



HAL
open science

Local Appearance Knowledge and Diffusion Wavelets Shape Variation Models for Muscle Segmentation

Salma Essafi, Georg Langs, Jean-Francois Deux, Guillaume Bassez, Alain
Rahmouni, Nikolaos Paragios

► **To cite this version:**

Salma Essafi, Georg Langs, Jean-Francois Deux, Guillaume Bassez, Alain Rahmouni, et al.. Local Appearance Knowledge and Diffusion Wavelets Shape Variation Models for Muscle Segmentation. [Research Report] 2009. inria-00362636v1

HAL Id: inria-00362636

<https://inria.hal.science/inria-00362636v1>

Submitted on 18 Feb 2009 (v1), last revised 6 Mar 2009 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Local Appearance Knowledge and Diffusion
Wavelets Shape Variation Models for Muscle
Segmentation*

Salma Essafi — Georg Langs — Jean-François Deux — Guillaume Bassez — Alain
Rahmouni — Nikos Paragios

N° 6821

January 2009

Thème BIO

*R*apport
de recherche



Local Appearance Knowledge and Diffusion Wavelets Shape Variation Models for Muscle Segmentation

Salma Essafi ^{*}, Georg Langs [†], Jean-François Deux [‡], Guillaume Bassez [‡],
Alain Rahmouni [‡], Nikos Paragios ^{*}

Thème BIO — Systèmes biologiques
Projet Galen

Rapport de recherche n° 6821 — January 2009 — 22 pages

Abstract: In this report, we present a novel prior knowledge representation of shape variation using diffusion wavelets and applied for medical image segmentation. One of the major advantage of our approach is that it can reflect arbitrary and continuous interdependencies in the training data. In contrast to state-of-the-art methods our framework during the learning stage optimizes the coefficients as well as the number and the position of landmarks using geometric (reconstructed surface) constraints. Saliency is encoded in the model and segmentation is expressed through the extraction of the corresponding features in a new data-set. The resulting paradigm supports hierarchies both in the model and the search space, can encode complex geometric and photometric dependencies of the structure of interest, and can deal with arbitrary topologies. In another hand, our report deals with a different model search methodology where we apply an approach related to active feature models; the location of landmarks is updated iteratively, using local features, and the canonical correlation analysis. We report promising results on two challenging medical data sets, that illustrate the potential of our method.

Key-words: MRI, Diffusion wavelet, Segmentation, Human skeletal muscle

^{*} Salma Essafi and Nikos Paragios are affiliated to Laboratoire MAS, Ecole Centrale Paris, Châtenay-Malabry, France and to Equipe GALEN, INRIA Saclay - Île-de-France, Orsay, France

[†] Georg Langs is affiliated to CIR lab, Department of Radiology, Medical University of Vienna, Vienna, Austria

[‡] Jean-François Deux, Alain Rahmouni and Guillaume Bassez are affiliated to Centre Hospitalier Universitaire Henri Mondor, Créteil, France

Une approche basée sur des ondelettes de diffusion pour la segmentation du muscle squelettique

Résumé : Dans ce rapport, nous présentons une approche basée sur des ondelettes de diffusion pour la segmentation du muscle squelettique. Contrairement aux méthodes de l'état de l'art, notre approche au cours de la phase d'apprentissage réussit à optimiser les coefficients des ondelettes, ainsi que leur nombre et leur position. Le modèle prend en charge aussi bien les hiérarchies dans l'espace de recherche, et peut aussi encoder les dépendances géométriques complexes et photométriques de la structure d'intérêt. En effet notre modélisation a l'avantage de traiter avec des topologies arbitraires. Les résultats expérimentaux effectués sur les données de mollet de 25 patients (dont 5 atteints de myopathies) montrent le potentiel de notre méthode pour segmenter le muscle en régions anatomiquement significatives.

Mots-clés : IRM, Ondelettes de diffusion, Segmentation, Muscle squelettique

Contents

1	Introduction	4
2	Shape Representation	7
2.1	Properties of Diffusion Wavelets	7
2.2	A Diffusion Operator Reflecting The Topology	8
2.3	Diffusion Wavelet Modeling of Shape Variations	8
2.4	Modeling Using the Orthomax Criterion	10
3	Data Search	12
3.1	Sparse Search	12
3.2	Appearance Feature Model	13
3.2.1	Appearance Representation Using Local Features	13
3.2.2	Search by canonical correlation analysis	13
4	Results	14
4.1	Experimental Set-Up	14
4.2	Experimental Validation	14
5	Conclusion	18

1 Introduction

Myopathies are diseases that affect the muscle system, and lead to a severe deterioration of the motoric abilities. These pathologies affect 4 to 6% of the population, i.e 25 to 30 million Europeans[15]. Diagnosis, as well as follow up for a given therapeutic strategy are often performed through biopsy. Magnetic resonance imaging (MRI) allows a non-invasive observation of the muscle fibers, their texture, and their global structure. It has the potential to replace biopsies (by e.g. diffusion tensor imaging (DTI)) and has the advantage of encompassing a section of the entire leg as opposed to small local samples. This enables the analysis of local properties, as well as understanding the global structural change of muscles affected by a disease. A crucial first step in this analysis is the accurate segmentation of individual muscles.

