

Coulomb shapes: using electrostatic forces for deformation-invariant shape representation

Davide Boscaïni, Ramūnas Girdziušas and Michael M. Bronstein

Institute of Computational Science, Faculty of Informatics, University of Lugano, Switzerland

Abstract

Canonical shape analysis is a popular method in deformable shape matching, trying to bring the shape into a canonical form that undoes its non-rigid deformations, thus reducing the problem of non-rigid matching into a rigid one. The canonization can be performed by measuring geodesic distances between all pairs of points on the shape and embedding them into a Euclidean space by means of multidimensional scaling (MDS), which reduces the intrinsic isometries of the shape into the extrinsic (Euclidean) isometries of the embedding space. A notable drawback of MDS-based canonical forms is their sensitivity to topological noise: different shape connectivity can affect dramatically the geodesic distances, resulting in a global distortion of the canonical form. In this paper, we propose a different shape canonization approach based on a physical model of electrostatic repulsion. We minimize the Coulomb energy subject to the local distance constraints between adjacent shape vertices. Our model naturally handles topological noise, allowing to ‘tear’ the shape at points of strong repulsion. Furthermore, the problem is computationally efficient, as it lends itself to fast multipole methods. We show experimental results in which our method compares favorably to MDS-based canonical forms.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations—Physically based modeling

1. Introduction

Canonical shape analysis is a popular approach in the study of deformable shapes and content-based retrieval. The main idea of this class of methods is, given a shape, to apply to it some ‘canonization’ process that allows to undo the shape deformation. The resulting canonical representation can be treated as a rigid shape, thus dramatically reducing the number of degrees of freedom in the problem, and effectively reducing non-rigid shape matching to a rigid one using e.g. iterative closest points (ICP) methods [BM92, CM92].

Several approaches for shape canonization have been studied. Elad and Kimmel [EK03] proposed measuring geodesic distances between all pairs of points on the shape and then using multidimensional scaling (MDS) to embed these distances into a three-dimensional Euclidean space approximately isometrically. Such a process maps all intrinsic isometries (deformations preserving the geodesic distances) of the shape into Euclidean ones (rotations, translations, and reflections). Given that the MDS step is computationally intensive, several numerical acceleration techniques were proposed, including multigrid [BBKY06] and vector extrapolation

[RBBK10] methods. The complexity of MDS can also be reduced by embedding a small set of surface points (landmarks) and interpolating the rest [Pla05]. The recent spectral MDS approach [AK13] follows this idea using the Laplace-Beltrami eigenbasis for interpolation.

As opposed to MDS-based methods using global structures (geodesic distances), other dimensionality reduction methods use only local information. For instance, Belkin and Niyogi [BN03] and Coifman and Lafon [CL06] proposed to perform the embedding in a low-dimensional space using the eigenfunctions of the Laplace-Beltrami operator. Such an embedding approximately preserves diffusion distances, which are isometry-invariant, similarly to geodesic distances. Though originally developed for high-dimensional data analysis, these algorithms work well when applied to shapes [Rus07], due to the ability of Laplace-Beltrami eigenfunctions to capture meaningful geometric properties of the underlying manifold [Lév06, VL08].

One major drawback of canonical forms representing intrinsic distances in a low-dimensional Euclidean space is that such a process usually incurs a metric distortion (em-



Figure 1: Shape connectivity can influence dramatically the geodesic distances. Shown is the geodesic distance from the white point on the index finger (hotter colors correspond to larger distances) on two poses of the hand, where in the right pose the fingers are glued together.

bedding error) which is data-dependent. Another drawback of canonical forms is their sensitivity to topological noise. Such noise results in connectivity change in the shape, which can alter dramatically the geodesic distances (see Figure 1, where a shortcut introduced by gluing the fingers of the hand ‘reroutes’ the distances). As a consequence, the canonical forms are globally distorted (see examples in Figure 2 and 6). It was shown that using other intrinsic distances (e.g. diffusion distances) instead of geodesic may alleviate the problem, but not resolve it completely [BBK⁺10].

