

Shape-based diagnosis of the aortic valve

Razvan Ioan Ionasec^{a,c}, Alexey Tsymbal^b, Dime Vitanovski^a, Bogdan Georgescu^a, S. Kevin Zhou^a, Nassir Navab^c and Dorin Comaniciu^a

^a Integrated Data Systems, Siemens Corporate Research, Princeton, USA

^b Siemens Corporate Technology, SE SCR 2, Erlangen, Germany

^c Computer Aided Medical Procedures, Technical University Munich, Germany

ABSTRACT

Disorders of the aortic valve represent a common cardiovascular disease and an important public-health problem worldwide. Pathological valves are currently determined from 2D images through elaborate qualitative evaluations and complex measurements, potentially inaccurate and tedious to acquire. This paper presents a novel diagnostic method, which identifies diseased valves based on 3D geometrical models constructed from volumetric data. A parametric model, which includes relevant anatomic landmarks as well as the aortic root and leaflets, represents the morphology of the aortic valve. Recently developed robust segmentation methods are applied to estimate the patient specific model parameters from end-diastolic cardiac CT volumes. A discriminative distance function, learned from equivalence constraints in the product space of shape coordinates, determines the corresponding pathology class based on the shape information encoded by the model. Experiments on a heterogeneous set of 63 patients affected by various diseases demonstrated the performance of our method with 94% correctly classified valves.

Keywords: physiological aortic valve modeling, computer-aided diagnosis, machine learning, discriminative distance function, case retrieval

1. DESCRIPTION OF PURPOSE

Valvular heart diseases (VHD) represent an important public-health problem notably common in elderly population.¹ Disorders of the aortic valve, which regulates the blood flow pumped into the body, present the highest prevalence among VHD and require elaborate clinical management. Furthermore, information on the morphology of the aortic valve has been mainly obtained from explanted valves or animal models.² There is still a lack of consensus concerning the optimal physiological description,³ while reported standard dimensions of the valve vary considerably. The diagnosis and evaluation of the aortic valve, crucial in the clinical workflow, is currently performed on 2D imaging using potentially inaccurate measurements.⁴ Inconclusive or discordant clinical findings enforce invasive investigation methods, significantly increasing procedures' risk and cost.

There are two main types of heart valve malfunction: stenosis and regurgitation. Valvular stenosis is defined as an obstruction of the blood flow caused by narrowing, stiffening, thickening, fusion or blockage of one or more valves. Valvular regurgitation is a condition in which blood leaks back in the wrong direction because the affected valve is not closing properly. Figure 1 demonstrates the appearance and morphological variations for valves affected with the same disease, for both stenotic and regurgitant valves. Both disorders can greatly interfere with the heart's function to pump an adequate amount of blood, causing serious health problems.

The use of computed tomography (CT) in assessing the aortic valve has recently received considerable attention.⁵⁻⁷ This modality captures the entire 3D structure of the aortic valve, offering accurate assessment preconditions. However, current practice requires the navigation of the data into sets of 2D planes for visualization and quantization. The potentially inaccurate results can lead to suboptimal clinical decisions or even erroneous diagnosis.

In this paper we introduce a novel diagnostic method, which has the potential to significantly improve the management of aortic valve disease. The inference is performed in a comprehensive feature space, which incorporates the complex morphologic information of the valve by applying robust statistical methods. The

Further author information: (Send correspondence to Razvan Ioan Ionasec: razvan.ionasec@siemens.com)

