Statistical Shape Clustering of Left Atrial Appendages

Jakob M. Slipsager[1]([✉]), Kristine A. Juhl[1], Per E. Sigvardsen[2],
Klaus F. Kofoed[2], Ole De Backer[2], Andy L. Olivares[3], Oscar Camara[3],
and Rasmus R. Paulsen[1]

[1] DTU Compute, Technical University of Denmark,
Kongens Lyngby, Denmark
jmsl@dtu.dk
[2] Department of Cardiology, Rigshospitalet, University of Copenhagen,
Copenhagen, Denmark
[3] Physense, Department of Information and Communication Technologies,
Universitat Pompeu Fabra,
Barcelona, Spain

Abstract. Fifteen percent of all strokes are caused by emboli formed in the left atrium (LA) in case of atrial fibrillation (AF). The most common site of thrombus formation is inside the left atrial appendage (LAA). The LAA is accounting for 70% to 90% of the thrombi formed in the LA in patients with non-valvular AF. Studies have shown there is a correlation between the LAA morphology and risk of ischemic stroke; Chicken Wing and Cauliflower LAA shapes are associated with lower and higher risk, respectively. These two LAA shape categories come from a popular classification in the medical domain, but it is subjective and based on qualitative shape parameters. In this paper, we describe a full framework for shape analysis and clustering of the LAA. Initially, we build a point distribution model to quantitatively describe the LAA shape variation based on 103 LAA surfaces segmented and reconstructed from multidetector computed tomography volumes. We are successfully able to determine point correspondence between LAA surfaces, by non-rigid volumetric registration of signed distance fields. To validate if LAA shapes are clustered, we employ an unsupervised clustering on the shape models parameters to estimate the natural number of clusters in our training set, where the number of shape clusters is estimated by validating the test log-likelihood of several Gaussian mixture models using two level cross-validation. We found that the LAAs surfaces basically formed two shape clusters broadly corresponding to the Chicken wing and non-Chicken Wing morphologies, which fits well with clinical knowledge.

Keywords: Left atrial appendage · Point distribution models · Clustering · Gaussian mixture models

1 Introduction

Atrial fibrillation (AF) causes a 5-fold increase in risk of ischemic stroke, being the cause for approximately 15% all strokes in the United States [7]. Around 70% to 90% of the cases, thrombi are formed inside the left atrial appendage (LAA) in patients with non-valvular AF [17]. The LAA is a complex tubular structure, with a high inter-patient variability, originating from the left atrium (LA). Studies have shown there is a correlation between LAA morphology and risk of ischemic stroke [4,8]. Di Biase et al. [4] reported that the popular named Chicken Wing morphology is associated with lower risk of stroke compared to non-Chicken Wing morphology. Several studies have focused on describing the varying LAA morphology, where the morphology is described by the LAA length, width, orfice/ostium size, and number of lobes. In a study based on 220 LAA obtained from necropsy studies, Ernt et al. [5] reported variation in LAA volumes ranging from 770 to 19,270 mm^3, minor orifice diameters ranging from 5 to 27 mm, major orifice diameters between 10 and 40 mm, and LAA lengths ranging between 16 and 51 mm.

The aim of this work was to quantitatively describe the LAA shape variation and clustering using a statistical shape model. We trained a point distribution model (PMD) based on LAA surfaces reconstructed from multidetector computed tomography (CT) images and later combined the trained PMD together with unsupervised clustering methods to examine the natural clustering of the LAA shapes.

2 Data and Preprocessing

The LAA surfaces are reconstructed from CT images, provided by the Department of Radiology, Rigshospitalet, University of Copenhagen. The data are acquired as part of the Copenhagen General Population Study [12], where participants are offered a research cardiac computed tomography angiography (CCTA) examination [6]. Participants are excluded from the examination if they, among other things, suffer from AF. The CCTA examinations are performed on a 320 detector CT scanner (Aquilion One, Toshiba, Medical Systems), with the scanner settings: Gantry rotation time 350 ms, detector collimation 0.5×320, X-ray tube voltage 100–120 kV, and X-ray tube current 280–500 mA. The acquired CT images have a matrix size $512 \times 512 \times 560$ and a voxel size $0.5 \times 0.5 \times 0.25$ mm.

