

# Spectral clustering based on local linear approximations

Ery Arias-Castro\*

*Department of Mathematics, University of California, San Diego*

*e-mail: [eariasca@ucsd.edu](mailto:eariasca@ucsd.edu)*

Guangliang Chen

*Department of Mathematics, Duke University*

*e-mail: [glchen@math.duke.edu](mailto:glchen@math.duke.edu)*

and

Gilad Lerman

*Department of Mathematics, University of Minnesota, Twin Cities*

*e-mail: [lerman@umn.edu](mailto:lerman@umn.edu)*

**Abstract:** In the context of clustering, we assume a generative model where each cluster is the result of sampling points in the neighborhood of an embedded smooth surface; the sample may be contaminated with outliers, which are modeled as points sampled in space away from the clusters. We consider a prototype for a higher-order spectral clustering method based on the residual from a local linear approximation. We obtain theoretical guarantees for this algorithm and show that, in terms of both separation and robustness to outliers, it outperforms the standard spectral clustering algorithm (based on pairwise distances) of Ng, Jordan and Weiss (NIPS '01). The optimal choice for some of the tuning parameters depends on the dimension and thickness of the clusters. We provide estimators that come close enough for our theoretical purposes. We also discuss the cases of clusters of mixed dimensions and of clusters that are generated from smoother surfaces. In our experiments, this algorithm is shown to outperform pairwise spectral clustering on both simulated and real data.

**AMS 2000 subject classifications:** 62H30, 62G20, 68T10.

**Keywords and phrases:** Spectral clustering, higher-order affinities, local linear approximation, local polynomial approximation, detection of clusters in point clouds, dimension estimation, nearest-neighbor search.

Received August 2010.

## 1. Introduction

In a number of modern applications, the data appear to cluster near some low-dimensional structures. In the particular setting of manifold learning [7, 17, 22, 47, 51], the data are assumed to lie near manifolds embedded in Euclidean space. When multiple manifolds are present, the foremost task is separating them,

---

\*Corresponding author.

meaning the recovery of the different components of the data associated with the different manifolds. Manifold clustering naturally occurs in the human visual cortex, which excels at grouping points into clusters of various shapes [21, 41]. It is also relevant for a number of modern applications. For example, in cosmology, galaxies seem to cluster forming various geometric structures such as one-dimensional filaments and two-dimensional walls [39, 52]. In motion segmentation, feature vectors extracted from moving objects and tracked along different views cluster along affine or algebraic surfaces [10, 23, 35, 53]. In face recognition, images of faces in fixed pose under varying illumination conditions cluster near low-dimensional affine subspaces [6, 19, 31], or along low-dimensional manifolds when introducing additional poses and camera views.

In the last few years several algorithms for multi-manifold clustering were introduced; we discuss them individually in Section 1.3.3. We focus here on spectral clustering methods, and in particular, study a prototypical multiway method relying on local linear approximations, with precursors appearing in [1, 2, 12, 27, 48]. We refer to this method as Higher-Order Spectral Clustering (HOSC). We establish theoretical guarantees for this method within a standard mathematical framework for multi-manifold clustering. Compared with all other algorithms we are aware of, HOSC is able to separate clusters that are much closer together; equivalently, HOSC is accurate under much lower sampling rate than any other algorithm we know of. Roughly speaking, a typical algorithm for multi-manifold clustering relies on local characteristics of the point cloud in a way that presupposes that all points, or at least the vast majority of the points, in a (small enough) neighborhood are from a single cluster, except in places like intersections of clusters. In contrast, though HOSC is also a local method, it can work with neighborhoods where two or more clusters coexist.

### 1.1. Higher-order spectral clustering (HOSC)

We introduce our higher-order spectral clustering algorithm in this section, tracing its origins to the spectral clustering algorithm of Ng et al. [42] and the spectral curvature clustering of Chen and Lerman [11, 12].

Spectral methods are based on building a neighborhood graph on the data points and partitioning the graph using its Laplacian [22, 34], which is closely related to the extraction of connected components. The version introduced by Ng et al. [42] is an emblematic example—we refer to this approach as SC. It uses an affinity based on pairwise distances. Given a scale parameter  $\epsilon > 0$  and a kernel  $\phi$ , define

$$\alpha(\mathbf{x}_1, \mathbf{x}_2) = \begin{cases} \phi(\|\mathbf{x}_1 - \mathbf{x}_2\|/\epsilon), & \mathbf{x}_1 \neq \mathbf{x}_2; \\ 0, & \mathbf{x}_1 = \mathbf{x}_2. \end{cases} \quad (1)$$

( $\|\cdot\|$  denotes the Euclidean norm.) Standard choices include the heat kernel  $\phi(s) = \exp(-s^2)$  and the simple kernel  $\phi(s) = \mathbf{1}_{\{|s|<1\}}$ . Let  $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^D$  denote the data points. SC starts by computing all pairwise affinities  $\mathbf{W} = (W_{ij})$ , with  $W_{ij} = \alpha(\mathbf{x}_i, \mathbf{x}_j)$ , for  $i, j = 1, \dots, N$ . It then computes the matrix

**Algorithm 1** Spectral Clustering (SC) [42]**Input:**

$\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$ : the data points  
 $\epsilon$ : the affinity scale  
 $K$ : the number of clusters

**Output:**

A partition of the data into  $K$  disjoint clusters

**Steps:**

- 1: Compute the affinity matrix  $\mathbf{W} = (W_{ij})$ , with  $W_{ij} = \alpha(\mathbf{x}_i, \mathbf{x}_j)$ .
- 2: Compute the  $\mathbf{Z} = (Z_{ij}) : Z_{ij} = W_{ij}/(D_i D_j)^{1/2}$ , where  $D_i = \sum_j W_{ij}$ .
- 3: Extract  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_K]$ , the top  $K$  eigenvectors of  $\mathbf{Z}$ .
- 4: Renormalize each *row* of  $\mathbf{U}$  to have unit norm, obtaining a matrix  $\mathbf{V}$ .
- 5: Apply  $K$ -means to the row vectors of  $\mathbf{V}$  in  $\mathbb{R}^K$  to find  $K$  clusters.
- 6: Accordingly group the original points into  $K$  disjoint clusters.

$\mathbf{Z} = (Z_{ij}) : Z_{ij} = W_{ij}/(D_i D_j)^{1/2}$ , where  $D_i = \sum_{1 \leq j \leq N} W_{ij}$  is the degree of the  $i$ th point in the graph with similarity matrix  $\mathbf{W}$ . Note that  $\mathbf{I} - \mathbf{Z}$  is the corresponding normalized Laplacian. Providing the algorithm with the number of clusters  $K$ , SC continues by extracting the top  $K$  eigenvectors of  $\mathbf{Z}$ , obtaining a matrix  $\mathbf{U} \in \mathbb{R}^{N \times K}$ , and after normalizing its rows, uses them to embed the data into  $\mathbb{R}^K$ . The algorithm concludes by applying  $K$ -means to the embedded points. See Algorithm 1 for a summary.

Spectral methods utilizing multiway affinities were proposed to better exploit additional structure present in the data. The spectral curvature clustering (SCC) algorithm of Chen and Lerman [11, 12] was designed for the case of hybrid linear modeling where the manifolds are assumed to be affine, a setting that arises in motion segmentation [35]. Assuming that the subspaces are all of dimension  $d$ —a parameter of the algorithm, SCC starts by computing the (polar) curvature of all  $(d+2)$ -tuples, creating an  $N^{\otimes(d+2)}$ -tensor. The tensor is then flattened into a matrix  $\mathbf{A}$  whose product with its transpose,  $\mathbf{W} = \mathbf{A}\mathbf{A}'$ , is used as an affinity matrix for the spectral algorithm SC. (In practice, the algorithm is randomized for computational tractability.) Kernel spectral curvature clustering (KSCC) [10] is a kernel version of SCC designed for the case of algebraic surfaces.

The SCC algorithm (and therefore KSCC) is not localized in space as it fits a parametric model that is global in nature. The method we study here may be seen as a localization of SCC, which is appropriate in our nonparametric setting since the manifolds resemble affine surfaces locally. This type of approach is mentioned in publications on affinity tensors [1, 2, 27, 48] and is studied here for the first time, to our knowledge. As discussed in Section 4, all reasonable variants have similar theoretical properties, so that we choose one of the simplest versions to ease the exposition. Concretely, we consider a multiway affinity that combines pairwise distances between nearest neighbors and the residual from the best  $d$ -dimensional local linear approximation. Formally, given a set of  $m \geq d+2$  points,  $\mathbf{x}_1, \dots, \mathbf{x}_m$ , define

$$\Lambda_d(\mathbf{x}_1, \dots, \mathbf{x}_m) = \min_{L \in \mathcal{A}_d} \max_{j=1, \dots, m} \text{dist}(\mathbf{x}_j, L), \quad (2)$$




FIG 1. The circle is of radius  $\epsilon/2$  and the band is of half-width  $\eta$ . Assuming we use the simple kernel, the  $m$ -tuple on the left has affinity  $\alpha_d$  equal to one, while the other two  $m$ -tuples have affinity equal to zero, the first one for having a diameter exceeding  $\epsilon$  and the second one for being ‘thicker’ than  $\eta$ .

where  $\text{dist}(\mathbf{x}, S) := \inf_{\mathbf{s} \in S} \|\mathbf{x} - \mathbf{s}\|$  for a subset  $S \subset \mathbb{R}^D$  and  $\mathcal{A}_d$  denotes the set of  $d$ -dimensional affine subspaces in  $\mathbb{R}^D$ . In other words,  $\Lambda_d(\mathbf{x}_1, \dots, \mathbf{x}_m)$  is the width of the thinnest tube (or band) around a  $d$ -dimensional affine subspace that contains  $\mathbf{x}_1, \dots, \mathbf{x}_m$ . (In our implementation, we use the mean-square error; see Section 3.) Given scale parameters  $\epsilon > \eta > 0$  and a kernel function  $\phi$ , define the following affinity:  $\alpha_d(\mathbf{x}_1, \dots, \mathbf{x}_m) = 0$  if  $\mathbf{x}_1, \dots, \mathbf{x}_m$  are not distinct; otherwise:

$$\alpha_d(\mathbf{x}_1, \dots, \mathbf{x}_m) = \phi\left(\frac{\text{diam}(\mathbf{x}_1, \dots, \mathbf{x}_m)}{\epsilon}\right) \cdot \phi\left(\frac{\Lambda_d(\mathbf{x}_1, \dots, \mathbf{x}_m)}{\eta}\right), \quad (3)$$

where  $\text{diam}(\mathbf{x}_1, \dots, \mathbf{x}_m)$  is the diameter of  $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ . See Figure 1 for an illustration.

Given data points  $\mathbf{x}_1, \dots, \mathbf{x}_N$  and approximation dimension  $d$ , we compute all  $m$ -way affinities, and then obtain pairwise similarities by clique expansion [2] (note that several other options are possible [12, 27, 48]):

$$W_{ij} = \sum_{i_1, \dots, i_{m-2}} \alpha_d(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_{m-2}}). \quad (4)$$

Though it is tempting to choose  $m$  equal to  $d + 2$ , a larger  $m$  allows for more tolerance to weak separation and small sampling rate. The down side is what appears to be an impractical computational burden, since the mere computation of  $\mathbf{W}$  in (4) requires order  $O(N^m)$  flops. In Section 1.4, we discuss how to reduce the computational complexity to  $O(N^{1+o(1)})$  flops, essentially without compromising performance.

Once the affinity matrix  $\mathbf{W}$  is computed, the SC algorithm is applied. We call the resulting procedure higher-order spectral clustering (HOSC), summarized in Algorithm 2. Note that HOSC is (essentially) equivalent to SC when  $\eta \geq \epsilon$ , and equivalent to SCC when  $\epsilon = \infty$ .

**Algorithm 2** Higher Order Spectral Clustering (HOSC)**Input:**

$\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$ : the data points  
 $d, m$ : the approximation dimension and affinity order  
 $\epsilon, \eta$ : the affinity scales  
 $K$ : the number of clusters

**Output:**

A partition of the data into  $K$  disjoint clusters

**Steps:**

- 1: Compute the affinity matrix  $\mathbf{W} = (W_{ij})$  according to (4).
- 2: Apply SC (Algorithm 1).

**1.2. Generative model**

It is time to introduce our framework. We assume a generative model where the clusters are the result of sampling points near surfaces embedded in an ambient Euclidean space, specifically, the  $D$ -dimensional unit hypercube  $(0, 1)^D$ . For a surface  $S \subset (0, 1)^D$  and  $\tau > 0$ , define its  $\tau$ -neighborhood as

$$B(S, \tau) = \{\mathbf{x} \in (0, 1)^D : \text{dist}(\mathbf{x}, S) < \tau\}.$$

The reach of  $S$  is the supremum over  $\tau > 0$  such that, for each  $\mathbf{x} \in B(S, \tau)$ , there is a unique point realizing  $\inf\{\|\mathbf{x} - \mathbf{s}\| : \mathbf{s} \in S\}$  [20]. It is well-known that, for  $C^2$  submanifolds, the reach bounds the radius of curvature from below [20, Lem. 4.17]. For a connection to computational geometry, the reach coincides with the condition number introduced in [43] for submanifolds without boundary. Let  $\text{vol}_d(S)$  denote the  $d$ -dimensional Hausdorff measure, and  $\partial S$  the boundary of  $S$  within  $(0, 1)^D$ . For an integer  $1 \leq d \leq D - 1$  and a constant  $\kappa \geq 1$ , let  $\mathcal{S}_d^2(\kappa)$  be the class of  $d$ -dimensional, connected,  $C^2$  submanifolds  $S \subset (0, 1)^D$  of  $1/\kappa \leq \text{diam}(S) \leq \kappa$  and  $\text{reach}(S) \geq 1/\kappa$ , and if  $S$  has a boundary,  $\partial S$  is a  $(d - 1)$ -dimensional  $C^2$  submanifold with  $\text{reach}(\partial S) \geq 1/\kappa$ . Given surfaces  $S_1, \dots, S_K \in \mathcal{S}_d^2(\kappa)$  and  $\tau < 1/\kappa$ , we generate clusters  $\mathcal{X}_1, \dots, \mathcal{X}_K$  by sampling  $N_k$  points uniformly at random in  $B(S_k, \tau)$ , the  $\tau$ -neighborhood of  $S_k$  in  $(0, 1)^D$ , for all  $k = 1, \dots, K$ . We call  $\tau$  the jitter level. Except for Section 2.3, where we allow for intersections, we assume that the surfaces are separated by a distance of at least  $\delta \geq 0$ , i.e.

$$\text{dist}(S_k, S_\ell) := \inf_{\mathbf{x} \in S_k} \inf_{\mathbf{y} \in S_\ell} \|\mathbf{x} - \mathbf{y}\| \geq \delta, \quad \forall k \neq \ell. \quad (5)$$

In that case, by the triangle inequality, the actual clusters are separated by at least  $\delta - 2\tau$ , i.e.

$$\text{dist}(\mathcal{X}_k, \mathcal{X}_\ell) \geq \delta - 2\tau.$$

We assume that the clusters are comparable in size by requiring that  $N_k \leq \zeta N_\ell$  for all  $k \neq \ell$ , for some finite constant  $\zeta$ . Let  $\mathbf{x}_1, \dots, \mathbf{x}_N$  denote the data points thus generated. See Figure 2 for an illustration.

Given data  $\mathcal{X} := \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ , we aim at recovering the clusters  $\mathcal{X}_1, \dots, \mathcal{X}_K$ . Formally, a clustering algorithm is a function taking data  $\mathcal{X}$ , and possibly other




FIG 2. This figure illustrates the generative model. Left: Three surfaces (here curves) with their  $\tau$ -neighborhood. The curves are separated by at least  $\delta$ . Right: Points sampled within the tubular neighborhoods of the surfaces.




FIG 3. This figure illustrates the generative model with outliers included in the data.

tuning parameters, and outputs a partition of  $\mathcal{X}$ . We say that it is ‘perfectly accurate’ if the output partition coincides with the original partition of  $\mathcal{X}$  into  $\mathcal{X}_1, \dots, \mathcal{X}_k$ . Our main focus is on relating the sample size  $N$  and the separation requirement in (5) (in order for HOSC to cluster correctly), and in particular we let  $\tau$  and  $\delta$  vary with  $N$ . This dependency is left implicit. In contrast, we assume that  $d, K, \zeta$  are fixed. Also, we assume that  $d, \tau, K$  are known throughout the paper (except for Section 2.1 where we consider their estimation). Though our setting is already quite general, we discuss some important extensions in Section 4.

We will also consider the situation where outliers may be present in the data. By outliers we mean points that were not sampled near any of the underlying surfaces. We consider a simple model where outliers are points sampled uniformly in  $(0, 1)^D \setminus \bigcup_k B(S_k, \delta_0)$  for some  $\delta_0 > 0$ , in general different from  $\delta$ . That is, outliers are at least a distance  $\delta_0$  away from the surfaces. We let  $N_0$  denote the number of outliers, while  $N$  still denotes the total number of data points, including outliers. See Figure 3 for an illustration.

### 1.3. Performance in terms of separation and robustness

#### 1.3.1. Performance of SC

A number of papers analyze SC under generative models similar to ours [3, 40, 44, 54], and the closely related method of extracting connected components of the neighborhood graph [3, 9, 36, 37]. The latter necessitates a compactly supported kernel  $\phi$  and may be implemented via a union-of-balls estimator for the support of the density [16]. Under the weaker (essentially Lipschitz) regularity assumption

$$C^{-1} \epsilon^d \leq \text{vol}_d(B(\mathbf{s}, \epsilon) \cap S) \leq C \epsilon^d, \quad \forall \epsilon \in (0, 1/C), \forall \mathbf{s} \in S, \quad (6)$$

Arias-Castro [3] shows that SC with a compactly supported kernel is accurate if

$$\delta - 2\tau \gg \text{sep}_N := \left( \frac{\log N}{N} \right)^{1/d} \vee \tau^{1-d/D} \left( \frac{\log N}{N} \right)^{1/D}. \quad (7)$$

( $a \vee b$  denotes the maximum of  $a$  and  $b$  and  $a_N \gg b_N$  if  $a_N/b_N \rightarrow \infty$  as  $N \rightarrow \infty$ ). With the heat kernel, the same result holds up to a  $\sqrt{\log N}$  multiplicative factor. See also [36, 37], which prove a similar result for the method of extracting connected components under stronger regularity assumptions. At the very least, (7) is necessary for the union-of-balls approach and for SC with a compactly supported kernel, because  $\text{sep}_N$  is the order of magnitude of the largest distance between a point and its closest neighbor from the same cluster [45]. Note that (6) is very natural in the context of clustering as it prevents  $S$  from being too narrow in some places and possibly confused with two or more disconnected surfaces. And, when  $C$  in (6) is large enough and  $\kappa$  is small enough, it is satisfied by any surface  $S$  belonging to  $\mathcal{S}_d^2(\kappa)$ . Indeed, such a surface resembles an affine subspace locally and (6) is obviously satisfied for an affine surface.

When outliers may be present in the data, as a preprocessing step, we identify as outliers data points with low connectivity in the graph with affinity matrix  $\mathbf{W}$ , and remove these points from the data before proceeding with clustering. (This is done between Steps 1 and 2 in Algorithm 1.) In the context of spectral clustering, this is very natural; see, e.g., [3, 12, 37]. Using the pairwise affinity (1), outliers are properly identified if  $\delta_0 - \tau$  satisfies the lower bound in (7) and if the sampling is dense enough, specifically [3],

$$N_k \geq (N^{d/D} \vee N\tau^{D-d}) \log(N), \quad \forall k = 1, \dots, K. \quad (8)$$

When the surfaces are only required to be of Lipschitz regularity as in (6), we are not aware of any method that can even detect the presence of clusters among outliers if the sampling is substantially sparser.

#### 1.3.2. Performance of HOSC

Methods using higher-order affinities are obviously more complex than methods based solely on pairwise affinities. Indeed, HOSC depends on more parameters and is computationally more demanding than SC. One, therefore, wonders




FIG 4. Left: data. Middle: output from SC. Right: output from HOSC. The sampling is much sparser than in the original paper of Ng et al. [42], which is why SC fails. This figure is part of Figure 13 in Section 3, which displays more numerical experiments.

whether this higher level of complexity is justified. We show that HOSC does improve on SC in terms of clustering performance, both in terms of required separation between clusters and in terms of robustness to outliers.

