Universitat
Pompeu Fabra
Barcelona

Enlargement, Subdivision and Individualization of Statistical Shape Models

Application to 3D Medical Image Segmentation

Marco Pereañez

TESI DOCTORAL UPF / 2017

Enlargement, Subdivision and Individualization of Statistical Shape Models

Application to 3D Medical Image Segmentation

By

MARCO PEREAÑEZ

A dissertation submitted to Universitat Pompeu Fabra in accordance with the requirements of the degree of DOCTOR OF PHILOSOPHY of the Department of Information and Communication Technologies.

Thesis Directors
Prof. Alejandro F. Frangi
Dr. Karim Lekadir

JUNE 2017

This work was carried out at the Center for Computational Imaging and Simulation Technologies in Biomedicine (CISTIB), Universitat Pompeu Fabra, Barcelona. This thesis was supported by a departmental scholarship from the Department of Information and Communications Technologies at Universitat Pompeu Fabra (DTIC-UPF).

Abstract

Statistical shape models (SSMs) are well-established tools for medical image segmentation and interpretation due to their ability to constrain the image analysis result within allowable limits as defined by a representative set of training examples. However, the construction of SSMs that are rich, i.e., that represent well the natural and complex variability of anatomical structures, is a difficult task, as this is often limited by the available delineated training shapes, which rarely exceed a few dozen. As a consequence, such models often generalize poorly to new groups of individuals, particularly given the increasing amounts of data nowadays routinely acquired for clinical practice and biomedical research. Computational methods are thus required to enrich these SSMs and increase the quality of the representation of the variability between individuals, with no additional datasets. This requirement is important as the delineation of new training shapes is always laborious, particularly for building SSMs of complex 3D shapes.

This thesis presents three original and complementary approaches to enhance the quality of the SSMs, that improve the accuracy of medical image segmentation in challenging applications. First, we enhance the statistical richness of SSMs by developing a technique capable of merging the shape representations and statistical properties of several pre-existing models with no original or additional raw data. Second, we enhance the geometrical quality of SSMs by developing a framework for modeling simultaneously both global and local characteristics of highly complex and/or

multi-part anatomical shapes. Last, we improve the specificity of SSMs for each dataset by integrating individual-specific non-imaging metadata such as demographic, clinical and behavioural variables into the SSM construction and image segmentation tasks. These techniques are demonstrated and validated by considering various imaging modalities such as magnetic resonance imaging (MRI) and computed tomography (CT), and different complex anatomies, including the human heart, brain and spine.

Resumen

Los modelos estadísticos de forma son herramientas bien establecidas para la segmentación e interpretación de imágenes médicas debido a su habilidad para restringir el resultado de la segmentación dentro de los límites permitidos usando como referencia un grupo representativo de ejemplos usados para entrenar el modelo. La construcción de modelos estadísticos de formas que son detalladas, porque representan bien la variabilidad compleja y natural de las estructuras anatómicas es una tarea dificil, frecuentemente limitada por la escasez de las delineaciones de formas disponibles para el entrenamiento del modelo. Estas raramente superan las pocas docenas. Como consecuencia, estos modelos a menudo no tienen un buen desempeÒo cuando se aplican a nuevos individuos, en particular dada la cantidad de datos que en la actualidad se obtienen para prácticas clínicas e investigación biomédica. Métodos computacionales son entonces requeridos para enriquecer estos modelos estadísticos de formas e incrementar la calidad de la representación de la variabilidad entre individuos y sin la disponibilidad de bases de datos adicionales para entrenamiento del modelo. Este requisito es importante ya que la delineación de nuevas formas de entrenamiento es siempre laboriosa en particular para construir modelos de formas complejas de tres dimensiones.