Shape models state of the art Model based segmentation and localization approaches are of prime interest, due to their ability to repeatably identify positions in the anatomy, and to learn and apply properties of large and representative populations. Of particular importance is the accurate representation of subtle local shape variation and the correct parameterization of the associated interdependencies between parts of the anatomical structures. This cannot be performed manually, but has to be learnt from the data, in order to avoid a bias that would compromise the model efficiency, and result accuracy. A second aspect that becomes prevalent when exploring large anatomical regions, is the highly heterogeneous nature of their appearance, but existing local consistency across patients.

State of the art segmentation methods mainly rely on a clearly defined topology, and an object boundary characterized by salient features (e.g. edges). An example for accurate segmentation achieved through deformable models combined with shape modeling as shown in [9]. Statistical Shape Models (SSM), proposed by Cootes [8], are used to capture and represent the shape variation in a training sets. Active Shape Model (ASM) [6], build a point distribution model, which allows for global scale analysis of shape variation by applying principal component analysis (PCA) to the positions of the boundary points, while using local appearance to fit the model to new data. These methods depend heavily on the accuracy of the inter-subject registration for group comparison and the parameterization of the shape.

The majority of work on anatomy segmentation is focused on brain, liver or heart data [13]. Muscle-compounds (see Fig. 1) pose a rather different and new challenge to segmentation algorithms because there is no prominent difference of tissue-properties between neighboring muscles. Border tissues in between muscles are only visible on specific locations, distributed in a very sparse and heterogeneous manner. Indeed muscle surfaces are only partially visible, while parts exhibit a structure, that can change dramatically between patients, or during the course of follow-up examinations. Related work on muscle segmentation includes simplex meshes [16], or atlas-based models which are deformed using a free-form deformation [15]. The main challenge of calf T1-MRI segmentation is the total absence of conventional image support. Edges between classes are almost invisible while at the same

time the global intensity characteristics of the muscle regions are almost identical. Consequently prior models based on conventional image-based search approaches will fail to separate the muscles regions. An alternative would be the explicit use of the image support during the construction of the model. In fact selective active shape models for example [7] use only a small portion of the landmarks during the search, while sparse models [14] aim to optimize shape representation by proposing a sparse representation that encodes sparsity and exploits geometric redundancy.

A promising line of research was initiated by a method [10] to build hierarchical active shape models of 2-D anatomical objects using 1-D wavelets, which are then used for shape based image segmentation. A further extension is performed in [23] [24], where spherical wavelet can characterize shape variation in a local fashion in both space and frequency, in contrast to spherical harmonics that have a global basis set. Indeed this representation aims to learn a shape prior of 3-D objects by applying PCA to the clusters of correlated diffusion wavelet coefficients, and to exploit this prior for image segmentation based on the diffusion wavelets presentation. A more recent work by Zhu & al. [26] presents an approach based on the diffusion wavelet however just applied for the matching of 3D shapes.

Therefore it would be useful to consider a framework in which, for a given data set, similarities are given only between very similar points, and where it would be possible to organize local information by diffusion into global parameterizations. Therefore we present a method that provide for building a generative model of shape variation, and use it for segmentation in new data, including the use of orthomax to obtain an optimal subdivision of the shape.




Figure 1: Calf MRI: (a) healthy (b) manual expert annotation and (c) zoom on AT, EDL and PL muscle groups.

Beyond pre-defined manifolds The shape model approach proposed in this report accounts for the systematic behavior of shape variation and image support in the anatomical structure. It obtains an accurate delineation of the partially visible muscle surfaces, that cannot be achieved with registration based methods, by accounting for local statistical modeling and appearance behavior. The shape representation is based on a finite set of landmarks, that are repeatedly identified and exhibit significant differentiation to the background on different examples of the anatomical structure.