Our main contribution in this paper is to establish a separation requirement for HOSC which is substantially weaker than (7) when the jitter  $\tau$  is small enough. Specifically, HOSC operates under the separation

$$\delta - 2\tau \gg (\tau \wedge \text{sep}_N) \vee \text{sep}_N^2, \quad (9)$$

where  $a \wedge b$  denotes the minimum of  $a$  and  $b$ , and  $\text{sep}_N$  is the separation required for SC with a compactly supported kernel, defined in (7). This is proved in Theorem 1 of Section 2. In particular, in the jitterless case (i.e.  $\tau = 0$ ), the magnitude of the separation required for HOSC is (roughly) the square of that for SC at the same sample size; equivalently, at a given separation, HOSC requires (roughly) the square root of the sample size needed by SC to correctly identify the clusters.

That HOSC requires less separation than SC is also observed numerically. In Figure 4 we compare the outputs of SC and HOSC on the emblematic example of concentric circles given in [42] (here with three circles). While the former fails completely, the latter is perfectly accurate. Indeed, SC requires that the majority of points in an  $\epsilon$ -ball around a given data point come from the cluster containing that point. In contrast, HOSC is able to properly operate in situations where the separation between clusters is so small, or the sampling rate is so low, that any such neighborhood is empty of data points except for the one point at the center. To further illustrate this point, consider the simplest possible setting consisting of two parallel line segments in dimension  $D = 2$ , separated by a distance  $\delta > 0$ , specifically,  $S_1 := \{(t, 0) : t \in [0, 1]\}$  and  $S_2 := \{(t, \delta) : t \in [0, 1]\}$ . Suppose  $N/2$  points are sampled uniformly on each of these line segments. It is well-known that the typical distance between a point on  $S_k$  and its nearest neighbor on  $S_k$  is of order  $O(1/N)$ ; see [45]. Hence, a method computing local statistics requires neighborhoods of radius at least of order  $1/N$ , for otherwise some neighborhoods are empty. From (9), HOSC is perfectly accurate when  $\delta = (\log N)^3/N^2$ , say. When the separation  $\delta$  is that small, typical ball of radius of order  $1/N$  around a data point contains about as many points from  $S_1$  as from  $S_2$  (thus SC cannot work). See Figure 5 for an illustration.



As a bonus, we also show that HOSC is able to resolve intersections in some (very) special cases, while SC is incapable of that. See Proposition 6 and also Figure 12.

To make HOSC robust to outliers, we do exactly as described above, identifying outliers as data points with low connectivity in the graph with affinity matrix  $\mathbf{W}$ , this time computed using the multiway affinity (3). The separation and sampling requirements are substantially weaker than (8), specifically,  $\delta_0 - \tau$  is required to satisfy the lower bound in (9) and the sampling

$$N_k \gg (N^{d/(2D-d)} \vee N\tau^{D-d}) \log(N), \quad \forall k = 1, \dots, K. \quad (10)$$

This is established in Proposition 5, and again, we are not aware of any method for detection that is reliable when the sampling is substantially sparser. For example, when  $\tau = 0$  and we are clustering curves ( $d = 1$ ) in the plane ( $D = 2$ ) (with background outliers), the sampling requirement in (8) is roughly  $N_k \gg N^{1/2} \log(N)$ , compared to  $N_k \gg N^{1/3} \log(N)$  in (10). In Figure 6 below we compare both SC and HOSC on outliers detection, using the data in Figure 4 but further corrupted with 33.3% outliers.




FIG 5. Clustering results obtained by SC (left) and HOSC (right) on a data set of two lines with small separation ( $\delta = 0.005$ ). 100 points are sampled from each line, equally spaced (at a distance 0.01). Note that the inter-point separation on the same cluster is twice as large as the separation between clusters. In this case, SC cannot separate the two lines correctly, as we have argued. In contrast, HOSC performs perfectly when clustering the data, which again agrees with the theory and our expectation. We have also tried increasing the separation  $\delta$  from 0.005 to 0.025, in which case both SC and HOSC perform correctly.




FIG 6. Left: data with outliers. Middle: outliers (black dots) detected by SC. Right: outliers (black dots) detected by HOSC. This figure is part of Figure 15 in Section 3, where more outliers-removal experiments are conducted.

### 1.3.3. Other methods

We focus on comparing HOSC and SC to make a strong point that higher-order methods may be preferred to simple pairwise methods when the underlying clusters are smooth and the jitter level is small. In fact, we believe that no method suggested in the literature is able to compete with HOSC in terms of separation requirements. We quickly argue why.

The algorithm of Kushnir et al. [32] is multiscale in nature and is rather complex, incorporating local information (density, dimension and principal directions) within a soft spectral clustering approach. In the context of semi-supervised learning, Goldberg et al. [25] introduce a spectral clustering method based on a local principal components analysis (PCA) to utilize the unlabeled points. Both methods rely on local PCA to estimate the local geometry of the data and they both operate by coarsening the data, eventually applying spectral clustering to a small subset of points acting as representative hubs for other points in their neighborhoods. They both implicitly require that, for the most part, the vast majority of data points in each neighborhood where the statistics are computed come from a single cluster. Souvenir and Pless [49] suggest an algorithm that starts with ISOMAP and then alternates in EM-fashion between the cluster assignment and the computation of the distances between points and clusters—this is done in a lower dimensional Euclidean space using an MDS embedding. Though this iterative method appears very challenging to be analyzed, it relies on pairwise distances computed as a preprocessing step to derive the geodesic distances, which implicitly assumes that the points in small enough neighborhoods are from the same manifold. Thus, like the SC algorithm, all these methods effectively rely on neighborhoods where only one cluster dominates. This is strong evidence that their separation requirements are at best similar to that of SC. The methods of Haro et al. [30] and Gionis et al. [24] are solely based on the local dimension and density, and are powerless when the underlying manifolds are of same dimension and sampled more or less uniformly, which is the focus of this paper. The method of Guo et al. [29] relies on minimizing an energy that, just as HOSC, incorporates the diameter and local curvature of  $m$ -tuples, with  $m = 3$  for curves and  $m = 4$  for surfaces in 3D, and the minimization is combinatorial over the cluster assignment. In principle, this method could be analyzed with the arguments we deploy here. That said, it seems computationally intractable.

### 1.4. Computational considerations

Thus it appears that HOSC is superior to SC and other methods in terms of separation between clusters and robustness to outliers, when the clusters are smooth and the jitter is small. But is HOSC even computationally tractable?

Assume  $K$  and  $D$  are fixed. The algorithm starts with building the neighborhood graph (i.e., computing the matrix  $\mathbf{W}$ ). This may be done by brute force in  $O(mN^m)$  flops. Clearly, this first step is prohibitive, in particular since

we recommend using a (moderately) large  $m$ . However, we may restrict computations to points within distance  $\epsilon$ , which essentially corresponds to using a compactly supported kernel  $\phi$ . Hence, we could apply a range search algorithm to reduce computations. Alternatively, at each point we may restrict computations to its  $\ell = \omega_N \log(N)$  nearest neighbors, with  $\omega_N \rightarrow \infty$ , or in a slightly different fashion, adapt the local scaling method proposed in [56] by replacing  $\epsilon$  in  $\alpha_d(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_m})$  by  $(\epsilon_{i_1} \cdots \epsilon_{i_m})^{1/m}$ , where  $\epsilon_i$  denotes the distance between  $\mathbf{x}_i$  and its  $\ell$ th nearest neighbor. The reason is that the central condition (12) effectively requires that the degree at each point be of order  $\log(N)^{m-1}$  (roughly), which is guaranteed if the  $\ell$ -nearest neighbors are included in the computations; see [3, 36] for rigorous arguments leading to that conclusion. In low dimensions,  $D = O(\log \log N)$ , a range search and  $\ell$ -nearest-neighbor search may be computed effectively with kd-trees in  $O(N \text{poly}(\log N))$  flops. In higher dimensions, it is essential to use methods that adapt to the intrinsic dimensionality of the data. Assuming that  $d$  is small, the method suggested in [8] has a similar computational complexity. Hence, the (approximate) affinity matrix  $\mathbf{W}$  can be computed in order  $O(N \text{poly}(\log N)) + O(N \cdot \ell^m)$ ; assuming  $m \leq \log(N)/(\omega_N \log \log(N))$ , this is of order  $O(N^{1+1/\omega_N})$ . This is within the possible choices for  $m$  in Theorem 1.

Assume we use the  $\ell$ -nearest-neighbor approximation to the neighborhood graph, with  $\ell = \omega_N \log(N)$ . Then computing  $\mathbf{Z}$  may be done in  $O(N^{1+1/\omega_N})$  flops, since the affinity matrix  $\mathbf{W}$  has at most  $\ell^m = O(N^{1/\omega_N})$  non-zero coefficients per row. Then extracting the leading  $K$  eigenvectors of  $\mathbf{Z}$  may be done in  $O(KN^{1+1/\omega_N})$  flops, using Lanczos-type algorithms [15]. Thus we may run the  $\ell$ -nearest neighbor version of HOSC in  $O(N^{1+1/\omega_N})$  flops, and it may be shown to perform comparably.

We actually implemented the  $\ell$ -nearest-neighbor variant of HOSC and tried it on a number of simulated datasets and a real dataset from motion segmentation. The results are presented in Section 3. The code is publicly available online [13].

### 1.5. Content

The rest of the paper is organized as follows. The main theoretical results are in Section 2 where we provide theoretical guarantees for HOSC, including in contexts where outliers are present or the underlying clusters intersect. We emphasize that HOSC is only able to separate intersecting clusters under very stringent assumptions. In the same section we also address the issue of estimating the parameters that need to be provided to HOSC. In theory at least, they may be chosen automatically. In Section 3 we implemented our own version of HOSC and report on some numerical experiments involving both simulated and real data. Section 4 discusses a number of important extensions, such as when the surfaces self-intersect or have boundaries, which are excluded from the main discussion for simplicity. We also discuss the case of manifolds of different intrinsic dimensions, suggesting an approach that runs HOSC multiple times with different  $d$ . And we describe a kernel version of HOSC that could take advantage

of higher degrees of smoothness. Other extensions are also mentioned, including the use of different kernels. The proofs are postponed to the Appendix.

## 2. Theoretical guarantees

Our main result provides conditions under which HOSC is perfectly accurate with probability tending to one in the framework introduced in Section 1.2. Throughout the paper, we state and prove our results when the surfaces have no boundary and for the simple kernel  $\phi(s) = \mathbf{1}_{\{|s|<1\}}$ , for convenience and ease of exposition. We discuss the case of surfaces with boundaries in Section 4.2 and the use of other kernels in Section 4.5.

**Theorem 1.** *Consider the generative model of Section 1.2. For  $\rho_N \rightarrow \infty$  slowly (e.g.,  $\rho_N = \log \log N$ ), assume the parameters of HOSC satisfy*

$$\log N \geq m \geq \frac{\log N}{\sqrt{\log \rho_N}}, \quad (11)$$

$$\epsilon \geq \left( \rho_N \frac{\log N}{N} \right)^{1/d} \vee \tau^{1-d/D} \left( \rho_N \frac{\log N}{N} \right)^{1/D}. \quad (12)$$

and

$$\eta \geq \epsilon \wedge (\tau + \rho_N \epsilon^2) \quad (13)$$

Assume that (5) holds with

$$\delta - 2\tau > \epsilon \wedge \rho_N \eta. \quad (14)$$

Under these conditions, when  $N$  is large enough, HOSC is perfectly accurate with probability at least  $1 - N^{-\rho_N}$ .

To relate this to the separation requirement stated in the Introduction, the condition (9) is obtained from (14) by choosing  $\epsilon$  and  $\eta$  equal to their respective lower bounds in (12) and (13).

We further comment on the theorem. First, the result holds if  $\rho_N = \rho$  and  $\rho$  is sufficiently large. We state and prove the result when  $\rho_N \rightarrow \infty$  as a matter of convenience. Also, by (11) and (14), the weakest separation requirement is achieved when  $m$  is at least of order slightly less than  $O(\log N)$  so that  $\rho_N$  is of order  $O(1)$ . However, as discussed in Section 1.4, the algorithm is not computationally tractable unless  $m = o(\log N)$ . This is another reason why we focus on the case where  $\rho_N \rightarrow \infty$ . Regarding the constraints (12)-(13) on  $\epsilon$  and  $\eta$ , they are there to guarantee that, with probability tending to one, each cluster is ‘strongly’ connected in the neighborhood graph. Note that the bound on  $\epsilon$  is essentially the same as that required by the pairwise spectral method SC [3, 36]. In turn, once each cluster is ‘strongly’ connected in the graph, clusters are assumed to be separated enough that they are ‘weakly’ connected in the graph. The lower bound (14) quantifies the required separation for that to happen.




FIG 7. Clustering results obtained by SC (left) and HOSC (right) on the data set of Figure 5, but with separation  $\delta = 0.025$  and jitter  $\tau = 0.01$ . In this example, neither SC nor HOSC can successfully separate the two lines. This example supports our claim that when the jitter is large enough (relative to separation), HOSC does not improve over SC and the two algorithms will output the same clustering.

Note that it is specific to the simple kernel. For example, the heat kernel would require a multiplicative factor proportional to  $\sqrt{\log N}$ .

So how does HOSC compare to SC? When the jitter is large enough that  $\tau \gg (\log(N)/N)^{1/d}$ , we have  $\eta \geq \epsilon$  and the local linear approximation contribution to (3) does not come into play. In that case, the two algorithms will output the same clustering (see Figure 7 for an example).

When the jitter is small enough that  $\tau \ll (\log(N)/N)^{1/d}$ , HOSC requires less separation, as demonstrated in Figure 5. Intuitively, in this regime the clusters are sampled densely enough relative to the thickness  $\tau$  that the smoothness of the underlying surfaces comes into focus and each cluster, as a point cloud, becomes locally well-approximated by a thin band. We provide some numerical experiments in Section 3 showing HOSC outperforming SC in various settings.

Thus, HOSC improves on SC only when the jitter is small. This condition is quite severe, though again, we do not know of any other method that can accurately cluster under the weak separation requirement displayed here, even in the jitterless case. It is possible that some form of scan statistic (i.e., matched filters) may be able to operate under the same separation requirement without needing the jitter to be small, however, we do not know how to compute it in our nonparametric setting—even in the case of hybrid linear modeling where the surfaces are affine, computing the scan statistic appears to be computationally intractable. At any rate, the separation required by HOSC is essentially optimal when  $\tau$  is of order  $O(N^{-1/d})$  or smaller. A quick argument for the case  $d = 1$  and  $D = 2$  goes as follows. Consider a line segment of length one and sample  $N$  points uniformly at random in its  $\tau$ -neighborhood, with  $\tau = O(1/N)$ . The claim is that this neighborhood contains an empty band of thickness of order slightly less than  $O(1/N^2)$ , and therefore cannot be distinguished from two parallel line segments. Indeed, such band of half-width  $\lambda$  inside that neighborhood is empty of sample points with probability  $(1 - \lambda/\tau)^N$ , which converges to 1 if  $N\lambda/\tau \rightarrow 0$ , and when  $\tau = O(1/N)$ , this is the case if  $\lambda = o(1/N^2)$ .

In regards to the choice of parameters, the recommended choices depend solely on  $(d, \tau, K)$ . These model characteristics are sometimes unavailable and we discuss their estimation in Section 2.1. Afterwards, we discuss issues such as outliers (Section 2.2) and intersection (Section 2.3).

### 2.1. Parameter estimation

In this section, we propose some methods to estimate the intrinsic dimension  $d$  of the data, the jitter  $\tau$  and the number of clusters  $K$ . Though we show that these methods are consistent in our setting, further numerical experiments are needed to determine their potential in practice.

Compared to SC, HOSC requires the specification of three additional parameters. This is no small issue in practice. In theory, however, we recommend choosing  $d$  and  $K$  consistent with their true values,  $\epsilon$  and  $\eta$  as functions of  $\tau$ , and  $m$  of order slightly less than  $\log(N)$ . The true unknowns are therefore  $(d, \tau, K)$ . We provide estimators for  $d$  and  $K$  that are consistent, and an estimator for  $\tau$  that is accurate enough for our purposes. Specifically, we estimate  $d$  and  $\tau$  using the correlation dimension [28] and an adaptation of our own design. The number of clusters  $K$  is estimated via the eigengap of the matrix  $\mathbf{Z}$ .

#### 2.1.1. The intrinsic dimension and the jitter level

A number of methods have been proposed to estimate the intrinsic dimensionality; we refer the reader to [33] and references therein. The correlation dimension, first introduced in [28], is perhaps the most relevant in our context, since surfaces may be close together. Define the pairwise correlation function

$$\text{Cor}(\epsilon) = \sum_i \sum_{j \neq i} \mathbf{1}_{\{\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon\}}.$$

The authors of [28] recommend plotting  $\log \text{Cor}(\epsilon)$  versus  $\log \epsilon$  and estimating the slope of the linear part. We use a slightly different estimator that allows us to estimate  $\tau$  too, if it is not too small. The idea is to regress  $\log \text{Cor}(\epsilon)$  on  $\log \epsilon$  and identify a kink in the curve. See Figure 8 for an illustration.

Though several (mostly ad hoc) methods have been proposed for finding kinks, we describe a simple method for which we can prove consistency. Fix  $\rho_N \rightarrow \infty$ , with  $\rho_N \ll \log N$ . Define

$$r_N = - \left\lceil \frac{\log \log(N) - \log N}{d \log \rho_N} \right\rceil - 2.$$

Let  $A_r = \log \text{Cor}(\rho_N^{-r})$ . If there is  $r \in \{3, \dots, r_N - 2D - 1\}$  such that

$$(A_r - A_{r+1}) / \log \rho_N > D - 1/2,$$

then let  $\hat{r} \geq 0$  be the smallest such  $r$ ; otherwise, let  $\hat{r} = r_N - 2D$ . Define  $\hat{\tau} = \rho_N^{-\hat{r}}$ ; and also  $\hat{d} = D$ , if  $\hat{r} = 3$ , and  $\hat{d}$  the closest integer to  $(A_3 - A_{\hat{r}}) / (\hat{r} \log \rho_N)$ , otherwise.




FIG 8. A correlation curve for a simulated data set of 240 points sampled from the  $\tau$ -neighborhood of three disjoint one-dimensional curves ( $d = 1$ ) in dimension ten ( $D = 10$ ) crossing all dimensions. The jitter is  $\tau = 0.01$ . We see that the linear part of the curve has slope (near) 1, which coincides with the intrinsic dimension of the curves. The kink appears near  $\hat{\tau} := \exp(-4.5) = 0.0111$ , a close approximation to  $\tau$ .

**Proposition 1.** Consider the generative model described in Section 1.2 with  $S_1, \dots, S_K \in \mathcal{S}_d^2(\kappa)$ . Assume that  $\tau \leq \rho_N^{-3}$  and, if there are  $N_0$  outliers, assume that  $N - N_0 \geq N/\rho_N$ . Then the following holds with probability at least  $1 - N^{-\sqrt{\rho_N}}$ : if  $\hat{r} < r_N - 2D$ , then  $\tau \in [\hat{\tau}/\rho_N, \rho_N \hat{\tau}]$ ; if  $\hat{r} = r_N - 2D$ , then  $\tau \leq \hat{\tau}$ ; moreover, if  $\hat{r} > 3$ ,  $\hat{d} = d$ .

In the context of Proposition 1, the only time that  $\hat{d}$  is inconsistent is when  $\tau$  is of order  $\rho_N^{-3}$  or larger, in which case  $\hat{d} = D$ ; this makes sense, since the region  $\bigcup_k B(S_k, \tau)$  is in fact  $D$ -dimensional if  $\tau$  is of order 1. Also,  $\hat{\tau}$  is within a  $\rho_N$  factor of  $\tau$  if  $\tau$  is not much smaller than  $(\log(N)/N)^{1/d}$ .

We now extend this method to deal with a smaller  $\tau$ . Consider what we just did. The quantity  $\text{Cor}(\epsilon)$  is the total degree of the  $\epsilon$ -neighborhood graph built in SC. Fixing  $(d, m)$ , we now consider the total degree of the  $\eta$ -neighborhood graph built in HOSC. Define the multiway correlation function

$$\text{Cor}_{d,m}(\epsilon, \eta) = \sum_i D_i^{1/(m-1)}.$$

Similarly, we shall regress  $\log \text{Cor}_{d,m}(\epsilon, \eta)$  on  $\log \eta$  and identify a kink in the curve (Figure 9 displays such a curve).

