



Particle-Based Shape Modeling for Arbitrary Regions-of-Interest

Hong Xu^(✉), Alan Morris, and Shireen Y. Elhabian

Scientific Computing and Imaging Institute, School of Computing,
University of Utah, Salt Lake City, UT, USA
{[hxu](mailto:hxu@sci.utah.edu),[amorris](mailto:amorris@sci.utah.edu),[shireen](mailto:shireen@sci.utah.edu)}@sci.utah.edu
<http://www.sci.utah.edu>

Abstract. Statistical Shape Modeling (SSM) is a quantitative method for analyzing morphological variations in anatomical structures. These analyses often necessitate building models on targeted anatomical regions of interest to focus on specific morphological features. We propose an extension to particle-based shape modeling (PSM), a widely used SSM framework, to allow shape modeling to arbitrary regions of interest. Existing methods to define regions of interest are computationally expensive and have topological limitations. To address these shortcomings, we use mesh fields to define free-form constraints, which allow for delimiting arbitrary regions of interest on shape surfaces. Furthermore, we add a quadratic penalty method to the model optimization to enable computationally efficient enforcement of any combination of cutting-plane and free-form constraints. We demonstrate the effectiveness of this method on a challenging synthetic dataset and two medical datasets.

1 Introduction

Statistical Shape Modeling (SSM) is a widespread method used to analyze shape variation across 3D anatomical samples within a population. These analyses are crucial in detecting common morphological pathologies and advancing the understanding of different diseases by studying the form-function relationships between anatomies [2, 5–8, 13, 16–18, 21]. While building SSMs, certain biomedical and clinical applications require a focus on specific anatomical regions of interest (ROIs) to tailor the analysis to precise morphological features (e.g. [1–4, 11, 14–16]). Such applications might require excluding certain surface aspects, modeling certain regions in isolation, or a mix of these. ROI definition without altering the input shape has been achieved using *constraints*, mathematical delimiters that restrict model construction to certain surface areas [11]. Our approach focuses on redesigning the constraint application method to improve its functionality, flexibility, and efficiency during SSM construction.

To construct such SSMs, two distinct families of shape representations can be used to allow for statistical analysis, *deformation fields* and *landmarks*. Whereas the former encodes *implicit* transformations between cohort samples and a pre-defined (or learned) atlas, the latter uses *explicit* landmark points spread on

shape surfaces that correspond across the population [19,20]. We focus on the latter approach given its extensive use due to its simplicity, computational efficiency, and interpretability for statistical analysis [19,22]. Although landmarks used to be manually placed on specific anatomical features of interest, the modern convention uses dense automatically-placed landmarks obtained through computational methods, such as minimum description length (MDL) [12], and particle-based shape modeling (PSM) [9,10]). We utilize PSM, an efficient and robust entropy-based optimization method that creates a system of dense landmarks or *particles*, which conform to all population shape surfaces while maintaining correspondence across them.

A previous attempt to constrain PSM particle distributions uses geometric primitives in the form of spheres or cutting planes to exclude regions [11]. This exclusion is achieved by projecting *virtual particles* onto these geometric primitives (represented as parametric constraints), relying on the entropy objective to repel landmark particles away from these areas. Such an approach has the advantage of not altering input surfaces, which can otherwise distort morphology or necessitate manual expert-driven reprocessing of data. However, it falters when arbitrary regions of interest cannot be expressed via geometric primitives, limiting the topologies to be modeled. It also exhibits poor scaling due to it requiring an additional set of projected virtual particles per constraint. Thereby, to address these shortcomings in the existing literature, we propose the use of the quadratic penalty method in the optimization to allow the simultaneous and scalable application of cutting-plane, spheres, other primitive constraints, as well as a proposed method of defining arbitrary surface constraints, or *free-form constraints* (FFCs). This method provides both flexibility in the definition of constraints to define ROIs and scalability with large-scale or heavily constrained populations without the need to reprocess data.