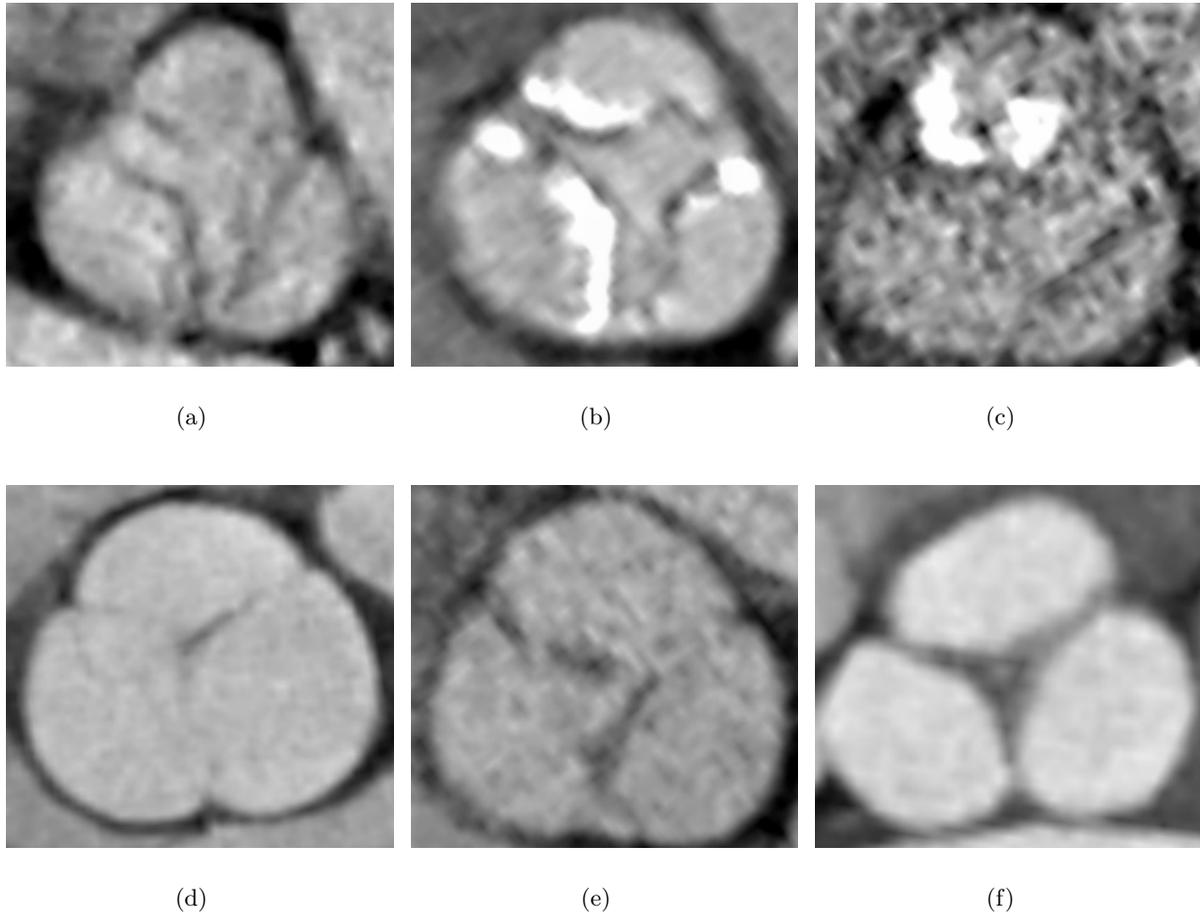


Figure 1. Examples of diseased valves from 4D CT in short axis view; (a),(b),(c) stenotic aortic valves and (d),(e),(f) regurgitant aortic valves.

aortic valve is represented through a 3D geometrical model as proposed by Ionasec *et al.*⁸ For each subject patient, the parameters of the model are estimated from computed tomography data applying recent machine learning techniques.^{8,9} In order to infer the pathology from the geometrical models we use the technique for learning discriminative distance functions,¹⁰ which was shown to perform well in domains with many irrelevant, weakly relevant and highly correlated features. Compared to the gold-standard,⁴ the proposed technique is fully reproducible as it always outputs the same results given the same input. Another important characteristic of the proposed approach is the integration of the entire 3D information in contrast to the state-of-the-art, which employs 2D data only. Experiments on a set of 63 patients demonstrated the accuracy of the system as well as significant anatomical regions for pathological valve.