Using the multiway correlation function, we then propose an estimator  $\hat{\tau}$  as follows. We assume that the method of Proposition 1 returned  $\hat{r} = r_N - 2D$ , for otherwise we know that  $\hat{\tau}$  is accurate. Choose  $d = \hat{d}$  and  $m \geq \log(N)(\log \rho_N)^2$ . Note that this is the only time we require  $m$  to be larger than  $\log N$ . Let  $B_s = \log \text{Cor}_{d,m}(\rho_N^{-\hat{r}}, \rho_N^{-\hat{r}-s})$ . If there is  $s \in \{0, \dots, \hat{r} - 1\}$  such that

$$(B_s - B_{s+1})/\log \rho_N > D - d - 1/2,$$

then let  $\hat{s}$  be the smallest one; otherwise, let  $\hat{s} = \hat{r}$ . We then redefine  $\hat{\tau}$  as  $\hat{\tau} = \rho_N^{-\hat{r}-\hat{s}+1}$ .




FIG 9. Correlation curves corresponding to SC (left) and HOSC (right) for the data set of Figure 8, but with a much smaller  $\tau = 1e - 4$ . We see that the pairwise correlation function works poorly in this case, while the multiway correlation curve has a kink near  $\hat{\tau} := \exp(-10.5) = 2.754e - 5$ , within a factor of  $\frac{1}{4}$  of the true  $\tau$ .

**Proposition 2.** *In the context of Proposition 1, assume that  $\hat{r} = r_N - 2D$ . Then redefining  $\hat{\tau}$  as done above, the following holds with probability at least  $1 - N^{-\sqrt{\rho_N}}$ : if  $\hat{s} < \hat{r}$ , then  $\tau \in [\hat{\tau}/\rho_N, \rho_N \hat{\tau}]$ ; if  $\hat{s} = \hat{r}$ , then  $\tau \leq \hat{\tau}$ .*

Now,  $\hat{\tau}$  comes close to  $\tau$  if  $\tau$  is not much smaller than  $(\log(N)/N)^{2/d}$ . Whether this is the case, or not, the statement of Theorem 1 applies with  $\hat{\tau}$  in place of  $\tau$  in (13).

Though our method works in theory, it is definitely asymptotic. In practice, we recommend using other approaches for determining the location of the kink and the slope of the linear part of the pairwise correlation function (in log-log scale). Robust regression methods with high break-down points, like least median of squares and least trimmed squares, worked well in several examples. We do not provide details here, as this is fairly standard, but the figures are quite evocative.

### 2.1.2. The number of clusters

HOSC depends on choosing the number of clusters  $K$  appropriately. A common approach consists in choosing  $K$  by inspecting the eigenvalues of  $\mathbf{Z}$ . We show that, properly tuned, this method is consistent within our model.

**Proposition 3.** *Compute the matrix  $\mathbf{Z}$  in HOSC with the same choice of parameters as in Theorem 1, except that knowledge of  $K$  is not needed. Set the number of clusters equal to the number of eigenvalues of  $\mathbf{Z}$  (counting multiplicity) exceeding  $1 - N^{-2}/\rho_N$ . Then with probability at least  $1 - N^{-\rho_N}$ , this method chooses the correct number of clusters.*

We implicitly assumed that  $d$  and  $\tau$  are known, or have been estimated as described in the previous section. The proof of Proposition 3 is parallel to that of [3, Prop. 4], this time using the estimate provided in part (A1) of the proof of Theorem 1. Details are omitted.






FIG 10. The top six eigenvalues of the weight matrix  $\mathbf{Z}$  obtained by HOSC in Step 2 for the same data used in Figure 8. Though in this example the clusters are well-separated, the eigengap is still very small (about 0.005).

Figure 10 illustrates a situation where the number of clusters is correctly chosen by inspection of the eigenvalues, more specifically, by counting the number of eigenvalue 1 in the spectrum of  $\mathbf{Z}$  (up to numerical error). This success is due to the fact that the clusters are well-separated, and even then, the eigengap is quite small.

We apply this strategy to more data later in Section 3, and show that it can correctly identify the parameter  $K$  in some cases (see Figure 14). In general we do not expect this method to work well when the data has large noise or intersecting clusters, though we do not know of any other method that works in theory under our very weak separation requirements.

## 2.2. When outliers are present

So far we have only considered the case where the data is devoid of outliers. We now assume that some outliers may be included in the data as described at the end of Section 1.2. As stated there, we label as outlier any data point with low degree in the neighborhood graph, as suggested in [3, 12, 37]. Specifically, we compute  $\mathbf{D}$  as in Step 2 of HOSC, and then label as outliers points  $\mathbf{x}_i$  with degree  $D_i$  below some threshold. Let  $\rho_N \rightarrow \infty$  slower than any power of  $N$ , e.g.,  $\rho_N = \log N$ . We propose two thresholds:

(O1) Identify as outliers points with degree:

$$D_i^{1/(m-1)} \leq \rho_N^{-1} \max_j D_j^{1/(m-1)}.$$

(O2) Identify as outliers points with degree:

$$D_i^{1/(m-1)} \leq \rho_N N e^d \eta^{D-d}.$$

Taking up the task of identifying outliers, only the separation between outliers and non-outliers is relevant, so that we do not require any separation between

the actual clusters. We first analyze the performance of (O1), which requires about the same separation between outliers and non-outliers as HOSC requires between points from different clusters in (14).

**Proposition 4.** *Consider the generative model described in Section 1.2. Assume that  $N - N_0 \geq N/\rho_N$  and that (11)-(13) hold. In terms of separation, assume that  $\delta_0 - \tau > \epsilon \wedge \rho_N \eta$ . Then with probability at least  $1 - N^{-\rho_N}$ , the procedure (O1) identifies outliers without error.*

We now analyze the performance of (O2), which requires a stronger separation between outliers and non-outliers, but operates under very weak sampling requirements.

**Proposition 5.** *Assume that  $m$  is as in (11), and*

$$\epsilon = (\rho_N \log(N)/N)^{1/(2D-d)}, \quad \eta = (\rho_N \log(N)/N)^{2/(2D-d)}. \quad (15)$$

*In terms of separation, assume that  $\delta_0 - \tau > \epsilon$ . In addition, suppose that*

$$N_k \geq \rho_N \log(N) N^{d/(2D-d)} \vee N \tau^{D-d}, \quad \forall k = 1, \dots, K. \quad (16)$$

*Then with probability at least  $1 - N^{-\rho_N}$ , the procedure (O2) identifies outliers without error.*

If  $\delta_0 = \tau$ , so that outliers are sampled everywhere but within the  $\tau$ -tubular regions of the underlying surfaces, then both (O1) and (O2) may miss some outliers within a short distance from some  $B(S_k, \tau)$ . Specifically, (O1) (resp. (O2)) may miss outliers within  $\epsilon \wedge \rho_N \eta$  (resp. within  $\epsilon$ ) from some  $B(S_k, \tau)$ . Using Weyl's tube formula [55], we see that there are order  $N_0(\epsilon \wedge \rho_N \eta)^{D-d}$  (resp.  $N_0 \epsilon^{D-d}$ ) such outliers, a small fraction of all outliers.

The sampling requirement (16) is weaker than the corresponding requirement for pairwise methods displayed in (8). In fact, (16) is only slightly stronger than what is required to just detect the presence of a cluster hidden in noise. We briefly explain this point. Instead of clustering, consider the task of detecting the presence of a cluster hidden among a large number of outliers. Formally, we observe the data  $\mathbf{x}_1, \dots, \mathbf{x}_N$ , and want to decide between the following two hypotheses: under the null, the points are independent, uniformly distributed in the unit hypercube  $(0, 1)^D$ ; under the alternative, there is a surface  $S_1 \in \mathcal{S}_d^2(\kappa)$  such that  $N_1$  points are sampled from  $B(S_1, \tau)$  as described in Section 1.2, while the rest of the points,  $N - N_1$  of them, are sampled from the unit hypercube  $(0, 1)^D$ , again uniformly. Assuming that the parameters  $d$  and  $\tau$  are known, it is shown in [4, 5] that the scan statistic is able to separate the null from the alternative if

$$N_1 \gg N^{d/(2D-d)} \vee N \tau^{D-d}. \quad (17)$$

We are not aware of a method that is able to solve this detection task at a substantially lower sampling rate, and (16) comes within a logarithmic factor from (17). We thus obtain the remarkable result that accurate clustering is possible within a log factor of the best (known) sampling rate that allows for accurate detection in the same setting.

### 2.3. When clusters intersect

We now consider the setting where the underlying surfaces may intersect. The additional conditions we introduce are implicit constraints on the dimension of, and the incidence angle at, the intersections. We suppose there is an integer  $0 \leq d_{\text{int}} \leq d - 1$  and a finite constant  $C > 0$  such that

$$\text{vol}_d(B(S_k \cap S_\ell, \epsilon) \cap S_k) \leq C\epsilon^{d-d_{\text{int}}}, \quad \forall \epsilon \in (0, 1/\kappa), \quad \forall k \neq \ell. \quad (18)$$

(The subscript  $\text{int}$  stands for ‘intersection’.) In addition, we assume that for some  $\theta_{\text{int}} \in (0, \pi/2]$ ,

$$\text{dist}(\mathbf{x}, S_\ell) \geq \delta \wedge \sin(\theta_{\text{int}}) \text{dist}(\mathbf{x}, S_k \cap S_\ell), \quad \forall \mathbf{x} \in S_k, \quad \forall k \neq \ell \text{ with } S_k \cap S_\ell \neq \emptyset. \quad (19)$$

(18) is slightly stronger than requiring that  $S_k \cap S_\ell$  has finite  $d_{\text{int}}$ -dimensional volume. If the surfaces are affine, it is equivalent to the condition  $\dim(S_k \cap S_\ell) \leq d_{\text{int}}, \forall k \neq \ell$ . (19), on the other hand, is a statement about the minimum angle at which any two surfaces intersect. For example, if the surfaces are affine within distance  $\delta$  of their intersection, then (19) is equivalent to their maximum (principal) angle being bounded from below by  $\theta_{\text{int}}$ . See Figure 11 for an illustration.

**Proposition 6.** *Consider the setting of Theorem 1, with (5) replaced by (19). In addition, assume that (18) holds. Define*

$$\gamma_N := N^2 \epsilon^d (\epsilon \wedge \rho_N \eta)^{d-d_{\text{int}}} (\sin \theta_{\text{int}})^{d_{\text{int}}-d}.$$

*Then there is a constant  $C > 0$  such that, with probability at least  $1 - C\gamma_N$ , HOSC is perfectly accurate.*




FIG 11. *Illustration of intersecting surfaces. Though the human eye easily distinguishes the two clusters, the clustering task is a lot harder for machine learning algorithms. The main issue is that there are too many data points at the intersection of the two tubular regions. However, in very special cases HOSC is able to separate intersecting clusters (see Figure 12 for such an example).*




FIG 12. *Left: data. Middle: output from HOSC. Right: Output from SC. This example shows that HOSC is able to separate intersecting curvilinear clusters when the incidence angle is perpendicular and there is no jitter ( $\tau = 0$ ). In particular, the conditions of Proposition 6 are satisfied. On the contrary, SC fails in this case.*

The most favorable case is when  $\tau = 0$  and  $\theta_{\text{int}} = \pi/2$ . Then with our choice of  $\epsilon$  and  $\eta$  in Theorem 1, assuming  $\rho_N$  increases slowly, e.g.,  $\rho_N \prec \log N$ , we have  $\gamma_N \rightarrow 0$  if  $2d_{\text{int}} < d$ , and partial results suggest this cannot be improved substantially. This constraint on the intersection of two surfaces is rather severe. Indeed, a typical intersection between two (smooth) surfaces of same dimension  $d$  is of dimension  $d - 1$ , and if so, only curves satisfy this condition. Figure 12 provides a numerical example showing the algorithm successfully separating two intersecting one-dimensional clusters. Thus, even with no jitter and the surfaces intersecting at right angle, HOSC is only able to separate intersecting clusters under exceptional circumstances. Moreover, even when the conditions of Proposition 6 are fulfilled, the probability of success is no longer exponentially small, but is at best of order  $(1/N)^{1-2d_{\text{int}}/d}$ . That said, SC does not seem able to properly deal with intersections at all (see also Figure 12). It essentially corresponds to taking  $\eta = \epsilon$  in HOSC, in which case  $\gamma_N$  never tends to zero.

Though the implications of Proposition 6 are rather limited, we do not know of any other clustering method which provably separates intersecting clusters under a similar generative model. This is a first small step towards finding such a method.

### 3. Software and numerical experiments

We include in this section a few experiments where a preliminary implementation of HOSC outperforms SC, to demonstrate that higher-order affinities can bring a significant improvement over pairwise affinities in the context of manifold clustering.

In our implementation of HOSC, we used the heat kernel  $\phi(s) = \exp(-s^2)$ . Following the discussion in Section 1.4, at each point we restrict the computations to its  $\ell$  nearest neighbors so that we practically remove the locality parameter  $\epsilon$  from the affinity function of (3) and obtain

$$\alpha_d(\mathbf{x}_1, \dots, \mathbf{x}_m) = \begin{cases} \phi(\Lambda_d(\mathbf{x}_1, \dots, \mathbf{x}_m)/\eta), & \text{if } \mathbf{x}_2, \dots, \mathbf{x}_m \in \ell\text{-NN}(\mathbf{x}_1) \text{ distinct;} \\ 0, & \text{otherwise,} \end{cases} \quad (20)$$

where  $\ell\text{-NN}(\mathbf{x}_1)$  is the set of the  $\ell$  nearest neighbors of  $\mathbf{x}_1$ . For computational ease, we used

$$\Lambda_d^{(2)}(\mathbf{x}_1, \dots, \mathbf{x}_m) = \min_{L \in \mathcal{A}_d} \sqrt{\frac{1}{m} \sum_{j=1}^m \text{dist}(\mathbf{x}_j, L)^2}, \quad (21)$$

which can be easily computed using the bottom  $m - d$  singular values of the  $m$  points. Note that, since  $\Lambda_d/\sqrt{m} \leq \Lambda_d^{(2)} \leq \Lambda_d$ , the results we obtained apply, with  $\eta$  changed by a  $\sqrt{m}$  factor, at most. (In the paper, the standard choice for  $\eta$  is a power of  $N$ , while  $m$  is of order at most  $\log N$ , so this factor is indeed negligible.) In practice, we always search a subinterval of  $[0, 1]$  for the best working  $\eta$  (e.g.,  $[.001, .1]$ ), based on the smallest variance of the corresponding clusters in the eigenspace (the row space of the matrix  $\mathbf{V}$ ), as suggested in [42]. When the given data contains outliers, the optimal choice of  $\eta$  is based on the largest gap between the means of the two sets of degrees (associated to the inliers and outliers), normalized by the maximum degree. The code is available online [13].

### 3.1. Synthetic data

We first generate five synthetic data sets in the unit cube  $(0, 1)^D$  ( $D = 2$  or  $3$ ), shown in Figure 13. In this experiment, the actual number of clusters (i.e.  $K$ ) and dimension of the underlying manifolds (i.e.  $d$ ) are assumed known to all algorithms. For HOSC, we fix  $\ell = 10, m = d + 2$ , and use the subinterval  $[0.001, 0.1]$  as the search interval of  $\eta$ . For SC, we considered two ways of tuning the scale parameter  $\epsilon$ : directly, by choosing a value in the interval  $[0.001, 0.25]$  (SC-NJW); and by the local scaling method of [56] (SC-LS), with the number of nearest neighbors  $\ell = 5, \dots, 15$ . The final choices of these parameters were also based on the same criterion as used by HOSC.

Figure 13 exhibits the clusters found by each algorithm when applied to the five data sets, respectively. Observe that HOSC succeeded in a number of difficult situations for SC, e.g., when the sampling is sparse, or when the separation is small at some locations.

We also plot the leading eigenvalues of the matrix  $\mathbf{Z}$  obtained by HOSC on each data set; see Figure 14. We see that in data sets 1, 2, 5, the number of eigenvalue 1 coincides with the true number of clusters, while in 3 and 4 there is some discrepancy between the  $K$ th eigenvalue and the number 1. Though we do not expect the eigengap method to work well in general, Figure 14 shows that it can be useful in some cases.

Figure 15 displays some experiments including outliers. We simply sampled points from the unit square  $(0, 1)^2$  uniformly at random and added them as outliers to the first three data sets in Figure 13, with percentages 33.3%, 60% and 60%, respectively. We applied SC and HOSC assuming knowledge of the proportion of outliers, and labeled points with smallest degrees as outliers. Choosing the threshold automatically remains a challenge; in particular, we did not test the theory.

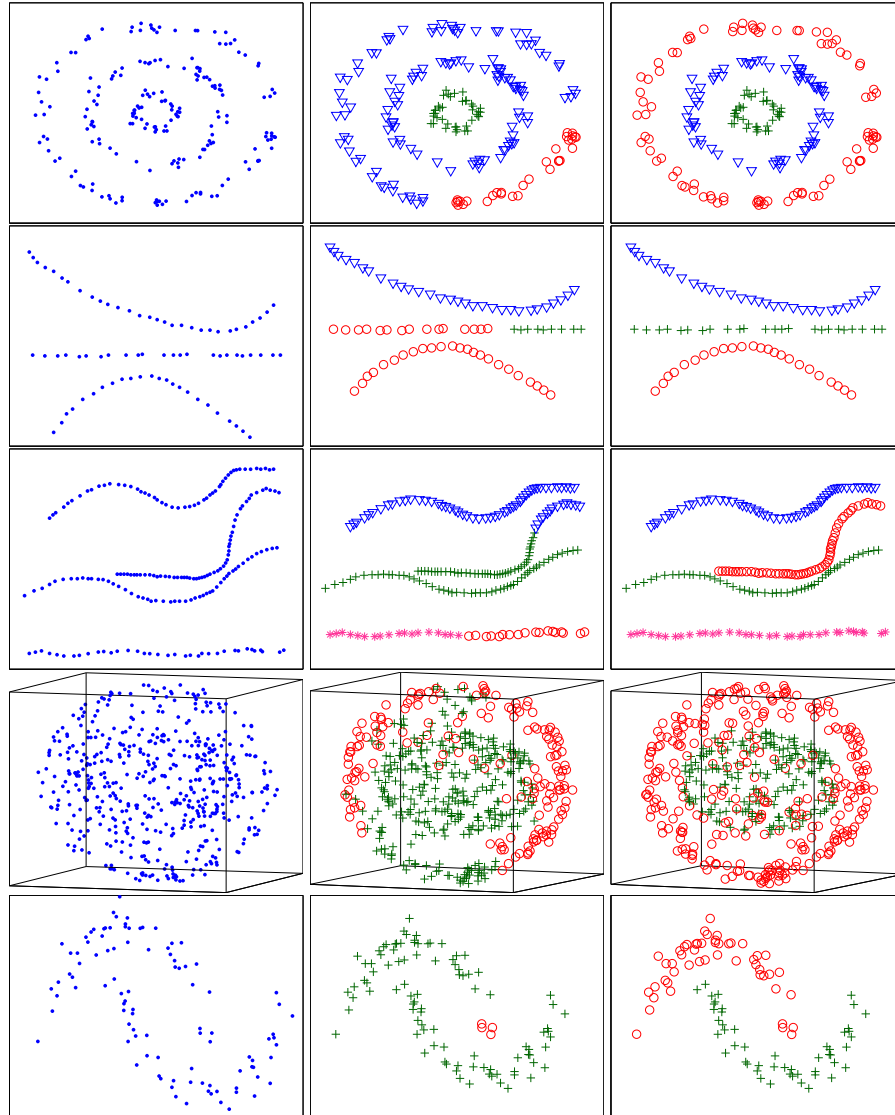


FIG 13. *Left column: data. (The third example shows a sphere containing an ellipsoid inside.) Middle column: best output from SC with the scale parameter chosen by both searching the interval  $[0.001, 0.25]$  and applying local scaling [56] with at most 15 nearest neighbors. Right column: output from HOSC. The optimal value of  $\eta$  is selected from the interval  $[0.001, 0.1]$ . We also tried the simple kernel instead of the heat kernel, and obtained same results except in data set 3.*

We observe that HOSC could successfully remove most of the true outliers, leaving out smooth structures in the data; in contrast, SC tended to keep isolated high-density regions, being insensitive to sparse smooth structures. A hundred replications of this experiment (i.e., fixing the clusters and adding randomly gen-

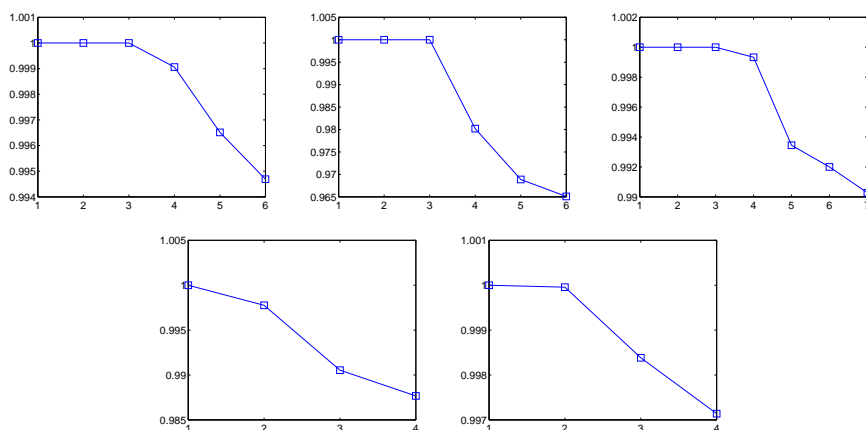


FIG 14. Top eigenvalues of the matrix  $\mathbf{Z}$  obtained by HOSC on each of the five data sets in Figure 13 (in same order).