In contrast to conventional shape modeling approaches that rely on a pre-defined topology, and parameterize the surface of an object with regard to an according manifold like a sphere, the method learns the appropriate topology from the training data, and uses an according shape representation based on diffusion wavelets to model its variation. For this we need two mechanisms: 1. A way to determine the intrinsic topology of a shape for which multiple examples are available, and 2. a means to represent the shape variation with regard to this topology, i.e., a way to exploit arbitrary topologies. In this report we present an approach, that tackles both aspects. First we learn the intrinsic topology of training data by means of shape maps. Then we represent the shape variation by diffusion wavelets. This representation can cope with arbitrary topologies that are defined by a diffusion kernel. In contrast to e.g. spherical wavelets [23] they are not confined to a spherical reference manifold. Therefore they can represent the variation of training examples in a much more efficient and natural way. Our work is motivated by the highly heterogeneous distribution of visual information, and the locally confined deformation behavior in calf MRI data (Fig. 1), where standard segmentation methods fail, due to the ambiguous texture and the sparse distribution of salient leg image information. The differentiation between individual muscles is very subtle, the distribution of reliable image information at the boundaries is un-even, and this parts can only be estimated from prior shape knowledge.

An innovative research direction lies on connecting computational harmonic analysis to a broad range of issues facing the need of exploiting high-dimensional data. Diffusion wavelets generalize standard wavelets, allowing for multi-scale analysis on general structures, such as manifolds, graphs and point clouds in Euclidean space. Indeed diffusion maps [3] and diffusion wavelets [4], both derived from the wavelet concept, provide a significant new tool for describing complicated architecture of high-dimensional data, since they encapsulate all the common advantages of wavelets. We apply our method to anatomical object segmentation, and more specifically to CT heart data and MRI calf muscle, and for the latter of interest in the study of Myopathies. We will present experimental results that underline, the superior efficiency that a model based on diffusion wavelets can achieve as opposed to existing approaches. We will also evaluate the effect of this model representation on the search performance.

The remainder of the report is organized as follows: in Sec. 2, the shape representation based on diffusion wavelets is presented, and in Sec. 3 we focus on the search in new data. In Sec. 4 we report experimental results and quantitative validation. Finally Sec. 5 concludes with a discussion.

2 Shape Representation

We parameterize the shape variation observed in the training data by means of diffusion wavelets. The *topology* of the structure that is modeled is defined by a diffusion kernel. It can be viewed as a generalization of standard parameterizations, e.g.: the diffusion kernel for a triangulated spherical surface would be the adjacency matrix weighted by the mutual distances. Defining the topology by a diffusion kernel instead of a fixed genus-0 manifold allows us to learn and incorporate more complex interaction patterns observed in the training data.

In the following we will first outline diffusion wavelets, then will explain the associated shape representation, and finally will note how the orthomax principle can be used to separate coherent sub-regions of the shape.

2.1 Properties of Diffusion Wavelets

Wavelets represent a robust mathematical tool for hierarchically decomposing functions. The theory is described extensively in [22]. Through wavelets decomposition, a function can be described in terms of a coarse overall shape, that is enriched by details that range from coarse to fine. They provide an elegant technique for representing the levels of detail regardless of the interest function type, e.g. an image, a curve, or a surface. Besides the major advantage of the wavelets remains on the compact support of basis functions in one hand, and on the inherently hierarchical representation since it is based on multi-resolution modeling of processes at different spatial and temporal scales. Thus with appropriately selected eigenfunctions or scaling functions of Markov matrices, we are able to optimally describe local variations in muscles shapes comparing to conventional approaches like PCA [4]. Diffusion Wavelet conveniently reshape many of the ideas of multi-scale analysis that have proven so useful in study of 1-D signals and 2-D images, making it possible to describe the complex nature of high-dimensional phenomena. Finally they provide an interesting alternative to global Fourier eigenfunctions for value function approximation.

The issue of finding meaningful structures and geometric descriptions of a data set is often tied to that of dimensionality reduction. Among the different techniques developed, particular attention has been paid to kernel methods. Their nonlinearity as well as their locality-preserving property are generally viewed as a major advantage over classical methods like principal component analysis and classical multidimensional scaling. As far as spectral kernel methods to achieve dimensional reduction are concerned one can refer to some recent approaches such as; Locally Linear Embedding (LLE) [25], Laplacian Eigenmaps [1] or still Hessian Eigen maps [12]. All these techniques minimize a quadratic distortion measure of the desired coordinates on the data, naturally leading to the eigenfunctions of Laplace-type operators as minimizers.

To get meaningful geometric descriptions of data sets, one should refer to diffusion process, where it's shown that eigenfunctions of the Markov matrix corresponding to the kernel of the diffusion process provide coordinates, called diffusion maps, for generating efficient representations of complex geometric structures.




Figure 2: Heart CT data: (a) shapes before alignment (b) after alignment and (c) Mean heart shape.

2.2 A Diffusion Operator Reflecting The Topology

Diffusion wavelets generalize standard wavelets, allowing for multi-scale analysis on general structures, such as manifolds, graphs and point clouds in Euclidean space. Diffusion wavelets [4], and the closely related diffusion maps [3], relate computational harmonic analysis to a broad range of dimensionality reduction questions. Both are derived from the wavelet concept, and provide a significant new tool for describing complicated architecture of high-dimensional data.