Esta tesis presenta tres propuestas originales y complementarias

para mejorar la calidad de los modelos estadísticos de formas que mejoran la precisión de la segmentación de la imagen médica en aplicaciones difíciles. Primero, proponemos mejorar la riqueza estadística de los modelos de formas por medio del desarrollo de una técnica capaz de unir la representación de forma y las propiedades estadísticas de muchos modelos pre-existentes sin datos originales adicionales. Segundo, proponemos mejorar la calidad geométrica de los modelos de formas estadísticas por medio de la creación de un marco de referencia para modelar simultáneamente las características globales y locales de alta complejidad y/o que constan de multiples partes anatomicas. Por último, proponemos mejorar la especificidad de los modelos estadísticos de formas para cada paciente mediante la integración de metadatos específicos al paciente y que no se derivan de la imagen, como por ejemplo, variables demográficas, conductuales y de entorno clínico, en la construcción de los modelos. Estas técnicas son demostradas y validadas mediante el uso de varias modalidades de imágenes tales como imágenes de resonancia magnética (MRI) y tomografía computarizada (CT); así como diversas formas anatómicas, incluyendo el corazón, el cerebro y la espina dorsal humanos.

TABLE OF CONTENTS

List of Tables

List of Figures

Figure **Page**

Acronyms

ASM	Active Shape Model
CAP	Cardiac Atlas Project
CBIR	content-based medical image retrieval
CMR	Cardiac Magnetic Resonance
CT	Computed Tomography
CVD	Cardio Vascular Disease
DETERMINE	Defibrillators to Reduce Risk by Magnetic Resonance Imaging Evaluation
ED	End Diastole
EF	Ejection Fraction
ES	End Systole
GRE	Gradient Recalled Echo
ICP	Iterative Closest Point
LV	Left Ventricle
MESA	Multi Ethnic Study of Atherosclerosis
MRI	Magnetic Resonance Imaging
PCA	Principal Component Analysis
PDM	Point Distribution Model
RMS	Root Mean Squared
RV	Right Ventricle
SSM	Statistical Shape Model

CHAPTER 1

GENERAL INTRODUCTION

1.1 Medical Image Segmentation in the Era of Personalized Medicine

Tomorrow's healthcare lies in early diagnosis and individually tailored treatments, i.e. delivering the right treatment to the right patient at the right time [1]. This revolutionary approach to medicine, the so-called personalized medicine, has been embraced as a vital societal challenge worldwide, with the active participation of researchers, care providers, industries, and policy makers [2]. The fundamental goal is to develop novel technological solutions for early patient diagnosis and prognosis, and personalized treatment and mon-

itoring.

Medical imaging plays an important role in this personalized medicine revolution because of its unique capability to capture *in vivo* the most subtle morphological and functional changes in diseased organs [3]. Recently, there has been an increasing interest in quantifying such changes through the development of quantitative imaging biomarkers and computational models of anatomy. On the one hand, compared to biochemical and histological biomarkers, imaging biomarkers have the advantage of remaining non-invasive and being spatially and temporally resolved [4]. Image-based models of anatomy constitute valuable tools for computer simulation of patient response to intervention [5].

One important pre-requisite to the estimation of imaging biomarkers and computational models is the segmentation of the boundaries of the organs and tissues of interest. There have been significant advances in medical image segmentation over the years and various approaches to the problem have been developed. However, in the era of personalized medicine, existing segmentation techniques are now faced with a significant increase in the medical images available in clinical practice. These statistics illustrate the growth and volume of data available today: In the UK alone, the total number of radiology tests performed between 2012-2013 was 41.1 million as reported by the National Health Service [6]. In the last 10 years, the volume of magnetic resonance imaging (MRI) and computed tomography (CT) scans has increased by 211%, and 167%, respectively. There is a need

for new scalable, reliable, and automatic methods for image segmentation and quantification that can address the much larger biological variability and complexity faced with in today's clinical practice and medical research.

1.2 Medical Image Segmentation using Statistical Shape Models

One natural approach to deal with the extraordinary increase in image data volumes is with statistical shape models (SSMs) [7]. With these methods, statistical models of anatomical shape can be constructed to encode the anatomical variability from a representative training sample of manually segmented organs. The image segmentation is then carried through specific algorithms designed to constrain the segmentation result within an allowable domain defined by the statistical model. One such algorithm is the Active Shape Models (ASMs) [8, 9], which comprises two steps; the training of a point distribution model (PDM), and the model-constrained image segmentation. These two stages can be summarized:

Training

1. *Shape Representation:* Given a set of manually delineated image volumes, define a set of points, $\mathbf{x} = \{x_1, y_1, z_1, \ldots, x_P, y_P, z_P\}$ that adequately describe the surface of the organ.