In this paper, we propose a different physical model for shape canonization referred to as *coulombization*, inspired by recent works on graph drawing [Ead84, Hu06, Hos12, GHN13]. Our approach suggests to consider the shape vertices as equally charged particles which tend to repulse by the Coulomb forces, being at the same time constrained by the material properties of the shape. The electrostatic equilibrium can be obtained by minimization of the Coulomb energy subject to the distance constraints imposed for each pair of adjacent nodes. Instead of simulating the complex Newtonian dynamics, we use the ‘force-directed’ gradient minimization [Hu06]. If the metric constraints are imposed exactly (i.e. the shape material is absolutely inelastic), such a canonical representation is isometric (without metric distortion). However, since closed polyhedral surfaces are known to be rigid, it is necessary to relax the metric constraints.

One way of doing it is to impose the metric constraints after a step of Coulomb energy minimization by projecting the position of the vertices on the spheres defined by the adjacent nodes and their initial (unnormalized) distance values.

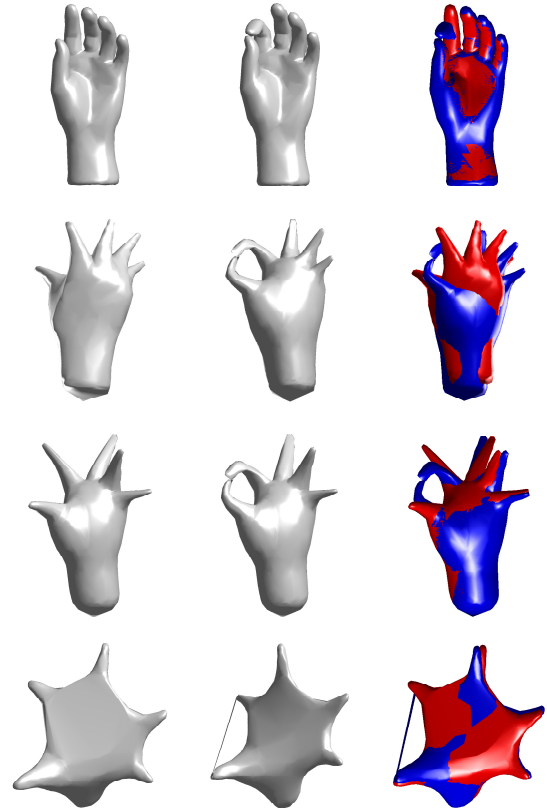


Figure 2: Illustration of the effect of topological noise on shape canonization. First row: human hand shape with touching (center) and separated (left) fingers. Second row: classical MDS canonical forms; third row: SMACOF MDS canonical forms; fourth rows: Coulomb shapes. Rightmost column: rigid alignment of the shapes shown in the first two columns using ICP.

This approach has been previously considered in the context of source localization [Gro59]. The principal consequence of this approach is that the constraints are not imposed precisely, but only ‘on average’. Thus, such a relaxation allows to overcome the rigidity of the polyhedra and let the surface to ‘tear’ under the influence of the repulsion forces (however, at the expense of metric distortion which stems from having only approximate metric constraints). Another important consequence concerns the fact that, with this approach, we can also handle efficiently topological artifacts, ruling out from this ‘average’ the local distance constraints that are violated the most due to repulsion forces.

From the computational standpoint, our problem has significant advantage over MDS-based approaches. First, Coulomb interactions are a classical case for the use of fast multipole methods (FMM) [GR87], leading to an $\mathcal{O}(n \log n)$ iteration complexity, where n is the number of vertices.

Furthermore, since we use only local (edge) distance constraints, there is no need to compute and store the $n \times n$ matrix of all pairwise distances.

The rest of the paper is organized as follows. In Section 2, we formulate our problem of shape coulombization. Section 3 shows experimental results comparing the proposed approach to MDS-based canonical forms. Finally, Section 4 concludes the paper.

2. Shape coulombization

We model a shape as a triangulated mesh (\mathbf{X}, E, T) , where E and T denote the edge and face structures, respectively, and \mathbf{X} is the $n \times 3$ matrix representing the 3D coordinates of the mesh vertices $\mathbf{x}_1, \dots, \mathbf{x}_n$. We assume that the shape is connected (i.e. has no disconnected components). We denote by d_{ij} , $(i, j) \in E$ the edge lengths, representing the discrete Riemannian metric of the shape. Since each vertex is connected to approximately a constant number of vertices (typically 5–6), the number of edges is $|E| = \mathcal{O}(n)$.