One hundred and five CT images with high contrast are randomly selected from the database (see Fig. 1a for an example of CT image). The raw CT-volumes are manually cropped, using Osirix, to only contain the tracer-enhanced regions with the LAA. After cropping, CT-volumes are blurred with a Gaussian filter kernel with standard deviation at 0.5 mm and the iso-surfaces of the inner part of the LAA is computed using the Marching Cubes algorithm [11] with a manually set iso-level in the range 150–250 Hounsfield Units. The selected iso-surface level varies, due to variations of the amount of tracer in the LAA. Image blurring and surface reconstruction are conducted using 3D Slicer [1]. A reconstructed LAA surface from the example CT image is shown in Fig. 1b.

(a) CT-image   (b) LAA surface
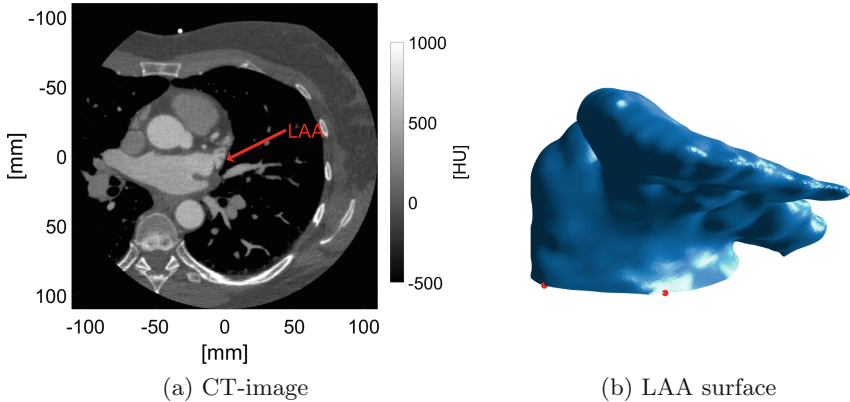
Fig. 1. (a) Slice from raw cardiac computed tomography (CT) image. (b) Left atrial appendage (LAA) surface reconstructed from the CT image shown in (a), where the red marks are manually placed landmarks. HU: Hounsfield Units (Color figure online)

3   Methods

The first goal of this work is to build a statistical shape model [3] to quantitatively describe the shape variation of the LAA. This model is created from a training set containing N Procrustes-aligned shapes; shapes in the training set are represented as a series of corresponding points.

3.1   Point Correspondence

Point correspondences between LAA surfaces are determined by registering a source surface $\mathcal{S}$ to each target surface $\mathcal{T}$ in the training set, such that each vertex is positioned on the same anatomical structures in both $\mathcal{S}$ and $\mathcal{T}$. Initially, $\mathcal{S}$ is aligned to $\mathcal{T}$ with a similarity transform by registration of four manually placed landmarks equally distributed in the LAA orifice (two out of the four landmarks are visible as the red marks in Fig. 1b). Furthermore, the registration is fine-tuned by an iterative close point (ICP) alignment [16]. The aligned source is now denoted $\mathcal{S}_{ICP}$. The surface registration of $\mathcal{S}_{ICP}$ and $\mathcal{T}$ is performed using a non-rigid volumetric registration algorithm. To be able to use the volumetric registration algorithm, $\mathcal{S}_{ICP}$ and $\mathcal{T}$ must be represented as volumes. We represent $\mathcal{S}_{ICP}$ and $\mathcal{T}$, as signed distance fields (SDF), where each voxel value in the SDF is equal to the signed Euclidean distance to the surface [13,15].

