

This is a repository copy of *Shape Simplification Through Graph Sparsification*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/114032/>

Proceedings Paper:

Escolano, Francisco, Curado, Manuel, Biasotti, Silvia et al. (1 more author) (2017) Shape Simplification Through Graph Sparsification. In: Graph-Based Representations in Pattern Recognition - 11th IAPR-TC-15 International Workshop, GbRPR 2017, Anacapri, Italy, May 16-18, 2017, Proceedings. Springer , pp. 13-22.

https://doi.org/10.1007/978-3-319-58961-9_2

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Shape Simplification Through Graph Sparsification

Francisco Escolano, Manuel Curado, Silvia Biasotti, and Edwin R. Hancock

Department of Computer Science and AI,
University of Alicante, 03690, Alicante Spain
`{sco,mcurado}@dccia.ua.es`

CNR-IMATI, Via de Marini, 6 (Torre di Francia), 16149 Genova, Italy
`silvia.biasotti@ge.imati.cnr.it`

Department of Computer Science,
University of York, York, YO10 5DD, UK
`erh@cs.york.ac.uk`

Abstract. In this paper, we draw on Spielman and Srivastava’s method for graph sparsification in order to simplify shape representations. The underlying principle of graph sparsification is to retain only the edges which are key to the preservation of desired properties. In this regard, sparsification by edge resistance allows us to preserve (to some extent) links between protrusions and the remainder of the shape (e.g. parts of a shape) while removing in-part edges. Applying this idea to alpha shapes (abstract representations which have a huge number of edges) opens up a way of introducing a hierarchy of the edge strength, thus being relevant for shape analysis and interpretation.

Keywords: Graph sparsification, Shape simplification, Alpha shapes

1 Introduction

1.1 Shape representations: Triangulations vs Alpha shapes

The traditional problem addressed by shape reconstruction is to recover a digital representation of a physical shape that has been scanned, where the scanned data contain a wide variety of defects or the representation of data acquired by different diagnostic equipments such as angiography, Computed Tomography (CT) and Magnetic Resonance (MR). To encode the data in a digital model different geometric representations have been explored in detail. The work reported in [?] organizes them into a spectrum with respect to the achieved trade-off between verbosity and complexity. Voxel grids are at one extreme of the spectrum, since they are the simplest, but the most verbose and less accurate representation. Although, in principle, the use of arbitrarily fine grids could achieve any level of approximation, the practical limit comes from constraints on the resolution. At the other end of the spectrum, the functional representations - using smooth functions to specify the continuous of points that make up the shape - provide

an accurate and complex representation. Piecewise linear representations are at the center of the spectrum.

The most popular representations in the piecewise class are the simplicial complexes [?], including triangular meshes that have become the *de-facto* standard in graphics accelerators [?] and tetrahedral meshes that are used to represent volumes and are used for the simulation of deformable models, such as organs or tissues. A generalization of the concept of triangulation are the so-called alpha (α -) shapes, that are families of piecewise linear simple curves in the Euclidean space associated with a dense and unorganized set of data points. An alpha shape is demarcated by a frontier, which is a linear approximation of the original shape. First introduced in the 2D plane by Edelsbrunner et al. [?], they were extended to 3D spaces [?] and higher dimensions [?]. In the case of 2D, an alpha shape consists of vertices, edges and triangles, while for 3D there are also tetrahedra. In our graph representation, we consider the 1-skeleton of both triangulations and alpha-shapes, i.e., the set of vertices and edges of the complex.




Fig. 1. From left to right: a point set, a triangulation and a sequence of three alpha shapes with increasing values of α .

Alpha shapes depend on the parameter α used as radius of spheres centered on the points that determine the connection among the neighbourhoods. A very small value will generate many isolated points and the alpha shape degenerates to the point cloud when $\alpha \rightarrow 0$. On the other hand, a large value of α will consider many points inside the spheres and therefore the size of the 1-skeleton considerably increases. The limit of the alpha shape when $\alpha \rightarrow \infty$ is the convex hull of the point cloud and the 1-skeleton of the alpha shape becomes the complete graph.