According to our physical model, we place at each vertex an equally-charged particle. Following Coulomb's law, the repulsion force between two charges at coordinates $\mathbf{x}_i, \mathbf{x}_j$ is inversely proportional to their Euclidean distance $\|\mathbf{x}_i - \mathbf{x}_j\|_2$. The overall 'stress' of a given configuration of points \mathbf{X} can be described by the *Coulomb energy*

$$\mathcal{E}(\mathbf{X}) = \sum_{i \neq j} \frac{1}{\|\mathbf{x}_i - \mathbf{x}_j\|_2}. \quad (1)$$

Our shape coulombization process consists of minimizing this energy subject to metric preservation constraints,

$$\mathbf{X}^* = \arg \min_{\mathbf{X}} \mathcal{E}(\mathbf{X}) \text{ s.t. } \|\mathbf{x}_i - \mathbf{x}_j\|_2 = d_{ij}, (i, j) \in E. \quad (2)$$

This problem has two forces that counterbalance each other. On the one hand, electrostatic repulsion forces try to draw the shape vertices apart as far as possible from each other. On the other hand, metric constraints do not allow the shape to stretch. The result is a canonical form that undoes the shape deformation similarly to MDS (Figure 2).

Note that the metric constraints in (2) guarantee that the resulting shape is isometric to the original one, hence there is no metric distortion typical to MDS-based method. At the same time, it is known that closed polyhedral shapes are rigid, with the exception of a few pathological examples such as the so-called Connelly sphere [Con79]. Therefore, problem (2) might be over-constrained.[†]

[†] More precisely, rigidity means that a polyhedron cannot be *continuously* bent isometrically. Some constrained optimization algorithms guarantee that all the iterations performed by the optimization are feasible, while other guarantee feasibility only of the solution. In the latter case, it is possible to obtain intermediate steps that are not isometric.

We see two possible remedies to this issue. First, we can reformulate the optimization problem (2) using approximate metric constraints allowing some elasticity, e.g.

$$\begin{aligned} \mathbf{X}^* &= \arg \min_{\mathbf{X}} \mathcal{E}(\mathbf{X}) \\ \text{s.t. } &\frac{1}{c} d_{ij} \leq \|\mathbf{x}_i - \mathbf{x}_j\|_2 \leq c d_{ij}, (i, j) \in E. \end{aligned} \quad (3)$$

where $c \geq 1$ is the *dilation constant*, controlling the allowed amount of metric distortion. A second alternative, which we discuss in the following, is adopting an optimization algorithm that imposes the metric constraints approximately.

2.1. Numerical optimization

We propose a simple approach to solving problem (2) using a projected gradient descent, consisting of alternating the (unconstrained) minimization of $\mathcal{E}(\mathbf{X})$ with a projection on the metric constraints. The gradient step at iteration t has the form

$$\mathbf{X}^{(t)} = \mathbf{X}^{(t-1)} - c^{(t-1)} \nabla \mathcal{E}(\mathbf{X}^{(t-1)}), \quad (4)$$

where c denotes the step size that can be fixed or determined by line search. In our experiments, we set the step size using the following heuristic,

$$c^{(t)} = 0.1 \frac{\max_{i,j=1,\dots,n} d_{ij}}{\max_{i=1,\dots,n} \|\nabla_{\mathbf{x}_i} \mathcal{E}(\mathbf{X}^{(t)})\|}.$$

The energy $\mathcal{E}(\mathbf{X})$ and its gradient $\nabla \mathcal{E}(\mathbf{X})$ can be evaluated with $\mathcal{O}(n \log n)$ complexity using the fast multipole method. In our experiments, we employed the code developed in [GG13,fmp].

After each gradient descent iteration, instead of enforcing the metric constraints strictly, we minimize their violation

$$\sum_{(i,j) \in E} (d_{ij} - \|\mathbf{x}_i - \mathbf{x}_j\|)^2. \quad (5)$$

Groginsky [Gro59] proposed solving this problem through the fixed point iteration

$$\mathbf{x}_i^{(t)} \leftarrow \frac{1}{v_i} \sum_{(i,j) \in E} \left(\mathbf{x}_j^{(t)} + d_{ij} \frac{\mathbf{x}_i^{(t)} - \mathbf{x}_j^{(t)}}{\|\mathbf{x}_i^{(t)} - \mathbf{x}_j^{(t)}\|_2} \right), \quad (6)$$

where v_i denotes the number of vertices connected to \mathbf{x}_i , which is assumed to be constant. Geometrically, equation (6) is nothing but the average of the projections of the vertex \mathbf{x}_i over the spheres of radii d_{ij} centered in \mathbf{x}_j . For the sake of brevity, we refer to this interpretation of (6) saying that we satisfy the constraints 'on average'. Moreover, under mild technical conditions it is possible to show that (6) is a convergent scheme [Gro59], [APV10]. The overall complexity of the constraint projection step is $\mathcal{O}(n)$.

