-\mathbf{x}_j\|\leq\varepsilon_n\}}$ in the definition of the graph cut, one considers edge weights $\eta(\|\mathbf{x}_i-\mathbf{x}_j\|/\varepsilon_n)$ where η is nonnegative, integrable and non-increasing, the results still hold, provided we change the surface tension σ_d with a surface tension associated to η defined as in [10].

Outline. In Subsection 1.3 we establish (3) and give an outline of the argument behind our main results. We present the proof of Theorem 1.1 in Section 2, while in Subsection 2.1 we show that the scaling is sharp (up to logarithmic corrections) in the sense that if $n^2 \varepsilon_n^{d+1} \to 0$ then $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ converges in probability to zero and hence does not converge almost surely to the (rescaled) relative perimeter. In Section 3 we prove the results on the asymptotic distribution of the error stated in Theorem 1.7. In Subsection 3.1 we study the type I and type II errors of the hypothesis test of Proposition 1.8. Finally, Appendix A deals with the bias estimate from Lemma 1.4.

1.2. **Discussion.** Here we discuss the connections between our work and related works in the literature. First we relate it to other estimators of perimeter based on a random sample. Then we contrast the type of the convergence and the scaling regimes considered in this paper with the ones needed for the convergence of graph-cut based machine learning algorithms for clustering and related tasks.

The problem of estimating the perimeter of a set, Ω , based on knowing which points of a random sample V_n belong to Ω , has been considered by a number of works. Cuevas, Fraiman, and Rodríguez-Casaet [9], considered estimators of the Minkowksi content, which agrees with perimeter for regular enough sets, but is a less general notion of the perimeter than the one we consider, (2). Their estimator is based on counting the vertices near the boundary (relative to a parameter ε_n), while we "count" edges of a graph. The error bound obtained was of order $O(n^{-1/2d})$. Pateiro-López and Rodríguez-Casal [24] consider a similar estimator and improve the bounds to $O((\ln n/n)^{1/(d+1)})$ under qualitative regularity conditions (rolling ball conditions). Cuevas, Fraiman, and Györfi [8], obtain the convergence of estimators similar to those of [9], when $\varepsilon_n \gg 1/n^{1/d}$ and under weaker conditions on the regularity of Ω , although still not in the full generality we consider in this paper.

Armendáriz, Cuevas, and Fraiman [4] consider a similar set-up to that in [9, 24] but with different sampling rates for Ω and Ω^c : let *n* denote the number of sample points in Ω and *k* the number of sample points in Ω^c . Under mild assumptions on the regularity of $\partial\Omega$, and under some conditions on *n* and *k* which include $1 \gg \varepsilon_{n,k} \gg (1/n^{1/3})$ and $k \gg (n/\varepsilon_{n,k})^{d/2}$ they obtain the asymptotic distribution of the error, under the same scaling in *n* and ε that we consider in Theorem 1.7. We notice that our estimator is different, and also that in [4] a very large number of points *k* in Ω^c is needed for the consistency to hold. This allows the authors to obtain the asymptotic distribution of the total error in any dimension, while in our setting we only obtain it in low dimensions. A further difference between their work and ours, is that we allow for a wider range in ε , namely $1 \gg \varepsilon_n \gg (1/n)^{2/(d+1)}$.

Jiménez and Yukich [17] give a different use to the point cloud and instead of considering a parameter ε to count points close to the boundary of the set or to define a geometric graph, they consider a new estimator based on the Delaunay triangulation induced by the cloud. They obtain results not only on estimating the perimeter of the set, but also integrals of functions over $\partial\Omega$.

Kothari, Nayyeri, and O'Donnell [21] and Neeman [23] consider an estimator essentially based on the following procedure: pick n random points uniformly distributed on D and to each point associate a random direction for a 'needle' based at the point with length of order $\sqrt{\varepsilon_n}$; then count how many of the needles touch the boundary of $\partial\Omega$. Their main motivation is to consider the perimeter estimation from the viewpoint of *property testing* as introduced in [19]. In that setting, the idea is to produce an algorithm that requires a small number of samples (essentially independent of the dimension d) in order to determine if a given set has a small perimeter or is 'far away' from a set that has small perimeter. The authors show completeness and soundness of the test they design (the notion of completeness and soundness is as in [21, 19]). The price to pay for the dimension independence of their statements is that of considering a different notion of testing; this is one of the main differences with our work since we consider the perimeter testing in the more classical framework of hypothesis testing (Proposition 1.8). We note that the completeness of [21] is analogous to the type I error, but the soundness is fundamentally different from estimating the type II error.

It is also worth mentioning the work of Belkin, Narayanan and Niyogi [5] where they consider an algorithm that requires as few samples as possible in order to estimate the perimeter of a convex body. Their results show that there is an algorithm that uses $O(d^4\gamma^{-2})$ samples to obtain an estimator for the true perimeter of the convex set, with an error of approximation of γ ; this statement holds with high probability. We notice the polynomial dependence on dimension in their estimates.

Let us now contrast the type of the convergence and the scaling regimes that we consider in this paper with the ones needed for the convergence of graph-cut based machine learning algorithms for clustering and related tasks.

Graph-cut based algorithms for tasks such as clustering have played an important role in machine learning [3, 7, 14, 15, 18, 25, 26, 27, 28]. Data clustering algorithms are called consistent if as the sample size $n \to \infty$ their outputs converge to a desired partitioning of the underlying measure being sampled. It is of interest to understand under what scaling of ε_n on n does the consistency hold. Here we showed that for a *fixed* set Ω the value of graph perimeter (and indeed of the graph-cut-based objective functionals such as Cheeger, ratio, or normalized cuts) converges to the perimeter of Ω ; this result holds even for rather sparse graphs (Theorem 1.6). In particular, partitioning such sparse graphs (as on Figure 2) does not provide almost any information about the clusters present in the data. We conclude that the convergence of the graph perimeter of a fixed set towards the continuum perimeter does not provide the needed information on the asymptotic properties of graph-cut based clustering algorithms. To obtain consistency of such algorithms one needs a stronger notion of convergence of graph based functionals towards continuum functionals. Recently the authors [10], and together with Laurent and Bresson [11] have developed the appropriate notion of convergence (based on Γ -convergence from the calculus of variations), and have applied it to consistency of Cheeger, ratio, sparsest, and normalized cut based point cloud clustering. More specifically, in [11] it is shown that consistency holds if $\frac{(\log(n))^{1/d}}{n^{1/d}} \ll \varepsilon \ll 1$ (for $d \ge 3$), in that regime the graphs are connected with high probability as $n \to \infty$.

1.3. Outline of the argument. In order to understand the asymptotic behavior of the graph perimeter, following [3] we first define a symmetric kernel $\phi_{\varepsilon} : D \times D \to (0, \infty)$ by

$$\phi_{\varepsilon}(x,y) = \frac{\mathbf{1}_{\{||x-y|| \le \varepsilon\}}}{\varepsilon^{d+1}} |\mathbf{1}_{\Omega}(x) - \mathbf{1}_{\Omega}(y)|.$$

Using the kernel ϕ_{ε_n} , we can then write $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ as

(15)
$$\operatorname{GPer}_{n,\varepsilon_n}(\Omega) = \frac{2}{n(n-1)} \sum_{i=1}^n \sum_{j=i+1}^n \phi_{\varepsilon_n}(\mathbf{x}_i, \mathbf{x}_j),$$

which is a *U-statistic* in the terminology of [16]. In order to understand the mean of this *U*-statistic, we define the *non-local perimeter* of the set $\Omega \subseteq D$ as the expression

(16)
$$\operatorname{Per}_{\varepsilon}(\Omega) := \int_{D} \int_{D} \phi_{\varepsilon}(x, y) \, \mathrm{d}x \mathrm{d}y = \frac{2}{\varepsilon^{d+1}} \int_{\Omega} \int_{D \setminus \Omega} \mathbf{1}_{\{\|x-y\| \le \varepsilon\}} \, \mathrm{d}x \mathrm{d}y.$$

A simple computation shows that $\operatorname{Per}_{\varepsilon_n}(\Omega)$ is the mean of $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$, that is,

(17)
$$\mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega)) = \operatorname{Per}_{\varepsilon_n}(\Omega).$$