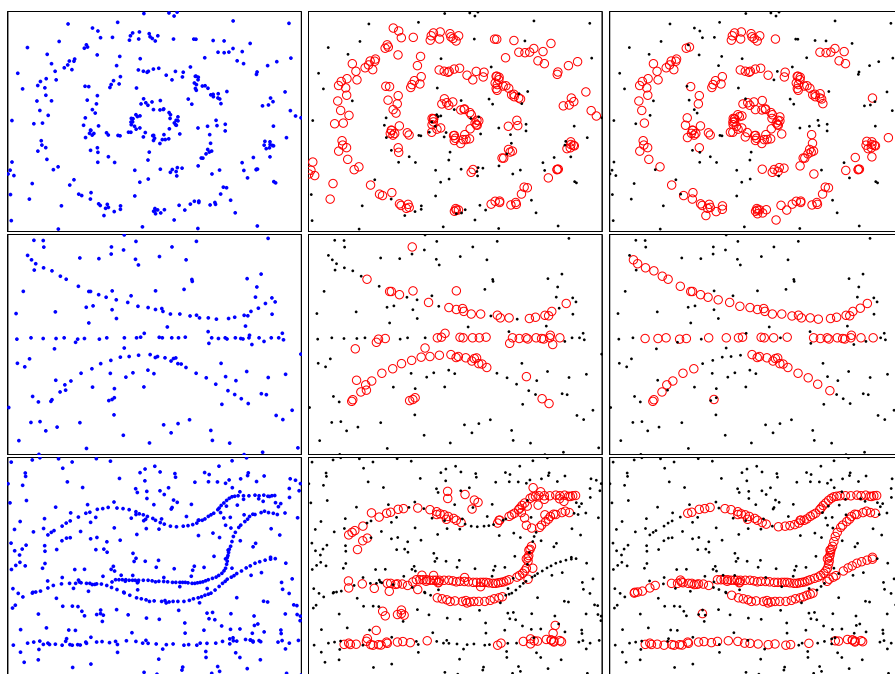


FIG 15. Outlier-removal experiments. Left column: data with outliers. The percentages of outliers are 33.3%, 60% and 60%, respectively. Middle: outliers (black dots) detected by pairwise spectral clustering (both SC-NJW and SC-LS, but only the better result is shown). Right: outliers (black dots) detected by HOSC. The use of the simple kernel (instead of the heat kernel) in HOSC gives very similar results.

erated outliers) show that the True Positive Rates (i.e., percentages of correctly identified outliers) for (SC, HOSC) are (58.1% vs 67.7%), (75.4% vs 86.8%) and (76.8% vs 88.0%), respectively.

### 3.2. Real data

We next compare SC and HOSC using the two-view motion data studied in [10, 46]. This data set contains 13 motion sequences: (1) *boxes*, (2) *carsnbus3*, (3) *deliveryvan*, (4) *desk*, (5) *lightbulb*, (6) *manycars*, (7) *man-in-office*, (8) *nrbooks3*, (9) *office*, (10) *parking-lot*, (11) *posters-checkerboard*, (12) *posters-keyboard*, and (13) *toys-on-table*; and each sequence consists of two image frames of a 3-D dynamic scene taken by a perspective camera (see Figure 16 for a few such sequences). Suppose that several feature points have been extracted from the moving objects in the two camera views of the scene. The task is to separate the trajectories of the feature points according to different motions. This application, which lies in the field of *structure from motion*, is one of the fundamental problems in computer vision.

Given a physical point  $\mathbf{x} \in \mathbb{R}^3$  and its image correspondences in the two views  $(x_1, y_1)', (x_2, y_2)' \in \mathbb{R}^2$ , one can always form a joint image sample  $\mathbf{y} = (x_1, y_1, x_2, y_2, 1)' \in \mathbb{R}^5$ . It is shown in [46] that, under perspective camera projection, all the joint image samples  $\mathbf{y}$  corresponding to different motions live on different manifolds in  $\mathbb{R}^5$ , some having dimension 2 and others having dimension 4. Exploratory analysis applied to these data suggests that the manifolds in this dataset mostly have dimension 2 (see Figure 17). Therefore, we will apply our algorithm (HOSC) with  $d = 2$  to these data sets in order to compare with pairwise spectral clustering (SC-NJW, SC-LS).




FIG 16. Three exemplary two-view motion sequences (arranged in columns): (4) desk, (6) manycars and (7) man-in-office. The true clusters are displayed in different colors and markers (the black dots are outliers).






FIG 17. The true clusters of the three sequences in Figure 16 (in same order), shown in top three principal dimensions. (The outliers have been removed from the data and thus are not displayed). These plots clearly indicate that the underlying manifolds are two dimensional.

We use the following parameter values for the two algorithms. In HOSC, we choose  $\ell = 20$ ,  $m = d + 2$ ,  $\eta \in [.0001, .1]$ , while in SC we try both searching the interval  $[.001, .5]$  (SC-NJW) and local scaling with at most 24 nearest neighbors (SC-LS).

The original data contains some outliers. In fact, 10 sequences out of the 13 are corrupted with outliers, with the largest percentage being about 32%. We first manually remove the outliers from those sequences and solely focus on the clustering aspects of the two algorithms. Next, we add outliers back and compare them regarding outliers removal. (Note that we need to provide both algorithms with the true percentage of outliers in each sequence.) By doing so we hope to evaluate the clustering and outliers removal aspects of an algorithm separately and thus in the most accurate way.

Table 1 presents the results from the experiments above. Observe that HOSC achieved excellent clustering results in all but two sequences, with zero error on

TABLE 1

The misclassification rates and the numbers of true outliers detected by HOSC, SC-NJW and SC-LS. In the clustering experiment, the outliers-free data is used; then the outliers are added back so that each of these algorithms can be applied to detect them. For SC-NJW, the tuning parameter is selected from the interval  $[.001, .5]$ ; for SC-LS, a maximum of 24 nearest neighbors are used; for HOSC, 20 nearest neighbors are used and the flatness parameter  $\eta$  is selected from the interval  $[.0001, .1]$

seq.	Data		Clustering Errors			# True Outliers Detected		
	#samples	#out.	SC-NJW	SC-LS	HOSC	SC-NJW	SC-LS	HOSC
1	115,121	2	0.85%	0.85%	0.85%	1	1	1
2	85,45,89	28	0%	0%	0%	24	24	24
3	62,192	0	30.3%	23.6%	30.3%	N/A	N/A	N/A
4	50,50,55	45	0.65%	2.58%	1.29%	35	30	37
5	51,121,33	0	0%	0%	0%	N/A	N/A	N/A
6	54,24,23,43	0	18.8%	0%	0%	N/A	N/A	N/A
7	16,57	34	19.2%	19.2%	0%	17	12	26
8	129,168,91	32	22.9%	17.8%	22.9%	12	17	23
9	76,109,74	48	0%	0%	0%	36	28	36
10	19,117	4	0%	47.8%	0%	0	0	1
11	100,99,81	99	0%	1.79%	0%	42	39	73
12	99,99,99	99	0.34%	0.34%	0%	80	43	91
13	49,42	35	33.0%	15.4%	2.20%	7	6	21




FIG 18. Clustering results of both HOSC and SC (left to right) on sequence (7). (The truth is shown in Figure 17, rightmost plot). In this example, HOSC correctly found the two clusters, using geometric information; in contrast, SC failed because it solely relies on pairwise distances.

eight sequences, one mistake on sequence (13), and two mistakes on each of (1) and (4). We remark that HOSC also outperformed the algorithms in [10, Table 1], in terms of clustering accuracy, but due to the main aim of this paper, we do not include those results in Table 1. In contrast, each of SC-NJW and SC-LS failed on at least five sequences (with over 15% misclassification rates), both containing the two bad sequences for HOSC. As a specific example, we display in Figure 18 the clusters obtained by both HOSC and SC on sequence (7), demonstrating again that higher order affinities can significantly improve over pairwise affinities in the case of manifold data. Regarding outliers removal, HOSC is also consistently better than SC-NJW and SC-LS (if not equally good).

## 4. Extensions

### 4.1. When the underlying surfaces self-intersect

In our generative model described in Section 1.2 we assume that the surfaces are submanifolds, implying that they do not self-intersect. This is really for convenience as there is essentially no additional difficulty arising from self-intersections. If we allow the surfaces to self-intersect, then we bound the maximum curvature (from above) and not the reach. We could, for example, consider surfaces of the form  $S = f(B_d(0, 1))$ , where  $f : B_d(0, 1) \rightarrow (0, 1)^D$  is locally bi-Lipschitz and has bounded second derivative. A similar model is considered in [38] in the context of set estimation. Clearly, proving that each cluster is connected in the neighborhood graph in this case is the same. The only issue is in situations where a surface comes within distance  $\epsilon$  from another surface at a location where the latter intersects itself. The geometry involved in such a situation is indeed complex. If we postulate that no such situation arises, then our results generalize immediately to this setting.




FIG 19. An example of a surface with a boundary coming close to another surface. This is a potentially problematic situation for HOSC as the points near the boundary of one surface and close to the other surface may be strongly connected to points from both clusters. Numerically, we show in Figure 13 such an example where HOSC is successful.

#### 4.2. When the underlying surfaces have boundaries

When the surfaces have boundaries, points near the boundary of a surface may be substantially connected with points on a nearby surface. See Figure 19 for an illustration. This is symptomatic of the fact that the algorithm is not able to resolve intersections in general, as discussed in Section 2.3, with the notable exception of clusters of dimension  $d = 1$ , as illustrated in the ‘two moons’ example of Figure 13.

If we require a stronger separation between the boundary of a surface and the other surfaces, specifically,

$$\text{dist}(\partial S_k, S_\ell) \geq \delta_{\ddagger}, \quad \forall k \neq \ell, \quad (22)$$

with  $\delta_{\ddagger} - 2\tau > \epsilon$ , no point near the boundary of a cluster is close to a point from a different cluster. (A corresponding requirement in the context of outliers would be that outliers be separated from the boundary of a cluster by at least  $\delta_{0,\ddagger}$ , with  $\delta_{0,\ddagger} - \tau > \epsilon$ .)

#### 4.3. When the data is of mixed dimensions

In a number of situations, the surfaces may be of different intrinsic dimensions. An important instance of that is the study of the distribution of galaxies in space, where the galaxies are seen to cluster along filaments ( $d = 1$ ) and walls ( $d = 2$ ) [39]. We propose a top-down approach, implementing HOSC for each dimension  $d$  starting at  $D - 1$  and ending at 1 (or between any known upper and lower bounds for  $d$ ).

At each step, the algorithm is run on each cluster obtained from the previous step, including the set of points identified as outliers. Indeed, when the dimension parameter of the algorithm is set larger than the dimension of the underlying surfaces, HOSC may not be able to properly separate clusters. For example, two parallel segments satisfying the separation requirement of Theorem 1 still belong to a same plane and HOSC with dimension parameter  $d = 2$

would not be able to separate the two line segments. Another reason for processing the outlier bin is the greater disparity in the degrees of the data points in the neighborhood graph often observed with clusters of different dimensions. At each step, the number of clusters is determined automatically according to the procedure described in Section 2.1, for such information is usually not available. The parameters  $\epsilon$  and  $\eta$  are chosen according to (15). Partial results suggest that, under some additional sampling conditions, this top-down procedure is accurate under weaker separation requirements than required by pairwise methods, which handle the case of mixed dimensions seamlessly [3]. The key is that an actual cluster  $\mathcal{X}_k$ , as defined in Section 1.2, is never cut into pieces. Indeed, properties (A1) and (A4) in the proof of Theorem 1, which guarantee the connectivity and regularity (in terms of comparable degrees) of the subgraph represented by  $\mathcal{X}_k$ , are easily seen to also be valid when the dimension parameter of the algorithm is set larger than  $d$ . (This observation might explain the success of the SCC algorithm of [12] in some mixed settings when using an upper bound on the intrinsic dimensions.)

#### 4.4. Clustering based on local polynomial approximations

For  $1 \leq d \leq D - 1$  and an integer  $r \geq 3$ , let  $\mathcal{S}_d^r(\kappa)$  be the subclass of  $\mathcal{S}_d^2(\kappa)$  of  $d$ -dimensional submanifolds  $S$  such that, for every  $\mathbf{x} \in S$  with tangent  $T_{\mathbf{x}}$ , the orthogonal projection  $S \cap B(\mathbf{x}, 1/\kappa) \rightarrow T_{\mathbf{x}}$  is a  $C^r$ -diffeomorphism with all partial derivatives of order up to  $r$  bounded in supnorm by  $\kappa$ . For example,  $\mathcal{S}_d^r(\kappa)$  includes a subclass of surfaces of the form  $S = f(B_d(0, 1))$ , where  $f : B_d(0, 1) \rightarrow (0, 1)^D$  is locally bi-Lipschitz and has its first  $r$  derivatives bounded. (We could also consider surfaces of intermediate, i.e., Hölder smoothness, a popular model in function and set estimation [18, 38].)

Given that surfaces in  $\mathcal{S}_d^r$  are well-approximated locally by polynomial surfaces, it is natural to choose an affinity based on the residual of the best  $d$ -dimensional polynomial approximation of degree at most  $r - 1$  to a set of points  $\mathbf{x}_1, \dots, \mathbf{x}_m$ . This may be implemented via the “kernel trick” with a polynomial kernel, as done in [10] for the special case of algebraic surfaces. The main difference with the case of  $C^2$  surfaces that we consider in the rest of the paper is the degree of approximation to a surface  $S \in \mathcal{S}_d^r$  by its osculating algebraic surface of order  $r - 1$ ; within a ball of radius  $\epsilon$ , it is of order  $O(\epsilon^r)$ .

Partial results suggest that, under similar conditions, the kernel version of HOSC with  $r$  known may be able to operate under a separation of the form (9), with the exponent  $2/d$  replaced by  $r/d$  and, in the presence of outliers, within a logarithmic factor of the best known sampling rate ratio achieved by any detection method [4, 5]:

$$\min_k N_k \geq N^{d/(rD-(r-1)d)} \vee N\tau^{D-d}. \quad (23)$$

Regarding the estimation of  $\tau$ , defining the correlation dimension using the underlying affinity defined here allows to estimate  $\tau$  accurately down to (essen-

tially)  $(\log(N)/N)^{r/d}$ , if the surfaces are all in  $\mathcal{S}_d^r(\kappa)$ . The arguments are parallel and we omit the details.

Thus, using the underlying affinity defined here may allow for higher accuracy, if the surfaces are smooth enough. However, this comes with a larger computational burden and at the expense of introducing a new parameter  $r$ , which would need to be estimated if unknown, and we do not know a good way to do that.

#### 4.5. Other extensions

The setting we considered in this paper, introduced in Section 1.2, was deliberately more constrained than needed for clarity of exposition. We list a few generalizations below, all straightforward extensions of our work.

- *Sampling.* Instead of the uniform distribution, we could use any other distribution with a density bounded away from 0 and  $\infty$ , or with fast decaying tails such as the normal distribution.
- *Kernel.* The rate of decay of the kernel  $\phi$  dictates the range of the affinity (3). Let  $\omega_N$  be a non-decreasing sequence such that  $N^{3m}\phi(\omega_N) \rightarrow 0$ . For a compactly supported kernel,  $\omega_N = \sup\{s : \phi(s) > 0\}$ , while for the heat kernel, we can take  $\omega_N = 2\sqrt{m \log N}$ . As we will take  $m \rightarrow \infty$ ,  $\phi$  is essentially supported in  $[0, \omega_N]$  so that points that are further than  $\omega_N \epsilon$  apart have basically zero affinity. Specifically, we use the following bounds:

$$\phi(1)\mathbf{1}_{\{|s|<1\}} \leq \phi(s) \leq \mathbf{1}_{\{|s|<\omega_N\}} + \phi(\omega_N).$$

The results are identical, except that statements of the form  $\delta - 2\tau > Z$  are replaced with  $\delta - 2\tau > \omega_N Z$ .

- *Measure of flatness.* As pointed out in the introduction, any reasonable measure of linear approximation could be used instead. Our choice was driven by convenience and simplicity.

## Appendix A: Preliminaries

We gather here some preliminary results. Recall that, for  $a, b \in \mathbb{R}$ ,  $a \vee b := \max(a, b)$ ;  $a \wedge b := \min(a, b)$ ;  $a_+ = a \vee 0$ . For  $(a_N), (b_N) \in \mathbb{R}^N$ ,  $a_N \prec b_N$  means  $a_N = O(b_N)$ ;  $a_N \asymp b_N$  means both  $a_N = O(b_N)$  and  $b_N = O(a_N)$ ;  $a_N \ll b_N$  means  $a_N = o(b_N)$ . For  $L \in \mathcal{A}_d$ ,  $P_L$  denotes the orthogonal projection onto  $L$ . The canonical basis of  $\mathbb{R}^D$  is denoted  $\mathbf{e}_1, \dots, \mathbf{e}_D$ .

### A.1. Large deviations bounds

The following result is a simple consequence of Hoeffding's or Bernstein's inequalities.

**Lemma 1** ([50], Lem. 5.3.7). *Let  $(X_i)_{i \geq 1}$  be independent random variables in  $[0, 1]$ .*

$$\text{If } 4a \leq \sum_i \mathbb{E}(X_i),$$

$$\mathbb{P}\left(\sum_i X_i \leq a\right) \leq \exp(-a).$$

$$\text{If } a \geq 8 \sum_i \mathbb{E}(X_i),$$

$$\mathbb{P}\left(\sum_i X_i \geq a\right) \leq \exp(-a).$$

### A.2. Some geometrical results

We start by quantifying how well a surface  $S \in \mathcal{S}^2(\kappa)$  is locally approximated by its tangent. Recall that, for an affine subspace  $T$ ,  $P_T$  denotes the orthogonal projection onto  $T$ . For any  $\mathbf{s} \in S$ , let  $T_{\mathbf{s}}$  denote the tangent of  $S$  at  $\mathbf{s}$ .