The main application of alpha shapes is the reconstruction of objects which have been sampled by points. How to determine the best value of α is not obvious and in practice α is found using a trial-and-error strategy. This leads the computation of quite large families of alpha shapes and the 1-skeleton increases as long as the value of α increases. Moreover, there are point-sets for which there is no unique α value, for instance because small α values capture local characteristics while larger ones determine large connectivity. For instance, this is the case when a point cloud is not uniformly sampled or the point cloud is supposed to represent either small or large features (for instance, it contains both thin and long handles like the examples shown in Fig. 1). Low density sampling requires a rather large radius to build a connected representation. But

a large value of α will unfortunately close some handles. In practice, a large value of α results in (among possibly other things) a closure of handles, connection of multiple components and joints (e.g., sharp turns) being destroyed. For this reasons and because the general size of the 1-skeleton, the approach proposed in this work is able to simplify connections without destroying the global topology of the alpha shape.

1.2 Contributions

Spielman and Srivastava [?] have developed an efficient method for graph sparsification based on edge resistance (which is proportional to the commute time of its end nodes). The method is based on the observation that the probability of an edge appearing in a random spanning tree of a graph is equal to its effective resistance. Drawing on Spielman and Teng’s approximately linear solver[?], they show how to efficiently compute resistance, and hence sample edges for the purposes of sparsification.

Herein, we present a unified view of resistance sparsification through sampling. In addition, we exploit such a sampling for retaining edges (both in triangulations and alpha spaces) that are key to the preservation of the topological properties of the input shape. Our experiments show high compression rates as the allowed error ϵ increases. However each shape is sensitive to a different value of ϵ . It is the persistence of a given edge as ϵ increases what will provide us with is the relative importance of a vertex. This characterisation is pivotal for subsequent tasks such as efficient shape matching and shape representation.

2 Graph Sparsification

2.1 Definition and Ingredients

Graph sparsification [?] is the principled study of how to significantly decrease the number of edges of an input graph G so that the output, H , preserves some of the structural properties of G .

Benczúr and Karger [?] showed that every cut in $G = (V, E)$ can be approximated in $H = (V, E')$, with $E' \subseteq E$, so that every cut in H has a value within $(1 \pm \epsilon)$ times its value in G . For instance, a K_n (complete) graph with n vertices and $O(n^2)$ edges can be approximated by a random d -regular graph, i.e. a graph with $O(dn)$ edges. This means that for every subset $S \subset V$ the ratio between the value of a cut in K_n and that of the same cut in the random d -regular graph H is n/d . This link between sparsification and random graphs is useful (to some extent). For instance, if an edge in G is included in H with probability p , we must set $p \gg 1/c$ where c is the value of the minimal cut. As a result, if we have m edges in G we can only have $O(m/c)$ edges in H .

This limitation leads to *non-uniform sampling*, i.e. to associate a different probability p_e to each edge $e \in E$. The edge e is included in E' with probability p_e and it is given a weight $1/p_e$ if it is included. This inverse weighting

ensures that the expected weight of e in H is unity.

The *choice of a suitable value of p_e* is the first step in graph sparsification. For cut sparsification, the choice of p_e relies on the *strong* connectivity c_e of e . The strong connectivity c_e is the maximum value of a cut in a connected component including e . This quantity is upper bounded by the standard connectivity of e (the minimal value of a cut separating its endpoints), but it is hard to find. However lower bounds $c'_e \leq c_e$ can be found through sparse certification (see details in [?]). In this way we have that $p_e = \rho/c'_e \geq \rho/c_e$, where ρ is the compression factor, is a good choice for p_e . The *compression factor* ρ has complexity $O(c(d+2)(\log n)/\epsilon^2)$ and it is in turn inversely proportional to the squared error ϵ^2 . The setting $p_e = \min\{1, \rho/c'_e\}$ then ensures the correctness of the approximation with probability $1 - n^{-d}$.

