Eurographics Italian Chapter Conference (2006)
G. Gallo and S. Battiato and F. Stanco (Editors)

Shape Comparison and Deformation Analysis in Biomedical Applications

S. Colantonio, D. Moroni, O. Salvetti

Istituto di Scienza e Tecnologie dell'Informazione, ISTI-CNR, Pisa, Italy

Abstract

In this paper, we present a method for comparing shapes and analyzing deformations of 3D image objects. This method is based on the definition of an object model based on multi-source 3D images.
Elective application cases of the proposed method consist of the analysis of deformable anatomical structures, useful to support medical diagnosis in routine clinical practice. In particular, preliminary aspects are discussed of a study case regarding heart dynamics.

Categories and Subject Descriptors (according to ACM CCS): I.4.6 [Image Processing and Computer Vision]: Pixel Classification; I.4.8 [Image Processing and Computer Vision]: Shape; I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations

1. Introduction

We introduce a general framework suitable to compare the shape and analyze the modification of periodically deforming 3D image objects.

The main idea is to define a *statistical dynamic reference model* of object classes, eventually imaged by multimodal data, which encodes morphological and functional properties and can be used to analyze different scenarios.

Multimodal data consist of images and signals acquired by different modalities that supply a faithful representation of the object and its non-rigid motion. A *dense* reconstruction is performed to identify the 3D shape of the object, which is then refined, exploiting the periodical feature of its motion, eventually through a registration procedure. The reconstructed volume is characterized by means of a set of properties suitable to assess the object *deformation pattern*, and through a *normalization* phase the object can be reduced to a reference model in order to compare different *states*.

The presented method is the current result of an ongoing research, devoted to analyze and assess the state of periodically deforming objects in the biomedical field. Actually, deformable structures are widespread in human anatomy (e.g., the heart, the lungs and the foot during gait), and their deformation modes offer useful clues in medical research, diagnosis and treatment. In particular, the study of the cardiac dynamics, by analyzing the left ventricle deformation, has been performed as a preliminary test and is here reported for showing the effectiveness of the method [CMO05, MCO05].

2. Problem definition

A deformable image object $\mathcal{O}$ is defined as an ordered list

$$\mathcal{O} = (\mathcal{O}_t) \quad t = 1, 2, \ldots \tag{1}$$

where $\mathcal{O}_t$ is a snapshot of the object at time t, and

$$\mathcal{O}_t = (V_t, P_t) \tag{2}$$

where $V_t = \{V_t^{\alpha}\}$ corresponds to a collection of smooth manifolds embedded in the background space $\Omega \in \mathbb{R}^d$, and P_t is a property function.

In this way, we assume a hypothesis quite common in computational anatomy [GM93], which is, to some extent, satisfied in practice, and implies that some differential geometric features, like curvatures, can be computed everywhere. Moreover, using a collection of

manifolds instead of a single manifold assures the possibility to describe multiple objects of different dimensionalities (curves, surfaces, volumes), and, at the same time, to treat manifolds with boundaries. To understand the usefulness of this assumption we may think to the problem of the left ventricle modeling: the structure of interest is the *myocardium*, which is contained between the *epicardium* and the *endocardium*, in such a case it is convenient to discriminate epicardial and endocardial surfaces and to characterize them using features different from those used to characterize the internal myocardial points.

Another assumption is that the transformations $V_t \rightarrow V_{t'}$ that rule the object deformation, are diffeomorphisms acting on an admissible shape space. This means that, fixed any reference time t_0, we can smoothly deform $\mathcal{O}_{t_0}$ to fit any other $\mathcal{O}_t$. In this way, we exclude tears or cracks of the object. Moreover, we avoid dealing with changes in topology, which would require to model shape *transition*, a task beyond our present scopes.

A further assumption is that the object deformation recurs at a regular interval of time T, that is

$$\exists\, T : \forall\, t: \mathcal{O}_t = \mathcal{O}_{t+T} \tag{3}$$

or, in other words, that the object motion is periodic. This requirement proves to be essential in our framework, since our analysis is based on a frequency domain approach. This assumption allows to consider a limited number of objects $\mathcal{O} = (\mathcal{O}_t)$, $t = 1, 2, \ldots, T$. In the following, we will refer to a *deformation cycle* as a complete deformation cycle of the object, starting from one reference instant and taking the structure back to its initial conditions.