**Lemma 2.** *For any  $S \in \mathcal{S}_d^2(\kappa)$  and  $\mathbf{s} \in S$ , the orthogonal projection onto  $T_{\mathbf{s}}$  is injective on  $B(\mathbf{s}, 1/(4\kappa)) \cap S$  and  $P_{T_{\mathbf{s}}}^{-1}$  has Lipschitz constant bounded by  $\sqrt{2}$  on its image, which contains  $B(\mathbf{s}, 1/(8\kappa)) \cap T_{\mathbf{s}}$ . Moreover,*

$$B(\mathbf{s}, \epsilon) \cap S \subset B(T_{\mathbf{s}}, \kappa\epsilon^2), \quad \forall \epsilon,$$

and

$$B(\mathbf{s}, \epsilon) \cap T_{\mathbf{s}} \subset B(S, 2\kappa\epsilon^2), \quad \forall \epsilon < 1/(8\kappa).$$

*Proof.* This sort of result is standard in differential geometry. We follow the exposition in [43]. We note that the manifold parameter  $\tau$  in [43], i.e., the inverse of the condition number, coincides with the manifold's reach. We thus fix here an  $S \in \mathcal{S}_d^2(\kappa)$  and denote  $\tau := \text{reach}(S)$ . Since  $1/\kappa$  is a lower bound on the reach for manifolds in  $\mathcal{S}_d^2(\kappa)$ , we have the inequality  $\tau \geq 1/\kappa$ .

Fix also a point  $\mathbf{s} \in S$ . Applying [43, Lem. 5.4], we obtain that  $P_{T_{\mathbf{s}}}$  is one-to-one on  $B(\mathbf{s}, \epsilon) \cap S$  for any  $\epsilon < \tau/2$ , in particular,  $\epsilon < 1/(2\kappa)$ . We obtain an estimate on the image of  $P_{T_{\mathbf{s}}}$  as follows. We note that [43, proof of Lem. 5.3] implies that

$$P_{T_{\mathbf{s}}}(B(\mathbf{s}, \epsilon) \cap S) \supseteq B(\mathbf{s}, \epsilon \cos \arcsin(\epsilon/(2\tau))) \cap T_{\mathbf{s}}. \quad (24)$$

Furthermore,

$$\text{if } \epsilon \leq 1/(4\kappa), \quad \cos \arcsin(\epsilon/(2\tau)) \geq \cos \arcsin(\kappa\epsilon/2) \geq \sqrt{63/64} > 1/2. \quad (25)$$

Combining (24) and (25), we conclude that

$$P_{T_{\mathbf{s}}}(B(\mathbf{s}, \epsilon) \cap S) \supseteq B(\mathbf{s}, \epsilon/2) \cap T_{\mathbf{s}}, \quad \forall \epsilon \leq 1/(4\kappa). \quad (26)$$

In particular, for  $\epsilon = 1/(4\kappa)$ , we obtain that the range of  $P_{T_{\mathbf{s}}}$  (when applied to  $B(\mathbf{s}, 1/(4\kappa)) \cap S$ ) contains the ball  $B(\mathbf{s}, 1/(8\kappa)) \cap T_{\mathbf{s}}$ .

Next, for any  $\mathbf{s}' \in B(\mathbf{s}, 1/(4\kappa)) \cap S$ , the derivative of the linear operator  $P_{T_{\mathbf{s}}}$  in the direction  $\mathbf{u}$ , a unit vector in  $T_{\mathbf{s}'}$ , is

$$\nabla_{\mathbf{u}}(P_{T_{\mathbf{s}}}) = (P_{T_{\mathbf{s}}}) \cdot \mathbf{u} = \cos \theta_1(T_{\mathbf{s}}, \text{span}\{\mathbf{u}\}) \geq \cos \theta_1(T_{\mathbf{s}}, T_{\mathbf{s}'}), \tag{27}$$

where  $\theta_1$  denotes the largest principal angle between the corresponding subspaces. In order to further bound from below the RHS of (27), we couple [43, Props. 6.2, 6.3] and use  $\tau \geq 1/\kappa$  to obtain that

$$\cos \theta_1(T_{\mathbf{s}}, T_{\mathbf{s}'}) \geq \sqrt{1 - 2\kappa\|\mathbf{s} - \mathbf{s}'\|}. \tag{28}$$

Combining (27) and (28) we conclude that  $P_{T_{\mathbf{s}}}^{-1}$  has Lipschitz constant bounded by  $\sqrt{2}$  in  $B(\mathbf{s}, 1/(4\kappa)) \cap T_{\mathbf{s}}$ .

For the inclusions, we use the fact that

$$\|P_{T_{\mathbf{s}}}(\mathbf{x}) - \mathbf{x}\| \leq (\kappa/2) \cdot \|\mathbf{s} - \mathbf{x}\|^2, \quad \forall \mathbf{x}, \mathbf{s} \in S, \tag{29}$$

which appears in [20, Th. 4.18(2)]. This immediately implies the first inclusion—which actually holds for any  $\epsilon > 0$  and with  $\kappa$  replaced by  $\kappa/2$ . The second inclusion follows by combining (26) with (29).  $\square$

Next, we estimate the volume of the intersection of the neighborhood of a surface and a ball centered at a point within that neighborhood.

**Lemma 3** ([3], Lem. 1). *For  $S$  satisfying (6),  $\mathbf{x} \in B(S, \tau)$  and  $\epsilon, \tau > 0$ ,*

$$\text{vol}_D(B(S, \tau) \cap B(\mathbf{x}, \epsilon)) \asymp \epsilon^d (\epsilon \wedge \tau)^{D-d}, \quad \text{vol}_D(B(S, \tau)) \asymp \tau^{D-d}.$$

The following result is on the approximation of a set of points in the neighborhood of a  $d$ -dimensional affine subspace by a  $d$ -dimensional affine subspace generated by a subset of  $d + 1$  points.

**Lemma 4.** *There is a constant  $C > 0$  depending only on  $d$  such that, if  $\mathbf{z}_1, \dots, \mathbf{z}_m \in B(L, \eta)$ , with  $L \in \mathcal{A}_d$  and  $m \geq d + 2$ , then there exists  $H \in \mathcal{A}_d$  generated by  $d + 1$  points among  $\mathbf{z}_1, \dots, \mathbf{z}_m$ , such that  $\mathbf{z}_1, \dots, \mathbf{z}_m \in B(H, C\eta)$ .*

*Proof.* For points  $\mathbf{a}_1, \dots, \mathbf{a}_k$ , let  $\text{aspan}\{\mathbf{a}_1, \dots, \mathbf{a}_k\}$  denote the affine subspace of minimum dimension passing through  $\mathbf{a}_1, \dots, \mathbf{a}_k$ . Let  $(i_1, i_2) \in \text{argmax}_{i,j} \|\mathbf{z}_i - \mathbf{z}_j\|$  and, for  $d \geq k \geq 3$ ,

$$i_k \in \text{argmax}_{i \neq i_1, \dots, i_{k-1}} \text{dist}(\mathbf{z}_i, \text{aspan}\{\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_{k-1}}\}).$$

Let  $A_k = \text{aspan}\{\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_{k+1}}\}$ , for  $d \geq k \geq 1$ . Define  $\lambda_1 = \|\mathbf{z}_{i_2} - \mathbf{z}_{i_1}\|$  and, for  $d \geq k \geq 2$ ,  $\lambda_k = \text{dist}(\mathbf{z}_{i_{k+1}}, \text{span}\{\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_k}\})$ . Also, let  $\mathbf{v}_1 = (\mathbf{z}_{i_2} - \mathbf{z}_{i_1})/\lambda_1$  and, for  $k \geq 2$ ,  $\mathbf{v}_k = (\mathbf{z}_{i_{k+1}} - P_{A_{k-1}}\mathbf{z}_{i_{k+1}})/\lambda_k$ . Without loss of generality, assume that  $\mathbf{z}_{i_1}$  is the origin, which allows us to identify a point  $\mathbf{z}$  with the vector  $\mathbf{z} - \mathbf{z}_{i_1}$ . Take  $\mathbf{z} \in \{\mathbf{z}_1, \dots, \mathbf{z}_m\}$  and express it as  $\mathbf{z} = a_1\mathbf{v}_1 + \dots + a_d\mathbf{v}_d + \mathbf{w}$ , with  $\mathbf{w} \perp A_d$ . We show that  $\|\mathbf{w}\| \leq C\eta$  for a constant  $C$  depending only on  $d$ , which implies that  $\mathbf{z} \in B(A_d, C\eta)$ . Let  $C_1 > 0$ , to be made sufficiently large later. By construction,

$A_1 \subset \dots \subset A_d$  and  $\lambda_1 \geq \dots \geq \lambda_d$  with  $\|P_{A_{k-1}^\perp} \mathbf{z}\| \leq \lambda_k$  for all  $k = 1, \dots, d$ . Consequently, if  $\lambda_d \leq C_1\eta$ , then  $\|\mathbf{w}\| \leq \|P_{A_{d-1}^\perp} \mathbf{z}\| \leq \lambda_d \leq C_1\eta$  and we are done. Therefore, assume that  $\lambda_d > C_1\eta$ . Define  $\mathbf{q}_k = P_L \mathbf{v}_k$ . We have

$$\|\mathbf{q}_k - \mathbf{v}_k\| = \|P_{L^\perp} \mathbf{v}_k\| = \frac{1}{\lambda_k} \|P_{L^\perp} P_{A_{k-1}^\perp} \mathbf{z}_{i_{k+1}}\| \leq \frac{1}{\lambda_k} \|P_{L^\perp} \mathbf{z}_{i_{k+1}}\| \leq \frac{\eta}{\lambda_k} \leq \frac{1}{C_1}.$$

Hence, for  $C_1$  large enough,  $\mathbf{q}_1, \dots, \mathbf{q}_d$  are linearly independent, and therefore span  $L$ . Suppose this is the case and define matrices  $\mathbf{V}$  with columns  $\mathbf{v}_1, \dots, \mathbf{v}_d$  and  $\mathbf{Q}$  with columns  $\mathbf{q}_1, \dots, \mathbf{q}_d$ . Then, by continuity, for  $C_1$  large enough we have

$$\|P_L - P_{A_d}\| = \|\mathbf{Q}(\mathbf{Q}^T \mathbf{Q})^{-1} \mathbf{Q}^T - \mathbf{V}\mathbf{V}^T\| \leq 1/2,$$

where  $\|\cdot\|$  here denotes the (Euclidean) operator norm. When  $C_1$  is that large, we have

$$\|P_L \mathbf{w}\| = \|(P_L - P_{A_d})\mathbf{w}\| \leq \frac{1}{2} \|\mathbf{w}\| \leq \frac{1}{2} (\|P_L \mathbf{w}\| + \|P_{L^\perp} \mathbf{w}\|),$$

so that  $\|P_L \mathbf{w}\| \leq \|P_{L^\perp} \mathbf{w}\|$ . Now, using the triangle inequality,

$$\|P_{L^\perp} \mathbf{w}\| \leq \|P_{L^\perp} \mathbf{z}\| + |a_1| \|P_{L^\perp} \mathbf{v}_1\| + \dots + |a_d| \|P_{L^\perp} \mathbf{v}_d\|. \tag{30}$$

Because  $\mathbf{z} \in B(L, \eta)$ , we have  $\|P_{L^\perp} \mathbf{z}\| \leq \eta$ . For the other terms, we have  $\|P_{L^\perp} \mathbf{v}_k\| \leq \eta/\lambda_k$  as before, and, using the fact that, by construction, the  $\mathbf{v}_1, \dots, \mathbf{v}_d$  are orthonormal with  $A_k = \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$  and  $\|P_{A_{k-1}^\perp} \mathbf{z}\| \leq \lambda_k$ , together with the Cauchy-Schwartz inequality, we also have

$$|a_k| = |\mathbf{v}_k^T \mathbf{z}| = \left| \mathbf{v}_k^T P_{A_{k-1}^\perp} \mathbf{z} \right| \leq \lambda_k.$$

Hence, the RHS in (30) is bounded by  $(d + 1)\eta$ , implying

$$\|\mathbf{w}\| \leq \|P_L \mathbf{w}\| + \|P_{L^\perp} \mathbf{w}\| \leq 2\|P_{L^\perp} \mathbf{w}\| \leq 2(d + 1)\eta.$$

We then let  $C = \max(C_1, 2d + 2)$ . □

Below we provide an upper bound on the volume of the three-way intersection of the neighborhood of a surface, a ball centered at a point on the surface and the neighborhood of an affine  $d$ -dimensional subspace passing through that point, in terms of the angle between this subspace and the tangent to the surface at that same point. The principal angles between linear subspaces  $L, L' \in \mathcal{A}_d$ , denoted by

$$\frac{\pi}{2} \geq \theta_1(L, L') \geq \dots \geq \theta_d(L, L') \geq 0,$$

are recursively defined as follows:

$$\cos \theta_r(L, L') = \min_{\mathbf{u} \in L} \min_{\mathbf{u}' \in L'} \mathbf{u}^T \mathbf{u}' = \mathbf{u}_r^T \mathbf{u}'_r,$$



subject to

$$\begin{aligned} \|\mathbf{u}\| &= \|\mathbf{u}'\| = 1; \\ \mathbf{u}^T \mathbf{u}_s &= 0, \forall s = 1, \dots, r - 1; \\ \mathbf{u}'^T \mathbf{u}'_s &= 0, \forall s = 1, \dots, r - 1. \end{aligned}$$

Note that the orthogonality constraints are void when  $r = 1$ . (Some authors use the reverse ordering, e.g, [26].)

**Lemma 5.** Consider a surface  $S \in \mathcal{S}_d^2(\kappa)$ . Suppose  $\epsilon \geq \eta \vee \tau$ ,  $\eta \geq \epsilon^2$  and  $\tau > 0$ . Let  $\Psi$  be the uniform distribution on  $B(S, \tau)$ . For  $\mathbf{s} \in S$ , let  $T_{\mathbf{s}}$  be the tangent space to  $S$  at  $\mathbf{s}$ . Then for  $L \in \mathcal{A}_d$  containing  $\mathbf{s}$ ,

$$\Psi(B(\mathbf{s}, \epsilon) \cap B(L, \eta)) \prec \epsilon^d (1 \wedge (\eta/\tau))^{D-d} \prod_{j=1}^d \left( 1 \wedge \frac{\eta \vee \tau}{\epsilon \theta_j(L, T_{\mathbf{s}})} \right).$$

*Proof.* Fix  $\mathbf{s} \in S$  and  $L \in \mathcal{A}_d$  containing  $\mathbf{s}$ , and let  $T := T_{\mathbf{s}}$  and  $\theta_j := \theta_j(L, T)$  for short. By definition,

$$\Psi(B(\mathbf{s}, \epsilon) \cap B(L, \eta)) = \frac{\text{vol}_D(B(S, \tau) \cap B(\mathbf{s}, \epsilon) \cap B(L, \eta))}{\text{vol}_D(B(S, \tau))}.$$

By Lemma 3, it suffices to show that

$$\text{vol}_D(B(S, \tau) \cap B(L, \eta) \cap B(\mathbf{s}, \epsilon)) \prec \epsilon^d (\eta \wedge \tau)^{D-d} \prod_{j=1}^d \left( 1 \wedge \frac{\eta \vee \tau}{\epsilon \theta_j} \right).$$

We divide the proof into two cases; though the proof is similar for both, the first case is simpler and allows us to introduce the main ideas with ease before generalizing to the second case.

*Case  $\epsilon^2 \leq \tau$ .* We use Lemma 2 and the fact that  $\tau \geq \epsilon^2$ , to get

$$B(S, \tau) \cap B(\mathbf{s}, \epsilon) \subset B(T, (1 + \kappa)\tau) \cap B(\mathbf{s}, \epsilon). \tag{31}$$

Ignoring the constant factor  $1 + \kappa$ , we bound

$$\text{vol}_D(B(T, \tau) \cap B(L, \eta) \cap B(\mathbf{s}, \epsilon)).$$

We may assume without loss of generality that  $\mathbf{s}$  is the origin and

$$T = \text{span}\{\mathbf{e}_1, \dots, \mathbf{e}_d\}, \text{ and}$$

$$L = \text{span}\{(\cos \theta_1)\mathbf{e}_1 + (\sin \theta_1)\mathbf{e}_{d+1}, \dots, (\cos \theta_d)\mathbf{e}_d + (\sin \theta_d)\mathbf{e}_{2d}\}.$$

Then

$$\begin{aligned} B(T, \tau) &= \{(z_1, \dots, z_D) : \sum_{j>d} z_j^2 \leq \tau^2\}; \\ B(L, \eta) &= \{(z_1, \dots, z_D) : \sum_{j \leq d} (z_j \sin \theta_j - z_{d+j} \cos \theta_j)^2 + \sum_{j>2d} z_j^2 \leq \eta^2\}; \\ B(\mathbf{s}, \epsilon) &= \{(z_1, \dots, z_D) : \sum_j z_j^2 \leq \epsilon^2\}. \end{aligned}$$

Take  $j \leq d$ ; since  $|z_{d+j}| \leq \tau$ , we have

$$|z_j \sin \theta_j - z_{d+j} \cos \theta_j| \leq \eta \quad \Rightarrow \quad |z_j| \leq 2(\eta \vee \tau) / \sin \theta_j \leq \pi(\eta \vee \tau) / \theta_j.$$

Therefore,

$$B(T, \tau) \cap B(L, \eta) \cap B(\mathbf{s}, \epsilon) \subset \prod_{j=1}^d \left[ -\epsilon \wedge \frac{\pi(\eta \vee \tau)}{\theta_j}, \epsilon \wedge \frac{\pi(\eta \vee \tau)}{\theta_j} \right] \times B_{D-d}(0, \eta \wedge \tau).$$

From that we obtain the desired bound.

*Case  $\tau \leq \epsilon^2$ .* The arguments here are a little different and we simply bound  $\text{vol}_D(B(S, \tau) \cap B(\mathbf{s}, \epsilon))$ . Assume that  $\epsilon < 1/(8\kappa)$ . Because  $P_T$  is contractile, we have

$$P_T(S \cap B(\mathbf{s}, \epsilon)) \subset T \cap B(\mathbf{s}, \epsilon) = B_d(0, \epsilon),$$

so that, by Lemma 2,

$$S \cap B(\mathbf{s}, \epsilon) \subset P_T^{-1}(B_d(0, \epsilon)),$$

where  $P_T^{-1} : B_d(0, \epsilon) \rightarrow S \cap B(\mathbf{s}, \epsilon)$ . Hence,

$$B(S, \tau) \cap B(\mathbf{s}, \epsilon) \subset \{(\mathbf{a}, \mathbf{b}) : \mathbf{a} \in B_d(0, \epsilon), \|\mathbf{b} - P_T^{-1}(\mathbf{a})\| \leq \tau\}.$$

And by direct integration, the set on the RHS has  $D$ -volume of order  $\epsilon^d \tau^{D-d}$  since  $P_T^{-1}$  is Lipschitz on  $B_d(0, \epsilon)$  by Lemma 2. □

A companion of the previous result, the following lemma provides a lower bound on the angle between the affine subspace and the tangent.

**Lemma 6.** *Let  $\epsilon, \eta > 0$ , and take  $S \in \mathcal{S}_d^2(\kappa)$ . Suppose  $L \in \mathcal{A}_d$  is such that  $B(L, \eta)$  contains  $\mathbf{s} \in S$  and  $\mathbf{y} \in B(\mathbf{s}, \epsilon)$ . Let  $T_{\mathbf{s}}$  the tangent to  $S$  at  $\mathbf{s}$ . Then*

$$\theta_1(L, T_{\mathbf{s}}) \geq \frac{\text{dist}(\mathbf{y}, S) - 2\kappa\epsilon^2 - \eta}{\epsilon + \eta}.$$

*Proof.* Let  $T$  denote  $T_{\mathbf{s}}$  for short, and let  $L'$  be the line passing through  $\mathbf{s}$  and  $P_L(\mathbf{y})$ . Since  $L' \subset L$ , we have  $\theta_1(L, T) \geq \theta_1(L', T)$ , and using the triangle inequality and the fact that  $\theta \geq \sin \theta$ , for  $\theta \geq 0$ , this is bounded below by

$$\frac{\text{dist}(P_L(\mathbf{y}), T)}{\text{dist}(P_L(\mathbf{y}), \mathbf{s})} \geq \frac{\text{dist}(\mathbf{y}, T) - \eta}{\text{dist}(\mathbf{s}, \mathbf{y}) + \eta}.$$

The denominator does not exceed  $\epsilon + \eta$ . For the numerator,

$$\text{dist}(\mathbf{y}, T) = \|P_T(\mathbf{y}) - \mathbf{y}\| \geq \text{dist}(\mathbf{y}, S) - \text{dist}(P_T(\mathbf{y}), S).$$

Since  $\|\mathbf{y} - \mathbf{s}\| \leq \epsilon$ , we have  $P_T(\mathbf{y}) \in T \cap B(\mathbf{s}, \epsilon)$ , so that  $\text{dist}(P_T(\mathbf{y}), S) \leq 2\kappa\epsilon^2$  by Lemma 2. Consequently, the numerator is bounded from below by  $\text{dist}(\mathbf{y}, S) - \kappa\epsilon^2 - \eta$ . □

Next is another result estimating some volume intersections. It is similar to Lemma 5, though the conditions are different.