2. METHODS

2.1. Modeling of the aortic valve

The morphology of the aortic valve is expressed by a 4D model estimated from computed tomography, as proposed by Ionasec *et al.*⁸ This representation includes key anatomical landmarks such as hinges, commissures, leaflet tips and coronary ostia as well as the aortic root and leaflets (see Figure 2). Non uniform rational B-splines (NURBS) together with physiological-driven topological and geometrical constraints are used to obtain a model capable to express a large spectrum of morphological and pathological variations.

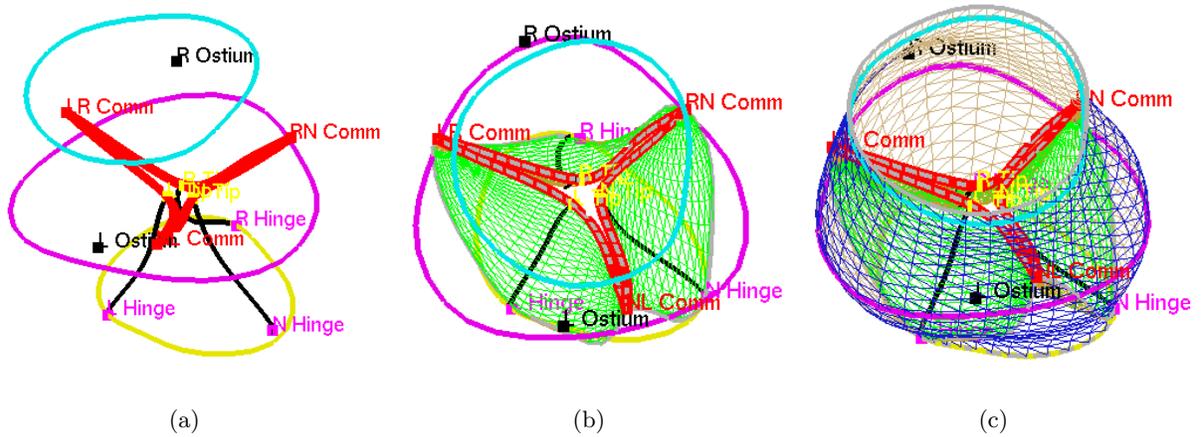


Figure 2. The proposed aortic valve model.

Recent advances in discriminative learning are applied to estimate the above presented model for each subject patient from end-diastolic 3D CT data.^{8,9} The parameter estimation is performed in an incremental fashion and includes three stages: rigid transformation estimation, landmark detection and boundary fitting. In the first stage the rigid parameters $\theta = (x, y, z, \alpha, \beta, \gamma, s_x, s_y, s_z)$ of the valve's location are estimated from the underlying volumetric data. A discriminative classifier H trained using the Probabilistic Boosting Tree (PBT)¹¹ in combination with Haar-like and steerable features⁹ incrementally scans the search space in order to determine the parameter $\hat{\theta}$ with maximum probability

$$\hat{\theta} = \underset{\theta \in \Theta}{\operatorname{argmax}} (p(\theta|I) = H(\theta|I)) \quad (1)$$

where I is the current volume, H is the trained rigid transformation detector, and $x, y, z, \alpha, \beta, \gamma, s_x, s_y, s_z$ are the translation, orientation and scale parameters.

In the second stage, the landmarks locations $(L_{comm}^R, L_{comm}^L, L_{comm}^N, L_{hinge}^R, L_{hinge}^L, L_{hinge}^N, L_{ostium}^R, L_{ostium}^L)$ relative to the estimate rigid $\hat{\theta}$ are detected using the same method. A thin-plate-spline (TPS) transformation¹² projects a prior computed mean model to the corresponding image location under consideration of the landmarks correspondence. A boundary detector, trained with the PBT and Steerable features,⁹ is used to evolve the boundary of the initial estimated model to its true location, as presented in Figure 3.