Additionally, Remark 4.3 in [10] establishes that the non-local perimeter $\operatorname{Per}_{\varepsilon_n}(\Omega)$ approaches a constant multiple of the relative perimeter of Ω as the parameter ε_n goes to zero. More precisely, if $\varepsilon_n \to 0$ as $n \to \infty$ then

(18)
$$\operatorname{Per}_{\varepsilon_n}(\Omega) \to \sigma_d \operatorname{Per}(\Omega)$$
 as $n \to \infty$

for σ_d the surface tension (4). This convergence also follows from the estimates in Appendix A in the special case that Ω has a smooth boundary. Combining (18) with (17) we conclude that if ε_n converges to zero as $n \to \infty$ then

(19)
$$\mathbb{E}(\operatorname{GPer}_{n,\varepsilon_n}(\Omega)) \to \sigma_d \operatorname{Per}(\Omega) \quad \text{as } n \to \infty.$$

Since the graph perimeter $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ is a *U*-statistic of order two we can use the general theory of *U*-statistics to obtain moment estimates for $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$. Let us first note that Hoeffding's decomposition theorem for *U*-statistics of order two (see [20]) implies that $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ can be written as:

(20)
$$\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \operatorname{Per}_{\varepsilon_n}(\Omega) = 2U_{n,1} + U_{n,2},$$

where $U_{n,1}$ is a U-statistic of order one (just a sum of centered independent random variables) and $U_{n,2}$ is a U-statistic of order two which is *canonical* or completely degenerate (see [20]). In order to define the variables $U_{n,1}$ and $U_{n,2}$, let us introduce the functions

$$\bar{\phi}_{\varepsilon}(x) := \int_{D} \phi_{\varepsilon}(x, z) \mathrm{d}z, \qquad \qquad x \in D,$$

(21)
$$g_{n,1}(x) := \bar{\phi}_{\varepsilon_n}(x) - \operatorname{Per}_{\varepsilon_n}(\Omega), \qquad x \in D,$$

$$g_{n,2}(x,y) := \phi_{\varepsilon_n}(x,y) - \phi_{\varepsilon_n}(x) - \phi_{\varepsilon_n}(y) + \operatorname{Per}_{\varepsilon_n}(\Omega), \quad x, y \in D$$

With the previous definitions, we can now define

(22)
$$U_{n,1} = \frac{1}{n} \sum_{i=1}^{n} g_{n,1}(\mathbf{x}_i),$$
$$U_{n,2} = \frac{2}{n(n-1)} \sum_{1 \le i < j \le n} g_{n,2}(\mathbf{x}_i, \mathbf{x}_j).$$

We remark that $\int_D g_{n,1}(z)dz = 0$ and that $\int_D g_{n,2}(x,z)dz = 0$ for all $x \in D$. Because of this, $U_{n,1}$ and $U_{n,2}$ are said to be canonical statistics of order one and two respectively (see [20]). Now, Bernstein's inequality [6] implies that

(23)
$$\mathbb{E}(|U_{n,1}|^p) \leq \frac{C_p}{n^p} \max\left(A_{n,1}^p, B_{n,1}^p\right),$$

where

(24)
$$A_{n,1} := ||g_{n,1}||_{\infty}, \quad B_{n,1} := \sqrt{n} ||g_{n,1}||_2.$$

and C_p is a universal constant. See also [12] for a slight generalization of the previous result.

On the other hand some of the moment estimates in [12] for canonical U-statistics of order two can be used to prove that

(25)
$$\mathbb{E}\left(|U_{n,2}|^p\right) \le \frac{C_p}{n^{2p}} \max\left(A_{n,2}^p, B_{n,2}^p, C_{n,2}^p\right),$$

where

(26)
$$A_{n,2} := ||g_{n,2}||_{\infty}, \quad B_{n,2} := n||g_{n,2}||_2, \quad (C_{n,2})^2 := n||\int_D g_{n,2}^2(\cdot, y) \, \mathrm{d}y||_{\infty}.$$

and C_p is a universal constant. From the decomposition (20) it follows that for $p \ge 1$

$$\mathbb{E}\left(|\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \operatorname{Per}_{\varepsilon_n}(\Omega)|^p\right) \le C_p(\mathbb{E}\left(|U_{n,1}|^p\right) + \mathbb{E}\left(|U_{n,2}|^p\right))$$

Thus in order to obtain the moment estimates for $\operatorname{GPer}_{n,\varepsilon_n}(\Omega)$ in Theorem 1.1, we focus on finding estimates for the quantities in (24) and (26).

Remark 1.10. The estimates on $U_{n,1}$ and $U_{n,2}$ exhibit a crossover in the nature when the parameter ε_n transitions between the sparse $(\frac{1}{n^{2/(d+1)}} \ll \varepsilon_n \ll \frac{1}{n^{1/d}})$ and the dense regime $(\frac{1}{n^{1/d}} \ll \varepsilon_n \ll 1)$. From the theory of U-statistics the crossover in the nature of the bounds, is connected to the different nature of the two components in the canonical decomposition for U-statistics. Under the dense regime the biggest source of error comes from the term $U_{n,1}$ while in the sparse regime the biggest source of error comes from the term $U_{n,1}$ and $U_{n,2}$ exhibit a different nature. In fact, one can think that $U_{n,1}$ is a global quantity as it is the sum of averages, while $U_{n,2}$ is simply the sum of pure interactions concentrating on the boundary of the set Ω . We believe that there is a deeper geometric and analytic reason for the scaling of the error that appears in the sparse regime $(\frac{1}{n^{2/(d+1)}} \ll \varepsilon_n \ll \frac{1}{n^{1/d}})$. Moreover, we believe that such understanding

would allow us to complete the study of the asymptotic distribution in Theorem 1.7 for the sparse regime. We expect such distribution to be of the Gaussian chaos type.

The bias estimates in Appendix A are obtained by a series of computations whose starting point is writing $\operatorname{Per}_{\varepsilon}(\Omega)$ in terms of an iterated integral, the outer one taken over the manifold $\partial\Omega$ and the inner one taken along the normal line to $\partial\Omega$ at an arbitrary point $x \in \partial\Omega$. Such computations show that the first order term of $\operatorname{Per}_{\varepsilon_n}(\Omega)$ on ε_n vanishes.

Finally, Theorem 1.7 is obtained by using the canonical decomposition of U-statistics and by noticing that in the dense regime $\frac{1}{n^{1/d}} \ll \varepsilon_n \ll 1$, the variable $U_{n,2}$ is negligible in relation to $U_{n,1}$. We make use of the CLT for triangular arrays after computing the variances of the involved variables.

2. Proof of Theorem 1.1

We first compute the moments of $U_{n,1}$ and so we start computing the quantities $A_{n,1}$ and $B_{n,1}$ from (24). Denote by T_{ε} the ε -tube around $\partial \Omega$, that is, consider the set

(27)
$$T_{\varepsilon} := \left\{ x \in \mathbb{R}^d : \operatorname{dist}(x, \partial \Omega) \le \varepsilon \right\}.$$

We also consider the half tubes T_{ε}^{-} and T_{ε}^{+} ,

(28)
$$T_{\varepsilon}^{-} := \{x \in \Omega : \operatorname{dist}(x, \partial \Omega) \le \varepsilon\}, \quad T_{\varepsilon}^{+} := \{x \in \Omega^{c} : \operatorname{dist}(x, \partial \Omega) \le \varepsilon\}.$$

With these definitions it is straightforward to check that

(29)
$$\bar{\phi}_{\varepsilon_n}(x) = \begin{cases} |B_d(x,\varepsilon_n) \cap \Omega| / \varepsilon_n^{d+1} & \text{if } x \in T_{\varepsilon_n}^+ \\ |B_d(x,\varepsilon_n) \cap \Omega^c| / \varepsilon_n^{d+1} & \text{if } x \in T_{\varepsilon_n}^- \\ 0 & \text{if } x \notin T_{\varepsilon_n} \end{cases}$$

Since $|B_d(x,\varepsilon_n) \cap \Omega|$ and $|B_d(x,\varepsilon_n) \cap \Omega^c|$ are bounded by $\alpha_d \varepsilon_n^d$, where α_d is the volume of the *d*-dimensional unit ball, we deduce that

$$A_{n,1} = O\left(\frac{1}{\varepsilon_n}\right).$$

In order to compute the quantity $B_{n,1}$ we use the following lemma, whose proof may be found in Appendix B.