   The non-rigid volumetric registration is conducted by solving the optimization given by:

$$\hat{\mathbf{T}}_\mu = \underset{\mathbf{T}_\mu}{\arg\min}\left(\mathcal{C}\left(\mathbf{T}_\mu; I_F, I_M\right)\right) \tag{1}$$

Here $\mathcal{C}$ is a cost-function, I_F is the fixed volume and I_M is the moving volume, where I_F and I_M are the SDF representation of $\mathcal{S}_{ICP}$ and $\mathcal{T}$ respectively. $\mathbf{T}_\mu$ is the non-rigid volumetric transformation that transform I_M to I_F.

The transformation is parameterised by a parameter-vector $\boldsymbol{\mu}$. In this work, we use a multi-level B-Spline transformation with five resolution levels. The cost function we are going to minimize is described by:

$$\mathcal{C} = \omega_1 MSD\left(\mu; I_F, I_M\right) + \omega_2 \mathcal{P}_{CP}(\mathbf{x}, \mathbf{y}) + \omega_3 \mathcal{P}_{BE}\left(\mu\right), \tag{2}$$

where MSD is the mean squared voxel value difference similarity measure, $\mathcal{P}_{CP}(\mathbf{x}, \mathbf{y})$ is penalizing large distances between landmarks and $\mathcal{P}_{BE}\left(\mu\right)$ is the bending energy penalty term. The weights: $\omega_1 = 1$, $\omega_2 = 0.15$ and $\omega_3 = 2$ are optimized using a grid search. The optimal transformation parameters are found using adaptive stochastic gradient descent [9] as optimizer, with 2048 random samples per iteration for a maximum of 500 iterations as implemented in the elastix library [10]. The estimated transformation determined between I_F and I_M is applied to $\mathcal{S}_{ICP}$ and the transformed surface is $\mathcal{S}_T$.

Since the volumetric registration is conducted on the SDF, it is not guaranteed that the zero level iso-surfaces fits perfect after the registration. This problem is solved using an approach originally described in [14], where vertices in $\mathcal{S}_T$ are propagated to $\mathcal{T}$ using Markov Random Field regularization of the correspondence vector field. After the vertices in $\mathcal{S}_T$ are propagated to $\mathcal{T}$ we have obtained a point correspondence surface $\mathcal{S}_{COR}$, where each vertex corresponds to a vertex in $\mathcal{T}$. The set of surfaces with point correspondence is used to construct a point distribution model using Procrustes alignment and principal component analysis (PCA) as described in [3].

3.2   Shape Clustering

To examine the natural shape clusters formed by our data set, we use the trained point distribution model to represent the surfaces by their PCA loadings and use the loadings to identify shape clusters. The PCA loadings $\mathbf{b}$ of a given surface is determined by:

$$\mathbf{b} = \mathbf{P}(\mathbf{x}' - \bar{\mathbf{x}}), \tag{3}$$

where $\mathbf{x}'$ is the input surface, $\bar{\mathbf{x}}$ is the Procrustes average shape of the N aligned $\mathcal{S}_{COR}$ and $\mathbf{P}$ is the set of the t first eigenvectors. We use the PCA loadings to estimate the natural number of shape clusters, by examining the log-likelihood (LLH) computed from multivariate Gaussian mixture models (GMM) fitted to the loadings. The probability density function of a GMM can be written as:

$$p(\mathbf{x}) = \sum_{i=k}^{K} \pi_k \mathcal{N}\left(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\right), \tag{4}$$

where $\mathbf{x}$ is the loadings, π_k is the mixing coefficient, K is the number of mixture components and $\mathcal{N}\left(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\right)$ is the multivariate Gaussian distribution with mean $\boldsymbol{\mu}_k$ and covariance matrix $\boldsymbol{\Sigma}_k$. From Eq. (4) the LLH function is given by [2]:

$$p(\mathbf{x}|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \sum_{i=1}^{N} \ln\left(\sum_{k=1}^{K} \pi_k \mathcal{N}(\mathbf{x}_i|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)\right) \tag{5}$$

In order to avoid over-fitting, the number of shape clusters is determined by using two level cross-validation. The first level performs leave-one-out cross-validation. Here the data are divided into $N-1$ training shapes and one test shape. The training set is used to train a GMM with K mixture components, while the test shape is used to validate the trained GMM, using the LLH as quality metric. This procedure is repeated until all N shapes have been used as the test shape, after which the mean test LLH is computed based on the N test LHH. The second cross-validation level iterates through $K = 1\ldots10$ mixture components, where the first level is conducted for every K. The number of shape clusters is equal to the number of mixture components, which results in the highest mean test LLH.