The above rationale leads to the second ingredient of sparsification, namely the *minimal number of samples* required to correctly sparsify the graph with high probability. For cut sparsification, we have that taking $O(n\rho)$, i.e. $O(n \log n/\epsilon^2)$, samples will suffice. This can be proved by means of the Chernoff bound, which is a standard information-theoretic tool for limiting the number of samples.

2.2 Spectral Formulation and Effective Resistances

An alternative approach to the the sparsification problem consists of enforcing the preservation of structural properties by bounding the quadratic form associated with the graph Laplacian of the sparsified graph H with respect to that of the input graph G (see the survey in [?]). Therefore, given $G = (V, E, w)$ we must obtain $H = (V, E', w')$ by taking $O(n \log n/\epsilon^2)$ independent samples, so that we satisfy (with probability at least $1/2$) the following constraint

$$\forall x \in \mathbb{R}^n : (1 - \epsilon) \leq x^T L_G x \leq x^T L_H x \leq (1 + \epsilon) x^T L_G , \quad (1)$$

where $\epsilon > 0$, $n = |V|$, and L_G, L_H are the respective Laplacian matrices of G and H . Recall that $L_G = D - W$ where D is the diagonal degree matrix and W is the weighted adjacency matrix, and that $x^T L_G x = \sum_{(u,v) \in E} (x(u) - x(v))^2 w_{uv}$ and similarly for L_H .

Since Laplacian matrices are Semidefinite Positive (SDP), which is denoted by $L_G \succeq 0$, we can reformulate Eq. 1 in terms of circumventing the hyper ellipsoid associated with L_H with that associated with L_G , i.e. one must satisfy

$$(1 - \epsilon)L_G \preceq L_H \preceq (1 + \epsilon)L_G, \text{ or equivalently } L_G \preceq L_H \preceq \kappa L_G , \quad (2)$$

with $\kappa = \frac{1+\epsilon}{1-\epsilon}$. This implies that all of the eigenvalues λ'_i of L_H satisfy $\lambda'_i \leq \kappa \lambda_i$, where λ_i is the corresponding eigenvalue of L_G . In addition, since Eq. 2 is invariant under rescaling, we have that

$$L_G^{-1/2} L_G L_G^{-1/2} \preceq L_G^{-1/2} L_H L_G^{-1/2} \preceq \kappa L_G^{-1/2} L_G L_G^{-1/2} \quad (3)$$

i.e.

$$I \preceq L_G^{-1/2} L_H L_G^{-1/2} \preceq \kappa I , \quad (4)$$

where I is the identity matrix and $L_G^{-1/2}L_H L_G^{-1/2}$ is the so called *relative Laplacian*. This leads to locating L_H so that the relative Laplacian is properly contained between I and κI . In this regard, the structure of L_H is determined by a weighted sum of outer products: $L_H = \sum_{e \in E} w'_e b_e b_e^T$, where w'_e are the unknown weights, $b_e = \delta_u - \delta_v = b_{uv}$, δ_u is the unit vector with a 1 at u and zeros elsewhere (similarly for v), and $e = (u, v)$ is the edge. In this regard, since $E' \subseteq E$, an edge of E not included in E' will have $w'_e = 0$. We define the random variables s_e (our unknowns) so that $w'_e = s_e w_e$ where $\mathbb{E}(s_e) = 1$ for all $e \in E$. Then, Eq. 4 can be rewritten as follows

$$I \preceq L_G^{-1/2} \left(\sum_{e \in E} s_e w_e L_G^{-1/2} b_e b_e^T \right) L_G^{-1/2} \preceq \kappa I. \quad (5)$$

It is well known that the Laplacian matrix L cannot be inverted since it contains the zero eigenvalue. Expressions including the inverse must be computed using the pseudo-inverse L^+ instead. The pseudo inverse plays a key role in defining the effective resistance across $e = (u, v)$ (the scaled commute time) R_e , which is given by

$$R_e = (\delta_u - \delta_v)^T L^+ (\delta_u - \delta_v) = b_e^T L^+ b_e. \quad (6)$$