2. *Correspondence Generation:* Propagate the aforementioned point configuration to all manually delineated samples, and construct a shape matrix by concatenating all shapes, $\mathbf{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$.

3. *Shape Alignment:* Since the objective is to model the shape of the organs, attributes that do not constitute shape are removed from samples, i.e., translations, rotations, and scaling depending on the application.

4. *Linear Decomposition:* Having the set of aligned shapes, one can then use matrix factorization techniques, such as, principal component analysis (PCA) [10], to extract the main modes of variation from the data. The resulting model is the mean $\bar{\mathbf{x}}$, and variance $\mathbf{C} = \mathbf{V}\mathbf{D}\mathbf{V}^T$, where, $\mathbf{V}$ and $\mathbf{D}$, are the principal components, and associated variance, respectively. With this model, new shapes are generated from the distribution by adding a linear combination of modes of variation to the mean shape, $\tilde{\mathbf{x}} = \bar{\mathbf{x}} + \mathbf{V}\mathbf{b}$, $\mathbf{b}_i \leq \pm 3\sqrt{\mathbf{D}_{i,i}}$, see Fig. 1.1

Image Segmentation

1. *Shape Initialization:* The first step in the image segmentation process is positioning an initial estimate of the shape, its location and orientation. Methods exist that localize [11], and adjust the pose [12] of the object of interest providing an initial estimation of the segmentation.

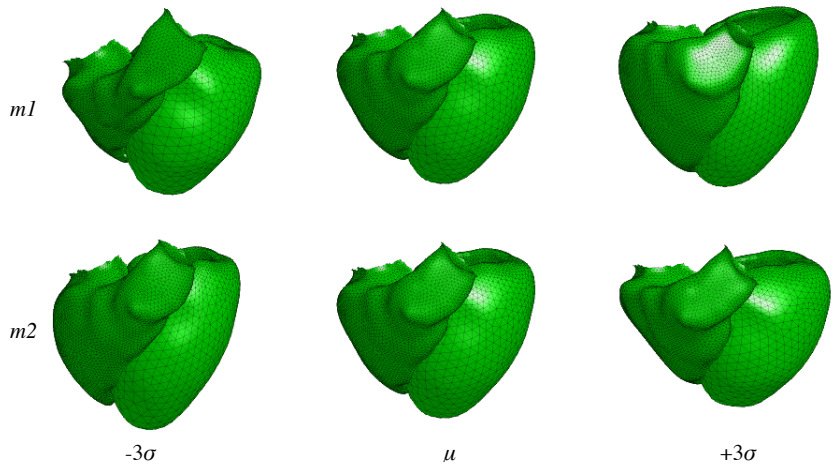

Figure 1.1: Bi-ventricular cardiac shapes generated by a SSM. Rows showing the first 2 modes of variation from the model. The center column showing the mean μ, and $\pm 3\sigma$ left and right, respectively.

2. *Feature Point Search:* Parting from the initial shape, the next step is the search of relevant features on the image. This is done for every point, by locally searching image patches that minimize the distance to a pre-trained model of image appearance, and updating the location of the points.

3. *Projection onto PDM:* Due to the noisy nature of the image, and false positive image feature detections, some points from the previous step will be incorrectly localized, and therefore, the resulting shape will not be a valid instance as defined by the model. To correct this, the resulting shape is projected onto the model space, and constrained to be within n standard deviations

from the mean. The resulting valid instance is then projected back to the image space.

4. *Iterate:* ASMs are an iterative procedure where step 2) is initialized with the current estimate of the shape from step 3). This procedure is repeated until no significant changes in shape are measured. The resulting shape is the output of the algorithm.

1.3 Statistical Shape Modeling: Current Challenges

Despite their many advantages and potential applications, SSMs are associated with limitations that reduce their quality and render them less robust when applied to large and variable databases, i.e.:

1. *Their construction is tedious and time consuming:* The construction of SSMs, is not a trivial task. It requires the collection of a sufficiently large and representative training sample, and landmark annotation with point correspondence of all training shapes. This is true with anatomical shapes, which cannot be easily annotated manually, such as for the heart.