Lemma 2.1. Let $p \ge 1$ and let $\Omega \subseteq D$, be a set with finite perimeter. Then, for all $\varepsilon > 0$ we have

$$\int_{D} \bar{\phi}_{\varepsilon}^{p}(x) \mathrm{d}x \leq \frac{\alpha_{d}^{p-1} \sigma_{d}}{\varepsilon^{p-1}} \operatorname{Per}(\Omega)$$

In particular, taking p = 1 in the previous expression, we obtain

(30)
$$\operatorname{Per}_{\varepsilon}(\Omega) \leq \sigma_d \operatorname{Per}(\Omega).$$

Using the previous lemma with p = 2 we deduce that $\int_D \bar{\phi}_{\varepsilon_n}^2(x) dx = O\left(\frac{1}{\varepsilon_n}\right)$, and since

$$\int_{D} g_{n,1}^{2}(x) \mathrm{d}x = \int_{D} \bar{\phi}_{\varepsilon_{n}}^{2}(x) \mathrm{d}x - \left(\operatorname{Per}_{\varepsilon_{n}}(\Omega)\right)^{2},$$

we conclude that

$$B_{n,1} = O\left(\sqrt{\frac{n}{\varepsilon_n}}\right).$$

From the previous computations, we deduce that

$$\mathbb{E}(|U_{n,1}|^p) \le C \max\left(\frac{1}{n^p \varepsilon_n^p}, \frac{1}{n^{p/2} \varepsilon_n^{p/2}}\right),$$

where C may depend on the set Ω and p. If $\frac{1}{n^{2/(d+1)}} \leq \varepsilon_n$, so that in particular $\frac{1}{n\varepsilon_n}$ is o(1), then

(31)
$$\mathbb{E}(|U_{n,1}|^p) \le \frac{C}{n^{p/2}\varepsilon_n^{p/2}}$$

Remark 2.2. Later on, in Lemma 3.1, we provide an explicit computation of $\int_D \bar{\phi}_{\varepsilon_n}^2(x) dx$ up to order $\frac{1}{\varepsilon_n}$, which is useful when studying the asymptotic distribution of a rescaled version of $U_{n,1}$.

Now we turn to the task of obtaining moment estimates for $U_{n,2}$. We estimate the quantities $A_{n,2}$, $B_{n,2}$ and $C_{n,2}$ from (26). Let us start by estimating $A_{n,2}$. Note that for any $(x, y) \in D \times D$, $\bar{\phi}_{\varepsilon_n}(x)$ and $\bar{\phi}_{\varepsilon_n}(y)$ are of order $\frac{1}{\varepsilon_n}$ and that $\operatorname{Per}_{\varepsilon_n}(\Omega)$ is of order one. Thus, it is clear from the definition of $g_{n,2}$ in (21) that

$$A_{n,2} = O\left(\frac{1}{\varepsilon_n^{d+1}}\right).$$

On the other hand, using $\phi_{\varepsilon_n}^2(x,y) = \frac{1}{\varepsilon_n^{d+1}} \phi_{\varepsilon_n}(x,y)$, we obtain that for every $x \in D$,

(32)

$$\int_{D} g_{n,2}^{2}(x,y) dy = \int_{D} \phi_{\varepsilon_{n}}^{2}(x,y) dy - \bar{\phi}_{\varepsilon_{n}}^{2}(x) + 2\theta_{n} \overline{\phi}_{n}(x) \\
- 2 \int_{D} \phi_{\varepsilon_{n}}(x,y) \bar{\phi}_{\varepsilon_{n}}(y) dy + \int_{D} \bar{\phi}_{\varepsilon_{n}}^{2}(y) dy - \theta_{n}^{2} \\
= \frac{1}{\varepsilon_{n}^{d+1}} \bar{\phi}_{\varepsilon_{n}}(x) - \bar{\phi}_{\varepsilon_{n}}^{2}(x) + 2\theta_{n} \overline{\phi}_{n}(x) \\
- 2 \int_{D} \phi_{\varepsilon_{n}}(x,y) \bar{\phi}_{\varepsilon_{n}}(y) dy + \int_{D} \bar{\phi}_{\varepsilon_{n}}^{2}(y) dy - \theta_{n}^{2} \\$$

where we are using $\theta_n := \operatorname{Per}_{\varepsilon_n}(\Omega)$. From this, it follows that

$$C_{n,2} = O\left(\sqrt{\frac{n}{\varepsilon_n^{d+2}}}\right).$$

Finally, upon integration of (32) and direct computations, we obtain

$$||g_{n,2}||_2^2 = \frac{\theta_n}{\varepsilon_n^{d+1}} - 2\int_D \bar{\phi}_{\varepsilon_n}^2(y) \mathrm{d}y + \theta_n^2$$

which implies that

$$B_{n,2} = O\left(\frac{n}{\varepsilon_n^{(d+1)/2}}\right).$$

Thus, from (25) we deduce that

$$\mathbb{E}(|U_{n,2}|^p) \le K_p \max\left(\frac{1}{n^{2p}\varepsilon_n^{p(d+1)}}, \frac{1}{n^p\varepsilon_n^{p(d+1)/2}}, \frac{1}{n^{3p/2}\varepsilon_n^{p(d+2)/2}}\right),$$

where K_p depends on the set Ω . Hence, if $\frac{1}{n^{2/(d+1)}} \leq \varepsilon_n$, we have

(33)
$$\mathbb{E}(|U_{n,2}|^p) \le \frac{K_p}{n^p \varepsilon_n^{p(d+1)/2}}$$

Combining (31) and (33) and using the canonical decomposition (20), we obtain (6).

2.1. Sharpness of the Rate in Theorem 1.1. A very simple argument shows that the rates for ε_n that guarantee the almost sure convergence of the graph perimeter to the actual perimeter in Theorem 1.1 are optimal in terms of scaling, up to logarithmic corrections.

In fact, suppose $n^2 \varepsilon_n^{d+1} = o(1)$ and let e_n denote the random variable that counts the number of edges that cross in the interface between Ω and its complement. In other words, we define

$$e_n := \varepsilon_n^{d+1} \sum_{i=1}^n \sum_{j=i+1}^n \phi_{\varepsilon_n}(\mathbf{x}_i, \mathbf{x}_j).$$

7 . 4

As a consequence, if Ω has finite perimeter then we have

(34)
$$\operatorname{GPer}_{\varepsilon_n,n}(\Omega) = \frac{2}{n(n-1)\varepsilon_n^{d+1}}e_n, \qquad \mathbb{E}(e_n) = \frac{n(n-1)\varepsilon_n^{d+1}}{2}\operatorname{Per}_{\varepsilon_n}(\Omega).$$

Note that e_n takes integer values in the range $\{0, 1, \dots, N\}$ for N = n(n-1)/2, so that

$$\mathbb{E}(e_n) = \sum_{k=1}^{N} k p_k^n \qquad p_k^n := \mathbb{P}(e_n = k).$$

The fact that $p_0^n + \cdots + p_N^n = 1$ implies

$$\mathbb{E}(e_n) = \sum_{k=1}^{N} k p_k^n \ge \sum_{k=1}^{N} p_k^n = (1 - p_0^n)$$

In particular, from (34) and (18) we deduce that if $n^2 \varepsilon_n^{d+1} \to 0$ and Ω has finite perimeter then

$$(1-p_0^n) \le \mathbb{E}(e_n) = o(1).$$

On the other hand, note that for any given $\gamma > 0$ it is true that $\operatorname{GPer}_{\varepsilon_n,n}(\Omega) > \gamma$ implies that $e_n \neq 0$. In turn

$$\mathbb{P}\left(\operatorname{GPer}_{\varepsilon_n}(\Omega) > \gamma\right) \le \mathbb{P}\left(e_n \neq 0\right) = 1 - p_0^n = o(1)$$

We conclude that if $n^2 \varepsilon_n^{d+1} \to 0$ then $\operatorname{GPer}_{\varepsilon_n,n}(\Omega)$ converges in probability to zero. Therefore, if Ω has a non-zero, finite perimeter then $\operatorname{GPer}_{\varepsilon_n,n}(\Omega)$ does not converge to $\sigma_d \operatorname{Per}(\Omega)$ in probability (nor almost surely, either).