All anatomical shapes are aligned by Procrustes alignment [21], which produces the series of volumes V_i^p , before calculating the mean shape \bar{V}^p . Results of this alignment on heart data can be seen in figure 2. After this registration, we can calculate the mean shape of the muscles/hearts as in the equation (2)

$$\bar{V}^p = \frac{1}{N} \sum_{i=1}^N V_i^p \quad (1)$$

$$S_i = V_i^p - \bar{V}^p \quad (2)$$

We build our diffusion operator T on a cloud of points embedded in a metric space. A matrix of graph weights for the point cloud is constructed by building a local bump function centered at each point and then the weight matrix is normalized through the symmetric Laplace-Beltrami to form the diffusion operator T .

2.3 Diffusion Wavelet Modeling of Shape Variations

Given this diffusion operator T defining the manifold, we use the corresponding hierarchical diffusion wavelets, to represent the shape variation.

First we build a hierarchical wavelet structure, the *diffusion wavelet tree*: We call upon a general multi resolution construction for efficiently computing, representing and compressing T^{2^j} , for $j > 0$. Indeed during the down-sampling process, and throughout a recursive sparse QR decomposition we aim to obtain the orthonormal bases of scaling functions, $\Phi = \{\phi_j\}$, the wavelets W_j , and compressed representation of T^{2^j} on ϕ_j , for j in the requested range.

Giving K maximum number of levels to compute, we will obtain after K steps in this fashion, a representation of T^{2j} onto a basis ϕ_j , with $1 \leq j \leq K$. For more calculation details please refer to [4].

Once the construction of the diffusion wavelet tree Φ is accomplished, we use it to represent the individual training shapes. In fact we calculate the diffusion wavelet coefficient Γ on the deviation from the mean of the aligned shapes, and obtain the following diffusion wavelet coefficient:

$$\Gamma_{S_i} = \Phi^{-1} \cdot S_i \quad (3)$$

Thus the shape representation of the muscles could be written:

$$V_i^p = \bar{V}^p + \Phi \Gamma_{S_i} \quad (4)$$

Once we have generated the diffusion wavelet coefficients for all training examples, we build our model as follows:

- i Consider the coefficient level by level,
- ii Computing covariance matrix of diffusion wavelet coefficient, and
- iii Apply PCA to each level of coefficients in order to obtain variation of the shape at a specific scale.

At this stage we have the diffusion wavelets coefficients which in the lowest level provide an overall approximation, whereas localized variations are captured hierarchically by the higher-level coefficients. Consequently we conduct our PCA study over each frequency level of the diffusion wavelets coefficients separately, instead of taking into the consideration all the points, with $\{\sigma_j\}_{j=1 \dots K}$ are the subset of eigenvectors corresponding to eigenvalues $\{\lambda_j\}_{j=1 \dots K}$ of the covariance matrix of the diffusion wavelets coefficients at each level j , ($1 \leq j \leq K$), which will call Γ_{level_j} (see eq. 5)

$$\Gamma_{level_j} = \{\Gamma_{S_i/level=j}\}_{i=1 \dots N} \quad (5)$$

and $\Gamma_{level_j}^*$ as the deviation from the mean of the latter; so at each level the coefficients would be expressed such as:

$$\Gamma_{level_j} = \bar{\Gamma}_{level_j} + \sigma_j (\sigma_j' \cdot \Gamma_{level_j}^*) \quad (6)$$

We can afterward gather together the coefficients of each shape at all level ($\Gamma_{S_i Rec}$), and consequently get the reconstruction of our initial data.

$$V_i^p = \bar{V}^p + \Phi \cdot \Gamma_{S_i Rec} \quad (7)$$

2.4 Modeling Using the Orthomax Criterion

In order to separate parts of the data with low correlation, we use the orthomax criterion [2] in the coefficient space of the wavelets. More especially we explore the varimax version [20] for optimizing sparsity corresponding to new variables being associated to localized modes of variation. Indeed orthomax rotations represents re parameterizations of the PCA space producing a simple basis.