Then, combining Eqs. 5 and 6 we obtain

$$I \preceq \sum_{e \in E} s_e w_e v_e v_e^T \preceq \kappa I, \quad (7)$$

where $v_e = L_G^{-1/2} b_e$, i.e., the squared norm of v_e is

$$\|v_e\|^2 = (L_G^{-1/2} b_e)^T (L_G^{-1/2} b_e) = (b_e^T L_G^{-1/2}) (L_G^{-1/2} b_e) = b_e^T L_G^+ b_e = R_e. \quad (8)$$

This squared norm allows us to treat $\sum_{e \in E} s_e w_e v_e v_e^T$ in Eq. 5 as a quadratic form quite close to the identity matrix I . This is extremely important since: i) the relative Laplacian relies on the effective resistances of G , and ii) we can pose the sparsification problem in terms of finding the sampling probabilities p_e so that the constraint in Eq. 5 is satisfied. To this end, Batson et al. [?] exploited the following fact:

$$\sum_{e \in E} \tilde{v}_e \tilde{v}_e^T = I, \quad (9)$$

where $\tilde{v}_e = w_e^{1/2} v_e$. This can be proved by using the $m \times n$ incidence matrix of G , i.e. B , with elements $B(e, v) = 1$ if v is e 's head, $B(e, v) = -1$ if v is e 's tail, and $B(e, v) = 0$ otherwise. Then the Laplacian matrix of G is given by $L_G = B^T W_e B$, where W_e is the diagonal $m \times m$ matrix where $W_e(e, e) = w_e$. Since the vectors $v_e = L_G^{-1/2} b_e$ rely on the columns of B^T , we have that vectors $\tilde{v}_e = v_e w_e^{1/2}$ are the columns of a $n \times m$ matrix $\tilde{V} = L_G^{-1/2} B^T W_e^{1/2}$. Then

$$\sum_{e \in E} \tilde{v}_e \tilde{v}_e^T = \tilde{V} \tilde{V}^T = L_G^{-1/2} B^T W_e^{1/2} W_e^{1/2} B L_G^{-1/2} = L_G^{-1/2} L_G L_G^{-1/2} = I.$$

In addition we have that

$$\|\tilde{v}_e\|^2 = (w_e^{1/2} L_G^{-1/2} b_e)^T (w_e^{1/2} L_G^{-1/2} b_e) = w_e (b_e^T L_G^+ b_e) = w_e R_e, \quad (10)$$

i.e. we obtain weighted effective resistances. The identity $\|\tilde{v}_e\|^2 = w_e R_e$ suggests to sample E with probabilities p_e proportional to $w_e R_e$.

Let y_1, y_2, \dots, y_q vectors drawn independently with replacement from the distribution

$$y = \frac{1}{\sqrt{p_e}} \tilde{v}_e \text{ with probability } p_e. \quad (11)$$

Then, the expectation of yy^T (which contains the effective resistances) is

$$\mathbb{E}[yy^T] = \sum_{e \in E} p_e \frac{1}{p_e} \tilde{v}_e \tilde{v}_e^T = I. \quad (12)$$

In addition, the shape of each of the q samples $y_i = \tilde{v}_e / \sqrt{p_e}$ leads to

$$\frac{1}{q} \sum_{i=1}^q y_i y_i^T = \frac{1}{q} \sum_{i=1}^q \#e \frac{\tilde{v}_e}{\sqrt{p_e}} \cdot \frac{\tilde{v}_e^T}{\sqrt{p_e}} = \frac{1}{q} \sum_{i=1}^q \#e \frac{\tilde{v}_e \tilde{v}_e^T}{p_e} = \sum_{e \in E} s_e \tilde{v}_e \tilde{v}_e^T, \quad (13)$$

where $\#e$ is the number of times that e is sampled, and $s_e = \#e / qp_e$. Then, we obtain

$$\frac{1}{q} \sum_{i=1}^q y_i y_i^T = \sum_{e \in E} s_e \tilde{v}_e \tilde{v}_e^T = \sum_{e \in E} s_e w_e v_e v_e^T, \quad (14)$$