2.2. Learning discriminative distance functions

The task we consider is to discriminate between healthy and diseased valves based on their geometrical models as explained in the previous section. Each case is represented with N 3D landmarks, which results in C parameters, of the respective model and is to be classified as healthy (-1) or diseased ($+1$):

$$\hat{y} = \underset{y \in \{-1, +1\}}{\operatorname{argmax}} (p(y|C)), \quad C = ((P_1^x, P_1^y, P_1^z), \dots, (P_N^x, P_N^y, P_N^z)) \quad (2)$$

Currently, in distance based classification, canonical distance functions are used most commonly. Today in many areas including computer vision the Euclidean distance is one of the most widely used, though it is well known that its use is justified only when the data distribution is Gaussian. During the last three decades, the importance of learning a proper distance function in machine learning has been gradually acknowledged.¹⁰ Different imaging applications have been considered, including image retrieval, object detection,¹³ motion estimation and

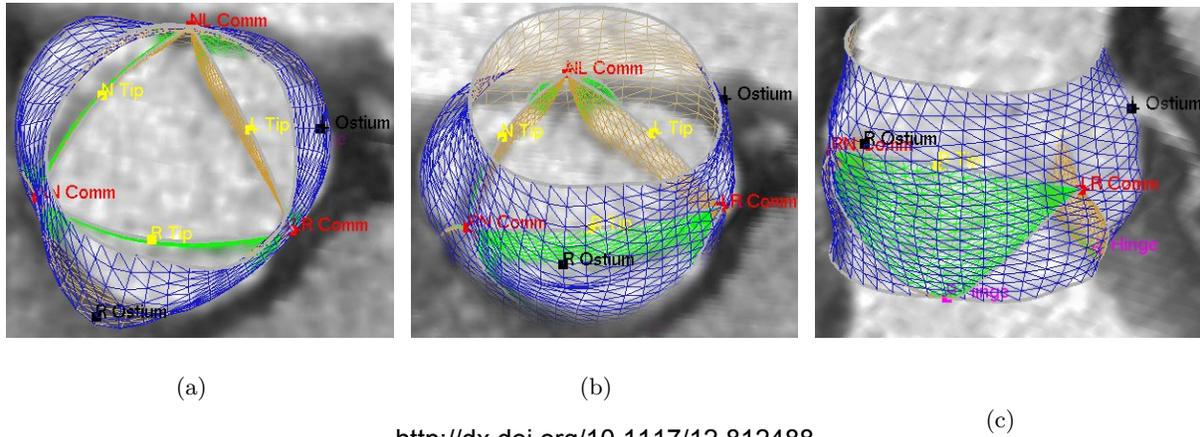


Figure 3. Estimation results of the aortic valve model from CT data.

image registration.¹⁴ However, to the best of our knowledge, we are the first to apply this to learning distance functions for meshes in the domain of cardiology.

There are several reasons that motivate the studies in the area of learning distance functions and their use in practice.¹⁰ First, it is easy to show that choosing an optimal distance function simplifies classifier learning. Next, learning distance functions breaks the learning process into two sequential stages (distance function learning followed by clustering or classification), where each stage requires search in a less complex functional space than in the immediate learning. Besides, it fosters the creation of more modular and thus more flexible systems, supporting component reuse. For example, the same distance function learning module can be used with various classification techniques and vice versa. Another important benefit is the opportunity for inductive transfer between similar tasks. While knowledge sharing at the level of creating a classifier is often non-trivial, one solution is to learn the same or similar distance function for each task. This approach is often used in computer vision applications.¹³

While historically the research on distance function learning has started from supervised learning of distance functions for k -nearest neighbor learning in the original "feature vector-object label" representation, today by far the most commonly used representation, especially in computer vision, is the one based on so called equivalence constraints.¹⁵ Interestingly, the distance learnt from labels is almost always metric, while learning from weak representations such as equivalence constraints usually provides more flexibility for learning arbitrary functions.