Let \mathbf{R} be an orthonormal rotation matrix in \mathbb{R}^k where $\mathbf{R}_{i,j}$ represents \mathbf{R} elements, and where k implies the number of eigenvectors with the largest eigenvalues λ_i . Besides σ denotes as previously in the report the $p \times N$ eigenvectors matrix. The orthogonal orthomax rotation matrix \mathbf{R} is calculated as follow:

$$\mathbf{R} = \arg \max_{\mathbf{R}} \left(\sum_{j=1}^k \sum_{i=1}^p (\sigma \mathbf{R})_{ij}^4 - \frac{1}{p} \sum_{j=1}^k \left(\sum_{i=1}^p (\sigma \mathbf{R})_{ij}^2 \right)^2 \right), \quad (8)$$




Figure 3: Data reconstruction through global PCA (a) and localized orthomax (b) rotated modes of the DW coefficients models. The surfaces representing the $\pm 3\sqrt{\lambda_i}$

Orthomax rotation effects are illustrated in Fig. 3, where one can easily notice that while the PCA modes demonstrate several spatially distributed effects within each mode, the varimax modes in the other hand show nicely isolated effects.

Furhtermore Fig.4 shows the 'flattening' of the eigenvalue spectrum carried out by the varimax rotation where the respective modes as well as variances are plotted. So this simple, yet powerful modification of PCA had enabled us to optimize sparsity leading to localized modes of variation, and which is more suitable for applications with sparse parameterizations as the pathologies we are working on.

In Fig.5 the resuming scheme of our pipeline model is depicted. In fact, once we succeeded to separate our data variations through the wavelet level space, we can then get the shape prior by projecting back the selected orthomax eigenvectors of the diffusion coefficients level into the right dimensions. So the equation 6 would be expressed with the orthomax components. Thus our pipeline allows to capture variations at every levels with-

out discriminating higher variations beyond lower ones. Our final shape reconstruction will be therefore defined as following:

$$\Gamma_{level}^{i*} = [\Gamma_{level_1}^{i*} \Gamma_{level_2}^{i*} \dots \Gamma_{level_K}^{i*}] \quad (9)$$

$$V_i^p = \bar{V}^p + \Phi \cdot \left(\bar{\Gamma}_{level_j} + \sigma_{j_{orthomax}} \left(\sigma_{j_{orthomax}}' \cdot \Gamma_{level}^{i*} \right) \right) \quad (10)$$




Figure 4: Comparison between PCA and Orthomax DW eigenvalues

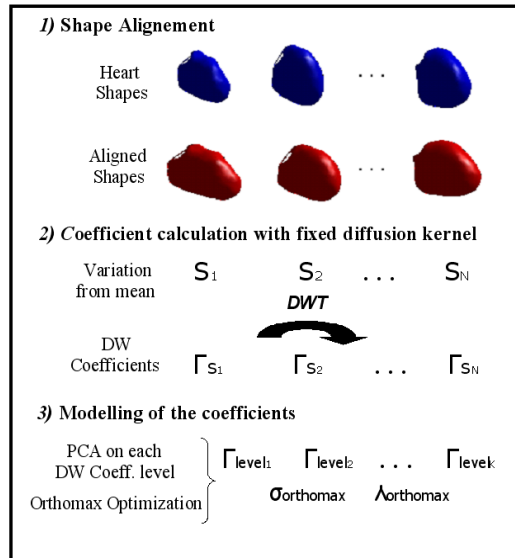


Figure 5: Scheme of Diffusion Wavelet Coefficient Process

3 Data Search

3.1 Sparse Search

Let us now consider a new anatomical data set in which we want to determine the position of the modeled object (e.g. a specific muscle/heart). The search with the model representation $\hat{\mathbf{V}}$, and appearance patch models $(b_i)_{i=1, \dots, N}$ for each landmark is performed in an iterative manner, starting from a coarse initialization obtained by atlas registration. Similar to a standard shape model inference approach, the positions of landmarks in new data are estimated by an energy minimization that involves both shape prior and appearance costs:

1. The landmark positions of $\hat{\mathbf{V}}$ are updated according to a local appearance model. For each landmark the position with highest probability of the corresponding local texture patch being consistent with regard to the learnt texture model b_i is chosen along a line orthogonal to the model surface. Let $b(\mathbf{x}_i^j)$ be the learned texture patch at the correct landmark position \mathbf{x}_i in the training example \mathbf{I}_j , and for landmark positions in a local neighborhood \mathcal{N} let $Q_i^j(\mathbf{x})$ be the correlation between the patch $b(\mathbf{x})$ and $b(\mathbf{x}_i^j)$ normalized within the neighborhood, i.e. $\int_{\mathbf{x} \in \mathcal{N}} Q_i^j(\mathbf{x}) = 1$, then the image support is

$$g_i = \text{mean}_{j=1, \dots, n} \left(\frac{Q_i^j(\mathbf{x})}{\int_{\mathbf{x} \in \mathcal{N} \setminus \mathbf{x}_i^j} Q_i^j(\mathbf{x})} \right). \quad (11)$$

That is, for a landmark in \mathcal{V} the image support is calculated from the local appearance behavior at the corresponding positions in the training set.