To image the object of interest, multimodal data can be considered. They should consist of synchronous signals and images, acquired by multiple modalities, and should give a faithful representation of the object itself. In particular, data should include at least a 3D image sequence $I(t)$ from which morphological and regional functional properties of $\mathcal{O}$ can be inferred. The representation may be enriched, when needed, by signal vectors conveying some other global information or constraints.

With the previous assumptions, a *statistical dynamic model* of an object of interest is constructed in "pseudo-ideal" or particularly interesting cases, and used as reference system for assessing the *deformation pattern* of real cases under examination.

Four steps are performed to obtain the aforementioned model:

- *Object reconstruction*: for each instant t, the V_t object collection of manifolds is identified and reconstructed in the 3D space by applying neural algorithms to the image sequence $I(t)$;
- *Object characterization*: morphological features and dynamic descriptors are extracted and coded in a property function P_t that for each element $\mathbf{x}$ of the manifolds V_t^{α} returns the property vector $(P_1, \ldots, P_m)$, where each P_i represents one of the selected features;
- *Object registration*: to eliminate acquisition inaccuracies and obtain the most likely reconstruction of the object, data acquired over multiple cycles are averaged through an inter-cycle registration. In this way, the mean shape and the mean property function are obtained;
- *Deformation pattern assessment*: suitable and significant shape descriptors are extracted and spatial distribution of the property functions are evaluated in order to obtain a description of the object dynamics. This description is then used in the normalization procedure: the quantities obtained in the reference cases are used to construct the dynamic model, and the values obtained in real cases are compared to them by means of an *energy-based* method [Mum98, MTY03].

In the following sections, these steps are described in more details.

2.1 Object Reconstruction

The 3D reconstruction of the object $\mathcal{O}$ is achieved via a *voxel* classification, that is by labeling each voxel in the image domain with semantic classes which describe voxel membership to the collection of manifolds, V_t. The classification is performed applying an advanced neural architecture to a set of extracted features.

The involved features can be divided into two classes. First, *low-level* features are considered: they are context-independent and do not require any knowledge and/or pre-processing. Some examples are voxel position, gray level value, gradients and other differentials, texture, and so forth. *Middle-level* features are also selected, since voxel classification can benefit from more accurate clues, specific of the problem at hand. In particular, if an intrinsic reference system can be individuated to describe the object shape, it can be used to define a relative voxel position. If, in addition, a priori information about the object shape is available, a reliable clue for detecting edges in the images is given by the gradient along the normal direction to the expected edge orientation.

Moreover, a *multiscale* approach is adopted: the features are computed on blurred images, supplying information about the behavior of the voxel neighborhood, which results in a more reliable classification.

The set of selected features are processed to accomplish the voxel classification by means of a *Multilevel Artificial Neural Network* (MANN), which assures different computational advantages [dBNPS03]. For each voxel $\mathbf{x}$, its computed features vector is divided into vectors $\Im_k(\mathbf{x})$, each one containing features of the same typology and/or correlated. Then each $\Im_k(\mathbf{x})$ is processed by a dedicated classifier based on an

unsupervised *Self Organizing Maps* (SOM) architecture. The set of parallel SOM modules constitutes the first level of the MANN which aims at clustering each portion of the feature vector into crisp classes, thus reducing the computational complexity. The output of this first level is then passed to a second and final level, consisting in a single *Error Back-Propagation* (EBP) module, which supplies voxel classification. Its output describes voxel membership to the various manifolds V_t^{α} in the collection V_t.

2.2 Object Characterization

The reconstructed object is further characterized by assigning significant properties to each voxel. We allow voxel belonging to two different manifolds V_t^{α} , V_t^{β} to be described by different sets of properties, but we require that points belonging to manifolds labeled with α, even if at different times, are described by a uniform sets of properties. More formally, for every index α, the disjoint union $\cup_t V_t^{\alpha}$ is equipped with a function

$$P^{\alpha}: \cup_t V_t^{\alpha} \rightarrow \mathbb{R}^{d(\alpha)}$$

having values in a suitable properties space.

This has two obvious advantages. Firstly, by allowing different properties spaces, the characterization can effectively deal with complex *structures*: indeed, for every 'substructure' we can select relevant properties that may be not meaningful or, worse, ill-defined for the whole structure. This stems, thus, for adaptability. In addition, time-uniformity of the description of every manifold labeled with 'α' guaranties ability to perform a time evolution analysis of the properties.