2 Method

The aforementioned automatic landmark placement methods take in a population of I -shapes $\mathcal{S} = \{\mathbf{S}_i\}_{i=1}^I$ (binary segmentations, meshes, or n -dimensional contours), and obtain particles $\mathcal{P} = \{\mathbf{P}_i\}_{i=1}^I$ where the i -th shape point distribution model (PDM) is denoted by J -particles $\mathbf{P}_i = [\mathbf{p}_{i,1}, \mathbf{p}_{i,2}, \dots, \mathbf{p}_{i,J}]$, where $\mathbf{p}_{i,j} \in \mathbb{R}^3$. Such particles are obtained by optimizing an objective $f(\mathcal{P})$, which give

$$f(\mathcal{P}) = H(\mathcal{P}) - \sum_{j=1}^J H(\mathbf{P}_i), \quad (1)$$

where H is an estimation of the differential entropy. The particles enable quantifying subtle differences and computing shape statistics (e.g., by performing the principal component analysis (PCA) on corresponding particles) by providing a population-specific anatomical mapping across the given cohort.

We constrain each shape \mathbf{S}_i by M_i -inequality constraints in the form $g_{i,m}(\mathbf{p}) \leq 0$, where $g_{i,m}(\mathbf{p})$ is a differentiable function. These parametric constraints can be in the form of cutting planes or spheres as showcased in [11] (by using the equations of planes or spheres), other parametric delimiters, or free-form constraints, which allow arbitrary surface region definition. These constraints limit the distribution of particles to regions that satisfy the inequality, a region more easily demarcated using parametric constraints in some anatomies, and/or free-form *surface-painting* in others.

In this section, we describe the use of a quadratic penalty method to allow efficient and simultaneous enforcement of an arbitrary number of parametric constraints, and the use of signed mesh vector fields to build free-form constraints that allow arbitrary surface region isolation. We will also showcase a friendly graphical interface to define these constraints.

2.1 Quadratic Penalty for Efficient Constrained PDM Construction

We define an extended objective function to express this constrained optimization problem in an unconstrained form. For each constraint function in the form $g_{i,m}(\mathbf{p}) \leq 0$, we add a quadratic penalty term $g_{i,m}^+(\mathbf{p}) = \max(0, g_{i,m}(\mathbf{p}))$ to the optimization objective, yielding

$$F(\mathcal{P}) = f(\mathcal{P}) + \sum_{i=1}^I \sum_{m=1}^{M_i} \sum_{j=1}^J g_{i,m}^+(\mathbf{p}_{i,j}). \quad (2)$$

We optimize this objective function using a Gauss-Seidel gradient descent scheme, with the second term preventing particles from violating constraints, hence restricting their movement exclusively to feasible regions. This method scales linearly with respect to the number of particles per shape, whereas the virtual particle model [11] scales quadratically.

2.2 Free-Form Constraints

We express free-form constraints in the same form $g_{i,m}(\mathbf{p}) \leq 0$ for each shape \mathbf{S}_i by attributing a distance and gradient field onto each vertex of a mesh \mathbf{M}_i . Any feasible region on the surface of \mathbf{M}_i can be delineated by a set of surface boundaries $\mathcal{B}_i = [\mathbf{B}_{i,1}, \mathbf{B}_{i,2}, \dots, \mathbf{B}_{i,B}]$, which are represented as vertex loops on the mesh surface. A distance field query for a particle \mathbf{p} , denoted $\mathbf{M}_i^d(\mathbf{p})$, provides the signed geodesic distance to the closest constraint boundary $\mathbf{B}_{i,*}$ from the projection of \mathbf{p} onto M_i , illustrated in Fig. 1 (b). Similarly, a gradient field query, denoted $\mathbf{M}_i^g(\mathbf{p})$, would provide the gradient direction, shown in Fig. 1 (c). Ultimately, the mesh \mathbf{M}_i together with its fields \mathbf{M}_i^d and \mathbf{M}_i^g , can approximate the distance and first-order gradients over near-surface points, effectively simulating a differentiable function $g_{i,m}(\mathbf{p})$. When integrated into the aforementioned penalty method 2.1, this approach can enforce arbitrary surface constraints.




Fig. 1. (a) Constrained particle distribution on a sphere, where yellow illustrates the feasible region of the constrained area where particles are allowed to be distributed, the gray is the infeasible region where if particles were to be there, they would be violating the constraint. (b) Distance field $\mathbf{M}_i^d(\mathbf{p})$ of signed geodesic distances to the surface of every mesh vertex. (c) Gradient field $\mathbf{M}_i^g(\mathbf{p})$ on the mesh surface at every mesh vertex represented using white arrows and the blue surface as the feasible region. (Color figure online)

2.3 Graphical Interface Tool

We include a graphical interface tool that can define cutting planes and FFCs and can roughly propagate these to all shapes in the population. Cutting planes are defined by prompting 3 points that the user can pick that are on the shape surface, and can be copy-pasted into all other shapes. FFCs are defined using a “painting” tool that can define included and excluded areas with an adjustable brush size. This tool allows precise and arbitrarily customizable definition of constraints. An FFC on a single shape can be propagated to others using deformation parameters computed from image registration. This functionality is also included. All the graphical interface functionality is illustrated in Fig. 2.