In order to identify shape appearance of the natural formed clusters, a new GMM is trained on the entire data set, where the number of mixture components is equal to the number of estimated shape clusters. We can now use the model to randomly sample PCA loadings within the different shape clusters and generate synthetic shapes base on the loadings by:

$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{Pb} \tag{6}$$

The synthetic shapes can be visualized to identify the different shape appearance of each cluster.

The GMMs are fitted to the training data by estimating a set of model parameters: π, $\boldsymbol{\mu}$, and $\boldsymbol{\Sigma}$, that maximize the LLH function. In this work, we estimate the parameters by the Expectation Maximization algorithm, with 100 random initialization and use the set of model parameters with highest training LLH.

4   Results

The point correspondence framework is applied to our 105 reconstructed LAA surfaces. We use the template surface shown in Fig. 2a as source. The template is the average shape of N Procrustes aligned $\mathcal{S}_{COR}$. The set of $\mathcal{S}_{COR}$ is computed as an initial registration of $\mathcal{S}$ and $\mathcal{T}$, where $\mathcal{S}$ is selected randomly from the pool of LAA surfaces.

We are able to determine point correspondences of the majority of the target surfaces (103 out of 105), with a median root mean square distance (RMS) between $\mathcal{T}$ and $\mathcal{S}_{COR}$ at 0.6 mm and a 75th percentile at 0.9 mm. The surfaces with RMS equal to the median and 75th percentile are shown in Fig. 2b and c, respectively. The figure shows $\mathcal{T}$, where the color scale indicate the distance between $\mathcal{T}$ and $\mathcal{S}_{COR}$. It is seen that $\mathcal{S}_{COR}$ matches $\mathcal{T}$ in most of the surface. It is also seen that the point correspondence framework are not able to find point correspondences in the most distal lobes of the LAA. A visual analysis of all $\mathcal{S}_{COR}$ shows that two of the surfaces have poor point correspondence and are therefore excluded from the training set, leaving 103 surfaces for the rest of the analysis.

(a) Template   (b) Median   (c) 75th percentile
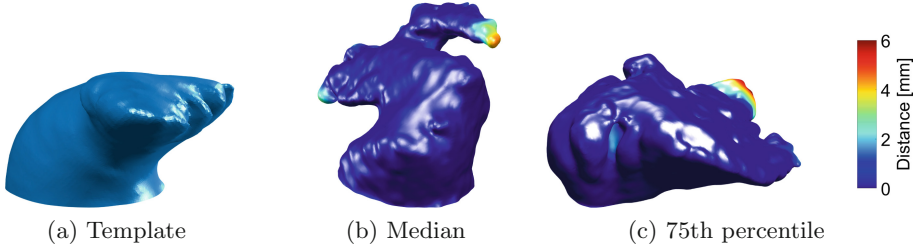
Fig. 2. (a): Template surface determined as the average shape of the N Procrustes-aligned point correspondence surfaces ($\mathcal{S}_{COR}$). (b) and (c) target surface, where the color scale indicate the distance to the $\mathcal{S}_{COR}$. (Color figure online)