2. The shape is constrained by either a local or global statistical shape model. In our work we use diffusion wavelet shape model. This procedure is iterated. During each iteration, the corresponding \mathbf{V}' is reconstructed to re-estimate the surface normals. After convergence the final reconstruction \mathbf{V}' is an estimate of the true muscle surface inferred from the data, and the prior model.

3.2 Appearance Feature Model

3.2.1 Appearance Representation Using Local Features

We employ a search framework related to Active Feature Models (AFM) proposed in [19]. The approach is related to Active Appearance Models (AAM) [5], with the main difference that AFMs describes appearance by means of local features, and infers model updates during search by means of canonical correlation analysis (CCA), which has advantages given noisy data. The benefit of the method is considerable in the presence of complex data, like muscle data, where large parts of the variation within the muscles have low relevance for landmark localization and a small training set has to compromise a decent representation of the texture in the model.

For the calf muscle segmentation, we utilize a Gabor jet with frequencies $\{0.3, 0.6, 0.9\}$ and directions $\{0, \pi/4, \pi/2, 3\pi/4, \}$ to describe the local texture at the landmark positions. During training we learn the relation between landmark displacements and the corresponding texture feature change by CCA. Still during this phase, model parameters are perturbed randomly generating a large number of displaced model instances. A functional relation is then learned from the resulting feature vectors describing local texture and the corresponding model parameter displacement by Canonical Correlation Analysis (CCA) [11].

3.2.2 Search by canonical correlation analysis


During search local features are extracted at the current landmark position estimates. Based on the relation learnt by CCA, according model parameter updates are performed . This results in an iterative search approach, that converges to the landmark positions, based on the local appearance, and constraint by the diffusion wavelet shape model.

4 Results

4.1 Experimental Set-Up

In order to evaluate the performance of the proposed method, we have considered two different data set:

- 25 calf muscles divided into two groups: 20 healthy control patients and 5 unhealthy cases. For each volume there are 90 slices of 4mm thickness, and with voxel spacing of 0.7812x0.7812x4 mm acquired with a 1.5 T Siemens scanner. Standard of reference annotation by experts for the Medial Gastrocnemius (MG) muscle, was available (Fig. 6.b). Correspondences for a set of 895 landmarks on the surfaces were obtained by an MDL based optimization [17]. We performed a leave-one-out cross validation on the whole data-sets, i.e., models were trained on 24 examples, and search was performed on the remaining example.
- 25 CT volumes of the heart, with an approximate voxel spacing of 1.5 mm, for which 90 anatomical standard of reference landmarks, and a set of 1451 control points for the left ventricle was available, in addition to the ground truth segmentation (Fig. 6.a) from experts concerning the diastole as well as the systole.



(a) Ground truth segmentation of papillary muscles



(b) *T-1* MRI slice supervised segmentation of a human calf

Figure 6: Standard reference segmentation of respectively the calf muscle and the left ventricle

4.2 Experimental Validation

Evaluating model reconstruction performance To assess the model representation of our framework, we model the shapes of the structures by diffusion wavelet shape models in a leave-one-out cross validation strategy. To evaluate the diffusion modeling approach, we assess two measures: 1. reconstruction error, and 2. search performance. We compared the reconstruction error for gaussian shape models, and the proposed diffusion wavelet modeling.

	Heart Data	Calf Data
<i>Gaussian model</i>	1.6154	2.1277
<i>DW Model with spatial kernel</i>	0.0755	0.1485
<i>DW Model with Shape Map kernel</i>	0.1100	0.1796

Table 1: Full Landmark Reconstruction Error (in voxel) with regard to three different shape models for heart and calf data sets.

We tried two different diffusion wavelet kernels: 1. just utilizing the spacial proximity of landmarks, and 2. a kernel based on the shape map distance of the landmarks[18], where shape maps were built with kernel size $k = 3$ and $\epsilon = 500$ for both data sets.

The goal is to see how far our model is able to detect the local shape variations based on different kernels. As a first comparison, we show in the Fig. 8 the reconstructed surfaces for our heart data set using projected wavelet coefficients on the set of principal components at level 1, where the surfaces are representing the $\pm 3\sqrt{\lambda_i}$ from left to right. Moreover the reconstruction error between the diffusion wavelet model and the reference model is calculated as an average surface error for all test shapes in table 1, where we also establish rate errors. In Fig 7 the reconstruction of the projected shape model (heart/muscle) is depicted, this reconstruction starts from the projection of the diffusion wavelet coefficients in respectively the first and the last level, and then extracting a new subset of coefficients from the eigenvectors that constitutes 99% of the variation in the correspondant level. To keep the variation in reasonable limits, the shape parameters are also restricted to $\pm 3\sqrt{\lambda_i}$. During the reconstruction experiments we were also able to check the effects of each eigenvector of any scale on the rate error.