Equivalence constraints are represented using triplets (C^1, C^2, y) , where C^1 and C^2 are coordinates in the original space (see equation 2) and $y \in \{+1, -1\}$ is a label indicating whether the two points are similar (from the same class or cluster) or dissimilar.¹⁰ Learning from these triples is often called learning in the product space (i.e. with pairs of points as input); see^{10, 14, 15} for examples. While learning in the product space is perhaps a more popular form of learning from equivalence constraints, yet another common alternative is to learn in the difference space, the space of vector differences; see¹⁶ for examples. The difference space is normally used with homogeneous high-dimensional data, such as pixel intensities or their PCA coefficients in imaging. While both representations demonstrate promising empirical results in different contexts, there is no understanding which representation is better and when.

Thus, learning from equivalence constraints usually results in a binary distance function, predicting only whether the two objects are similar or dissimilar. Though often, in order to avoid this limitation, the signed margin of margin-based classifiers such as SVM and boosting is used as the required distance function.

There are two essential reasons that motivate the use of equivalence constraints in learning distance functions; their availability in some learning contexts and the fact that they are a natural input for optimal distance function learning.¹⁰ It can be shown that the optimal distance function for classification is of the form $p(y_i \neq y_j | x_i, x_j)$.

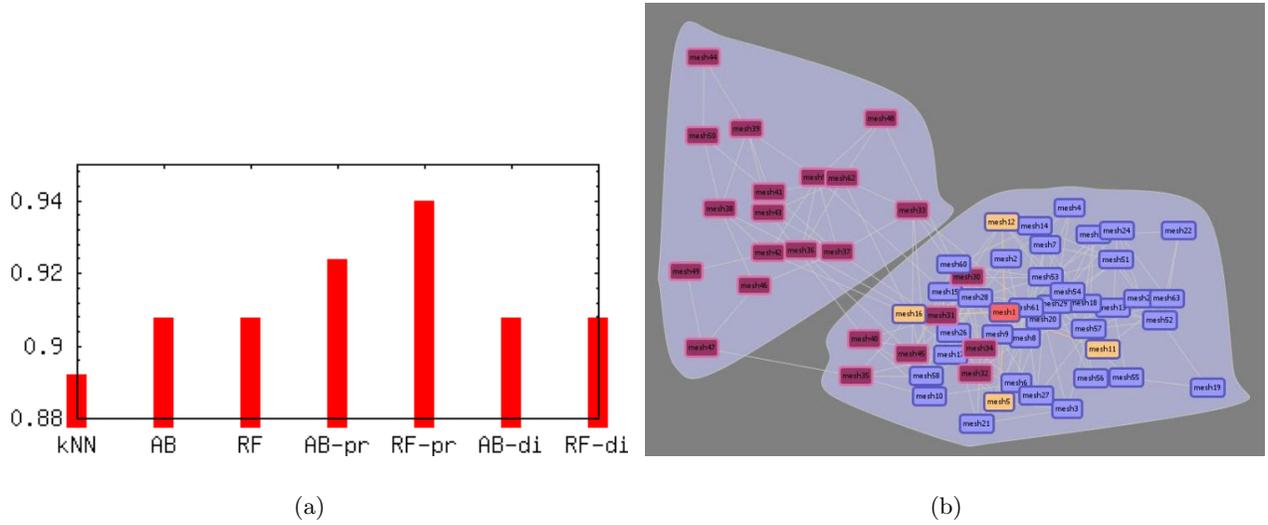


Figure 4. (a) The disease detection accuracy for different learning techniques. (b) Relative neighborhood graph, healthy (blue) and diseased (red), current and most similar patients (orange and yellow, respectively).

Under the i.i.d. assumption the optimal distance measure can be expressed in terms of generative models $p(x|y)$ for each class as follows:¹³

$$p(y_i \neq y_j | x_i, x_j) = \sum_y p(y|x_i)(1 - p(y|x_j)) \quad (3)$$

Because of this characteristics of equivalence constraints, supervised learning algorithms, which receive canonical labelled points as input, often turn them into equivalence constraints and operate directly on them.¹⁰

Clearly, the most popular learning algorithm in the area of distance function learning from weak representations is boosting.¹⁷ In our experiments we also use Random Forests which demonstrated promising results in our preliminary experiments with public benchmark datasets and which are normally faster than boosting.