The optimization process alternates between Coulomb energy minimization (one or a few steps of (4)) and metric constraint projection (one or a few steps of (6)). Varying the number of Coulomb steps and the projection steps in this

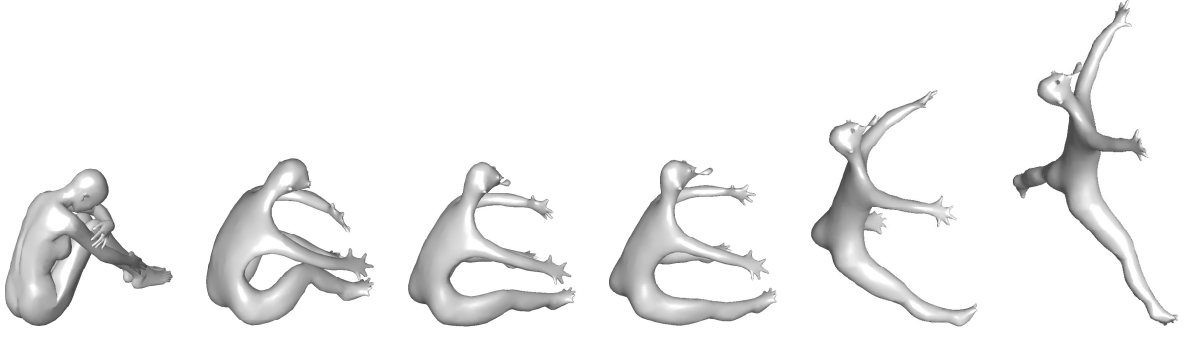


Figure 3: Coulombization of a female shape (leftmost) after (left to right) 300, 600, 900, 3000, and 10000 iterations.

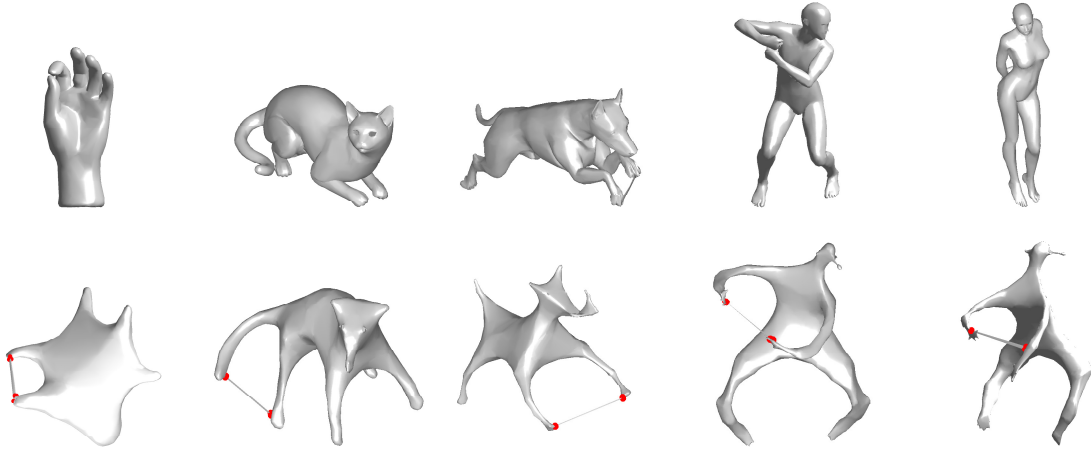


Figure 4: Robustness of Coulomb shapes to topological noise. First row: shapes with point-wise connectivity artifacts; second row: the corresponding coulombization. Shown in red are vertices that violate the constraints the most, and at which the surface is allowed to ‘tear’.

alternating schemes allows controlling the resulting metric distortion (Figure 5). Figure 3 shows a few iterations of the proposed coulombization process. Similarly to MDS methods, there is no theoretical guarantee of global convergence. However, having a good initialization (the original shape vertex coordinates), we observe good convergence behavior.