3. PROOF OF THEOREM 1.7

The proof of Theorem 1.7 relies on the following lemma, whose proof may be found in Appendix C.

Lemma 3.1. Assume that $\varepsilon_n \to 0$ as $n \to \infty$, and that $\Omega \subset D$ is an open set with smooth boundary so that $dist(\Omega, \partial D) > 0$. Then

(35)
$$\operatorname{Var}(g_{n,1}(X_1)) = \frac{C_d \operatorname{Per}(\Omega)}{\varepsilon_n} + O(1),$$

where C_d is given by (9).

Now we turn our attention to the proof of Theorem 1.7.

Proof of Theorem 1.7. Note that from (20), (21) and (22) we have

$$\sqrt{\frac{n\varepsilon_n}{4C_d \operatorname{Per}(\Omega)}} \left(\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \operatorname{Per}_{\varepsilon_n}(\Omega) \right) = \sqrt{\frac{\varepsilon_n}{nC_d \operatorname{Per}(\Omega)}} \sum_{i=1}^n g_{n,1}(\mathbf{x}_i) + \sqrt{\frac{n\varepsilon_n}{4C_d \operatorname{Per}(\Omega)}} U_{n,2}.$$

From the moment estimates (33), we deduce that

(36)
$$\sqrt{\frac{n\varepsilon_n}{4C_d \operatorname{Per}(\Omega)}} U_{n,2} \xrightarrow{P} 0.$$

On the other hand, we note that from (35)

$$\frac{\operatorname{std}(g_{n,1}(\mathbf{x}_1))\sqrt{\varepsilon_n}}{\sqrt{C_d\operatorname{Per}(\Omega)}}\to 1,\quad \text{as }n\to\infty,$$

where $\operatorname{std}(g_{n,1}(\mathbf{x}_1))$ is the standard deviation of $g_{n,1}(\mathbf{x}_1)$. Lyupanov's condition which is sufficient to allow us to use the central limit theorem for triangular arrays is easily checked from (2.1). We deduce that

$$\sqrt{\frac{\varepsilon_n}{nC_d \operatorname{Per}(\Omega)}} \sum_{i=1}^n g_{n,1}(X_i) \stackrel{w}{\longrightarrow} N(0,1).$$

Combining with (36) and the Slutsky's theorem, we obtain the desired result. Finally, to obtain the last statement in the theorem, we note that from the bias estimates in Lemma 1.4,

$$\sqrt{n\varepsilon_n} |\operatorname{Per}_{\varepsilon_n}(A) - \sigma_\eta \operatorname{Per}(A)| = O(n^{1/2}\varepsilon_n^{5/2})$$

Under the condition (10), we conclude that $\sqrt{n\varepsilon_n} |\operatorname{Per}_{\varepsilon_n}(A) - \sigma_\eta \operatorname{Per}(A)| \to 0$. This implies (11).

3.1. Application to Perimeter Testing. Here we prove Proposition 1.8. We assume that $\Omega \subseteq D$ is an open set with smooth boundary such that $\overline{\Omega} \subset D$. Note that under the null hypothesis, if $l_n > Z_{\alpha}$, then,

(37)
$$Z_{\alpha} < l_n \le \sqrt{\frac{n\varepsilon_n}{4C_d \operatorname{Per}(\Omega)}} \left(\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \operatorname{Per}_{\varepsilon_n}(\Omega) \right),$$

where we used that $\operatorname{Per}_{\varepsilon_n}(\Omega) < \sigma_d \operatorname{Per}(\Omega)$ by Lemma 2.1. Thus, using Theorem 1.7, we deduce that asymptotically, the type I error of our test is

$$\mathbb{P}_{H_0}\left(l_n > Z_\alpha\right) \le \mathbb{P}\left(\sqrt{\frac{n\varepsilon_n}{4C_d \operatorname{Per}(\Omega)}} \left(\operatorname{GPer}_{n,\varepsilon_n}(\Omega) - \operatorname{Per}_{\varepsilon_n}(\Omega)\right) > Z_\alpha\right) \approx \mathbb{P}\left(N(0,1) > Z_\alpha\right) = \alpha.$$

Where \approx arrises because of the difference of the asymptotic distribution and its discrete approximation.

In order to compute the type II error of our test, suppose that $Per(\Omega) = \rho'$ where $\rho' > \rho$. In that case,

(38)

$$\mathbb{P}_{H_A} \left(l_n \leq Z_\alpha \right) = \mathbb{P}_{H_A} \left(\sqrt{\frac{n\varepsilon_n}{4C_d\rho}} (\sigma_d \rho - \operatorname{Per}_{n,\varepsilon_n}(\Omega)) \geq -Z_\alpha \right) \\ = \mathbb{P}_{H_A} \left(\sqrt{\frac{n\varepsilon_n}{4C_d\rho}} (\operatorname{Per}_{\varepsilon_n}(\Omega) - \operatorname{Per}_{n,\varepsilon_n}(\Omega)) \geq -Z_\alpha + \sqrt{\frac{n\varepsilon_n}{4C_d\rho}} (\operatorname{Per}_{\varepsilon_n}(\Omega) - \sigma_d \rho) \right),$$

Now recall that $\lim_{n\to\infty} \operatorname{Per}_{\varepsilon_n}(\Omega) = \sigma_d \operatorname{Per}(\Omega) = \sigma_d \rho' > \sigma_d \rho$. In particular, we deduce that

$$\lim_{n \to \infty} \sqrt{\frac{n\varepsilon_n}{4C_d\rho}} (\operatorname{Per}_{\varepsilon_n}(\Omega) - \sigma_d \rho) = +\infty.$$

Thus, for large enough n,

$$-Z_{\alpha} + \sqrt{\frac{n\varepsilon_n}{4C_d\rho}} (\operatorname{Per}_{\varepsilon_n}(\Omega) - \sigma_d \rho) \ge \frac{1}{2} \sqrt{\frac{n\varepsilon_n}{4C_d\rho}} (\sigma_d \rho' - \sigma_d \rho).$$

Hence, for large enough n,

$$\mathbb{P}_{H_A}\left(l_n \leq Z_\alpha\right) \leq \mathbb{P}_{H_A}\left(\sqrt{\frac{n\varepsilon_n}{4C_d\rho}}(\operatorname{Per}_{\varepsilon_n}(\Omega) - \operatorname{Per}_{n,\varepsilon_n}(\Omega)) \geq \frac{1}{2}\sqrt{\frac{n\varepsilon_n}{4C_d\rho}}(\sigma_d\rho' - \sigma_d\rho).\right)$$

Using the moment estimates from Theorem 1.1, and Markov's inequality, we deduce that

$$\mathbb{P}_{H_A}(l_n \le Z_\alpha) = O\left(\frac{1}{\sqrt{n\varepsilon_n}}\right).$$

That is, the type II error is of order $\frac{1}{\sqrt{n\varepsilon_n}}$.

APPENDIX A. PROOF OF LEMMA 1.4

Since $\Omega \subset D$ and Ω has smooth boundary the relative perimeter of Ω with respect to D in the generalized sense (2) simply corresponds to the usual perimeter of $\partial\Omega$ in the sense that

$$\operatorname{Per}(\Omega) = \int_{\partial \Omega} \, \mathrm{d}\mathcal{H}^{d-1} = \mathcal{H}^{d-1}(\partial \Omega).$$

Additionally, for all $\varepsilon \leq \delta := \operatorname{dist}(\Omega, \partial D)$ we have that

$$\operatorname{Per}_{\varepsilon}(\Omega) = \frac{2}{\varepsilon^{d+1}} \int_{\Omega} |B_d(x,\varepsilon) \cap \Omega^c| \, \mathrm{d}x,$$

where $B_d(x, r)$ denotes the ball of radius r in \mathbb{R}^d centered at x and Ω^c denotes the complement of Ω in all of space. Moreover, since $\partial\Omega$ is a compact smooth manifold, we can assume without the loss of generality (by taking ε small enough) that for every $x \in T_{\varepsilon}$ there is a unique point P(x) in $\partial\Omega$ closest to x. Furthermore, we can assume that the map P is smooth. We may further write

$$\operatorname{Per}_{\varepsilon}(\Omega) = \frac{2}{\varepsilon^{d+1}} \int_{T_{\varepsilon}^{-}} |B_d(x,\varepsilon) \cap \Omega^c| \, \mathrm{d}x,$$

where T_{ε}^{-} is defined in (28). This reformulation makes it natural to write the previous integral as an iterated integral; the outer integral is taken over the manifold $\partial\Omega$ and the inner integral is taken along the normal line to $\partial\Omega$ at an arbitrary point x along the boundary.