We implemented both product and difference space in our framework. The signed margin of models constructed using AdaBoost and Random Forests is used as the required distance function.

3. RESULTS

The performance of the proposed method was evaluated on cardiac 4D CT data from 63 patients, with significant variation in resolution (0.28mm to 2.0mm) and capture ranges (80 to 350 slices with sizes from 512x 512 to 153x153 pixels). The evaluation set includes 40 healthy, 5 regurgitant, 3 bicuspid and 15 stenotic aortic valves, summarizing a total of 127 volumes associated with manual annotations.

The model fitting performance is evaluated using a point to mesh error measure, three-fold cross validation and the annotated models as the ground-truth. Table 1 demonstrates the average accuracy of the proposed method of 1.73mm at a computation time of 2.46 seconds per volume.

The results of the leave-one-out cross validation are summarized in Figure 4(a). The figure demonstrates the classification accuracy of k -Nearest Neighbors (kNN), AdaBoost (AB) and Random Forests (RF), as well as AdaBoost and Random Forests¹⁸ in the product and difference spaces ($AB-pr$, $RF-pr$, $AB-di$ and $RF-di$). Each cross-validation run was performed on 200 pre-selected features determined using the Gain-Ratio feature filter. Learning from equivalence constraints improves the accuracy in comparison with learning from the canonical space. For this task the product space (with RF) leads to the best accuracy, predicting the correct class

	Mean (mm)	Std dev	Median (mm)	90% (mm)
Rigid	2.97	3.36	2.01	4.01
Landmark	2.27	3.39	1.38	2.80
Boundary	1.73	4.03	0.92	1.76

Table 1. Error of each stage in the model estimation framework.