2.2. Handling topological noise

The structure of our optimization scheme allows to naturally handle topological artifacts. The main idea is to down-weight or eventually disconnect pairs of points that experience strong tension, allowing the mesh to ‘tear’ in a point-wise manner. Due to the local nature of the topological noise, instead of searching for coordinates $\mathbf{x}_i, \mathbf{x}_j$ that match with d_{ij} in a uniform way, we want to let violate the constraints for a small amount of vertices. If equation (5) reminds us the L^2 norm,

$$\sum_{(i,j) \in E} |d_{ij} - \|\mathbf{x}_i - \mathbf{x}_j\||. \quad (7)$$

led to a sparsity-inducing L^1 norm.

As noted in [APV10], (7) could be solved replacing (6) by a weighted projection

$$\mathbf{x}_i^{(t)} \leftarrow \sum_{(i,j) \in E} w_{ij}^{(t)} \left(\mathbf{x}_j^{(t)} + d_{ij} \frac{\mathbf{x}_i^{(t)} - \mathbf{x}_j^{(t)}}{\|\mathbf{x}_i^{(t)} - \mathbf{x}_j^{(t)}\|_2} \right), \quad (8)$$

where

$$w_{ij}^{(t)} = \frac{(\|\mathbf{x}_i^{(t)} - \mathbf{x}_j^{(t)}\|_2 - d_{ij})^{-1}}{\sum_{(i,j) \in E} (\|\mathbf{x}_i^{(t)} - \mathbf{x}_j^{(t)}\|_2 - d_{ij})^{-1}}. \quad (9)$$

In order to avoid division by zero, the distances are always lower-bounded by a fixed threshold 10^{-16} . The weights are updated in parallel as in the Jacobi method for linear systems. Such weighting relaxes the metric constraints for edges that tend to stretch significantly, allowing for some elasticity at such points. Furthermore, after an initial stage of repulsion (typically 10 – 100 iterations), we select a small number of edges (e.g. 1% of $|E|$) with the largest deviations

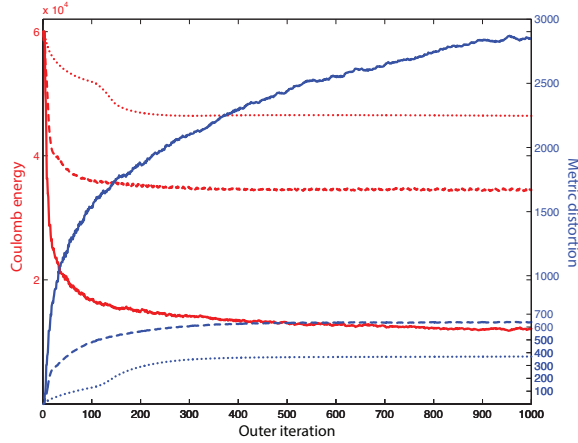


Figure 5: Typical convergence behavior of the shape coulombization algorithm. Shown is Coulomb energy (red, left scale) and the average metric constraint violation (blue, right scale) vs outer iteration number for different number of projection steps (solid: 1, dashed: 10, dotted: 100).

of the Euclidean distances from their given edge lengths d_{ij} and completely remove them from the adjacency set. This amounts to ‘tearing’ the surface.

3. Results

In this section, we show examples of shape coulombization using our method, implemented as described in the previous section. For our experiments, we used shape from the TOSCA dataset [BBK08], sampled at approximately 3.5K vertices. Topological noise was introduced by adding faces to connect between almost-touching vertices with large geodesic and small Euclidean distances. As baseline, we used canonical forms implemented with two MDS methods: classical scaling [YH38] and SMACOF [dL77] (for details on these algorithms, we refer the reader to [BG05]). Geodesic distances were computed using the fast marching algorithm [KS98].

Figure 6 shows a sample of the TOSCA shapes [BBK08] and their normalizations with the coulombization and MDS canonical forms. One can clearly see that our approach achieves a similar effect in undoing the deformation, while being drastically less sensitive to topological noise: in fact, its effect on the resulting Coulomb shapes is practically unnoticeable. On the other hand, canonical forms manifest severe global distortions as the result of topological noise. This difference is highlighted in Figure 2, where we align the resulting canonical forms using ICP.