Three types of properties are considered:

- intensity based properties;
- local shape descriptors;
- local dynamic behaviour descriptors.

Examples of properties of the first type are gray level value, gradients, textures and so on. They are extracted form the image sequence $I(t)$ – the one which leads us to object reconstruction. If data collected from other imaging modalities are available, after performing registration, we can fuse this information to further annotate the object (for example, in the case of the heart, information regarding perfusion and metabolism, obtained e.g. by means of PET imaging, can be referred to the reconstructed myocardium).

Geometric based properties, belonging to the second type, are extracted directly from the collection of manifolds, V_t, and are essential to describe locally the shape of the object. Again, we may distinguish between context independent features (automatically computable for every manifold of a given dimensionality) and problem-specific properties.

Finally, the local dynamic behavior may be described by properties borrowed from continuous mechanics; they, however, require motion estimation that we haven't pursued yet.

2.3 Object registration

Periodically deforming objects offer the opportunity to average data acquired on multiple deformation cycles. In other words, data available for each instant time of the deformation cycle can be fused in order to obtain an average shape and properties function of the characterized object. Of course, averaging should be performed between physically corresponding points, hence a registration step is essential.

More precisely, object averaging is realized considering a basic strategy on a given shape population. Elements of this population are compared up to a subgroup of irrelevant alignment, for instance by taking rigid transformation as irrelevant group of alignment. This means that a manifold V_t^{κ} and a rotated version $V_t^{\kappa'}$ of it are considered equivalent. So, formally, shapes are understood as equivalent classes of manifolds.

To average a given population, it is of course necessary to be able to compare different shapes belonging to it. After having performed registration by irrelevant transformations, residual discrepancies have to be evaluated. There are several methods for evaluating residual discrepancies, but, loosely speaking, we can say that any such method consists in a way of defining the 'work' to be done to deform one shape into another.

Here, however, we assume that the involved deformations are small: it is then possible to treat the space of involved deformation as a linear space and to perform usual algebraic manipulations. In particular, the centroid of the shapes population is defined and may be considered as the shape that minimizes residual discrepancies.

More in detail, a sketch (2D case) of the approach is given in Figure 1. Given a shape population, we randomly choose a characterized object. Then the following steps are performed:

1. We use rigid transformation to align the various shapes in the population to the reference one. The rigid transformation involved is considered as 'irrelevant' (i.e. characterized objects are considered equivalent up to rigid transformation).
2. Then objects are non-rigidly registered using free form deformation (a 'universal' class of non-rigid matching), thus dissolving differences in shape. We displayed also the displacement fields converting the reference shape in any member of the population.
3. Their (warped) properties functions are then averaged. This gives the averaged object 'seen' in the coordinates of the reference one.

4. The mean deformation is computed (using, as already stated, a linear approximation, valid in case of small deformations): this is the transformation we have to apply to the shape of the reference object to shift it to the shapes population centroid.

5. The shapes population centroid is computed and the average properties function is referred to this domain.

It can be shown that the choice of the reference objects introduce a bias in the final result. Indeed the algorithm fails to find the actual shapes population centroid, due to concatenation error of (standard pairwise) registration algorithms: registering '*a*' to '*b*' and '*b*' to '*c*' and then composing the transformations leads to a transformation that generally differs from the direct registration from '*a*' to '*c*'.

It is possible to overcome these difficulties in two ways. First, using a refined registration algorithm that 'consistently' registers the population (extremely low concatenation error); this is not theoretically difficult but results in an algorithm which is computationally intensive.

Otherwise, we can use an iterative approach based on an iterative version of the previous algorithm. Simply, we let the output of the algorithm at the n-th step to take the role of the reference characterized structure for the next step. The procedure ends when there is not appreciable difference between the results of the *n* and *n*+1 steps.

The final result is $\bar{\mathcal{C}} = \{\bar{V}_t, \bar{P}_t\}$, t=1, …, T ($\bar{\mathcal{C}}$ $\bar{V}$, and $\bar{P}$ average on the period).

Experimental results show that the final result is not biased by the initial reference structure choice.