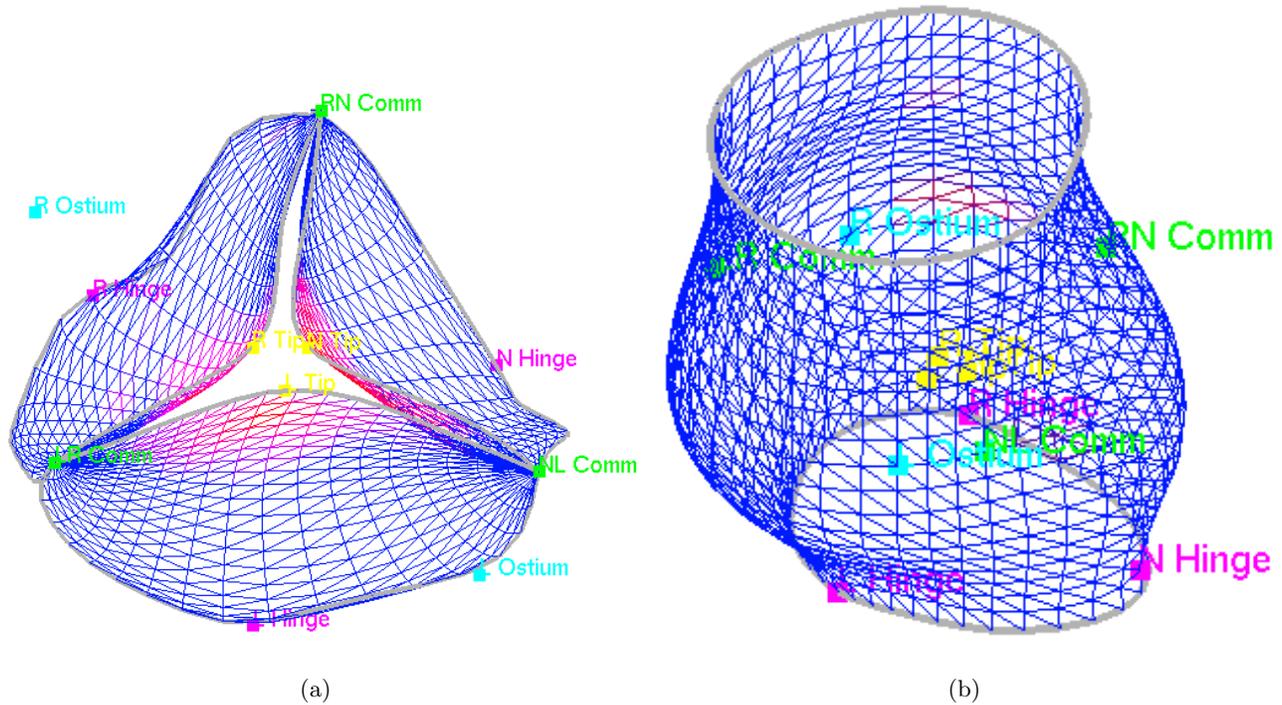


Figure 5. (a) The disease detection accuracy for different learning techniques. (b) and (c) Discriminate anatomical regions for aortic valve disease; blue - less discriminate, red - more discriminate.

(pathological or healthy) in 93.7% cases. Learning distance functions from equivalence constraints helps to combine the power of boosting and random forests with the transparency of case retrieval.

In order to better understand the correlation between morphology and pathology we investigated the computed Gain-Ratio values for each coordinate of the geometrical model. The most relevant regions for diseased valves are the free parts of the leaflets (Figure 5(a)), around the tips, as well as the right coronary ostium region (Figure 5(b)). This demonstrates, in concordance with the current pathological knowledge, the significantly higher sensibility of the aortic leaflets.

The learnt inter-valve similarity can be efficiently visualized using neighbourhood graphs for better decision support and knowledge discovery.¹⁹ Presumably the most useful kind of neighbourhood graphs is the so-called relative neighbourhood graph (RNG). In RNG, two nodes A and B (corresponding to valves in our case) are connected with an edge if there is no node C such that the distances from both C to A and C to B are smaller than the distance from A to B. RNGs nicely visualize the underlying similarity and patient clusterings; they are always connected and are planar or close to planar. Figure 4(b) demonstrates that clustering the RNG for the aortic valve similarity learnt with RF in the product space nearly perfectly divides the set of valves into healthy (blue) and diseased (red). The neighbourhood graphs are easy to read; they allow to visualize additional features, textual, nominal, real-valued or even image thumbnails, at the nodes, remind the currently popular

in the semantic web community social networks, and form a nice alternative to the more orthodox and better known to clinicians (in particular, geneticists) heatmaps and recently introduced treemaps. With the help of neighbourhood graphs difficult cases and clear outliers in terms of the classification context studied can easily be spotted, that may help with the tasks of better understanding of the problem domain, validation and refinement of the learnt distance function, and ultimately, decision support.

4. CONCLUSION

In this paper we proposed a novel method for shape-based diagnosis of the aortic valve. Recent machine learning techniques are applied to estimate the parameters of an individual 3D morphological model from cardiac CT data. A robust algorithm, which learns a discriminative distance function in the shape's product space, is applied to classify the pathology of the subject valve. Our method integrates the entire three-dimensional geometrical expression of subject patient in an automatic classification framework offering a reliable and reproducible diagnosis solution, which has the potential to advance the clinical management of aortic valve diseases.

ACKNOWLEDGMENTS

The authors would like to thank Medical University of South Carolina for generously providing the images. Razvan Ioan Ionasec would like to thank "The Dinu Patriciu Foundation" for supporting his doctoral program.