To make this idea precise, we first let $\mathbf{N}(x)$ denote the outer unit normal to $\partial\Omega$ at $x \in \partial\Omega$ and then consider the transformation $(x,t) \in \partial\Omega \times (0,1) \mapsto x - t\varepsilon \mathbf{N}(x)$ for all ε sufficiently small. The Jacobian of this transformation equals $\varepsilon \det(I + t\varepsilon \mathbf{S}_x)$, where \mathbf{S}_x denotes the shape operator (or second fundamental form) of $\partial\Omega$ at x, see [13] for instance. For all ε sufficiently small, we may therefore conclude that

$$\frac{1}{\varepsilon} \int_{T_{\varepsilon}^{-}} |B_d(x,\varepsilon) \cap \Omega^c| \, \mathrm{d}x = \int_{\partial \Omega} \left(\int_0^1 |B_d(x-t\varepsilon \mathbf{N}(x),\varepsilon) \cap \Omega^c| \det(I+t\varepsilon \mathbf{S}_x) \, \mathrm{d}t \right) \, \mathrm{d}\mathcal{H}^{d-1}(x).$$

As a consequence, we also have that

(39)
$$\operatorname{Per}_{\varepsilon}(\Omega) = \frac{2}{\varepsilon^d} \int_{\partial\Omega} \left(\int_0^1 |B_d(x - t\varepsilon \mathbf{N}(x), \varepsilon) \cap \Omega^c| \det(I + t\varepsilon \mathbf{S}_x) \, \mathrm{d}t \right) \, \mathrm{d}\mathcal{H}^{d-1}(x).$$

With the expression (39) in hand, we may now proceed to establish (8) by expanding $\operatorname{Per}_{\varepsilon}(\Omega)$ in terms of ε and appealing to some elementary computations that show that the first order term in ε vanishes.

For a fixed $x \in \partial \Omega$, we first wish to understand the behavior of the function

$$g_x(\varepsilon) := \frac{1}{\varepsilon^d} \left(\int_0^1 |B_d(x - t\varepsilon \mathbf{N}(x), \varepsilon) \cap \Omega^c| \det(I + t\varepsilon \mathbf{S}_x) \, \mathrm{d}t \right)$$

for ε in a neighborhood of zero. Without loss of generality, we may assume that x = 0, that $\mathbf{N}(x) = e_d$ and that around x the boundary $\partial\Omega$ coincides with the graph $\hat{x} = (x_1, \ldots, x_{d-1}) \mapsto (\hat{x}, f(\hat{x})) \in \mathbb{R}^d$ of a smooth function $f(\hat{x})$ that satisfies both f(0) = 0 and $\nabla f(0) = 0$ simultaneously. By symmetry of the shape operator \mathbf{S}_x , there exists an orthonormal basis for \mathbb{R}^{d-1} (where we identify \mathbb{R}^{d-1} with the hyperplane $\{(\hat{x}, x_d) : x_d = 0\}$) consisting of eigenvectors of the shape operator. We let v_1, \ldots, v_{d-1} denote the eigenvectors of \mathbf{S}_x and $\kappa_1, \ldots, \kappa_{d-1}$ the corresponding eigenvalues (also known as principal curvatures). In particular, whenever $\|\hat{x}\| \leq \varepsilon$ we have that

(40)
$$f(\hat{x}) = \frac{1}{2} \sum_{i=1}^{d-1} \kappa_i \langle \hat{x}, v_i \rangle^2 + O(\varepsilon^3)$$

where curvatures $\kappa_i = \kappa_i(x)$ and the $O(\varepsilon^3)$ error term can be uniformly bounded.

With these reductions in place, we first define $u(\hat{y}):=\sqrt{\varepsilon^2-\|\hat{y}\|^2}$ and then let

$$h(\hat{y}, t; \varepsilon) := \begin{cases} 2u(\hat{y}) & \text{if } f(\hat{y}) + \varepsilon t < -u(\hat{y}) \\ u(\hat{y}) - \varepsilon t - f(\hat{y}) & \text{if } -u(\hat{y}) \le f(\hat{y}) + \varepsilon t \le u(\hat{y}) \\ 0 & \text{otherwise.} \end{cases}$$

A direct calculation then shows that

(41)
$$|B_d(x - t\varepsilon \mathbf{N}(x), \varepsilon) \cap \Omega^c| = \int_{B_{d-1}(0,\varepsilon)} h(\hat{y}, t; \varepsilon) \, \mathrm{d}\hat{y}$$

and an application of (40) shows that $h(\hat{y},t;\varepsilon)=2\sqrt{\varepsilon^2-\|\hat{y}\|^2}$ only if

$$\|\hat{y}\|^2 = \varepsilon^2 - O(\varepsilon^4)$$
 and $u(\hat{y}) = O(\varepsilon^2)$

It therefore follows that

$$\int_{B_{d-1}(0,\varepsilon)\cap\{f(\hat{y})+\varepsilon t<-u(\hat{y})\}} h(\hat{y},t;\varepsilon) \, \mathrm{d}\hat{y} \le O(\varepsilon^2) \int_{B_{d-1}(0,\varepsilon)\cap\{\|\hat{y}\|\ge\sqrt{\varepsilon^2-O(\varepsilon^4)}\}} \, \mathrm{d}\hat{y} = O(\varepsilon^{d+3}).$$

We then let A_t^{ε} denote the set $A_t^{\varepsilon} := \{\hat{y} \in B_{d-1}(0, \varepsilon) : -u(\hat{y}) \le f(\hat{y}) + \varepsilon t \le u(\hat{y})\}$ and use the previous estimate in (41) to uncover

(42)
$$|B_d(x - t\varepsilon \mathbf{N}(x), \varepsilon) \cap \Omega^c| = \int_{B_{d-1}(0,\varepsilon) \cap A_t^\varepsilon} (u(\hat{y}) - \varepsilon t - f(\hat{y})) \, \mathrm{d}\hat{y} + O(\varepsilon^{d+3}).$$

We may then note that

$$\det(I + \varepsilon t \mathbf{S}_x) = (1 + t\varepsilon \kappa_1) \dots (1 + t\varepsilon \kappa_{d-1}) = 1 + t\varepsilon H_x + O(\varepsilon^2),$$

where $H_x := \sum_{i=1}^{d-1} \kappa_i$ represents the mean curvature. Using this fact in (42) then yields

$$g_x(\varepsilon) = \frac{1}{\varepsilon^d} \int_0^1 \left(\int_{B_{d-1}(0,\varepsilon) \cap A_t^\varepsilon} u(\hat{y}) - \varepsilon t - f(\hat{y}) \, \mathrm{d}\hat{y} \right) (1 + t\varepsilon H_x) \, \mathrm{d}t + O(\varepsilon^2).$$

Now let $f^{\varepsilon}(z) := \frac{1}{\varepsilon} f(\varepsilon z)$ and define the corresponding subset C_t^{ε} of $(0,1) \times B_{d-1}(0,1)$ as

$$C_t^{\varepsilon} := \left\{ (t, z) \in (0, 1) \times B_{d-1}(0, 1) : -\sqrt{1 - \|z\|^2} \le f^{\varepsilon}(z) + t \le \sqrt{1 - \|z\|^2} \right\},\$$

then make the change of variables $\hat{y} = \varepsilon z$ to see that

$$g_x(\varepsilon) = \int_{C_t^{\varepsilon}} \left(\sqrt{1 - \|z\|^2} - t - f^{\varepsilon}(z) \right) (1 + t\varepsilon H_x) \, \mathrm{d}z \mathrm{d}t + O(\varepsilon^2).$$

Recalling (40) shows that

(43)
$$f^{\varepsilon}(z) = \frac{\varepsilon}{2} \sum_{i=1}^{d-1} \kappa_i \langle z, v_i \rangle^2 + O(\varepsilon^2),$$

which then allows us to obtain an expansion of $g_x(\varepsilon)$ in terms of ε according to the relation

(44)
$$g_x(\varepsilon) = \int_{C_t^{\varepsilon}} \left(\sqrt{1 - \|z\|^2} - t\right) dt dz + \varepsilon \int_{C_t^{\varepsilon}} \left(tH_x(\sqrt{1 - \|z\|^2} - t) - \frac{1}{2}\sum_{i=1}^{d-1} \kappa_i \langle z, v_i \rangle^2\right) dt dz + O(\varepsilon^2).$$

The bias estimate (8) then directly follows after computing each of these terms individually.