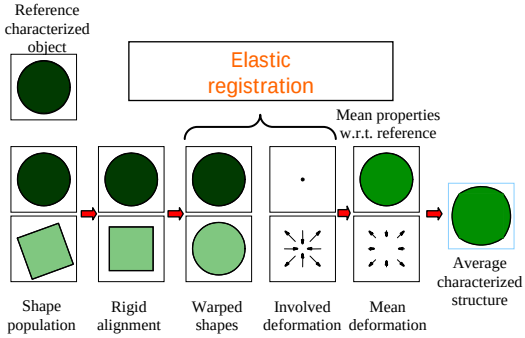

Figure 1. *A sketch of the approach adopted for the registration procedure.*

2.4 Pattern deformation assessment

The averaged shapes obtained in the previous step can be used to assess the dynamic behavior of the object and identify its deformation pattern.

However, the voxelwise characterization of the reconstructed objects is not suited for state assessment. Indeed, the given description of the whole objects (collection of manifolds described by functions) has a dimensionality far too high to make the problem computationally feasible. Moreover, it would be essential to be able to compare anatomical structures belonging to different patients and, at the moment, the idea is to use a deformable model (given for example by mass-spring models [dBLS03]) and to normalize every instance of anatomical structure to that model: in this way anatomical structures (belonging to the same family) are uniformly described and can be then compared.

Combining these two issues, we should look for a new set of 'more intrinsic' features $\mathfrak{I}_t$ that should be enough simple and, at the same time, capturing essential information about the objects.

To obtain these new kinds of features, global information about the objects can be extracted from the properties function, without introducing any model. For example, one may consider the 'property spectrum' by which we mean the probability density functions (PDF) of a given property $P_s(\bullet, t)$. This consists in a function defined on a 'shape independent' domain (the real line), capturing how the property is globally distributed; thus, comparison of different property spectra is directly feasible; however, to reduce dimensionality, it is effective and convenient to compute the momenta of the PDF (mean, variance, …).

However, properties spectrum does not convey any information at all about regional distribution of the property. In practical situation, this is a drawback which cannot be ignored: for example, a small 'highly abnormal' region may not affect appreciably the PDF, but its clinical relevance is, usually, not negligible. Hence, spatial distribution of properties has to be analyzed; in some cases, approaches which do not need a refined model of the object (e.g., Gaussian image, spherical harmonics or Gabor spherical wavelets) may be suitable. However, in general one should define a model of the objects (whose primitives -elementary bricks- are regions, patches or landmarks) and then propagate it to the set of instances to be analyzed by using matching techniques. Then, we may consider the average of a property on regions or patches (or the value in a landmark) as a good feature, since comparisons between averages on homologous regions can be immediately performed.

Following this recipe, a vector of features with the desired properties is obtained for each phase of the cycle. The deforming object is then described by the dynamics of the temporal sequence of feature vectors obtained at different phases of the deformation cycle.

At this point, we use again our assumptions to perform a fruitful feature transformation. The smoothness assumption implies that the objects have mainly low frequency excited deformation modes. We extend this slightly assuming that this holds true also for the features lists $(\mathfrak{I}_t)_{t=1,2,\ldots,T}$. We assume that the fundamental frequency of the motion is also the main component of each feature tracked on time.

With this assumption, an obvious choice is given by the Fourier transform, followed by a low pass filter, which supplies a new features vector Θ.

The evaluation of the above mentioned parameters $\Im_t$ at each phase t implicitly codifies information regarding object dynamics. Actually, we avoid defining a complex model of the object kinematics and exploit its periodic characteristic by constructing a rich representation of each phase of the deformation cycle.

Once obtained the values of $\Im_t$ and the average shapes in the reference cases, that is, once built the reference model, the same four steps procedure can be performed in real cases, obtaining a set of actual values $\Im_t$*. The normalizing operation consists in computing the *distance* of these values from the model. In particular, the shape comparison can be performed using a method based on the deformation energies, and, more precisely, on the *geodesic energies*, which define, via functional norms, a metric for infinitesimal deformations [KSMJ04, HM05].

3. Results

An elective case study for the presented methodology is cardiac analysis, whose clinical relevance can be hardly overestimated. We restrict our analysis to the left ventricle (LV) that, pumping oxygenated blood around the body, is the part of the heart for which contraction abnormalities are more clinically significant.

The proposed methodology is, of course, not universal, in the sense that there are some intrinsic limitations that prevent it to be potentially applied in any scenario. Indeed, our analysis is limited to a single deformation cycle and so only pathologies that affect every deformation cycle can be considered. Moreover, we require that physiological and (selected) pathological states induce different feature dynamics. This requirement is not too restrictive; actually, it is well known that many pathologies are correlated to abnormal shape patterns at end systole.