3 Results

We demonstrate our results by integrating our method into an open-source implementation of the particle-based shape modeling (PSM) framework, ShapeWorks [9], and produce SSMs from three datasets. The first is a synthetic dataset of ellipsoids that vary between values of 10, 20, 30, and 40, in each of their three major axes, totaling 64 ellipsoids. These ellipsoids are constrained by a free-form boundary that divides each ellipsoid into upper and lower halves by a full period of a sine wave projected onto the surface, providing a challenging but uniformly delimited population of shapes. Figure 3 shows a few examples and the modes of variation from the SSM. The constraints have the desired effect, and the modes of variation meet expectations as they mimic the variation in the three major axes.

The second is a dataset of 25 computerized tomography (CT) femurs, where the region of interest is the proximal femur sans the lesser trochanter (femoral head, neck, and greater trochanter). For each shape, we use a cutting plane

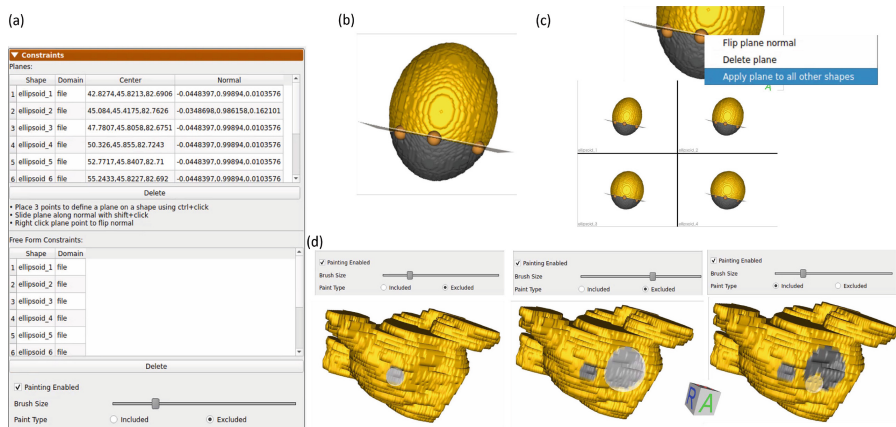


Fig. 2. (a) The constraint panel shows the constraints that have been defined and the tools to define the constraints. (b) Cutting-plane constraints are defined by ctrl-clicking 3 points on the shape surface. (c) Constraints can be flipped or applied to all other shapes via the right-click menu. (d) FFCs are defined with a painting tool with different brush sizes and options to customize included and excluded areas. We show how the painting of excluded areas of different sizes applied to a segmentation of a left atrium.

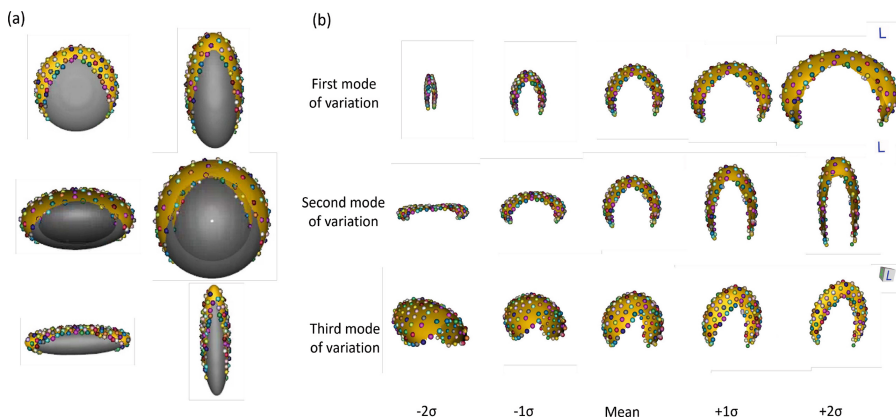


Fig. 3. (a) Sample ellipsoids from the dataset with feasible regions in yellow and restricted regions in grey. (b) The first three modes of variation in the dataset, which show the variation in corresponding major axes. (Color figure online)

constraint to exclude the shaft and a free-form constraint to exclude the lesser trochanter. Figure 4 illustrates a few examples and the first two modes of variation. We observe that a cutting plane allows a more straightforward exclusion of the shaft whilst the FFC precisely excludes the lesser trochanter. The constraints restrict the movement of particles to the feasible region as expected, and the modes of variation meet expectation as well.