**Lemma 7.** Consider a surface  $S \in \mathcal{S}_d^2(\kappa)$ . Let  $\Psi$  be the uniform distribution on  $B(S, \tau)$ . Then for  $\epsilon \geq \eta$  and  $\tau > 0$ ,

$$\sup_{\mathbf{y}, L} \Psi(B(\mathbf{y}, \epsilon) \cap B(L, \eta)) \prec \epsilon^d (1 \wedge (\eta/\tau))^{D-d},$$

where the supremum is over  $\mathbf{y} \in \mathbb{R}^D$  and  $L \in \mathcal{A}_d$ , and the implicit constants depend only on  $\kappa, d$ . Also, for  $\epsilon \geq 10\eta$ ,  $\eta \geq 10\kappa\epsilon^2$  and  $\tau > 0$ , and any  $\mathbf{x} \in B(S, \tau)$ ,

$$\sup_L \Psi(B(\mathbf{x}, \epsilon) \cap B(L, \eta)) \succ \epsilon^d (1 \wedge (\eta/\tau))^{D-d}.$$

*Proof.* The proof is similar to that of Lemma 5. We divide the proof into two parts.

*Upper bound.* Let  $\mathbf{x} \in B(S, \tau) \cap B(\mathbf{y}, \epsilon) \cap B(L, \eta)$ . When  $\eta \geq \tau$ , we use

$$B(S, \tau) \cap B(\mathbf{y}, \epsilon) \cap B(L, \eta) \subset B(S, \tau) \cap B(\mathbf{x}, 2\epsilon),$$

while, when  $\eta \leq \tau$ , we use

$$B(S, \tau) \cap B(\mathbf{y}, \epsilon) \cap B(L, \eta) \subset B(L, \eta) \cap B(\mathbf{x}, 2\epsilon).$$

In both cases, we conclude with Lemma 3.

*Lower bound.* Let  $\mathbf{s}$  be the point on  $S$  closest to  $\mathbf{x}$ , with tangent subspace  $T$ . When  $\eta \geq 2\tau + 4\kappa\epsilon^2$ , take as  $L$  the translate of  $T$  passing through  $\mathbf{x}$  and use Lemma 2 to get

$$B(S, \tau) \cap B(\mathbf{x}, \epsilon) \subset B(T, \tau + \kappa(\tau + \epsilon)^2) \subset B(L, \eta),$$

and therefore

$$B(S, \tau) \cap B(\mathbf{x}, \epsilon) \cap B(L, \eta) \supset B(S, \tau) \cap B(\mathbf{x}, \epsilon).$$

We then use Lemma 3. Now, suppose  $\eta \leq 2\tau + 4\kappa\epsilon^2$  and notice that, since  $\eta \geq 10\kappa\epsilon^2$ , we have  $\tau \geq 3\kappa\epsilon^2$ . First, assume that  $\epsilon \geq 10\tau$ . We use Lemma 2 to get

$$B(S, \tau) \cap B(\mathbf{x}, \epsilon) \supset B(T, \tau - 2\kappa\epsilon^2) \cap B(\mathbf{s}, \epsilon) \cap B(\mathbf{x}, \epsilon),$$

and therefore,

$$B(S, \tau) \cap B(\mathbf{x}, \epsilon) \cap B(L, \eta) \supset B(T, \tau - 2\kappa\epsilon^2) \cap B(L, \eta) \cap B(\mathbf{s}, \epsilon) \cap B(\mathbf{x}, \epsilon). \quad (32)$$

Without loss of generality, assume that  $\mathbf{x}$  is the origin,  $L = \text{span}\{\mathbf{e}_1, \dots, \mathbf{e}_d\}$ . Since the volume is least when  $\|\mathbf{x} - \mathbf{s}\| = \tau$ , assume that  $\mathbf{s} = \tau\mathbf{e}_{d+1}$  (seen as a point in space). Define  $\nu = (\eta + 2\kappa\epsilon^2)/2$  and note that  $\nu \leq \eta \wedge (2\tau)$  by the

conditions on  $\eta$  and  $\tau$ . Then

$$\begin{aligned} B(T, \tau - 2\kappa\epsilon^2) \cap B(L, \eta) &\supset \{(z_1, \dots, z_D) : \sum_{j>d+1} z_j^2 + (z_{d+1} - \nu)^2 \leq (\eta/3)^2\}; \\ B(\mathbf{s}, \epsilon) &= \{(z_1, \dots, z_D) : \sum_{j \neq d+1} z_j^2 + (z_{d+1} - \tau)^2 \leq \epsilon^2\}; \\ B(\mathbf{x}, \epsilon) &= \{(z_1, \dots, z_D) : \sum_j z_j^2 \leq \epsilon^2\}. \end{aligned}$$

By the conditions imposed on  $\epsilon, \eta, \tau$ , the RHS in (32) contains

$$B_d(0, \epsilon/10) \times [\eta/4, 3\eta/4] \times B_{D-d-1}(0, \eta/10).$$

Therefore the result. Finally assume that  $\tau \geq \epsilon/10$  and take  $L$  passing through  $\mathbf{x}$  and  $\mathbf{z} = (1 - \lambda)\mathbf{x} + \lambda\mathbf{s}$ , where  $\lambda = \epsilon/(2\tau)$ . We have  $\|\mathbf{z} - \mathbf{x}\| \leq \epsilon/2$  and  $\|\mathbf{z} - \mathbf{s}\| \leq \tau - \epsilon/2$ , so that  $B(\mathbf{z}, \epsilon/2) \subset B(S, \tau) \cap B(\mathbf{x}, \epsilon)$  by the triangle inequality. Hence,

$$B(S, \tau) \cap B(L, \eta) \cap B(\mathbf{x}, \epsilon) \supset B(L, \eta) \cap B(\mathbf{z}, \epsilon/2).$$

We then conclude with Lemma 3. □

**Lemma 8.** *Let  $\Psi$  be the uniform distribution on a measurable subset  $A \subset \mathbb{R}^D$  of positive  $D$ -volume. Then for  $\epsilon \geq \eta$ ,*

$$\sup_{\mathbf{y}, L} \Psi(B(\mathbf{y}, \epsilon) \cap B(L, \eta)) \prec \epsilon^d \eta^{D-d},$$

where the supremum is over  $\mathbf{y} \in \mathbb{R}^D$  and  $L \in \mathcal{A}_d$ , and the implicit constant depends only on  $d$  and  $\text{vol}_D(A)$ .

*Proof.* The proof is parallel to (and simpler than) that of Lemma 7. We omit details. □

### A.3. A perturbation bound

In the proof of Theorem 1, we follow the strategy outlined in [42] based on verifying the following conditions (where (A4) has been simplified). Let  $I_k = \{i : \mathbf{x}_i \in \mathcal{X}_k\}$  and let  $\mathring{\mathbf{W}}_k$  denote the matrix with coefficients indexed by  $i, j \in I_k$  and defined as

$$\mathring{W}_{ij} = \sum_{i_1, \dots, i_{m-2} \in I_k} \alpha_d(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_{m-2}}), \quad \mathring{D}_i = \sum_{j \in I_k} \mathring{W}_{ij}.$$

Let  $\mathring{W}_{ij} = 0$  if  $i \in I_k, j \in I_\ell$ , with  $k \neq \ell$ . Those are the coefficients of  $\mathbf{W}$  and  $\mathbf{D}$  under infinite separation, i.e., assuming  $\delta = \infty$ . (In fact  $\delta > \epsilon + 2\tau$  is enough since we use the simple kernel.)

(A1) For all  $k$ , the second largest eigenvalue of  $\mathring{\mathbf{W}}_k$  is bounded above by  $1 - \gamma$ .

(A2) For all  $k, \ell$ , with  $k \neq \ell$ ,

$$\sum_{i \in I_k} \sum_{j \in I_\ell} \frac{W_{ij}^2}{\mathring{D}_i \mathring{D}_j} \leq \nu_1.$$

(A3) For all  $k$  and all  $i \in I_k$ ,

$$\frac{1}{\mathring{D}_i} \sum_{j \notin I_k} W_{ij} \leq \nu_2 \left( \sum_{s,t \in I_k} \frac{W_{st}^2}{\mathring{D}_s \mathring{D}_t} \right)^{-1/2}.$$

(A4) For all  $k$  and all  $i, j \in I_k$ ,  $\mathring{D}_i \leq Q \mathring{D}_j$ .

The following result is a slightly modified version of [42, Th. 2], stated and proved in [3, Th. 7]. See also [11, Th. 4.5]. Recall the matrix  $\mathbf{V}$  defined in Algorithm 1.

**Theorem 2.** Let  $\mathbf{v}_1, \dots, \mathbf{v}_N$  denote the row vectors of  $\mathbf{V}$ . Under (A1)-(A4), there is an orthonormal set  $\{\mathbf{r}_1, \dots, \mathbf{r}_K\} \subset \mathbb{R}^K$  such that,

$$\frac{1}{N} \sum_{k=1}^K \sum_{i \in I_k} \|\mathbf{v}_i - \mathbf{r}_k\|^2 \leq 4Q\gamma^{-2}(K^2\nu_1 + K\nu_2^2).$$

### Appendix B: Main proofs

For a set  $A$ , its cardinality is denoted by  $\#A$ . Throughout the paper,  $C$  denotes a generic constant that does not depend on the sample size  $N$  and satisfies  $C \geq 1$ .

#### B.1. Proof of Theorem 1

Given Theorem 2, we turn to proving that the four conditions (A1)-(A4) hold with probability tending to one with  $\nu_1 = \nu_2^2 = (\rho_N/\zeta)^{-m/2}$ ,  $\gamma > C^{-m}N^{-2}$  and  $Q \leq C^m$  for some constant  $C > 0$ . Since  $m \log(\rho_N/\zeta) \gg \log N$ , this implies

$$\max_{i=1, \dots, N} \min_{k=1, \dots, K} \|\mathbf{v}_i - \mathbf{r}_k\| \rightarrow 0.$$

Therefore, since the  $\mathbf{r}_k$ 's are themselves orthonormal, the  $K$ -means algorithm with near-orthogonal initialization outputs the perfect clustering.

We restrict ourselves to the case where  $\tau \leq (\rho_N^2 \log(N)/N)^{1/d}$ , for otherwise  $\eta \geq \epsilon$  and HOSC is essentially SC, studied in [3]. With that bound on  $\tau$ , (12) reduces to  $\epsilon \geq (\rho_N^2 \log(N)/N)^{1/d}$ . By the same token, we assume that  $\eta \leq \epsilon$ , so that  $\epsilon \geq \eta \geq \tau + \rho_N \epsilon^2$ .

To verify conditions (A2), (A3) and (A4) we need to estimate the degree of each vertex under infinite separation and the edge weights under finite separation. We start with the case of infinite separation.

**Proposition 7.** *With probability at least  $1 - N^{-\rho_N^2/(K\zeta)}$ ,*

$$\mathbf{1}_{\{\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon/2\}} N_k \epsilon^d \prec \mathring{W}_{ij}^{1/(m-2)} \prec \mathbf{1}_{\{\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon\}} N_k \epsilon^d; \tag{33}$$

and also,

$$\mathring{D}_i^{1/(m-1)} \asymp N_k \epsilon^d, \tag{34}$$

uniformly over  $i, j \in I_k$  and  $k = 1, \dots, K$ .

*Proof.* Within a cluster, the linear approximation factor in (3) is a function of the proximity factor. This is due to Lemma 2. Formally, let  $\mathring{G}_{i,\epsilon}$  denote the degree of  $\mathbf{x}_i$  in the neighborhood graph built by SC, i.e.

$$\mathring{G}_{i,\epsilon} = \#\{j \in I_k, j \neq i : \mathbf{x}_j \in B(\mathbf{x}_i, \epsilon)\},$$

Then Proposition 7 is a direct consequence of Lemma 9, which relates  $\mathring{G}_{i,\epsilon}$  to  $\mathring{W}_{ij}$  and  $\mathring{D}_i$ , and of Proposition 8, which estimates  $\mathring{G}_{i,\epsilon}$ .  $\square$

**Lemma 9.** *We have*

$$\mathbf{1}_{\{\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon/2\}} (\mathring{G}_{i,\epsilon/2} - 1)^{\{m-2\}} \leq \mathring{W}_{ij} \leq \mathbf{1}_{\{\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon\}} (\mathring{G}_{i,\epsilon} - 1)^{\{m-2\}},$$

and,

$$\mathring{G}_{i,\epsilon/2} (\mathring{G}_{i,\epsilon/2} - 1)^{\{m-2\}} \leq \mathring{D}_i \leq \mathring{G}_{i,\epsilon} (\mathring{G}_{i,\epsilon} - 1)^{\{m-2\}},$$

where  $r^{\{m\}} = r(r-1) \cdots (r-m+1)$ .

Note that  $r^{\{m\}} \leq r^m$ , and  $r^{\{m\}} \geq (r/3)^m$  for  $r \geq m$ .

*Proof.* We focus on the first expression, as the second expression is obtained by summing the first one over  $j \in I_k, j \neq i$ , where  $k$  is such that  $i \in I_k$ . Therefore, fix  $i, j \in I_k$ . The upper bound on  $\mathring{W}_{ij}$  comes from the fact that

$$\text{diam}(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_1, \dots, \mathbf{x}_{m-2}) \leq \epsilon \quad \Rightarrow \quad \mathbf{x}_1, \dots, \mathbf{x}_{m-2} \in B(\mathbf{x}_i, \epsilon).$$

The lower bound comes from

$$\mathbf{x}_1, \dots, \mathbf{x}_{m-2} \in B(\mathbf{x}_i, \epsilon/2) \quad \Rightarrow \quad \text{diam}(\mathbf{x}_i, \mathbf{x}_1, \dots, \mathbf{x}_{m-2}) \leq \epsilon,$$

and the fact that,

$$\mathbf{x}_1, \dots, \mathbf{x}_{m-2} \in B(S_k, \tau) \cap B(\mathbf{x}_i, \epsilon/2) \quad \Rightarrow \quad \mathbf{x}_1, \dots, \mathbf{x}_{m-2} \in B(T_{\mathbf{s}_i}, \eta),$$

where  $\mathbf{s}_i$  is the point on  $S_k$  closest to  $\mathbf{x}_i$ . Indeed, take  $\mathbf{x} \in B(S_k, \tau) \cap B(\mathbf{x}_i, \epsilon/2)$  and let  $\mathbf{s} \in S_k$  such that  $\|\mathbf{x} - \mathbf{s}\| \leq \tau$ . By the triangle inequality,  $\|\mathbf{s} - \mathbf{s}_i\| \leq \epsilon/2 + 2\tau$ , so that, by Lemma 2,  $\mathbf{s} \in B(T_{\mathbf{s}_i}, \kappa(\epsilon/2 + 2\tau)^2)$ . Therefore,  $\mathbf{x} \in B(T_{\mathbf{s}_i}, \kappa(\epsilon/2 + 2\tau)^2 + \tau)$ . We then conclude with the fact that  $\tau \leq \epsilon$  and  $\eta \geq \tau + \rho_N \epsilon^2$ , with  $\rho_N \rightarrow \infty$ .  $\square$

Note that  $N \leq K\zeta N_k$ , which together with (12) implies

$$N_k \epsilon^d (1 \wedge (\epsilon/\tau))^{D-d} \geq \rho_N^2 / (K\zeta) \log N, \quad \forall k = 1, \dots, K. \tag{35}$$

The following bound on  $\mathring{G}_{i,\epsilon}$  is slightly more general than needed at this point.

**Proposition 8.** Assume that (35) holds. Then with probability at least  $1 - N^{-\rho_N^2/(K\zeta)}$ ,

$$\mathring{G}_{i,\epsilon} \asymp N_k \epsilon^d (1 \wedge (\epsilon/\tau))^{D-d}, \tag{36}$$

uniformly over  $i \in I_k$  and  $k = 1, \dots, K$ .

*Proof.* This is done in the proof of [3, Eq. (A4)] and we repeat the arguments here for future reference. Let  $\Psi_k$  denote the uniform distribution on  $B(S_k, \tau)$ . By definition, for any (measurable) set  $A$ ,

$$\Psi_k(A) = \frac{\text{vol}_D(A \cap B(S_k, \tau))}{\text{vol}_D(B(S_k, \tau))}. \tag{37}$$

Since  $\mathring{G}_{i,\epsilon}$  is the sum of independent Bernoulli random variables, by Lemma 1, it suffices to bound it in expectation. Using Lemma 3, we have

$$\mathbb{E} \left( \mathring{G}_{i,\epsilon} \right) = N_k \Psi_k(B(\mathbf{x}_i, \epsilon)) \asymp N_k \epsilon^d (1 \wedge (\epsilon/\tau))^{D-d}.$$

Applying Lemma 1 and (35), we then get

$$\mathbb{P} \left( \mathring{G}_{i,\epsilon} > 16 \mathbb{E} \mathring{G}_{i,\epsilon} \right) \vee \mathbb{P} \left( \mathring{G}_{i,\epsilon} < \mathbb{E} \mathring{G}_{i,\epsilon} / 8 \right) \leq N^{-2(\rho_N^2/(K\zeta))}.$$

We then apply the union bound and use the fact that  $N \cdot N^{-2(\rho_N^2/(K\zeta))} \leq N^{-\rho_N^2/(K\zeta)}$ , since  $\rho_N^2 \rightarrow \infty$ . □

We now turn to bounding the size of the edge weights  $W_{ij}$  under finite separation. We do so by comparing them with the edge weights under infinite separation.

**Proposition 9.** With probability at least  $1 - N^{-\rho_N}$ ,

$$(W_{ij} - \mathring{W}_{ij})^{1/(m-2)} \prec \mathbf{1}_{\{\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon\}} N \epsilon^d / \rho_N. \tag{38}$$

uniformly over  $i \in I_k, j \in I_\ell$  and  $k, \ell = 1, \dots, K$ .

*Proof.* If  $k = \ell$ ,  $W_{ij} - \mathring{W}_{ij}$  is the sum of  $\alpha_d(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_{m-2}})$  over (distinct)  $i_1, \dots, i_{m-2}$  that are not all in  $I_k$ . When  $k \neq \ell$ ,  $\mathring{W}_{ij} = 0$  and  $W_{ij}$  is again the same sum except this time over all (distinct)  $i_1, \dots, i_{m-2}$ . Both situations are similar and we focus on the latter. We assume that  $\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon$ , for otherwise the bound is trivially satisfied. Note that this implies that  $\rho_N \eta \leq \delta - 2\tau \leq \epsilon$ .