i.e. a proper sampling process leads to the relative Laplacian. This is ensured insofar $\frac{1}{q} \sum_{i=1}^q y_i y_i^T$ and $\mathbb{E}yy^T$ conform the Chernoff bound for matrices [?]:

$$\mathbb{E} \left[\left\| \frac{1}{q} \sum_{i=1}^q y_i y_i^T - \mathbb{E}[yy^T] \right\| \right] \leq \min \left(CM \sqrt{\frac{\log q}{q}}, 1 \right), \quad (15)$$

where $\|\mathbb{E}yy^T\| < 1$ and $\sup_y \|y\| \leq M$. The first norm condition is verified since $\mathbb{E}[yy^T] = I$. For verifying the second norm condition we must set the link between $w_e R_e$ (weighted effective resistances) and p_e (sampling probabilities). In order to do so, Spielman and Srivastava [?] exploit the fact that $\sum_e w_e R_e = n-1$. Therefore, we may set

$$p_e = \frac{w_e R_e}{n-1} \text{ so that } \|y\| = \frac{1}{\sqrt{p_e}} \sqrt{w_e R_e} = \sqrt{\frac{n-1}{w_e R_e}} \sqrt{w_e R_e} = \sqrt{n-1}. \quad (16)$$

Therefore, taking $q = 9C^2 n \log n / \epsilon^2$ yields

$$\mathbb{E} \left[\left\| \frac{1}{q} \sum_{i=1}^q y_i y_i^T - \mathbb{E}[yy^T] \right\| \right] \leq C \sqrt{\epsilon^2 \frac{\log(9C^2 n \log n / \epsilon^2)(n-1)}{9C^2 n \log n}} \leq \epsilon/2, \quad (17)$$

for n large enough and $\epsilon \geq 1/\sqrt{n}$.

Summarising, the resistance-based sparsifier [?] consists in five steps:

1. Given the input graph $G = (V, E, w)$, estimate the effective resistances R_e for each $e \in E$.
2. Set an error tolerance ϵ . Set $E' = \emptyset$, $w' = \emptyset$, define $H = (V, E', w')$ and set $\#e = 0$ for all edges in E .
3. Make $q = 9C^2 n \log n / \epsilon^2$ independent samples (with replacement) with probability $p_e \propto w_e R_e$. Each sample is associated with an edge e .
4. If e is selected from a cumulative sum test, then increment $\#e$ and add e to E' with weight $1/p_e$.
5. For all $e \in E'$ set $w'_e = \frac{\#e}{qp_e}$.

Finally, the computation R_e can be accomplished using exact spectral methods [?]. However, this step takes $O(n^3)$ steps and the eigenvalues are ill conditioned if the graph G has several connected components. This is why Spielman and Srivastava [?] propose to approximate the computation of effective resistances by exploiting the Achlioptas version [?] of the Johnson-Lindenstrauss (JL) Lemma. This lemma states that if we project the original vectors (for instance those belonging to the effective resistance embedding) onto a subspace spanned by $O(\log n)$ random vectors, the distances between the projected vectors and the original ones are preserved, and then to some extent are given by ϵ .

3 Experiments

We have performed several experiments on the reduction of the 1-skeleton of both triangulations and alpha shapes. As previously mentioned, triangulations are the standard *de-facto* representations of the surface of 3D objects.

Triangulations are sets of triangles and vertices and are fully described by their 1-skeleton. All vertices of a triangulation have the same importance. For instance, it is not possible to distinguish peaks, pits or passes from other structures. Moreover, connections are all represented without any relations with their importance (for instance from shape outliers or dense regions). For this reason, it is necessary to derive more abstract, high level shapes. In this sense, sparsification can act as a tool able to determine a hierarchy between the vertex connections. It may therefore determine a relative importance of the vertices.