We begin by considering the first term in the expansion, i.e.

$$\mathbf{I} := \int_{C_t^{\varepsilon}} \left(\sqrt{1 - \|z\|^2} - t \right) \, \mathrm{d}t \mathrm{d}z.$$

Given $\varepsilon > 0$ and $z \in B_{d-1}(0,1)$ define $c(z) := \max\{-\sqrt{1-\|z\|^2} - f^{\varepsilon}(z), 0\}$ and $C(z) := \min\{\sqrt{1-\|z\|^2} - f^{\varepsilon}(z), 1\}$, so that we may easily write

$$\mathbf{I} = \int_{B_{d-1}(0,1)} (C(z) - c(z)) \left(\sqrt{1 - \|z\|^2} - \frac{C(z) + c(z)}{2} \right) \, \mathrm{d}z.$$

As the set where $c(z) \neq 0$ has measure at most $O(\varepsilon^2)$, we easily conclude that

$$\mathbf{I} = \int_{B_{d-1}(0,1)} C(z) \left(\sqrt{1 - \|z\|^2} - \frac{C(z)}{2} \right) \, \mathrm{d}z + O(\varepsilon^2).$$

If C(z) = 1 then $\sqrt{1 - \|z\|^2} - \frac{C(z)}{2} = \frac{1}{2}(1 - \|z\|^2) + O(\varepsilon^2)$ as well. In any case, it follows that

(45)
$$\mathbf{I} = \frac{1}{2} \int_{B_{d-1}(0,1)} (1 - ||z||^2) \, \mathrm{d}z + O(\varepsilon^2) = \frac{\sigma_d}{2} + O(\varepsilon^2).$$

We now proceed to compute the second term in the expansion

$$\begin{aligned} \text{II} &:= H_x \int_{C_t^{\varepsilon}} \left(t \sqrt{1 - \|z\|^2} - t^2 \right) \, \mathrm{d}t \mathrm{d}z \\ &= H_x \int_{B_{d-1}(0,1)} C^2(z) \left(\frac{\sqrt{1 - \|z\|}}{2} - \frac{C(z)}{3} \right) \, \mathrm{d}z + O(\varepsilon^2) \end{aligned}$$

and the third term in the expansion

$$\text{III} := \frac{1}{2} \sum_{i=1}^{d-1} \kappa_i \int_{C_t^{\varepsilon}} \langle z, v_i \rangle^2 \, \mathrm{d}t \mathrm{d}z = \frac{1}{2} \sum_{i=1}^{d-1} \kappa_i \int_{B_{d-1}(0,1)} \langle z, v_i \rangle^2 C(z) \, \mathrm{d}z + O(\varepsilon^2)$$

in a similar fashion. We always have $C(z) = \sqrt{1 - ||z||^2} + O(\varepsilon)$, so that (46)

$$II = \frac{H_x}{6} \int_{B_{d-1}(0,1)} (1 - ||z||^2)^{3/2} dz + O(\varepsilon) = \frac{H_x \operatorname{vol}(\mathcal{S}^{d-2})}{6} \int_0^1 (1 - r^2)^{3/2} r^{d-2} dr + O(\varepsilon)$$

The third term follows similarly by appealing to spherical coordinates, in that we have

$$III = \frac{1}{2} \sum_{i=1}^{d-1} \kappa_i \int_{B_{d-1}(0,1)} \sqrt{1 - \|z\|^2} \langle z, v_i \rangle^2 \, \mathrm{d}z + O(\varepsilon)$$
$$= \frac{H_x \mathrm{vol}(\mathcal{S}^{d-2})}{2(d-1)} \int_0^1 \sqrt{1 - r^2} r^d \, \mathrm{d}r + O(\varepsilon) = \mathrm{II} + O(\varepsilon)$$

thanks to an integration by parts in the final term. We therefore have that $I = \sigma_d/2 + O(\varepsilon^2)$ and $II - III = O(\varepsilon)$, so that $g_x(\varepsilon) = \sigma_d/2 + O(\varepsilon^2)$ and

$$\operatorname{Per}_{\varepsilon}(\Omega) = 2 \int_{\partial \Omega} g_x(\varepsilon) \, \mathrm{d}\mathcal{H}^{d-1} = \sigma_d \operatorname{Per}(\Omega) + O(\varepsilon^2)$$

as desired.

We may also show that when Ω is a fixed ball, say $\Omega = B_d(x_c, \frac{1}{3})$ for $x_c \in \mathbb{R}^d$ the center point of D, that the absolute value of the difference between $\operatorname{Per}_{\varepsilon}(\Omega)$ and $\sigma_d \operatorname{Per}(\Omega)$ remains bounded from below by $c\varepsilon^2$ for c > 0 some positive constant. The proof proceeds similarly to the proof of the bias estimate above. In particular, this shows that the bound in Lemma 1.4 is optimal in terms of scaling for general sets with smooth boundary.

APPENDIX B. PROOF OF LEMMA 2.1

The proof follows the same argument used in the proof of Theorem 4.1 in [10] or Theorem 6.2 in [1]. We assume that $dist(\Omega, \partial D) > 0$. Such assumption implies that the perimeter of Ω with respect to D, that is $Per(\Omega)$ defined in (2), is equal to the perimeter of Ω with respect to \mathbb{R}^d . We remark that a slight modification of the argument we present below proves the result in the general case and hence we omit the details (see for example the proof of Theorem 4.1 in [10]).

First we prove that for any function $u : \mathbb{R}^d \to [0,1]$ with $u \in W^{1,1}(\mathbb{R}^d) \cap C^{\infty}(\mathbb{R}^d)$ and for all $\varepsilon > 0$ we have

(47)
$$\int_{\mathbb{R}^d} \left(\int_{\mathbb{R}^d} \frac{\mathbf{1}_{\|x-y\| \le \varepsilon}}{\varepsilon^{d+1}} |u(y) - u(x)| \mathrm{d}y \right)^p \mathrm{d}x \le \frac{\alpha_d^{p-1} \sigma_d}{\varepsilon^{p-1}} \int_{\mathbb{R}^d} \|\nabla u(x)\| \mathrm{d}x,$$

Inequality (47) follows from

$$\begin{split} \int_{\mathbb{R}^d} \left(\int_{\mathbb{R}^d} \frac{\mathbf{1}_{\|x-y\| \leq \varepsilon}}{\varepsilon^{d+1}} |u(y) - u(x)| \mathrm{d}y \right)^p \mathrm{d}x &= \frac{1}{\varepsilon^p} \int_{\mathbb{R}^d} \left(\int_{B_d(0,1)} |u(x+\varepsilon h) - u(x)| \mathrm{d}h \right)^p \mathrm{d}x \\ &\leq \frac{\alpha_d^{p-1}}{\varepsilon^p} \int_{\mathbb{R}^d} \int_{B_d(0,1)} |u(x+\varepsilon h) - u(x)|^p \mathrm{d}h \, \mathrm{d}x \\ &\leq \frac{\alpha_d^{p-1}}{\varepsilon^p} \int_{\mathbb{R}^d} \int_{B_d(0,1)} |u(x+\varepsilon h) - u(x)| \mathrm{d}h \, \mathrm{d}x \\ &= \frac{\alpha_d^{p-1}}{\varepsilon^{p-1}} \int_{\mathbb{R}^d} \int_{B_d(0,1)} \left\| \int_0^1 \nabla u(x+t\varepsilon h) \cdot h \mathrm{d}t \right\| \mathrm{d}h \, \mathrm{d}x \\ &\leq \frac{\alpha_d^{p-1}}{\varepsilon^{p-1}} \int_{\mathbb{R}^d} \int_{B_d(0,1)} \int_0^1 \|\nabla u(x+t\varepsilon h) \cdot h\| \mathrm{d}t \, \mathrm{d}h \, \mathrm{d}x \\ &= \frac{\alpha_d^{p-1}}{\varepsilon^{p-1}} \int_0^1 \int_{B_d(0,1)} \int_{\mathbb{R}^d} \|\nabla u(x) \cdot h\| \mathrm{d}x \, \mathrm{d}h \, \mathrm{d}t \\ &= \frac{\alpha_d^{p-1}}{\varepsilon^{p-1}} \int_0^1 \int_{\mathbb{R}^d} \|\nabla u(x)\| \int_{B_d(0,1)} \left\| \frac{\nabla u(x)}{\|\nabla u(x)\|} \cdot h \right\| \, \mathrm{d}h \, \mathrm{d}x \, \mathrm{d}t \\ &= \frac{\alpha_d^{p-1}}{\varepsilon^{p-1}} \int_0^1 \int_{\mathbb{R}^d} \|\nabla u(x)\| \, dx \end{split}$$