The LV structure is modelled as a 3D manifold (the myocardium) with boundary. The boundary has two connected components which are the surfaces corresponding to epicardium and endocardium.

The structure is represented by a multimedia object consisting in a 3D MRI sequence together with a synchronous ECG.

We describe henceforth how the steps of the methodology are applied.

To perform reconstruction we used a pre-processing phase devoted to the automatic localization of the left ventricle cavity (LVC).

First, *fuzzy c-means* is used to cluster the image into homogeneous regions (see Figure 2).

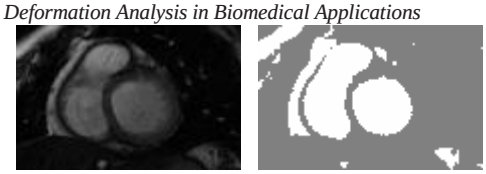

Figure 2. *An example of an input slice of the 3D MRI scan, and the result of the 2-clusters fuzzy c-means algorithm.*

Connected components are then labelled and tracked over time as shown in Figure 3. This allows extracting a set of features describing the behaviour in time of geometric properties like area and eccentricity (more precisely we used the first entries in the periodogram of the geometric property) [MCO05]. On the base of learned examples, we modelled the PDF of features vector as a multidimensional normal distribution. For any new case, then, we used maximum likelihood estimator to select the LVC among candidate regions.

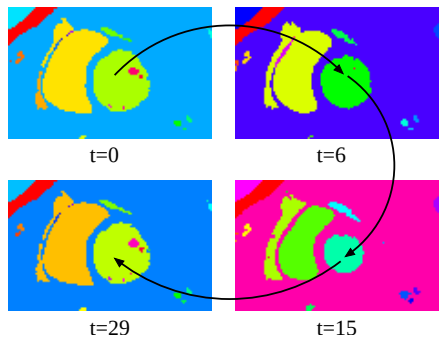

Figure 3. *An example of the tracking procedure applied to the left ventricle cavity.*

The located LVC is used to define a *Hybrid spherical/cylindrical reference system* (HRS) that is suitable to model the LV, which approximately resembles a *bullet-shaped* structure. Indeed, an isoradial surface {r = cost} is bullet-shaped in HRS (Figure 4).

More in detail, to define the HRS, we consider as the z axis of the reference system the long axis l of the LV, computed as the principal inertia axis of the 3D LVC. Then, we select a *switching point* O, along l, such that points lying above O are given cylindrical coordinates, while points lying beneath are described by spherical coordinates. More precisely, once fixed O, the plane H through O and orthogonal to l divides the 3D space in two half-spaces H^+ above H and H^- beneath H. Then, points $P \in H^+$ are described by cylindrical coordinates (r,φ,z), whereas spherical coordinates (r,φ,θ) w.r.t. O are associated to points $P \in H^-$.

The selection of point O is based on the request that the field of unit vectors pointing in the direction of increasing radial coordinate should be, as far as possible, orthogonal to the computed LVC borders. O is computed as follows (see Figure 5):

$O = \text{argmin} \sum_j \alpha_j$

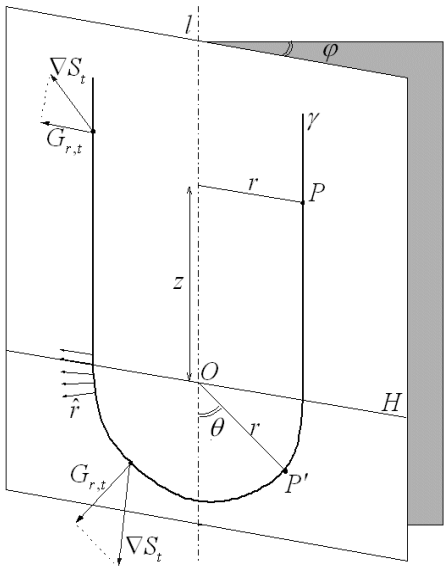

Figure 4. *The hybrid reference system. O: switching point from cylindrical to spherical coordinates. l: LVC axis.* $P=(r,\phi,z)$*: point in the cylindrical system.* $P'=(r,\phi,\theta)$*: point in the spherical system.* $\hat{r}$ *: unit vector field in the radial direction.* $G_{r,t} = \nabla S_t \cdot \hat{r}$ *: gradient radial component. The surface* $\gamma=\{r=cost\}$ *represents expected heart surfaces orientation.*

The HRS is used to extract the following features for voxel classification:

- Position w.r.t. HRS
- Intensity
- Mean intensity (computed applying Gaussian filters)
- Gradient norm
- Radial gradient

In particular radial gradient (that is d(I)/dr, the component of gradient in the direction of increasing radial coordinate) is motivated by the expected orientation of epicardial and endocardial surfaces.