The point distribution model is trained on the 103 Procrustes-aligned $\mathcal{S}_{COR}$ and we choose to represent the shapes using their first five PCA loadings. The first five PCA loadings are used, since the remaining 98 PCA loadings each describes only a small fraction (less than 5 %) of the total shape variation in the studied data. Ten GMMs, with $K = 1 \ldots 10$ mixture components, are trained on the PCA loadings and the test LLH is computed from each GMM using cross-validation. The mean test LLH and mean train LLH are shown in Fig. 3 for each validated GMM. It can be seen that, according to the LLH test, a GMM with two mixture components gets the best validation performance. This means that the studied dataset of LAA most likely form two different shape clusters.
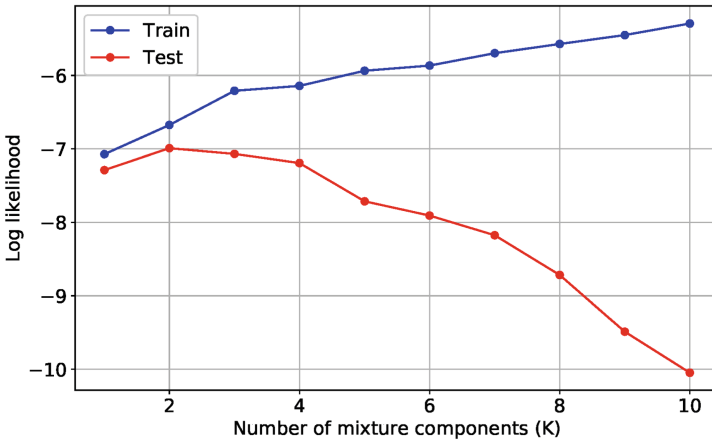

Fig. 3. Training and testing log likelihood (LLH) after training Gaussian mixture models (GMMs) on 103 PCA loadings using two level leave-one-out cross-validation.

In order to identify the shape appearance of the clusters, we train a new GMM, with two mixture components, on the entire data set. We generate four synthetic shapes by sampling PCA loadings from mixture component one and two of the new GMM, which can be visualized in Figs. 4 and 5. It can be observed in Fig. 4 that surfaces sampled from cluster one have similar LAA morphology, with an obvious bend in the primary lobe, a particular characteristic of Chicken Wing morphologies. It can also be appreciated the variability within the cluster in terms of LAA orifice characteristics and volumes. On the other hand, surfaces samples from cluster two, illustrated in Fig. 5, do not present a bending of the primary lobe, but a wider one with several secondary lobes. These particular characteristics are typical of non-Chicken Wing LAA morphologies such as Cauliflower ones.
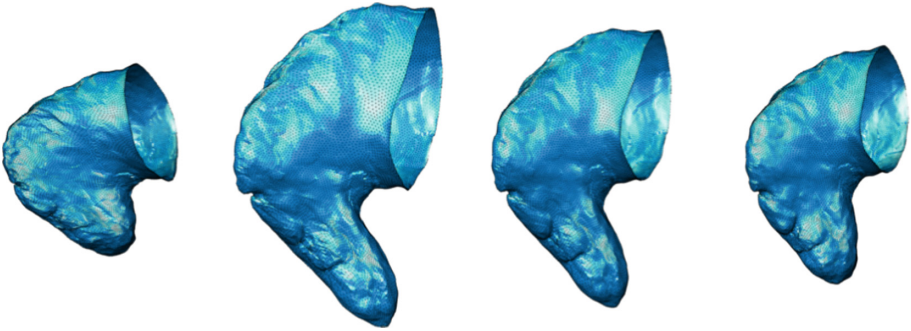
Fig. 4. Four synthetic shapes generated by sampling PCA loadings from mixture component one of a Gaussian mixture model (GMM) with two components.
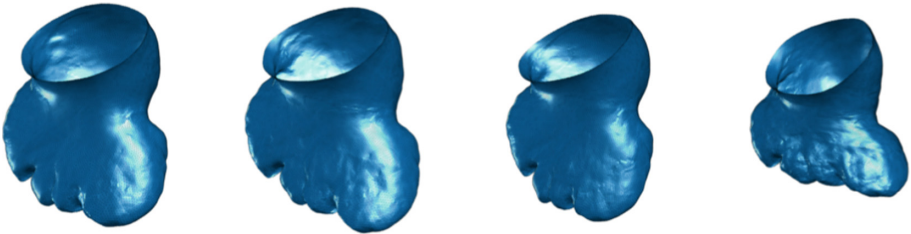
Fig. 5. Four synthetic shapes generated by sampling PCA loadings from mixture component two of a Gaussian mixture model (GMM) with two components.