Alpha shapes provide a family of shape *representations* that is very useful when performing shape reconstruction. The reason for this is that they connect vertices with all neighbourhoods that are enclosed in a ball of radius alpha. In general, alpha shapes generalize triangulations and their importance is mainly theoretical. In our experiments on 3D point clouds, triangulations represent the external boundary connections, while alpha shapes encode spatial (volumetric) relationships.

Figure 2 shows five triangulations used in our experiments. These 3D models correspond to an abstract shape, a cactus, a deer, a cup and a cow model, respectively. Most of these models contain features that can be considered to be at a small scale (for instance the small handles in the abstract shape, the details of the cow and deer models, etc) or to a larger scale, such as the handles and

the elongated parts (in the cactus, the deer and the cup models). The results in Table 1 report the number of edges when sparsification is performed and how they vary when the value of the ϵ parameter increases¹.




Fig. 2. Examples of triangulations used in our experiments.

Table 1. Statistics on the number of edges of the 1-skeleton of some triangulations when the parameter ϵ increases.

Triangulation	$\epsilon = 0$	$\epsilon = 0.25$	$\epsilon = 0.75$	$\epsilon = 1.25$	$\epsilon = 1.75$	$\epsilon = 2.25$
Model in Fig. 2(a)	3906	3905	3837	3090	2098	1438
Model in Fig. 2(b)	4623	4622	4550	3643	2543	1782
Model in Fig. 2(c)	15012	15011	14757	11871	80623	5506
Model in Fig. 2(d)	18837	18836	18504	17146	10392	7290
Model in Fig. 2(e)	21759	21757	21415	17201	11677	7989

Similarly Fig. 3 shows the 1-skeletons of five alpha shapes that were constructed over various point clouds, also varying the α value. These correspond to two different versions of the abstract shape already shown in 2(a), two alpha shapes of the deer model in 2(c) and an alpha shape from the cow point set that correspond to 2(e). The choice of these alpha shapes is motivated by the presence of small and larger handles and features that alpha shapes have difficulty capturing with a single choice of the parameter α , as previously discussed. The results in Table 2 report the number of edges of the 1-skeleton of the alpha shape when the value of the ϵ parameter increases. From these experiment, we think that with sparsification would be possible to overcome the limitations of alpha shapes in the sense that we hope that it will be possible to commence from a quite large value of the parameters α and then to remove the redundant edges by using sparsification, thus implementing a connected, progressive, geometrical-topological peeling of the shape.

Finally, Fig. 4 shows the potential degradation of the topological properties of the simplified shape as ϵ increases. For the abstract shape (models (a),(b) in Fig. 3) we observe that the shape of the graph is preserved up to $\epsilon = 1.2$. However, for $\epsilon = 1.25$ the representation collapses to the most important con-

¹ In this paper, the parameter ϵ controls the number of samples needed by the process, whereas the weight for choosing the edges is given by effective resistances.




Fig. 3. Examples of alpha shapes used in our experiments.

nected component. This is partially due to the fact that the link between the original resistances R_e ($\epsilon = 0$) and their sampled counterparts R'_e is governed by $R'_e = (1 \pm \epsilon)^2 R_e$ according to the JL lemma, if we do not compute them by spectral means. In addition, as ϵ increases we reduce the number of samples $q = 9C^2 n \log n / \epsilon^2$ ($C = 1$ in this paper). This leads to an increment of entropy, which in turn flattens the importance of certain key edges. Therefore, the critical value of ϵ is larger for shapes with an increasing number of nodes. For instance, for the deer alpha shapes we have that the critical value of ϵ is in the range $[1.4, 1.45]$ whereas for the cow alpha shape we have that it is in the range $[1.4, 1.5]$. For triangulations the values are similar but larger. For the blob the critical value is close to 1.4, and for the remaining ones is the range $[1.4, 1.5]$.




Fig. 4. Degradation of topological properties as ϵ increases. 2D projections of the blob alpha shape. From left to right: $\epsilon = 0$, $\epsilon = 0.75$, $\epsilon = 1.0$ and $\epsilon = 1.25$.

Table 2. Statistics on the number of edges of the 1-skeleton of some alpha shapes when the parameter ϵ increases.