Figure 7 shows a confusion matrix of shapes from four classes (male, female, dog, cat) with four near-isometric deformations and one containing topological noise. These shape classes are challenging due to similarity between the

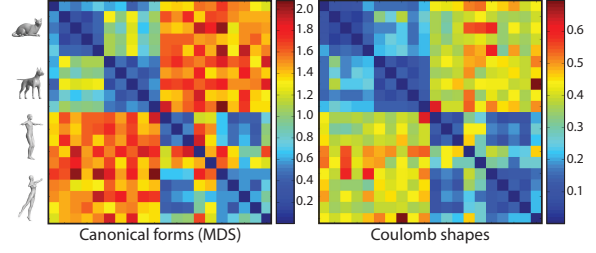


Figure 7: Distances between shapes computed using MDS-based canonical forms (left) and the proposed approach (right).

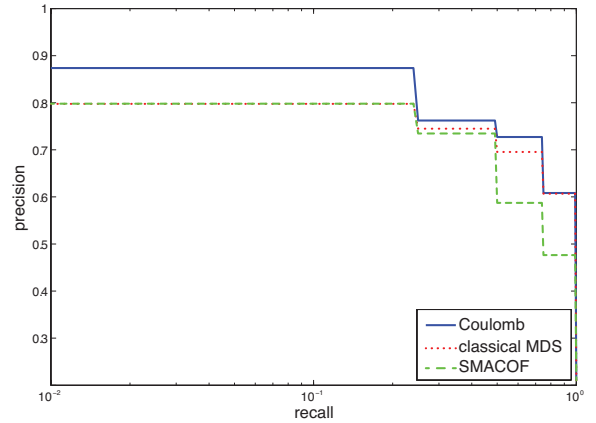


Figure 8: Precision-recall curve of MDS-based canonical forms and the proposed approach.

bipeds and the quadrupeds. Distances between shapes were computed by first canonizing the shape using the proposed approach followed by ICP; the distance between shapes is measured as the maximum distance between corresponding points after rigid alignment. Figure 8 shows the corresponding precision-recall curve. For comparison, we show the performance of MDS-based canonical forms.

3.1. Complexity

The theoretical complexity of MDS-based canonical forms includes the computation of geodesic distances ($\mathcal{O}(n \log n)$ using fast marching) and the SMACOF algorithm, which has an $\mathcal{O}(n^2)$ iteration complexity. On the other hand, the computation of Coulomb shapes involves precomputation of the edge lengths ($\mathcal{O}(n)$); the iteration complexity comprises the Coulomb energy gradient computation ($\mathcal{O}(n \log n)$ using fast multipole method) and projection ($\mathcal{O}(n)$). Table 1 shows the typical timing (averaged over 100 runs) for MDS and Coulomb canonical forms.