Evaluating search performance To evaluate the search behavior we compared the proposed method with a standard shape model search in an active shape model manner, based on an even sampling of the object surface, and gradients in the volumes, and a sparse shape model proposed in [14] The latter uses a similar appearance model, and allows for the assessment of the effect of replacing the multivariate Gaussian landmark model, with the diffusion wavelet shape model. The error measure evaluated is the mean distance of the model landmarks between standard of reference and segmentation result (Tab.2). The latter gives also an indication of the displacement along the surface, which is relevant if the result is used for navigation. Models were initialized with minimal overlap to the target shape, and the accuracy of the final result was quantified by the mean landmark error between ground truth annotation and search result. As far as quantitative comparison is concerned, (Tab.2) results clearly show how the diffusion wavelet model outstands the sparse model and that for both anatomical data sets, with for example 6.97 voxels for DW model over 8.72 error voxels for the sparse model through the challenging calf data sets.




Figure 7: Diffusion Wavelets Model reconstruction. First row for Heart results and second for Calf muscle. Data, green: standard of reference segmentation, red: reconstruction result for a. first scale and b. last wavelet scale.




Figure 8: Reconstructed surfaces for Heart CT data using projected wavelet coefficients on the set of principal components that represent 99% of the total variance at level 1. The surfaces representing the $\pm 3\text{sqrt}(\lambda_i)$ from left to right. The figure is dedicated to the axial view.

Results The diffusion wavelet model was able to recover the shape with superior accuracy. In the muscle data the standard search approach failed due to the ambiguous texture in large regions of the target shape. In Fig. 10 examples for standard, sparse model [14], and muLtiscale diffusion wavelet based search are depicted. In fact one of the important points that distinguishes our methodology from robust ASMs, is that we learn the distribution

	Heart Data	Calf Data
<i>Standard Model</i>	18.35	33.04
<i>Sparse Model</i>	5.81	8.72
<i>Diffusion Wavelet Model</i>	4.32	6.97

Table 2: Heart/Calf Search for Segmentation. Landmark Reconstruction Error (in voxel) comparison between three different search models.

of both image and shape information during the training phase to optimally exploit the anatomical properties of the data. This is not the case with of robust ASMs which for a given sampling they consider a subset of the control points according to the observed image support.

AFM Search Regarding the muscle data our approach was able to recover the shape with superior accuracy, in the meanwhile the standard search approach failed due to the ambiguous texture in large regions of the target shape. In Fig 9 one can visualize the boxplots of the dice volume overlap with the expert labels for the 25 patients concerning the standard and the active feature model search. It’s interesting to note that with the help of image support based on local texture descriptors, the method performs better for muscle segmentation due to a restriction to more relevant information being used for regression and fitting.




Figure 9: Results of the segmentation using the DW-AFM model. (a) Boxplots of the Dice Similarity Coefficients Measure and (b) landmark error (voxel) after finishing the search phase over the whole data set, with (1) our current approach and (2) standard model.




Figure 10: Model search result for Heart muscle(upper row) and Calf muscles(down row). data, green: standard of reference segmentation, red: search results for (a, d, g) standard gradient search approach, and uniform sampling, (b, e, h) sparse shape models and (c, f, i) diffusion wavelet model.

5 Conclusion

We present a multiscale shape prior based on diffusion wavelets shape representation and segmentation. Diffusion wavelet shape models are able to take advantage of the subtle inter dependencies in training data, by clustering coefficients based on correlation, and representing the topology of the structure by a diffusion kernel, instead of a fixed pre defined manifold. Furthermore we are using the orthomax criterion which is established to be suitable for building sparse representations -as it's the case of the current pathologies studied in the report- with relative ease leading to localized modes of variation. We have shown that in the context of anatomical structures, the diffusion wavelet transformation is able to accurately and efficiently detect the locations and spatial scales of shape variations. The

validation on detecting patterns on two complex data sets shows promising results indicating the advantage of using a learnt model parameterization. Also we want to point out that contrary to the standard AAM's model, we are introducing a substantially more flexible and powerful shape model (hierarchical diffusion wavelets vs. global PCA on point coordinates), a local appearance description and according search paradigm. As for the work of Spherical wavelets by Nain & al [23], we are extending the representative power: complete freedom of object topology defined by a diffusion kernel in our case, vs. a sphere in the case of [23] and the orthomax criterion to allow for a 'soft' subdivision, vs. n-Cuts in Nain & al. work. Future work includes the integration of model learning approaches that learn the locations in non-annotated data in a weakly- or unsupervised manner. The prime goal remains to achieve the segmentation accuracy necessary for a clinical application in myopathy diagnosis and follow up monitoring, to allow for a non-invasive analysis of the muscle structure.