Fig. 4. (a) Example of defined constraints. The feasible region is shown in yellow and the constrained region in grey. (b) The first two modes of variation in the dataset. Notice that particles are excluded from the lesser trochanter. (Color figure online)

The third dataset comprises 21 segmentation of left atria models obtained from MRIs. The pulmonary veins represent the area of greatest variation both in anatomical structure (e.g. number of veins, common veins, etc.) as well as greatest variability in segmentation by expert observers (e.g. length into vein to segment). While the position of veins may be important from a shape modeling perspective, their exact shape is not particularly relevant to LA shape morphology. Thus, we paint a free-form constraint exclusion area around the veins. Figure 5 showcases some examples of the shape and the first three modes of variation. The models meet expectations.

4 Conclusion

We demonstrate a flexible and more scalable approach to define regions of interest in fully-groomed shapes for landmark-based statistical shape modeling by allowing arbitrary definition of surface constraints via FFCs and incorporating mixed constraint types into the optimization. This significantly improves the usability of PSM methods, obviating the need for reprocessing datasets. Future work includes the automatic propagation of constraints to the entire cohort given manual definitions on certain representative shapes.

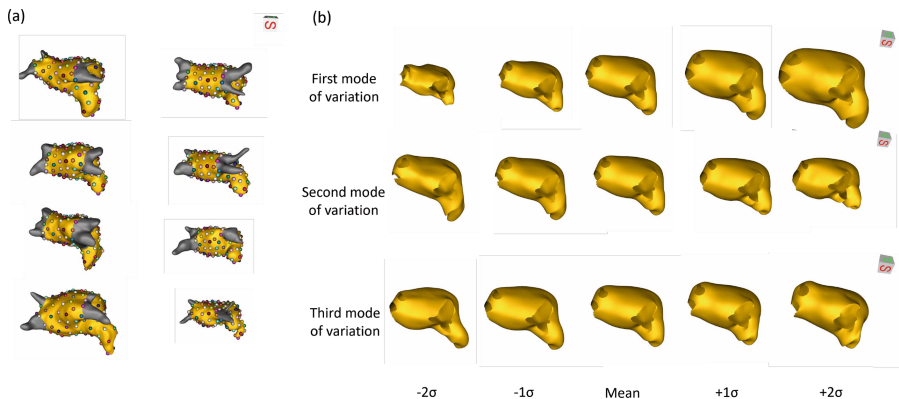


Fig. 5. (a) Example of defined constraints where the left atrium (in yellow) segmentations have the pulmonary veins excluded (in grey). (b) The first three modes of variation in the dataset. Notice how the pulmonary vein areas remain hollow. (Color figure online)

References

1. Atkins, P.R., et al.: Prediction of femoral head coverage from articulated statistical shape models of patients with developmental dysplasia of the hip. *J. Orthop. Res.* **40**(9), 2113–2126 (2022). <https://doi.org/10.1002/jor.25227>
2. Atkins, P.R., et al.: Quantitative comparison of cortical bone thickness using correspondence-based shape modeling in patients with cam femoroacetabular impingement. *J. Orthop. Res.* **35**(8), 1743–1753 (2017)
3. Atkins, P.R., et al.: Which two-dimensional radiographic measurements of cam femoroacetabular impingement best describe the three-dimensional shape of the proximal femur? *Clin. Orthop. Relat. Res.* **477**(1), 242 (2019)
4. Audenaert, E.A., Pattyn, C., Steenackers, G., De Roeck, J., Vandermeulen, D., Claes, P.: Statistical shape modeling of skeletal anatomy for sex discrimination: Their training size, sexual dimorphism, and asymmetry. *Front. in Bioeng. Biotechnol.* **7** (2019). DOI: <https://doi.org/10.3389/fbioe.2019.00302>, <https://www.frontiersin.org/articles/10.3389/fbioe.2019.00302>
5. Bhalodia, R., Dvoracek, L.A., Ayyash, A.M., Kavan, L., Whitaker, R., Goldstein, J.A.: Quantifying the severity of metopic craniosynostosis: a pilot study application of machine learning in craniofacial surgery. *J. Craniofac. Surg.* **31**(3), 697–701 (2020). <https://doi.org/10.1097/SCS.0000000000006215>
6. Bruse, J.L.: A statistical shape modelling framework to extract 3D shape biomarkers from medical imaging data: assessing arch morphology of repaired coarctation of the aorta. *BMC Med. Imaging* **16**, 1–19 (2016)
7. Carriere, N., et al.: Apathy in Parkinson’s disease is associated with nucleus accumbens atrophy: a magnetic resonance imaging shape analysis. *Mov. Disord.* **29**(7), 897–903 (2014)
8. Cates, J., et al.: Computational shape models characterize shape change of the left atrium in atrial fibrillation. *Clin. Med. Insights: Cardiol.* **8s1**, CMC.S15710 (2014). <https://doi.org/10.4137/CMC.S15710>