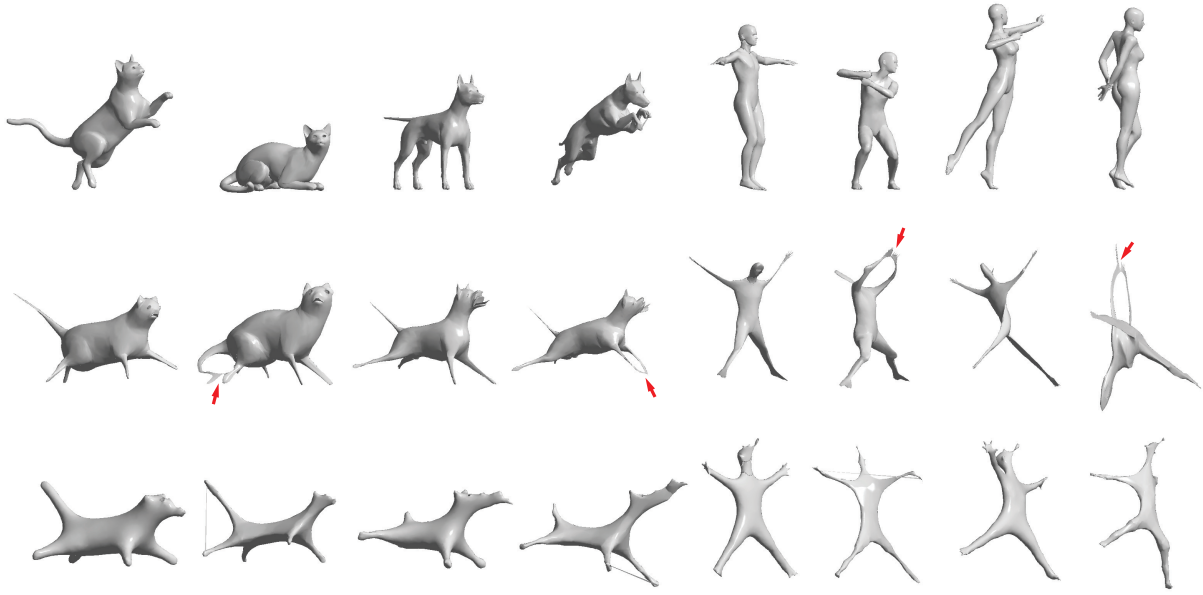


Figure 6: Examples of shapes (top row), their MDS canonical forms (computed using SMACOF, middle row) and the proposed coulombization (bottom row). Shapes in odd columns have different connectivity (the cat’s tail is glued to the body, the dog’s legs are glued together, etc). These topological artifacts result in significant distortions of the MDS canonical forms, marked with red arrows. On the other hand, the proposed approach handles the topological noise gracefully, allowing the surface to ‘tear’ at these points without introducing global distortions.

Table 1: Complexity comparison (time in sec. of the pre-computation stage and one iteration of the optimization) of SMACOF MDS-based canonical forms and the proposed approach (using a single Coulomb energy minimization step and a single metric projection step) for shapes with different number of vertices n .

Shape	n	MDS		Coulomb	
		Prec.	Iter.	Prec.	Iter.
<i>swissroll</i>	561	0.14	0.03	0.003	0.01
<i>cat</i>	3505	10.24	1.59	0.036	0.045
<i>michael</i>	22738	647.6	277.5	0.13	0.21

4. Conclusion

We studied the use of electrostatic model for shape canonization. Compared to MDS-based canonical forms, our approach has several important advantages. First, we avoid the computation and storage of the full distance matrix, using only local information (edge lengths). This leads to linear $\mathcal{O}(n)$ storage complexity. Second, Coulomb forces are very well suited for the use of fast multipole methods, which allow to achieve iteration complexity of $\mathcal{O}(n \log n)$. Unlike some other use cases of fast multi-body simulations in the embedding literature [vdM13, YPK13], rigorous gradient approximation estimates are available as the model is precisely

expressed with the spatially decaying Coulomb potential. Third, the averaged spherical projections provide a simple way to impose metric constraints, and allow to naturally handle topological noise. Our experiments show that the proposed method performs better than previous approaches in the presence of such artifacts.

Though in this paper we treated shapes represented as triangular meshes, in fact we do not use the triangle structures, but only the edges. This means that our approach should also be applicable to point clouds, which we plan to study in future works. Another important future research direction is how to exactly control the embedding error, which is a potentially big advantage of the proposed approach.

Acknowledgment

We would like to thank Randolph Schärfig for helping to generate meshes with topological artifacts and Davide Eynard for helping to vectorize our Matlab code. This research was supported by ERC Starting Grant no. 307047 (COMET).

References

- [AK13] AFLALO Y., KIMMEL R.: Spectral multidimensional scaling. *Proc. of the National Academy of Science* (2013). 1
- [APV10] AGARWAL A., PHILLIPS J. M., VENKATASUBRAMANIAN S.: Universal Multi-Dimensional Scaling. In *Proc. of*