Define

$$G_{i,\epsilon} = \#\{j \neq i : \mathbf{x}_j \in B(\mathbf{x}_i, \epsilon)\}, \tag{39}$$

which is the equivalent of  $\mathring{G}_{i,\epsilon}$  under finite separation, as well as

$$H_{i,\epsilon,\eta}(L) = \#\{j \neq i : \mathbf{x}_j \in B(\mathbf{x}_i, \epsilon) \cap B(L, \eta)\},$$

and

$$H_{i,j,\epsilon,\eta}^* = \max_M H_{i,\epsilon,\eta}(L_M),$$

where the maximum is over all  $M \subset \{1, \dots, N\}$ , of size  $|M| = d + 1$  such that  $\mathbf{x}_j \in B(L_M, \eta)$ . Then Proposition 9 is a direct consequence of Lemma 10, which relates  $G_{i,\epsilon}$  and  $H_{i,j,\epsilon,\eta}^*$  to  $W_{ij}$ , and of Propositions 10 and 11, which bound  $G_{i,\epsilon}$  and  $H_{i,j,\epsilon,\eta}^*$ , respectively.  $\square$

**Lemma 10.** *There is a constant  $C > 0$  such that*

$$W_{ij} \leq (G_{i,\epsilon} + 1)^{d+1} (H_{i,j,\epsilon,C\eta}^*)^{\{m-d-1\}}. \tag{40}$$

*Proof.* By definition of the affinity (3) and the triangle inequality, we have

$$W_{ij} \leq \sum_M \mathbf{1}_{\{\exists L \in L_d: \mathbf{x}_n \in B(\mathbf{x}_i, \epsilon) \cap B(L, \eta), \forall n \in M \cup \{i, j\}\}},$$

where the sum is over  $M \subset \{1, \dots, N\}$  such that  $|M| = m - 2$  and  $i, j \notin M$ . For a subset  $M \subset \{1, \dots, N\}$ , of size  $|M| = d + 1$ , let  $L_M$  denote the affine subspace spanned by  $\{\mathbf{x}_n, n \in M\}$ . By Lemma 4, we may limit ourselves to subspaces  $L$  that are generated by  $d + 1$  data points, obtaining

$$\begin{aligned} W_{ij} &\leq \sum_M \mathbf{1}_{\{\mathbf{x}_n \in B(\mathbf{x}_i, \epsilon), \forall n \in M\}} \\ &\quad \times \sum_{M'} \mathbf{1}_{\{\mathbf{x}_n \in B(\mathbf{x}_i, \epsilon) \cap B(L_M, C\eta), \forall n \in M' \cup \{i, j\}\}}, \end{aligned} \tag{41}$$

where  $M$  is any subset of  $\{1, \dots, N\}$  of size  $d + 1$ , and  $M'$  is any subset of  $\{1, \dots, N\} \setminus (M \cup \{i, j\})$  such that  $M' \cup M \cup \{i, j\}$  is of size  $m$ . Such an  $M$  is of size at most  $m - d - 1$  and does not contain  $i$  or  $j$ . For any  $M$ ,  $B(\mathbf{x}_i, \epsilon) \cap B(L_M, C\eta)$  contains at most  $H_{i,j,\epsilon,C\eta}^*$  data points other than  $\mathbf{x}_i$  and  $\mathbf{x}_j$ , so that the second sum is bounded by  $(H_{i,j,\epsilon,C\eta}^*)^{m-d-1}$  independently of  $M$ . Similarly,  $B(\mathbf{x}_i, \epsilon)$  contains at most  $G_{i,\epsilon} + 1$  points, so the first sum is bounded by  $(G_{i,\epsilon} + 1)^{d+1}$ . The result follows.  $\square$

**Proposition 10.** *Assume that (35) holds. Then with probability at least  $1 - N^{-\rho_N^2/(K\zeta)}$ ,*

$$G_{i,\epsilon} \prec N\epsilon^d (1 \wedge (\epsilon/\tau))^{D-d}, \tag{42}$$

*uniformly over  $i = 1, \dots, N$ .*

*Proof.* We have

$$\mathbb{E}(G_{i,\epsilon}) = \sum_{\ell} N_{\ell} \Psi_{\ell}(B(\mathbf{x}_i, \epsilon)).$$

Now, by Lemma 3, for all  $\ell$  such that  $\text{dist}(\mathbf{x}_i, S_{\ell}) \leq \epsilon + \tau$ ,

$$\Psi_{\ell}(B(\mathbf{x}_i, \epsilon)) \prec \epsilon^d (1 \wedge (\epsilon/\tau))^{D-d}.$$

Hence,

$$\mathbb{E}(G_{i,\epsilon}) \prec N\epsilon^d (1 \wedge (\epsilon/\tau))^{D-d}.$$

We then use Lemma 1 and (35).  $\square$



**Proposition 11.** *With probability at least  $1 - N^{-\rho_N}$ ,*

$$H_{i,j,\epsilon,\eta}^* \prec \frac{N\epsilon^d}{\rho_N}, \quad (43)$$

*uniformly over  $i \in I_k, j \in I_\ell$  and  $k \neq \ell$  in  $\{1, \dots, K\}$ .*

*Proof.* For  $L \in \mathcal{A}_d$ ,  $H_{i,\epsilon,\eta}(L)$  is a sum of independent Bernoulli random variables, with expectation

$$\mathbb{E}(H_{i,\epsilon,\eta}(L)) = \sum_{\ell} N_{\ell} \Psi_{\ell}(B(\mathbf{x}_i, \epsilon) \cap B(L, \eta)).$$

Take  $\ell$  such that  $B(S_{\ell}, \tau) \cap B(\mathbf{x}_i, \epsilon) \cap B(L, \eta) \neq \emptyset$ , and let  $\mathbf{x}$  be in that set and  $\mathbf{s}$  be the point on  $S_{\ell}$  closest to  $\mathbf{x}$ . Then by the triangle inequality and the fact that  $\epsilon \geq \eta \geq \tau$ ,

$$B(S_{\ell}, \tau) \cap B(\mathbf{x}_i, \epsilon) \cap B(L, \eta) \subset B(S_{\ell}, \tau) \cap B(\mathbf{s}, 3\epsilon) \cap B(L_{\mathbf{s}}, 3\eta),$$

where  $L_{\mathbf{s}}$  is the translate of  $L$  passing through  $\mathbf{s}$ . Therefore,

$$\Psi_{\ell}(B(\mathbf{x}_i, \epsilon) \cap B(L, \eta)) \leq \mathbf{1}_{\{\text{dist}(\mathbf{x}_i, S_{\ell}) \leq \epsilon + \tau\}} \cdot \sup_{\mathbf{s} \in S_{\ell}} \Psi_{\ell}(B(\mathbf{s}, 3\epsilon) \cap B(L_{\mathbf{s}}, 3\eta)).$$

Our focus is on  $L$  such that  $\mathbf{x}_i, \mathbf{x}_j \in B(L, \eta)$ , which transfers as  $\mathbf{x}_i, \mathbf{x}_j \in B(L_{\mathbf{s}}, 3\eta)$  by the triangle inequality. Since  $\mathbf{x}_i$  and  $\mathbf{x}_j$  belong to different clusters, for a given  $\ell$ , at least one of them does not belong to  $\mathcal{X}_{\ell}$ . Hence, by Lemma 6 and the fact that  $\delta \gg \eta \geq \tau + \kappa\epsilon^2$ ,  $\theta_1(L, T_{\mathbf{s}}) \succ \delta/\epsilon$  uniformly over  $\mathbf{s} \in S_{\ell}$  and  $\ell$ . (Remember that  $\theta_1(L, T)$  denotes the largest principal angle between  $L$  and  $T$ .) Together with Lemma 5, we thus get

$$\Psi_{\ell}(B(\mathbf{s}, 3\epsilon) \cap B(L_{\mathbf{s}}, 3\eta)) \leq C\epsilon^d(\eta/\delta).$$

Hence, by the fact that  $\delta \geq \rho_N\eta$ , we have

$$\mathbb{E}(H_{i,\epsilon,\eta}(L)) \leq CN\epsilon^d(\eta/\delta) \leq CN\epsilon^d/\rho_N.$$

With Lemma 1 and (12), we then get

$$\sup_L \mathbb{P}(H_{i,\epsilon,\eta}(L) > 16CN\epsilon^d/\rho_N) \leq N^{-2\rho_N}.$$

Hence, by the union bound,

$$\mathbb{P}(H_{i,j,\epsilon,\eta}^* > 16CN\epsilon^d/\rho_N) \leq N^{d+1} \cdot N^{-2\rho_N}. \quad (44)$$

The right hand side is bounded by  $N^{-\rho_N}$  eventually.  $\square$

We now turn to verifying (A1)-(A4).

- Verifying (A4): (34) immediately implies (A4) with  $Q = C^m$  for some constant  $C > 0$ .

- Verifying (A3): Take  $k = 1, \dots, K$ . By (33), (34) and (38),

$$\sum_{i,j \in I_k} \frac{W_{ij}^2}{\mathring{D}_i \mathring{D}_j} \prec \epsilon^{-2}(1 + (\rho_N/\zeta)^{-2(m-2)}) \prec \epsilon^{-2},$$

and also,

$$\frac{1}{\mathring{D}_i} \sum_{j \notin I_k} W_{ij} \prec (N/N_k) \epsilon^{-d} (\rho_N/\zeta)^{-(m-1)}.$$

Since  $N/N_k \leq N$ ,  $\epsilon \succ N^{-1/d}$  and  $m \log(\rho_N/\zeta) \gg \log N$ , we may take  $\nu_2 = (\rho_N/\zeta)^{-m/2}$ .

- Verifying (A2): Take  $k, \ell = 1, \dots, K$ , with  $k \neq \ell$ . Then by (33), (34) and (38),

$$\sum_{i \in I_k} \sum_{j \in I_\ell} \frac{W_{ij}^2}{\mathring{D}_i \mathring{D}_j} \prec \epsilon^{-2d} (\rho_N/\zeta)^{-2(m-2)}.$$

Since  $\epsilon \succ N^{-1/d}$  and  $m \log(\rho_N/\zeta) \gg \log N$ , we may take  $\nu_1 = (\rho_N/\zeta)^{-m}$ .

- Verifying (A1): As suggested in [42], we approach this through a lower bound on the Cheeger constant. Let  $\mathring{\mathbf{Z}}_k$  be the matrix obtained from  $\mathring{\mathbf{W}}_k$  following SC. That  $\mathring{\mathbf{Z}}_k$  has eigenvalue 1 with multiplicity 1 results from the graph being fully connected [14]. The Cheeger constant of  $\mathring{\mathbf{W}}_k$  is defined as:

$$h_k = \min_{|I| \leq N_k/2} \frac{\sum_{i \in I} \sum_{j \in I_k \setminus I} \mathring{W}_{ij}}{\sum_{i \in I} \mathring{D}_i},$$

where the minimum is over all subsets  $I \subset I_k$  of size  $|I| \leq N_k/2$ . The spectral gap of  $\mathring{\mathbf{Z}}_k$  is then at least  $h_k^2/2$ . By (33)-(34), there is a constant  $C > 0$  such that,

$$h_k \geq C^{-m} (N_k \epsilon^d)^{-1} \min_{|I| \leq N_k/2} \frac{\sum_{i \in I} \sum_{j \in I_k \setminus I} \mathbf{1}_{\{\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon/2\}}}{|I|}.$$

From here, the proof is identical to that of [3, Eq. (A1)], which bounds the minimum from below by  $1/N_k$ , so that  $h_k \geq C^{-m} N_k^{-1}$ .

**B.2. Proof of Proposition 6**

From the proof of Theorem 1, it suffices to verify that (A2) and (A3) still hold under the conditions of Proposition 6, and in view of (19), we may focus on  $W_{ij}$  for  $i \in I_k$  and  $j \in I_\ell$ , with  $k \neq \ell$ , such that  $\|\mathbf{x}_i - \mathbf{x}_j\| \leq \epsilon$  and with  $\mathbf{x}_j$  close to an intersection, specifically, for some  $p \neq \ell$ ,

$$\text{dist}(\mathbf{x}_j, S_\ell \cap S_p) \leq \nu, \quad \text{where } \nu := (\sin \theta_{\text{int}})^{-1} (\epsilon \wedge \rho_N \eta).$$

In fact, we show that, under the conditions of Proposition 6, with probability at least  $1 - \gamma_N$ , there is no such a pair of points  $(\mathbf{x}_i, \mathbf{x}_j)$ . For fixed  $(k, \ell, p)$ , the probability that  $\mathbf{x}_i \sim \Psi_k$  and  $\mathbf{x}_j \sim \Psi_\ell$  satisfy these conditions is

$$\mathbb{E} \left( \Psi_k(B(\mathbf{x}_j, \epsilon)) \mathbf{1}_{\{\mathbf{x}_j \in B(S_\ell \cap S_p, \nu)\}} \right), \tag{45}$$

after integrating over  $\mathbf{x}_i$ . By Lemma 3,

$$\Psi_k(B(\mathbf{x}_j, \epsilon)) \prec \epsilon^d.$$

where the implicit constant depends only on  $\kappa, d$ . Moreover, by condition (18),

$$\Psi_\ell(B(S_\ell \cap S_p, \nu)) \prec \nu^{d-d_{\text{int}}}.$$

Therefore, using the union bound, the probability that there is such a pair of points is of order not exceeding

$$\sum_{k, \ell} N_k N_\ell \cdot \epsilon^d \nu^{d-d_{\text{int}}} = N^2 \epsilon^d \nu^{d-d_{\text{int}}} = \gamma_N \rightarrow 0.$$

**B.3. Proof of Propositions 4 and 5**

Without loss of generality, we assume that  $\delta_0$  is small and that  $\eta \leq \epsilon/10$ . Let  $\Psi_0$  be the uniform distribution on  $(0, 1)^D \setminus \bigcup_k B(S_k, \delta_0)$ . By Lemma 3, this set has  $D$ -volume of order  $1 - O(K\delta_0^{D-d})$ , with  $K\delta_0^{D-d}$  small since  $K$  is fixed. Therefore, for  $A \subset (0, 1)^D$ ,

$$\Psi_0(A) \asymp \text{vol}_D \left( A \setminus \bigcup_k B(S_k, \delta_0) \right).$$

Let  $I_0 \subset \{1, \dots, N\}$  index the outliers and let  $N_0$  be the number of outliers.

In view of how the procedures (O1) and (O2) work, we need to bound the degrees of non-outliers from below and the degrees of outliers from above. The following lower bound holds

$$N_k \epsilon^d (1 \wedge (\eta/\tau))^{D-d} \geq (\rho_N / (K\zeta)) \log N, \quad \forall k = 1, \dots, K. \tag{46}$$

For (O1), it comes from (11)-(12) and the fact that, for all  $k \neq 0$ ,  $N_k \geq N / (K\zeta\rho_N)$ , since  $N \leq K\zeta N_k + N_0$ , implying  $N_k \geq (N - N_0) / (K\zeta)$ , and  $N - N_0 \geq N / \rho_N$  in our assumptions. For (O2), it comes from (15) and (16) (and the inequality holds with  $\rho_N$  in place of  $\rho_N / (K\zeta)$ ). In the same vein,

$$N_k (1 \wedge (\eta/\tau))^{D-d} \gg N \eta^{D-d}, \quad \forall k = 1, \dots, K. \tag{47}$$

We prove a result that is more general than what we need now.

**Proposition 12.** Assume (46) and (47). Then with probability at least  $1 - N^{-\rho_N/(K\zeta)}$ ,

$$N_k \epsilon^d (1 \wedge (\eta/\tau))^{D-d} \prec D_i^{1/(m-1)} \prec N_k \epsilon^d (1 \wedge (\eta/\tau))^{(D-d)(1-\frac{d+1}{m-1})}, \quad (48)$$

uniformly over  $i \in I_k$ ,  $k \neq 0$ ; and also,

$$\begin{aligned} D_i^{1/(m-1)} &\prec (N - N_0) \epsilon^d (1 \wedge (\eta/\tau))^{(D-d)(1-\frac{d+1}{m-1})} \xi^{1-\frac{d+1}{m-1}} \mathbf{1}_{\{\delta_0 \leq \epsilon + \tau\}} \\ &\quad + N \epsilon^d \eta^{(D-d)(1-\frac{d+1}{m-1})}, \end{aligned} \quad (49)$$

uniformly over  $i \in I_0$ , where  $\xi = 1$  if  $\tau \geq \epsilon$ , and  $\xi = 1 \wedge (\eta/\delta_0)$ , otherwise.

*Proof.* Define

$$H_{i,\epsilon,\eta} = \max_{L \in \mathcal{A}_d} H_{i,\epsilon,\eta}(L).$$

Proposition 12 is a direct consequence of Lemma 11 which relates  $D_i$  to  $G_{i,\epsilon}$  (defined in Section B.1) and  $H_{i,\epsilon,\eta}$ , and of Propositions 13 and 14 (together with (47)), which bound  $G_{i,\epsilon}$  and  $H_{i,\epsilon,\eta}$ , respectively.  $\square$

**Lemma 11.** There is a constant  $C > 0$  such that

$$H_{i,\epsilon/2,\eta}^{\{m-1\}} \leq D_i \leq G_{i,\epsilon}^{\{d+1\}} (H_{i,\epsilon,C\eta}^*)^{\{m-d-2\}}. \quad (50)$$

*Proof.* We get the upper bound by following the arguments in the proof of (38). For the lower bound, we simply have

$$\begin{aligned} D_i &\geq \sum_{M:|M|=m-1} \mathbf{1}_{\{\exists L \in L_d: \mathbf{x}_j \in B(\mathbf{x}_i, \epsilon/2) \cap B(L, \eta), \forall j \in M\}} \\ &\geq H_{i,\epsilon/2,\eta}^{\{m-1\}}. \end{aligned}$$

$\square$

The bounds for  $G_{i,\epsilon}$  and  $H_{i,\epsilon,\eta}$  that follow are more general than needed at this point. In particular, the case of large  $\tau$  will only be useful in Section C.

**Proposition 13.** Assume (46) holds with  $\epsilon$  in place of  $\eta$ . Then with probability at least  $1 - N^{-\rho_N/(K\zeta)}$ ,

$$N_k \epsilon^d (1 \wedge (\epsilon/\tau))^{D-d} \prec G_{i,\epsilon} \prec N_k \epsilon^d (1 \wedge (\epsilon/\tau))^{D-d} + N_0 \epsilon^D. \quad (51)$$

uniformly over  $i \in I_k$  and  $k = 1, \dots, K$ . Also,

$$G_{i,\epsilon} \prec (N - N_0) \epsilon^d (1 \wedge (\epsilon/\tau))^{D-d} \mathbf{1}_{\{\delta_0 \leq \epsilon + \tau\}} + N \epsilon^D. \quad (52)$$

uniformly over  $i \in I_0$

*Proof.* The proof is similar to that of Proposition 8. We bound  $G_{i,\epsilon}$  in expectation. Suppose  $i \in I_k$  with  $k \neq 0$ . Then by Lemma 3

$$\mathbb{E}(G_{i,\epsilon}) \geq N_k \Psi_k B(\mathbf{x}_i, \epsilon) \asymp N_k \epsilon^d (1 \wedge (\epsilon/\tau))^{D-d}.$$

For the upper bound, by Lemma 3 and the simple bound

$$\Psi_0(B(\mathbf{x}_i, \epsilon)) \prec \epsilon^D,$$

we have

$$\mathbb{E}(G_{i,\epsilon}) = \sum_{\ell} N_{\ell} \Psi_{\ell}(B(\mathbf{x}_i, \epsilon)) \prec (N - N_0) \epsilon^d (1 \wedge (\epsilon/\tau))^{D-d} + N_0 \epsilon^D,$$

with  $N - N_0 \leq (K\zeta)N_k$  for any  $k \neq 0$ . As in as in Proposition 8, we then use Lemma 1 together with (46) and the union bound, to conclude the proof of (51). The proof of (52) is identical, except that, when  $\delta_0 > \tau + \epsilon$ , we have  $\Psi_{\ell}(B(\mathbf{x}_i, \epsilon)) = 0$  if  $\ell \neq 0$  and  $i \in I_0$ .  $\square$

**Proposition 14.** *If (46) holds, then with probability at least  $1 - N^{-\rho N/(K\zeta)}$ ,*

$$H_{i,\epsilon/2,\eta} \succ N_k \epsilon^d (1 \wedge (\eta/\tau))^{D-d}, \quad H_{i,\epsilon/2,\eta}^* \prec N_k \epsilon^d (1 \wedge (\eta/\tau))^{D-d} + N_0 \epsilon^d \eta^{D-d}, \quad (53)$$

uniformly over  $i \in I_k$  and  $k \neq 0$ ; and also,

$$H_{i,\epsilon/2,\eta}^* \prec (N - N_0) \epsilon^d (1 \wedge (\eta/\tau))^{D-d} \xi \mathbf{1}_{\{\delta_0 \leq \epsilon + \tau\}} + N \epsilon^d \eta^{D-d}. \quad (54)$$

uniformly over  $i \in I_0$

*Proof.* First assume that  $i \in I_k$  with  $k \neq 0$ . For the lower bound in (53), let  $L$  be a subspace such that

$$\Psi_k(B(\mathbf{x}_i, \epsilon) \cap B(L, \eta)) \succ \epsilon^d (1 \wedge (\eta/\tau))^{D-d},$$

which exists by the lower bound in Lemma 7. We have  $H_{i,\epsilon,\eta} \geq H_{i,\epsilon,\eta}(L)$ , and the term on the right hand side is a sum of independent Bernoulli random variables with expectation

$$\mathbb{E}(H_{i,\epsilon,\eta}(L)) = N_k \Psi_k(B(\mathbf{x}_i, \epsilon) \cap B(L, \eta)) \succ N_k \epsilon^d (1 \wedge (\eta/\tau))^{D-d}.$$

We then apply Lemma 1, using (46), and the union bound. For the upper bound in (53), the arguments are the same as in the proof of (43), except for the following bound in expectation, valid for any  $L \in \mathcal{A}_d$ ,

$$\begin{aligned} \mathbb{E}(H_{i,\epsilon,\eta}(L)) &= \sum_{\ell} N_{\ell} \Psi_{\ell}(B(\mathbf{x}_i, \epsilon) \cap B(L, \eta)) \\ &\prec (N - N_0) \epsilon^d (1 \wedge (\eta/\tau))^{D-d} + N_0 \epsilon^d \eta^{D-d}, \end{aligned}$$

by Lemmas 7 and 8.