where in the first equation we used the change of variables $h = \frac{x-y}{\varepsilon}$, in the first inequality we used Jensen's inequality and in the second inequality the fact that u takes values in [0, 1].

Now, for any set $\Omega \subseteq D$ as in the statement, we can find a sequence of functions $\{u_k\}_{k\in\mathbb{N}}$ with $u_k: \mathbb{R}^d \to [0,1]$, $u_k \in W^{1,1}(\mathbb{R}^d) \cap C^{\infty}(\mathbb{R}^d)$ and such that

(48)
$$u_k \xrightarrow{L^1(\mathbb{R}^d)} \mathbf{1}_{\Omega}, \quad \lim_{k \to \infty} \int_{\mathbb{R}^d} \|\nabla u_k(x)\| \mathrm{d}x = \operatorname{Per}(\Omega)$$

Such sequence can be obtained for example with the aid of standard mollifiers (see Theorem 13.9 in [22]). It follows from (47) and from (48) that

$$\int_{\mathbb{R}^d} \left(\int_{\mathbb{R}^d} \frac{\mathbf{1}_{\|x-y\| \le \varepsilon}}{\varepsilon^{d+1}} |\mathbf{1}_{\Omega}(y) - \mathbf{1}_{\Omega}(x)| \mathrm{d}y \right)^p \mathrm{d}x \le \frac{\alpha_d^{p-1} \sigma_d}{\varepsilon^{p-1}} \operatorname{Per}(\Omega).$$

Finally, notice that

$$\int_{D} \bar{\phi}^{p}_{\varepsilon}(x) \mathrm{d}x = \int_{\mathbb{R}^{d}} \left(\int_{\mathbb{R}^{d}} \frac{\mathbf{1}_{\|x-y\| \leq \varepsilon}}{\varepsilon^{d+1}} |\mathbf{1}_{\Omega}(y) - \mathbf{1}_{\Omega}(x)| \mathrm{d}y \right)^{p} \mathrm{d}x \leq \frac{\alpha_{d}^{p-1} \sigma_{d}}{\varepsilon^{p-1}} \operatorname{Per}(\Omega)$$

APPENDIX C. PROOF OF LEMMA 3.1

The proof is based on similar computations to the ones in Appendix A and thus we simply highlight the main ideas. First of all note that

$$\operatorname{Var}\left(g_{n,1}(X_1)\right) = \int_D \bar{\phi}_{\varepsilon_n}^2(x) dx - \left(\operatorname{Per}_{\varepsilon_n}(\Omega)\right)^2.$$

Since $\operatorname{Per}_{\varepsilon_n}(\Omega) = O(1)$, our task reduces to computing the integral in the above expression. It follows from (29) that for all ε_n small enough (so that $T_{\varepsilon_n} \subseteq D$),

$$\int_{D} \bar{\phi}_{\varepsilon_{n}}^{2}(x) \mathrm{d}x = \frac{1}{\varepsilon_{n}^{2(d+1)}} \int_{T_{\varepsilon_{n}}^{-}} |B_{d}(x,\varepsilon_{n}) \cap \Omega^{c}|^{2} \mathrm{d}x + \frac{1}{\varepsilon_{n}^{2(d+1)}} \int_{T_{\varepsilon_{n}}^{+}} |B_{d}(x,\varepsilon_{n}) \cap \Omega|^{2} \mathrm{d}x$$

We compute the first of the integrals from the above expression. As in the proof of Lemma 1.4, we write

$$\frac{1}{\varepsilon_n^{2(d+1)}} \int_{T_{\varepsilon_n}} |B_d(x,\varepsilon_n) \cap \Omega^c|^2 \, \mathrm{d}x$$

= $\frac{1}{\varepsilon_n} \int_{\partial\Omega} \left(\frac{1}{\varepsilon_n^{2d}} \int_0^1 |B_d(x-t\varepsilon_n \mathbf{N}(x),\varepsilon_n) \cap \Omega^c|^2 \det(I+t\varepsilon_n \mathbf{S}_x) \, \mathrm{d}t \right) \, \mathrm{d}\mathcal{H}^{d-1}(x).$

For $x \in \partial \Omega$ we study the expression

$$h_{\varepsilon_n}(x) := \int_0^1 \frac{|B_d(x - t\varepsilon_n \mathbf{N}(x), \varepsilon_n) \cap \Omega^c|^2}{\varepsilon_n^{2d}} \det(I + t\varepsilon_n \mathbf{S}_x) \mathrm{d}t$$

Note that det $(I + t\varepsilon_n \mathbf{S}_x) = (1 + t\varepsilon_n \kappa_1) \cdots (1 + t\varepsilon_n \kappa_{d-1}) = 1 + O(\varepsilon_n)$, where $\kappa_1, \ldots, \kappa_{d-1}$ are the principal curvatures (eigenvalues of the shape operator S_x). Hence,

$$h_{\varepsilon_n}(x) = \int_0^1 \frac{|B_d(x - t\varepsilon_n \mathbf{N}(x), \varepsilon_n) \cap \Omega^c|^2}{\varepsilon_n^{2d}} dt + O(\varepsilon_n).$$

Without loss of generality, we may assume that x = 0, that $N(x) = e_d$ and that around x the boundary $\partial\Omega$ coincides with the graph $\hat{x} = (x_1, \dots, x_{d-1}) \mapsto (\hat{x}, f(\hat{x})) \in \mathbb{R}^d$ of a smooth function $f(\hat{x})$ that satisfies both f(0) = 0 and $\nabla f(0) = 0$ simultaneously, and we denote by v_1, \ldots, v_{d-1} the eigenvectors of \mathbf{S}_x (just as in Appendix A). Then, $f: B_{d-1}(0, \varepsilon_n) \to \mathbb{R}$, satisfies $f(\hat{x}) = \sum_{i=1}^{d-1} \kappa_i \langle \hat{x}, v_i \rangle^2 + O(\varepsilon_n^3)$. Now for fixed $t \in [0, 1]$ we define H_{1-t} to be the hyperplane

$$H_{1-t} := \{ x = (\hat{x}, x_d) : x_d \ge 1 - t \},\$$

and we let

$$A_t := |B_d(0,1) \cap H_{1-t}|.$$

With these definitions we can write

(49)
$$|B_{d}(x - t\varepsilon_{n}\mathbf{N}(x), \varepsilon_{n}) \cap \Omega^{c}| = \varepsilon_{n}^{d}A_{t} + \sum_{i=1}^{d-1} \kappa_{i} \int_{B_{d-1}(0,\varepsilon_{n}\sqrt{1-t})} \langle \hat{x}, v_{i} \rangle^{2} d\hat{x} + O(\varepsilon_{n}^{d+2})$$
$$= \varepsilon_{n}^{d}A_{t} + \frac{\sum_{i=1}^{d-1} \kappa_{i}}{d-1} \int_{B_{d-1}(0,\varepsilon_{n}\sqrt{1-t})} \|\hat{x}\|^{2} d\hat{x} + O(\varepsilon_{n}^{d+2})$$
$$= \varepsilon_{n}^{d}A_{t} + O(\varepsilon_{n}^{d+1}),$$

where the second equality holds due to symmetry, and where the last equality follows after computing $\int_{B_{d-1}(0,\varepsilon_n\sqrt{1-t})} \|\hat{x}\|^2 d\hat{x}$ using polar coordinates. From the above, it follows that

$$\frac{|B_d(x - t\varepsilon_n \mathbf{N}(x), \varepsilon_n) \cap \Omega^c|^2}{\varepsilon_n^{2d}} = A_t^2 + O(\varepsilon_n).$$