- the *International Conference on Knowledge Discovery and Data Mining* (2010), pp. 1149–1158. 3, 4
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 5
- [BBK*10] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R., MAHMOUDI M., SAPIRO G.: A Gromov-Hausdorff Framework with Diffusion Geometry for Topologically-Robust Non-Rigid Shape Matching. *International Journal of Computer Vision* 89, 2–3 (2010), 266–286. 2
- [BBKY06] BRONSTEIN M. M., BRONSTEIN A. M., KIMMEL R., YAVNEH I.: Multigrid multidimensional scaling. *Numerical Linear Algebra with Applications* 13 (2006), 149–171. 1
- [BG05] BORG I., GROENEN P. J. F.: *Modern Multidimensional Scaling: Theory and Applications*. Springer, 2005. 5
- [BM92] BESL P. J., MCKAY N. D.: A Method for Registration of 3-D Shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 14, 2 (1992), 239–256. 1
- [BN03] BELKIN M., NIYOGI P.: Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural computation* 15, 6 (2003), 1373–1396. 1
- [CL06] COIFMAN R. R., LAFON S.: Diffusion maps. *Applied and Computational Harmonic Analysis* 21, 1 (2006), 5–30. 1
- [CM92] CHEN Y., MEDIONI G.: Object modelling by registration of multiple range images. *Image and Vision Computing* 10, 3 (1992), 145–155. 1
- [Con79] CONNELLY R.: The Rigidity of Polyhedral Surfaces. *Mathematics Magazine* 52, 5 (1979), 275–283. 3
- [dL77] DE LEEUW J.: Applications of Convex Analysis to Multidimensional Scaling. In *Recent developments in statistics* (1977), pp. 133–145. 5
- [Ead84] EADES P.: A heuristic for graph drawing. *Congressus Numerantium* 42 (1984), 149–160. 2
- [EK03] ELAD A., KIMMEL R.: On Bending Invariant Signatures for Surfaces. *Trans. on Pattern Analysis and Machine Intelligence* 25, 10 (2003), 1285–1295. 1
- [fmp] FMPS (Fast Multi-Particle Scattering) software. <http://www.cims.nyu.edu/cmcl/fmps/fmps.html>. 3
- [GG13] GIMBUTAS Z., GREENGARD L.: Fast multi-particle scattering: a hybrid solver for the Maxwell equations in microstructured materials. *Journal of Computational Physics* 233, 1 (2013), 22–32. 3
- [GHN13] GANSNER E. R., HU Y., NORTH S.: A Maxent-Stress Model for Graph Layout. 927–940. 2
- [GR87] GREENGARD L., ROKHLIN V.: A Fast Algorithm for Particle Simulations. *Journal of Computational Physics* 73, 2 (1987), 325–348. 2
- [Gro59] GROGINSKY H. L.: Position Estimation Using Only Multiple Simultaneous Range Measurements. *IRE Trans. on Aeronautical and Navigational Electronics* 6, 3 (1959), 178–187. 2, 3
- [Hos12] HOSOBÉ H.: Numerical Optimization-Based Graph Drawing Revisited. In *IEEE Pacific Visualization Symposium* (2012), pp. 80–88. 2
- [Hu06] HU Y. H.: Efficient, High-Quality Force-Directed Graph Drawing. *The Mathematica Journal* 10, 1 (2006), 37–71. 2
- [KS98] KIMMEL R., SETHIAN J. A.: Computing Geodesic Paths on Manifolds. *Proc. of the National Academy of Sciences* 95, 15 (1998), 8431–8435. 5
- [Lév06] LÉVY B.: Laplace-beltrami eigenfunctions: towards an algorithm that “understands” geometry. In *IEEE International Conference on Shape Modeling and Applications* (2006). 1
- [Pla05] PLATT J. C.: FastMap, MetricMap, and Landmark MDS are all Nyström algorithms. In *Proc. of the International Workshop on Artificial Intelligence and Statistics* (2005), pp. 261–268. 1
- [RBBK10] ROSMAN G., BRONSTEIN M. M., BRONSTEIN A. M., KIMMEL R.: Nonlinear Dimensionality Reduction by Topologically Constrained Isometric Embedding. *International Journal of Computer Vision*, 89 (2010), 56–68. 1
- [Rus07] RUSTAMOV R. M.: Laplace-Beltrami Eigenfunctions for Deformation Invariant Shape Representation. In *Eurographics Symposium on Geometry Processing* (2007), pp. 225–233. 1
- [vdM13] VAN DER MAATEN L. J. P.: Barnes-Hut-SNE. *arXiv:1301.3342* (2013). 6
- [VL08] VALLET B., LÉVY B.: Spectral Geometry Processing with Manifold Harmonics. *Computer Graphics Forum* (2008), 251–260. 1
- [YH38] YOUNG G., HOUSEHOLDER A. S.: Discussion of a set of points in terms of their mutual distances. *Psychometrika* 3, 1 (1938), 19–22. 5
- [YPK13] YANG Z., PELTONEN J., KASKI S.: Scalable Optimization of Neighbor Embedding for Visualization. In *Proc. of the International Conference on Machine Learning* (2013), pp. 127–135. 6