Now, assume that  $i \in I_0$ . Again, the arguments are the same as in the proof of (43), except that the bounds in expectation are different. Specifically, if  $\delta_0 > \epsilon + \tau$ , then  $\Psi_{\ell}(B(\mathbf{x}_i, \epsilon) \cap B(L, \eta)) = 0$ ,  $\forall \ell \neq 0$ , so that, by Lemma 8, for any  $L \in \mathcal{A}_d$ ,

$$\mathbb{E}(H_{i,\epsilon,\eta}(L)) \prec N_0 \epsilon^d \eta^{D-d}.$$

Otherwise,

$$\mathbb{E}(H_{i,\epsilon,\eta}(L)) \prec (N - N_0) \epsilon^d (1 \wedge (\eta/\tau))^{D-d} \xi + N_0 \epsilon^d \eta^{D-d}.$$

$\square$

We are now in a position to prove Propositions 4 and 5. We first consider (O1). By (48) and (49), and the fact that  $\tau \leq \eta \leq \rho_N^{-3/(D-d)}$ , we have

$$\max_i D_i^{1/(m-1)} \prec (N - N_0)\epsilon^d \prec (N/\rho_N)\epsilon^d.$$

On the one hand, by (48),  $D_i^{1/(m-1)} \succ N_k\epsilon^d \succ (N/\rho_N)\epsilon^d$ , uniformly over  $i \in I_k$ ,  $\forall k \neq 0$ . Hence, since  $\rho_N \rightarrow \infty$ , no non-outlier is identified as an outlier. On the other hand, by (49), for any  $i \in I_0$ ,

$$D_i^{1/(m-1)} \prec N\epsilon^d(\xi^{1-\frac{d+1}{m-1}} + \eta^{D-d-\frac{d+1}{m-1}}) \ll N\epsilon^d/\rho_N^2,$$

since  $\xi \prec \eta/\delta_0 \prec \rho_N^{-3}$  and  $\eta \leq \epsilon \leq \rho_N^{-3/(D-d)}$ . Hence, all outliers are identified as such.

We now consider (O2). On the one hand, by (48) and (16), and the expression for  $\epsilon$  and  $\eta$ , we have

$$D_i^{1/(m-1)} \succ N_k\epsilon^d(1 \wedge (\eta/\tau))^{D-d} \succ \rho_N^3 \log N \succ \rho_N^2 N\epsilon^d \eta^{D-d},$$

uniformly over  $k \neq 0$  and  $i \in I_k$ . Hence, no non-outlier is identified as an outlier. On the other hand, by (49), for any  $i \in I_0$ ,

$$D_i^{1/(m-1)} \prec N\epsilon^d \eta^{D-d-\frac{d+1}{m-1}} \prec N\epsilon^d \eta^{D-d},$$

which comes from  $m \gg \log(N)/\log(\rho_N)$ . Hence, all outliers are identified as such.

## Appendix C: Proofs for the estimation of parameters

### C.1. Proof of Proposition 1

Recalling the definition of  $G_{i,\epsilon}$  in (39), we have

$$\text{Cor}(\epsilon) = \sum_i G_{i,\epsilon}.$$

Let  $\epsilon_r = \rho_N^{-r}$  and let  $r_0$  be the integer defined by  $\epsilon_{r_0+1} < \tau \leq \epsilon_{r_0}$ . Define

$$r_N^* := ((1 - d/D)r_0 + (d/D)r_N) \wedge r_N,$$

and note that, for  $r \leq r_N^*$ , (46) with  $\epsilon$  in place of  $\eta$  is satisfied for  $\epsilon_r$ . As there are only order  $\log N$  such  $r$ 's, Proposition 13 and the union bound imply that, with probability at least  $1 - \log(N)N^{-\rho_N/(K\zeta)}$ ,

$$(N/\rho_N)^2 \epsilon_r^d (1 \wedge (\epsilon_r/\tau))^{D-d} \prec \text{Cor}(\epsilon_r) \prec (N/\rho_N)^2 \epsilon_r^d (1 \wedge (\epsilon_r/\tau))^{D-d},$$

uniformly over  $r \leq r_N^*$ . Note that we used the fact that  $N^2 \epsilon_r^D \ll (N/\rho_N)^2 \epsilon_r^d (1 \wedge (\epsilon_r/\tau))^{D-d}$ , which holds since  $r, r_0 \geq 3$ . When this is the case,

$$A_r = \begin{cases} 2 \log N - dr \log \rho_N + O(1), & r \leq r_0; \\ 2 \log N - Dr \log \rho_N - (D-d) \log \tau + O(1), & r > r_0. \end{cases}$$

In particular, for  $r \leq r_N^*$ ,

$$\frac{A_r - A_{r+1}}{\log \rho_N} = \begin{cases} d + o(1), & r \leq r_0 - 1; \\ D + o(1), & r \geq r_0 + 1. \end{cases}$$

From the first part, we see that  $\hat{r} \geq r_0 \wedge (r_N - \lceil 2D/d \rceil)$ , since  $d \leq D - 1$  and  $\rho_N \rightarrow \infty$ . To use the second part, note that  $r_0 + 2 \leq r_N^*$  if, and only if,  $r_0 \leq r_N - \lceil 2D/d \rceil$ . If this is the case,  $\hat{r} \leq r_0 + 1$ . From this follows the statement in Proposition 1.

### C.2. Proof of Proposition 2

We follow the proof of Proposition 1. We assume that  $\hat{d} = d$ , which happens with probability tending to one. Let  $\eta_s = \rho_N^{-\hat{r}-s}$  and  $s_0 = r_0 - \hat{r}$ . Define

$$s_N^* := ((2Dd + d - 2)/(D - d) + s_0) \wedge (\hat{r} - 1),$$

and note that, for  $s \leq s_N^*$ , (46) is satisfied for  $\epsilon_{\hat{r}}$  and  $\eta_s$ . Indeed, using the fact that  $\epsilon_{\hat{r}} \geq (\log(N)/N)^{1/d} \rho_N^{2D+1}$  and  $\tau \leq \rho_N^{-r_0}$ , we get

$$\begin{aligned} N_k \epsilon_{\hat{r}}^d (1 \wedge (\eta_s/\tau))^{D-d} &\geq (N/(K\zeta)\rho_N)(\log(N)/N)\rho_N^{(2D+1)d} (1 \wedge \rho_N^{(s_0-s)(D-d)}) \\ &= \rho_N \log(N) \cdot \rho_N^{-2+(2D+1)d-(D-d)(s-s_0)_+}, \end{aligned}$$

and the exponent in  $\rho_N$  is non-negative by the upper bound on  $s$ . As there are only order  $\log N$  such  $s$ 's, Proposition 12 and the union bound imply that, with probability at least  $1 - \log(N)N^{-\rho_N/(K\zeta)}$ ,

$$\begin{aligned} \text{Cor}(\epsilon_{\hat{r}}, \eta_s) &\succ (N/\rho_N)^2 \zeta^{-1} \epsilon_{\hat{r}}^d (1 \wedge (\eta_s/\tau))^{D-d}, \\ \text{Cor}(\epsilon_{\hat{r}}, \eta_s) &\prec (N/\rho_N)^2 \epsilon_{\hat{r}}^d (1 \wedge (\eta_s/\tau))^{D-d-(d+1)/(m-1)}, \end{aligned}$$

uniformly over  $s \leq s_N^*$ . Note that we used the fact that

$$N^2 \epsilon_{\hat{r}}^d \eta_s^{D-d} \ll (N/\rho_N)^2 \epsilon_{\hat{r}}^d (1 \wedge (\eta_s/\tau))^{D-d}.$$

When this is the case,

$$B_s = 2 \log N - d\hat{r} \log \rho_N + O(1),$$

when  $s \leq s_0$ , and

$$B_s = 2 \log N - D\hat{r} \log \rho_N + (-(D-d) + O(1/m))(s \log \rho_N + \log \tau) + O(1),$$

when  $s > s_0$ . In particular, for  $s \leq s_N^*$ ,

$$\frac{B_s - B_{s+1}}{\log \rho_N} = \begin{cases} o(1), & s \leq s_0 - 1; \\ D - d + o(1), & s = s_0 + 1. \end{cases}$$

From here the arguments are parallel to those used in Proposition 1.

## Acknowledgements

GC was at the University of Minnesota, Twin Cities, for part of the project. The authors would like to thank the Institute for Mathematics and its Applications (IMA), in particular Doug Arnold and Fadil Santosa, for holding a stimulating workshop on multi-manifold modeling that GL co-organized, and EAC and GL participated in. The authors also thank Jason Lee for providing the last dataset in Figure 13. Finally, the authors are grateful to two anonymous referees and Associate Editor for providing constructive feedback and criticism, which helped improve the presentation of the paper. This work was partially supported by grants from the National Science Foundation (DMS-06-12608, DMS-09-15160, DMS-09-15064) and a grant from the Office of Naval Research (N00014-09-1-0258).

## References

- [1] S. AGARWAL, K. BRANSON, AND S. BELONGIE. Higher order learning with graphs. In *Proceedings of the 23rd International Conference on Machine Learning (ICML '06)*, volume 148, pages 17–24, 2006.
- [2] S. AGARWAL, J. LIM, L. ZELNIK-MANOR, P. PERONA, D. KRIEGMAN, AND S. BELONGIE. Beyond pairwise clustering. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '05)*, volume 2, pages 838–845, 2005.
- [3] E. ARIAS-CASTRO. Clustering based on pairwise distances when the data is of mixed dimensions. *IEEE Trans. Inform. Theory*, 57(3):1692–1706, 2011. [MR2815843](#)
- [4] E. ARIAS-CASTRO, D. L. DONOHO, X. HUO, AND C. A. TOVEY. Connect the dots: how many random points can a regular curve pass through? *Adv. in Appl. Probab.*, 37(3):571–603, 2005. [MR2156550](#)
- [5] E. ARIAS-CASTRO, B. EFROS, AND O. LEVI. Networks of polynomial pieces with application to the analysis of point clouds and images. *J. Approx. Theory*, 162(1):94–130, 2010. [MR2565828](#)
- [6] R. BASRI AND D. JACOBS. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(2):218–233, 2003.
- [7] M. BELKIN AND P. NIYOGI. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(16):1373–1396, 2003.
- [8] A. BEYGEZIMER, S. KAKADE, AND J. LANGFORD. Cover trees for nearest neighbor. In *Proceedings of the 23rd International Conference on Machine Learning (ICML '06)*, pages 97–104, 2006.
- [9] M. R. BRITO, E. L. CHÁVEZ, A. J. QUIROZ, AND J. E. YUKICH. Connectivity of the mutual  $k$ -nearest-neighbor graph in clustering and outlier detection. *Statist. Probab. Lett.*, 35(1):33–42, 1997. [MR1467706](#)
- [10] G. CHEN, S. ATEV, AND G. LERMAN. Kernel spectral curvature clustering (KSCC). In *Dynamical Vision Workshop, IEEE 12th International Conference on Computer Vision*, pages 765–772, Kyoto, Japan, 2009.



- [11] G. CHEN AND G. LERMAN. Foundations of a multi-way spectral clustering framework for hybrid linear modeling. *Found. Comput. Math.*, 9(5):517–558, 2009. [MR2534403](#)
- [12] G. CHEN AND G. LERMAN. Spectral curvature clustering (SCC). *Int. J. Comput. Vision*, 81(3):317–330, 2009.
- [13] G. CHEN, G. LERMAN, AND E. ARIAS-CASTRO. Higher order spectral clustering (hosc) algorithm. Matlab code. Current version available at <http://www.math.duke.edu/~glchen/hosc.html>.
- [14] F. R. K. CHUNG. *Spectral graph theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 1997. [MR1421568](#)
- [15] J. K. CULLUM AND R. A. WILLOUGHBY. *Lanczos Algorithms for Large Symmetric Eigenvalue Computations, Vol. 1: Theory*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2002. [MR1948689](#)
- [16] L. DEVROYE AND G. L. WISE. Detection of abnormal behavior via non-parametric estimation of the support. *SIAM J. Appl. Math.*, 38(3):480–488, 1980. [MR0579432](#)
- [17] D. L. DONOHO AND C. GRIMES. Image manifolds which are isometric to euclidean space. *J. Math. Imaging Vis.*, 23(1):5–24, 2005. [MR2208906](#)
- [18] R. M. DUDLEY. Metric entropy of some classes of sets with differentiable boundaries. *J. Approx. Theory*, 10:227–236, 1974. [MR0358168](#)
- [19] R. EPSTEIN, P. HALLINAN, AND A. YUILLE.  $5 \pm 2$  eigenimages suffice: An empirical investigation of low-dimensional lighting models. In *IEEE Workshop on Physics-based Modeling in Computer Vision*, pages 108–116, June 1995.
- [20] H. FEDERER. Curvature measures. *Trans. Amer. Math. Soc.*, 93:418–491, 1959. [MR0110078](#)
- [21] D. J. FIELD, A. HAYES, AND R. F. HESS. Contour integration by the human visual system: Evidence for a local ‘association field’. *Vision Research*, 33(2):173–193, 1993.
- [22] M. FILIPPONE, F. CAMASTRA, F. MASULLI, AND S. ROVETTA. A survey of kernel and spectral methods for clustering. *Pattern Recogn.*, 41(1):176–190, 2008.
- [23] Z. FU, W. HU, AND T. TAN. Similarity based vehicle trajectory clustering and anomaly detection. In *Proceedings of the IEEE International Conference on Image Processing (ICIP '05)*, volume 2, pages 602–605, 2005.
- [24] A. GIONIS, A. HINNEBURG, S. PAPADIMITRIOU, AND P. TSAPARAS. Dimension induced clustering. In *Proceedings of the eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining (KDD '05)*, pages 51–60, New York, NY, USA, 2005.
- [25] A. GOLDBERG, X. ZHU, A. SINGH, Z. XU, AND R. NOWAK. Multi-manifold semi-supervised learning. In *Twelfth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2009.

- [26] G. H. GOLUB AND C. F. VAN LOAN. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, third edition, 1996. [MR1417720](#)
- [27] V. GOVINDU. A tensor decomposition for geometric grouping and segmentation. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '05)*, volume 1, pages 1150–1157, June 2005.
- [28] P. GRASSBERGER AND I. PROCACCIA. Measuring the strangeness of strange attractors. *Physica D*, 9:189–208, 1983. [MR0732572](#)
- [29] Q. GUO, H. LI, W. CHEN, I.-F. SHEN, AND J. PARKKINEN. Manifold clustering via energy minimization. In *ICMLA '07: Proceedings of the Sixth International Conference on Machine Learning and Applications*, pages 375–380, Washington, DC, USA, 2007. IEEE Computer Society.
- [30] G. HARO, G. RANDALL, AND G. SAPIRO. Stratification learning: Detecting mixed density and dimensionality in high dimensional point clouds. *Advances in Neural Information Processing Systems (NIPS)*, 19:553, 2007.
- [31] J. HO, M. YANG, J. LIM, K. LEE, AND D. KRIEGMAN. Clustering appearances of objects under varying illumination conditions. In *Proceedings of International Conference on Computer Vision and Pattern Recognition (CVPR '03)*, volume 1, pages 11–18, 2003.
- [32] D. KUSHNIR, M. GALUN, AND A. BRANDT. Fast multiscale clustering and manifold identification. *Pattern Recogn.*, 39(10):1876–1891, 2006.
- [33] E. LEVINA AND P. BICKEL. Maximum likelihood estimation of intrinsic dimension. In *Advances in Neural Information Processing Systems (NIPS)*, volume 17, pages 777–784. MIT Press, Cambridge, Massachusetts, 2005.
- [34] U. LUXBURG. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. [MR2409803](#)
- [35] Y. MA, A. Y. YANG, H. DERKSEN, AND R. FOSSUM. Estimation of subspace arrangements with applications in modeling and segmenting mixed data. *SIAM Review*, 50(3):413–458, 2008. [MR2429444](#)
- [36] M. MAIER, M. HEIN, AND U. VON LUXBURG. Cluster identification in nearest-neighbor graphs. In *Algorithmic Learning Theory*, pages 196–210. Springer, 2007.
- [37] M. MAIER, M. HEIN, AND U. VON LUXBURG. Optimal construction of k-nearest-neighbor graphs for identifying noisy clusters. *Theor. Comput. Sci.*, 410(19):1749–1764, 2009. [MR2514706](#)
- [38] E. MAMMEN AND A. B. TSYBAKOV. Asymptotical minimax recovery of sets with smooth boundaries. *Ann. Statist.*, 23(2):502–524, 1995. [MR1332579](#)
- [39] V. MARTÍNEZ AND E. SAAR. *Statistics of the Galaxy Distribution*. Chapman and Hall/CRC press, Boca Raton, 2002.
- [40] H. NARAYANAN, M. BELKIN, AND P. NIYOGI. On the relation between low density separation, spectral clustering and graph cuts. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems (NIPS)*, volume 19. MIT Press, Cambridge, MA, 2007.

- [41] H. NEUMANN, A. YAZDANBAKHS, AND E. MINGOLLA. Seeing surfaces: The brain's vision of the world. *Physics of Life Reviews*, 4(3):189–222, 2007.
- [42] A. Y. NG, M. I. JORDAN, AND Y. WEISS. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems (NIPS)*, volume 14, pages 849–856, 2002.
- [43] P. NIYOGI, S. SMALE, AND S. WEINBERGER. Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.*, 39(1):419–441, 2008. [MR2383768](#)
- [44] B. PELLETIER AND P. PUDLO. Operator norm convergence of spectral clustering on level sets. *Journal of Machine Learning Research*, 12:385–416, 2011. [MR2783172](#)
- [45] M. PENROSE. *Random Geometric Graphs*, volume 5 of *Oxford Studies in Probability*. Oxford University Press, Oxford, 2003. [MR1986198](#)
- [46] S. RAO, A. YANG, S. SASTRY, AND Y. MA. Robust algebraic segmentation of mixed rigid-body and planar motions from two views. *International Journal of Computer Vision*, 88(3):425–446, 2010. [MR2719048](#)
- [47] S. ROWEIS AND L. SAUL. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.
- [48] A. SHASHUA, R. ZASS, AND T. HAZAN. Multi-way clustering using supersymmetric non-negative tensor factorization. In *Proceedings of the European Conference on Computer Vision (ECCV '06)*, volume 4, pages 595–608, 2006.
- [49] R. SOUVENIR AND R. PLESS. Manifold clustering. In *IEEE International Conference on Computer Vision (ICCV '05)*, volume 1, pages 648–653, 2005.
- [50] M. TALAGRAND. *The Generic Chaining*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2005. [MR2133757](#)
- [51] J. B. TENENBAUM, V. DE SILVA, AND J. C. LANGFORD. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- [52] R. VALDARNINI. Detection of non-random patterns in cosmological gravitational clustering. *Astronomy & Astrophysics*, 366:376–386, 2001.
- [53] R. VIDAL AND Y. MA. A unified algebraic approach to 2-D and 3-D motion segmentation and estimation. *Journal of Mathematical Imaging and Vision*, 25(3):403–421, 2006. [MR2283615](#)
- [54] U. VON LUXBURG, M. BELKIN, AND O. BOUSQUET. Consistency of spectral clustering. *Ann. Statist.*, 36(2):555–586, 2008. [MR2396807](#)
- [55] H. WEYL. On the volume of tubes. *Amer. J. Math.*, 61(2):461–472, 1939. [MR1507388](#)
- [56] L. ZELNIK-MANOR AND P. PERONA. Self-tuning spectral clustering. In *Advances in Neural Information Processing Systems (NIPS)*, volume 17, pages 1601–1608, 2004.