We conclude that

$$h_{\varepsilon_n}(x) = \int_0^1 A_t^2 dt + O(\varepsilon_n)$$

Therefore,

$$\frac{1}{\varepsilon_n^{2(d+1)}} \int_{T_{\varepsilon_n}^-} |B_d(x,\varepsilon_n) \cap \Omega^c|^2 \, \mathrm{d}x = \frac{1}{\varepsilon_n} \int_0^1 A_t^2 \mathrm{d}t \operatorname{Per}(\Omega) + O(1).$$

Analogously, we can obtain a similar expression for $\frac{1}{\varepsilon_n^{2(d+1)}} \int_{T_{\varepsilon_n}^+} |B_d(x,\varepsilon_n) \cap \Omega|^2 dx$ and from this we deduce (35).

Acknowledgements. DS and NGT are grateful to NSF for its support (grants DMS-1211760 and DMS-1516677). JvB was supported by NSF grant DMS 1312344/DMS 1521138. The authors are grateful to ICERM, where part of the research was done during the research cluster: *Geometric analysis methods for graph algorithms*. Furthermore they are grateful to Ki-Net (NSF Research Network Grant RNMS11-07444) for opportunities provided. The authors would like to thank the Center for Nonlinear Analysis of the Carnegie Mellon University for its support.

REFERENCES

- G. ALBERTI AND G. BELLETTINI, A non-local anisotropic model for phase transitions: asymptotic behaviour of rescaled energies, European J. Appl. Math., 9 (1998), pp. 261–284.
- [2] L. AMBROSIO, N. FUSCO, AND D. PALLARA, Functions of bounded variation and free discontinuity problems, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, 2000.
- [3] E. ARIAS-CASTRO, B. PELLETIER, AND P. PUDLO, *The normalized graph cut and Cheeger constant: from discrete to continuous*, Adv. in Appl. Probab., 44 (2012), pp. 907–937.
- [4] I. ARMENDÁRIZ, A. CUEVAS, AND R. FRAIMAN, Nonparametric estimation of boundary measures and related functionals: asymptotic results, Adv. in Appl. Probab., 41 (2009), pp. 311–322.
- [5] M. BELKIN, H. NARAYANAN, AND P. NIYOGI, *Heat flow and a faster algorithm to compute the surface area of a convex body*, in FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06), Washington, DC, USA, 2006, IEEE Computer Society, pp. 47–56.
- [6] S. BERNSTEIN, On a modification of Chebyshev's inequality and of the error formula of Laplace, Ann. Sci. Inst. Sav. Ukraine, Sect. Math, 1 (1924), pp. 38–49.
- [7] X. BRESSON, T. LAURENT, D. UMINSKY, AND J. H. VON BRECHT, Convergence and energy landscape for cheeger cut clustering, in Advances in Neural Information Processing Systems (NIPS), P. L. Bartlett, F. C. N. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, eds., 2012, pp. 1394–1402.
- [8] A. CUEVAS, R. FRAIMAN, AND L. GYÖRFI, *Towards a universally consistent estimator of the Minkowski content*, ESAIM Probab. Stat., 17 (2013), pp. 359–369.
- [9] A. CUEVAS, R. FRAIMAN, AND A. RODRÍGUEZ-CASAL, A nonparametric approach to the estimation of lengths and surface areas, Ann. Statist., 35 (2007), pp. 1031–1051.
- [10] N. GARCÍA TRILLOS AND D. SLEPČEV, Continuum limit of total variation on point clouds, Archive for Rational Mechanics and Analysis, (2015), pp. 1–49. online first.
- [11] N. GARCÍA TRILLOS, D. SLEPČEV, J. VON BRECHT, T. LAURENT, AND X. BRESSON, Consistency of Cheeger and ratio graph cuts. to appear in Journal of Machine Learning Research, 2015.
- [12] E. GINÉ, R. LATAŁA, AND J. ZINN, *Exponential and moment inequalities for U-statistics*, in High dimensional probability, II (Seattle, WA, 1999), vol. 47 of Progr. Probab., Birkhäuser Boston, Boston, MA, 2000, pp. 13–38.
- [13] A. GRAY, *Tubes*, Progress in Mathematics, Birkhäuser Basel, 2004.
- [14] L. HAGEN AND A. KAHNG, New spectral methods for ratio cut partitioning and clustering., IEEE Trans. Computer-Aided Design, 11 (1992), pp. 1074 –1085.
- [15] M. HEIN AND S. SETZER, Beyond spectral clustering tight relaxations of balanced graph cuts, in Advances in Neural Information Processing Systems (NIPS), J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, eds., 2011, pp. 2366–2374.
- [16] W. HOEFFDING, Probability inequalities for sums of bounded random variables, Journal of the American statistical association, 58 (1963), pp. 13–30.
- [17] R. JIMÉNEZ AND J. E. YUKICH, Nonparametric estimation of surface integrals, Ann. Statist., 39 (2011), pp. 232–260.

- [18] R. KANNAN, S. VEMPALA, AND A. VETTA, On clusterings: Good, bad and spectral, Journal of the ACM (JACM), 51 (2004), pp. 497–515.
- [19] M. KEARNS AND D. RON, Testing problems with sublearning sample complexity, Journal of Computer and System Sciences, 61 (2000), pp. 428 – 456.
- [20] V. S. KOROLJUK AND Y. V. BOROVSKICH, *Theory of U-statistics*, vol. 273 of Mathematics and its Applications, Kluwer Academic Publishers Group, Dordrecht, 1994. Translated from the 1989 Russian original by P. V. Malyshev and D. V. Malyshev and revised by the authors.
- [21] P. KOTHARI, A. NAYYERI, R. O'DONNELL, AND C. WU, *Testing surface area*, in Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '14, SIAM, 2014, pp. 1204–1214.
- [22] G. LEONI, *A first course in Sobolev spaces*, vol. 105 of Graduate Studies in Mathematics, American Mathematical Society, Providence, RI, 2009.
- [23] J. NEEMAN, *Testing surface area with arbitrary accuracy*, in Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC '14, New York, NY, USA, 2014, ACM, pp. 393–397.
- [24] B. PATEIRO-LÓPEZ AND A. RODRÍGUEZ-CASAL, Length and surface area estimation under smoothness restrictions, Adv. in Appl. Probab., 40 (2008), pp. 348–358.
- [25] J. SHI AND J. MALIK, Normalized cuts and image segmentation, Pattern Analysis and Machine Intelligence, IEEE Transactions on, 22 (2000), pp. 888–905.
- [26] A. SZLAM AND X. BRESSON, *Total variation and cheeger cuts.*, in ICML, J. Faenkranz and T. Joachims, eds., Omnipress, 2010, pp. 1039–1046.
- [27] U. VON LUXBURG, A tutorial on spectral clustering, Statistics and computing, 17 (2007), pp. 395–416.
- [28] Y.-C. WEI AND C.-K. CHENG, Towards efficient hierarchical designs by ratio cut partitioning, in Computer-Aided Design, 1989. ICCAD-89. Digest of Technical Papers., 1989 IEEE International Conference on, IEEE, 1989, pp. 298–301.

DIVISION OF APPLIED MATHEMATICS, BROWN UNIVERSITY, PROVIDENCE, RI, 02912, USA. *E-mail address*: nicolas_garcia_trillos@brown.edu

DEPARTMENT OF MATHEMATICAL SCIENCES, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PA, 15213, USA.

E-mail address: slepcev@math.cmu.edu

DEPARTMENT OF MATHEMATICS AND STATISTICS, CALIFORNIA STATE UNIVERSITY, LONG BEACH, LONG BEACH, CA 90840, USA.

E-mail address: James.vonBrecht@csulb.edu