REFERENCES

1. Nkomo, V., Gardin, J., Skelton, T., Gottdiener, J., Scott, C., and Enriquez-Sarano, M., "Burden of valvular heart diseases: a population-based study," *Lancet* **368**(9540), 1005–1011 (2006).
2. Dagum, P., Green, G., Nistal, F., Daughters, G., Timek, T., Foppiano, L., Bolger, A., Ingels, N., and Miller, D., "Deformational dynamics of the aortic root: modes and physiologic determinants," *Circulation* **100**(19 Suppl), II54–II62 (1999).
3. Anderson, R., "The surgical anatomy of the aortic root," *Multimedia Manual of Cardiothoracic Surgery (MMCTS)* (2006). doi:10.1510/mmcts.2006.002527.
4. Vahanian, A., Baumgartner, H., Bax, J., Butchart, E., Dion, R., Filippatos, G., Flachskampf, F., Hall, R., Iung, B., Kasprzak, J., Nataf, P., Tornos, P., Torracca, L., and Wenink, A., "Guidelines on the management of valvular heart disease: The task force on the management of valvular heart disease of the European Society of Cardiology," *European heart journal* **28**(2), 230–268 (2007).
5. Ionasec, R. I., Georgescu, B., Comaniciu, D., Vogt, S., Schoepf, U., and Gassner, E., "Patient specific 4d aortic root models derived from volumetric image data sets," in [*Radiological Society of North American (RSNA)*], (November 2008).
6. Gassner, E., Ionasec, R. I., Georgescu, B., Vogt, S., Schoepf, U., and Comaniciu, D., "Performance of a dynamic aortic valve model for quantification of the opening area at cardiac mdct . comparison to manual planimetry," in [*Radiological Society of North American (RSNA)*], (November 2008).
7. Pouleur, A., le Polain de Waroux, J., Pasquet, A., Vanoverschelde, J., and Gerber, B., "Aortic valve area assessment: multidetector ct compared with cine mr imaging and transthoracic and transesophageal echocardiography," *Radiology* **244**(3), 745–754 (2007).
8. Ionasec, R. I., Georgescu, B., Gassner, E., Vogt, S., Kutter, O., Scheuring, M., Navab, N., and Comaniciu, D., "Dynamic model-driven quantification and visual evaluation of the aortic valve from 4d ct," in [*International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*], (September 2008).
9. Zheng, Y., Barbu, A., Georgescu, B., Scheuring, M., and Comaniciu, D., "Fast automatic heart chamber segmentation from 3d ct data using marginal space learning and steerable features," *Proc. International Conference on Computer Vision* (2007).
10. Bar-Hillel, A. and Weinshall, D., "Learning distance functions by coding similarity," in [*24th Int. Conf. on Machine Learning*], **227**, 1031–1038 (2007).
11. Tu, Z., "Probabilistic boosting-tree: Learning discriminative methods for classification, recognition, and clustering," in [*ICCV 2005*], **2**, 1589–1596 (2005).

12. Duchon, J., "Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces," *RAIRO Analyse Numerique* **10**, 5–12 (1976).
13. Mahamud, S. and Hebert, M., "The optimal distance measure for object detection," in [*Proc. of Int. Conf. on Computer Vision and Pattern Recognition*], **1**, 248–255 (2003).
14. Zhou, S., Shao, J., Georgescu, B., and Comaniciu, D., "Boostmotion: Boosting a discriminative similarity function for motion estimation," in [*Proc. CVPR*], (2006).
15. Hertz, T., *Learning distance functions: algorithms and applications*, PhD thesis, The Hebrew University of Jerusalem (2006).
16. Yu, J., Amores, J., Sebe, N., and Tian, Q., "Toward robust distance metric analysis for similarity estimation," in [*Int. Conf. on Computer Vision and Pattern Recognition*], **1**(316-322) (2006).
17. Freund, Y. and Schapire, R. E., "A decision-theoretic generalization of on-line learning and an application to boosting," *Journal of Computer and System Sciences* **55**(1), 119–139 (1997).
18. Breiman, L., "Random forests," in [*Machine Learning*], 5–32 (2001).
19. Tsymbal, A., Huber, M., Zillner, S., Hauer, T., and Zhou, K., "Visualizing patient similarity in clinical decision support," in [*Alexander Hinneburg*], 304–311 (2007).