9. Cates, J., Elhabian, S., Whitaker, R.: ShapeWorks. In: Statistical Shape and Deformation Analysis, pp. 257–298. Elsevier (2017). <https://doi.org/10.1016/B978-0-12-810493-4.00012-2>
10. Cates, J., Fletcher, P.T., Styner, M., Shenton, M., Whitaker, R.: Shape modeling and analysis with entropy-based particle systems. In: Karssemeijer, N., Lelieveldt, B. (eds.) Information Processing in Medical Imaging: 20th International Conference, IPMI 2007, Kerkrade, The Netherlands, July 2-6, 2007. Proceedings, pp. 333–345. Springer, Berlin, Heidelberg (2007). https://doi.org/10.1007/978-3-540-73273-0_28
11. Datar, M., Cates, J., Fletcher, P.T., Gouttard, S., Gerig, G., Whitaker, R.: Particle based shape regression of open surfaces with applications to developmental neuroimaging. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009, pp. 167–174. Springer, Berlin, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04271-3_21
12. Davies, R.H., Twining, C.J., Cootes, T.F., Waterton, J.C., Taylor, C.J.: A minimum description length approach to statistical shape modeling. *IEEE Trans. Med. Imaging* **21**(5), 525–537 (2002)
13. Harris, M.D., Datar, M., Whitaker, R.T., Jurrus, E.R., Peters, C.L., Anderson, A.E.: Statistical shape modeling of cam femoroacetabular impingement. *J. Orthop. Res.* **31**(10), 1620–1626 (2013). <https://doi.org/10.1002/jor.22389>
14. Jacxsens, M., et al.: Thinking outside the glenohumeral box: Hierarchical shape variation of the periarticular anatomy of the scapula using statistical shape modeling. *J. Orthop. Res.* **38**(10), 2272–2279 (2020). <https://doi.org/10.1002/jor.24589>
15. Jacxsens, M., Elhabian, S.Y., Brady, S.E., Chalmers, P.N., Tashjian, R.Z., Henninger, H.B.: Coracoacromial morphology: a contributor to recurrent traumatic anterior glenohumeral instability? *J. Shoulder Elbow Surg.* **28**(7), 1316–1325 (2019)
16. Lenz, A.L.: Statistical shape modeling of the talocrural joint using a hybrid multi-articulation joint approach. *Sci. Rep.* **11**(1), (2021). <https://doi.org/10.1038/s41598-021-86567-7>
17. Merle, C., et al.: How many different types of femora are there in primary hip osteoarthritis? an active shape modeling study. *J. Orthop. Res.* **32**(3), 413–422 (2014)
18. Merle, C., et al.: High variability of acetabular offset in primary hip osteoarthritis influences acetabular reaming—a computed tomography-based anatomic study. *J. Arthroplasty* **34**(8), 1808–1814 (2019)
19. Sarkalkan, N., Weinans, H., Zadpoor, A.A.: Statistical shape and appearance models of bones. *Bone* **60**, 129–140 (2014)
20. Thompson, D.W., et al.: On growth and form. On growth and form. (1942)
21. van Buuren, M., et al.: Statistical shape modeling of the hip and the association with hip osteoarthritis: a systematic review. *Osteoarthritis and Cartilage* **29**(5), 607–618 (2021). <https://doi.org/10.1016/j.joca.2020.12.003>, <https://www.sciencedirect.com/science/article/pii/S106345842031219X>
22. Zachow, S.: Computational planning in facial surgery. *Facial Plast. Surg.* **31**(05), 446–462 (2015)