Using the 2-level ANN, voxels are classified on the basis of their features vector as belonging or not to epi- and endocardial surfaces.

More in detail, the set of extracted features is divided into two vectors $\Im_1$, $\Im_2$ containing respectively *position*, *intensity* and *mean intensity*, and *position*, *gradient norm* and *radial gradient*. The position w.r.t. HRS is replicated in both vectors because it reveals salient for clustering both features sub-sets. Then, the first level of the MANN consists of two SOM modules, which have been defined as 2D lattice of neurons and dimensioned experimentally, controlling the asymptotic behaviour of the number of excited neurons versus the non-excited ones, when increasing the number of total neurons [dBPO04]. A 8×8 lattice SOM was then trained, according to Kohonen's training algorithm [Koh97], for clustering the features vector$\Im_1$, while $\Im_2$ was processed by a l0×10 lattice SOM.

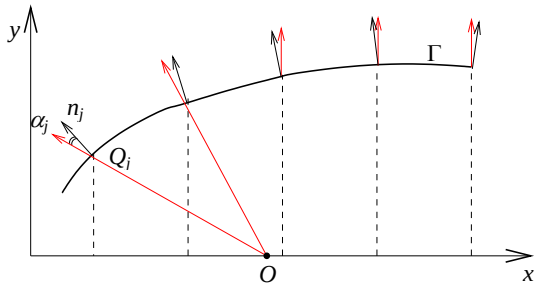

Figure 5. *Selection of the switching point O. The x axis corresponds to the long axis l of the LVCs; along the y axis, the mean radius of the LVCs is reported;* Γ *is the curve obtained interpolating the points* Q_j *identified by the intersection between the LVC borders and the plane orthogonal to l; n is the normal vector to* Γ *in* Q_j*;* α_j *is the angle between the normal vector* n_j *and the vector* OQ_j*.*

A single EBP module has been trained to combine the results of the first level and supply the final response of the MANN. Different architectures have been tested, finding the best performance for a network with only one hidden layer of 15 units, trained according to the *Resilient Back-Propagation* algorithm [RB93].

Voxel classification, obtained by means of the MANN, gives a representation of them as a point cloud so that a visualization of the voxelwise reconstruction can be obtained by triangulation. Characterization of the reconstructed structure is obtained annotating every voxel with intensity, Gaussian and mean curvature, *wall thickness* and HRS properties. In particular, Gaussian and mean curvature have been included as shape descriptors whereas wall thickness, which is a classical cardiac parameter, is one example of problem-specific property: it is defined as the thickness of the myocardium along a coordinate ray and it is expected to increase during contraction (myocardium, being almost water, is, with good approximation, incompressible).

This characterization is translated in a more amenable form by computing properties spectrum and regional features. In computing spectrum, coordinates w.r.t. HRS have been disregarded, with the exception of radial coordinate; intensity has also been excluded. For any property only mean and variance have been considered.

For computing regional features, so far, we used a popular model of the LV (see [FNV01] for a review of 3D-cardiac modelling). In 2D, as shown in Figure 6, it is defined by the intersections of cardiac surfaces with a pencil of equally spaced rays.

The 3D version is obtained by stacking the 2D construction along the axis of the LV.

4. Discussion and conclusions

A framework for the shape comparison and deformation analysis has been introduced for the study of periodically deforming objects. Using reference cases, a statistical dynamic model is constructed by coding the dynamics of the object in a rich representation of object shape and functional properties. The deformation pattern and dynamic behaviour of real objects are then obtained by normalizing them to the reference model.