5   Conclusion

In this work we have presented a full framework for the extraction and quantification of shape clusters of left atrial appendages and demonstrated that the two primary shape clusters broadly correspond to the main LAA morphological categories in standard clinical classification, Chicken Wing and non-Chicken Wing LAA shapes. The framework enables future statistical inference on the relation between LAA shape characteristics and stroke risk.

References

1. 3D Slicer. https://www.slicer.org/
2. Bishop, C.M.: Pattern Recognition and Machine Learning, 1st edn. Springer, New York (2006)
3. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models-their training and application. Comput. Vis. Image Underst. **61**(1), 38–59 (1995)
4. Di Biase, L., et al.: Does the left atrial appendage morphology correlate with the risk of stroke in patients with atrial fibrillation?: results from a multicenter study. J. Am. Coll. Cardiol. **60**(6), 531–538 (2012)
5. Ernst, G., et al.: Morphology of the left atrial appendage. Anat. Rec. **242**(4), 553–561 (1995)
6. Fuchs, A., et al.: Normal values of left ventricular mass and cardiac chamber volumes assessed by 320-detector computed tomography angiography in the copenhagen general population study. Eur. Heart J.-Cardiovasc. Imaging **17**(9), 1009–1017 (2016)
7. Go, A.S., et al.: Prevalence of diagnosed atrial fibrillation in adults: national implications for rhythm management and stroke prevention: the anticoagulation and risk factors in atrial fibrillation (ATRIA) study. JAMA **285**(18), 2370–2375 (2001)
8. Khurram, I.M., et al.: Relationship between left atrial appendage morphology and stroke in patients with atrial fibrillation. Heart Rhythm **10**(12), 1843–1849 (2013)
9. Klein, S., Pluim, J.P., Staring, M., Viergever, M.A.: Adaptive stochastic gradient descent optimisation for image registration. Int. J. Comput. Vis. **81**(3), 227 (2009)
10. Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P.: Elastix: a toolbox for intensity-based medical image registration. IEEE Trans. Med. Imaging **29**(1), 196–205 (2010)
11. Lorensen, W.E., Cline, H.E.: Marching cubes: a high resolution 3D surface construction algorithm. ACM SIGGRAPH Comput. Graph. **21**, 163–169 (1987)
12. Nordestgaard, B.G., et al.: The effect of elevated body mass index on ischemic heart disease risk: causal estimates from a mendelian randomisation approach. PLoS Med. **9**(5), e1001212 (2012)
13. Paulsen, R.R., Baerentzen, J.A., Larsen, R.: Markov random field surface reconstruction. IEEE Trans. Vis. Comput. Graph. **16**(4), 636–646 (2010)
14. Paulsen, R.R., Hilger, K.B.: Shape modelling using markov random field restoration of point correspondences. In: Taylor, C., Noble, J.A. (eds.) IPMI 2003. LNCS, vol. 2732, pp. 1–12. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45087-0_1
15. Paulsen, R.R., Marstal, K.K., Laugesen, S., Harder, S.: Creating ultra dense point correspondence over the entire human head. In: Sharma, P., Bianchi, F.M. (eds.) SCIA 2017. LNCS, vol. 10270, pp. 438–447. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59129-2_37
16. Rusinkiewicz, S., Levoy, M.: Efficient variants of the ICP algorithm. In: 2001 Proceedings of the Third International Conference on 3-D Digital Imaging and Modeling, pp. 145–152. IEEE (2001)
17. Wunderlich, N.C., Beigel, R., Swaans, M.J., Ho, S.Y., Siegel, R.J.: Percutaneous interventions for left atrial appendage exclusion: options, assessment, and imaging using 2D and 3D echocardiography. JACC Cardiovasc. Imaging **8**(4), 472–488 (2015)