Alpha-shape	α	$\epsilon = 0$	$\epsilon = 0.25$	$\epsilon = 0.75$	$\epsilon = 1.25$	$\epsilon = 1.75$	$\epsilon = 2.25$
Model in Fig. 3(a)	3	8492	8491	6960	4167	2453	1621
Model in Fig. 3(b)	10	9526	9525	7643	4316	2563	1663
Model in Fig. 3(c)	1	39224	39223	33083	19476	11596	7531
Model in Fig. 3(d)	10	39707	39705	33416	19440	11664	7502
Model in Fig. 3(e)	10	55598	55596	48044	28769	17304	11235

4 Conclusions

In this paper, we have shown that graph sparsification leads to a principled way of simplifying shapes. Experiments on both triangulations and alpha shapes show promising preliminary results. In particular, it introduces a hierarchy (and

therefore a priority queue) of the edge strength. It is relevant for shape analysis and interpretation. We plan to further develop these ideas, in particular, in relation to the filtrations induced by the theory of topological persistence [?,?].

Acknowledgments. F. Escolano and M. Curado are funded by Project TIN2015-69077-P of the Spanish Government. E. Hancock is

References

1. Achlioptas, D.: Database-friendly random projections. In Buneman, P., ed.: Proceedings of the Twentieth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, PODS, May 21-23, 2001, Santa Barbara, California, USA, ACM (2001)
2. Batson, J.D., Spielman, D.A., Srivastava, N.: Twice-ramanujan sparsifiers. *SIAM J. Comput.* **41**(6) (2012) 1704–1721
3. Batson, J.D., Spielman, D.A., Srivastava, N., Teng, S.: Spectral sparsification of graphs: theory and algorithms. *Commun. ACM* **56**(8) (2013) 87–94
4. Benczúr, A.A., Karger, D.R.: Approximating s - t minimum cuts in $\tilde{O}(n^2)$ time. In: Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, STOC, Philadelphia, Pennsylvania, USA, May 22-24, 1996. (1996) 47–55
5. Berger, M., Tagliasacchi, A., Seversky, L.M., Alliez, P., Guennebaud, G., Levine, J.A., Sharf, A., Silva, C.T.: A survey of surface reconstruction from point clouds. *Computer Graphics Forum* (2016) n/a–n/a
6. Edelsbrunner, H.: Alpha shapes - a survey. In: Tessellations in the Sciences: Virtues, Techniques and Applications of Geometric Tilings. Springer Verlag (2011)
7. Edelsbrunner, H., Harer, J.: Computational Topology: An Introduction. American Mathematical Society (2010)
8. Edelsbrunner, H., Kirkpatrick, D., Seidel, R.: On the shape of a set of points in the plane. *IEEE Transactions on Information Theory* **29**(4) (Jul 1983) 551–559
9. Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. *Discrete Comput. Geom.* **28** (2002) 511–533
10. Edelsbrunner, H., Mücke, E.P.: Three-dimensional alpha shapes. *ACM Trans. Graph.* **13**(1) (January 1994) 43–72
11. Naylor, B., Bajaj, C., Edelsbrunner, H., Kaufman, A., Rossignac, J.: Computational representations of geometry. SIGGRAPH’96 course notes. (1996)
12. Paoluzzi, A., Bernardini, F., Cattani, C., Ferrucci, V.: Dimension-independent modeling with simplicial complexes. *ACM Trans. Graph.* **12**(1) (January 1993) 56–102
13. Qiu, H., Hancock, E.R.: Clustering and embedding using commute times. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(11) (2007) 1873–1890
14. Rudelson, M., Vershynin, R.: Sampling from large matrices: An approach through geometric functional analysis. *J. ACM* **54**(4) (July 2007)
15. Spielman, D.A., Srivastava, N.: Graph sparsification by effective resistances. *SIAM J. Comput.* **40**(6) (2011) 1913–1926
16. Spielman, D.A., Teng, S.: Spectral sparsification of graphs. *SIAM J. Comput.* **40**(4) (2011) 981–1025