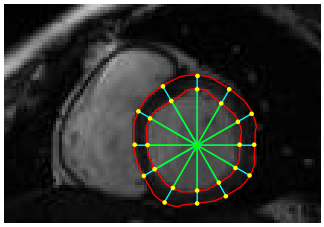

Figure 6. *The pencil of equally spaced rays used to compute the characterizing features.*

The framework consists in several modules performing a) object reconstruction, b) object characterization, c) object registration and d) pattern deformation assessment. Solutions to specific tasks proposed in each module are, to a large extent, independent and may be combined with other methods, thus broadening the potential application field of the framework.

The elective case studies are represented by the analysis of deformable anatomical structures, such as the heart. Actually, for demonstrating the effectiveness of the proposed framework, we have shown the preliminary results of the study of the cardiac dynamics, and in particular of the left ventricle. The next step will be to employ the obtained results for defining a general method to assess the state of the deformable object, and, in particular, the physio-pathological states of the left ventricle.

Acknowledgments

This work has been carried on within the ISTI-CNR project “SIMCAR”.

Authors would like to thank the Institute of Clinical Physiology – CNR, Pisa, for providing study case.

This work has been partially supported by European Project Network of Excellence MUSCLE – FP6-507752 (Multimedia Understanding through Semantics, Computation and Learning), and CNR – Finmeccanica Agreement.

References

[CMO05] COLANTONIO S., MORONI D., SALVETTI O.: A methodological approach to the study of periodically deforming anatomical structures. In *Proc. International Conference on Advanced Information and Telemedicine Technologies*, AITTH-2005 (Nov. 2005), vol. 1, pp. 32-36.

[dBLS03] DI BONA S., LUTZEMBERGER L., SALVETTI O.: A simulation model for analyzing brain structures deformations. *Physics in Medicine and Biology* (2003), Vol. 48 n. 24, pp. 4001-4022.

[dBNPS03] DI BONA S., NIEMANN H., PIERI G., SALVETTI O.: Brain volumes characterisation using hierarchical neural networks. *Artificial Intelligence in Medicine* (2003), vol. 28, pp. 307-322.

[dBPO04] DI BONO M.G., PIERI G., SALVETTI O.:A tool for system monitoring based on artificial neural networks. *WSEAS Transaction on System* (2004), vol. 3 n. 2, pp. 746-751.

[FNV01] FRANGI A.F., NIESSES W.J., VIERGEVER M.A.: Three-dimensional modelling for functional analysis of cardiac images: a review. *IEEE Transaction on Medical Imaging* (2001), vol. 20 n. 1, pp. 2-25.

[GM93] GRENANDER U., MILLER M.I.: Computational anatomy: an emerging discipline. Q. Appl. Math. LVI (1998), vol. 4, pp. 617-694.

[HM05] HOLM D.D., MARSDEN J.E.: Momentum maps and measure-valued solutions (Peakons, filaments and sheets) for the EPDiff Equation. In *The Breast of Symplectic and Poisson Geometry* (2005), A Festshrift for Alan Weinstein, 203-235, Progr. Math., 232, J.E. Marsden and T.S. Ratiu, Editors.

[Koh97] KOHONEN T.: Self-Organizing Maps. *Springer Series in Information Sciences* (1997), second ed. vol. 30, Berlin Springer.

[KSMJ04] KLASSEN E., SRIVASTAVA A., WASHNGTON M., SHANTANU H.J.: Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Trans. on PAMI* (2004), vol. 26, nr. 3, pp. 372-383.

[MCO05] MORONI D., COLANTONIO S., SALVETTI O.: MRI Left Ventricle Segmentation and Reconstruction for the Study of the Heart Dynamics. In *Proc. 5th IEEE International Symposium on Signal Processing and Information Technology, ISSPIT-05* (Dec. 2005).

[MTY03] MILLER M.I., TROUVÉ A., YOUNES L.: The metric spaces, Euler equations, and normal geodesic image motions of computational anatomy. In *Proc. IEEE International Conference on Image Processing* (2003), vol. 2, pp. 635-638.

[Mum98] Mumford D.: Pattern theory and vision. *Questions Mathematiques En Traitement du Signal et de l’Image* (1998), Intitut Henri Poincaré, Paris, chap. 3, pp. 7-13.

[RB93] RIEDMILLER M., BRAUN H.: A direct adaptive method for faster backpropagation learning: The RPROP algorithm. In *H. Ruspini, editor, Proceeding of the IEEE International Conference on Neural Networks* (ICNN) (1993); pp. 586-591