Acknowledgments

This work was supported by Association Française contre les Myopathies (AFM: <http://www.afm-france.org>) under the DTI-MUSCLE project.

References

- [1] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*. MIT Press, 15:1373–1396, 2003.
- [2] M. Browne. An overview of analytic rotation in exploratory factor analysis. *Multivariate Behavioral Research*, 36(1):111–150, 2001.
- [3] R. R. Coifman and S. Lafon. Diffusion maps. *Appl. Comput. Harmon. Anal.*, 21:5–30, 2006.
- [4] R. R. Coifman and M. Maggioni. Diffusion wavelets. *Appl. Comput. Harmon. Anal.*, 21:53–94, 2006.
- [5] T. F. Cootes, G. Edwards, and C. Taylor. A comparative evaluation of active appearance model algorithms. In *Proceedings of BMVC'98*, volume 2, pages 680–689, 1998.
- [6] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE TPAMI*, 23(6):681–685, 2001.
- [7] T. F. Cootes, A. Hill, C. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12(6):355–366, 1994.
- [8] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models-their training and application. *Computer Vision and Image Understanding*, 61:38–59, 1995.
- [9] D. Cremers and M. Rousson. *Deformable Models. Theory and Biomaterial Applications*, chapter Efficient Kernel Density Estimation Of Shape And Intensity Priors For Level Set Segmentation, pages 447–460. Jasjit S. Suri and Aly A. Farag, 2007.
- [10] C. Davatzikos, X. Tao, and S. Dinggang. Hierarchical active shape models, using the wavelet transform. *IEEE Transactions on Medical Imaging*, 22:414–423, 2003.
- [11] R. Donner, M. Reiter, G. Langs, P. Peloschek, and H. Bischof. Fast active appearance model search using canonical correlation analysis. *IEEE TPAMI*, 28(10):1690 – 1694, October 2006.
- [12] D. L. Donoho and C. Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *PNAS*, 10:5591–5596, 2003.
- [13] J. Duncan and N. Ayache. Medical image analysis: progress over two decades and the challenges ahead. *TPAMI*, 22:85–106, 2000.
- [14] S. Essafi, G. Langs, and N. Paragios. Sparsity, redundancy and optimal image support towards knowledge-based segmentation. In *Proc. of IEEE CVPR'08*, 2008.
- [15] J. W. Fernandez and P. J. Hunter. An anatomically based patient-specific finite element model of patella articulation: towards a diagnostic tool. *Bio. and Model. in Mechbio.*, 4:20–38, 2005.
- [16] B. Gilles, L. Moccozet, and N. Magnenat-Thalmann. Anatomical modelling of the musculoskeletal system from mr. *MICCAI*, pages 289–296, October 2006.

-
- [17] G. Langs, R. Donner, P. Peloschek, and H. Bischof. Robust autonomous model learning from 2d and 3d data sets. In *Proc. of MICCAI'07*, 2007.
 - [18] G. Langs and N. Paragios. Modeling the structure of multivariate manifolds: Shape maps. In *Proc. of IEEE CVPR'08*, 2008.
 - [19] G. Langs, P. Peloschek, R. Donner, M. Reiter, and H. Bischof. Active feature models. In *Proceedings of the International Conference on Pattern Recognition ICPR 06.*, volume 1, pages 417–420, 2006.
 - [20] K. E. Leung and J. G. Bosch. Localized shape variations for classifying wall motion in echocardiograms. In *Medical Image Computing and Computer-Assisted Intervention MICCAI*, volume 4791, pages 52–59, 2007.
 - [21] B. Luo and E. R. Hancock. Iterative procrustes alignment with the em algorithm iterative procrustes alignment with the em algorithm. *Image and Vision Computing - Elsevier*, 2002.
 - [22] Y. Meyer. Wavelets - Algorithms and applications. *Applied Mathematics*, 1993.
 - [23] D. Nain, S. Haker, A. Bobick, and A. Tannenbaum. Shape-driven 3d segmentation using spherical wavelets. In R. Larsen, M. Nielsen, and J. Sporring, editors, *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2006*, volume 4190 of *LNCS*, pages 66–74. Springer, 2006.
 - [24] D. Nain, S. Haker, A. Bobick, and A. Tannenbaum. Multiscale 3-d shape representation and segmentation using spherical wavelets. *IEEE Trans Med Imaging*, 26(4):598–618, 2007.
 - [25] S. Roweis and L. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.
 - [26] K. Zhu, Y. Wong, W. Lu, and J. Fuh. A diffusion wavelet approach for 3-d model matching. *Computer-Aided Design*, 41(1):28–36, 2009.



Unité de recherche INRIA Futurs
Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur

INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399