2. *They do not generalize well for highly variable shapes:* The dimensionality of the data is typically much higher than the number of samples. As a result, the available number of modes of variation is inevitably low compared to the number of landmarks

that represent the shape. This means that a sub-space spanned by a permissive number of eigenvectors cannot always generalize well to unseen shapes. The statistical model over-constrains the image segmentation process, which can lead to incorrect segmentations. See Fig. 1.2.

3. *They do not handle well complex multi-part structures:* SSMs rely on the assumption that after removing non-shape related attributes such as pose and scale from shapes, the resulting points follow a Gaussian distribution. Such assumption becomes largely invalid when dealing with articulated or highly complex geometries. An example is the complex variability found in the spine due to potential differences in forward/lateral flexion, extension, and rotation [13]. Here, the relative location of parts/components can vary in non-Gaussian manner, making the image segmentation task even more challenging than for single object organs.

4. *They are biased towards global shape variability:* Standard SSMs built using eigendecomposition encode well the overall global geometry of the anatomical shape, but give less weight to the variation of finer geometrical patterns. In some applications, this could lead to unsatisfactory segmentation of fine anatomical details (see Fig. 1.2), affecting later analysis, such as biomechanical simulations. In spine modeling, for instance, the global geometry of the vertebral bodies can be well represented with SSMs, but more complex regions such as the processes, which interact with the ligaments and muscles [14], are more problematic.

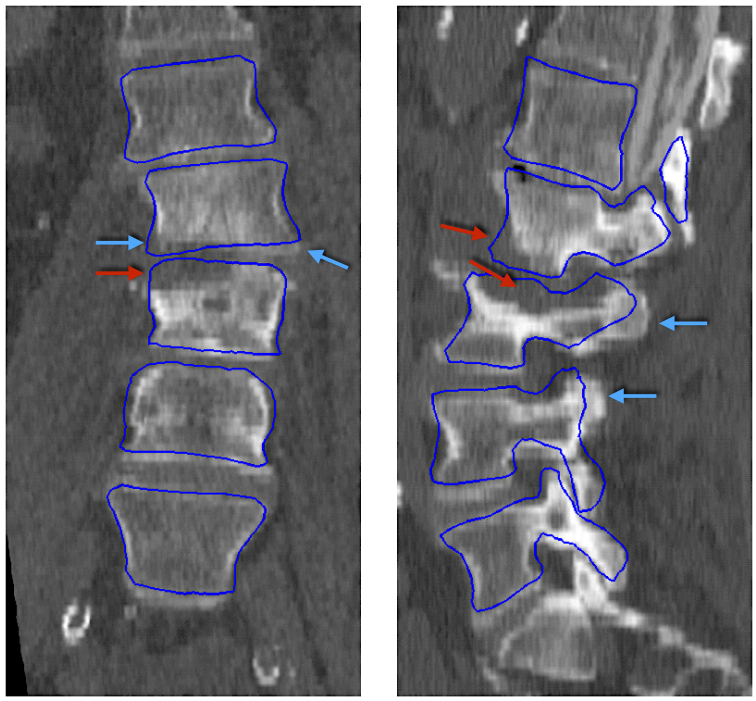

Figure 1.2: An example of an automatic ASM segmentation of the lumbar spine (L1-L5). Blue arrows on both figures show areas where the low generalization ability of the underling PDM fails to capture local details such as on the spinous process. The red arrows show the inability of the PDM to correctly capture the shape of a fractured vertebra (L3 vertebra).

There is a need for new methods to complement the standard construction of SSMs to enhance their quality and richness, and to enable handling anatomical shapes with high geometrical complexity and biological variability, such as for the heart, spine, or brain. Sev-

eral studies have attempted to enhance the quality of SSMs through different strategies. For instance, Lötjönen et al. [15] have proposed enlarging model variability by creating "new" training instances artificially from known shapes, by deforming the real shapes non-rigidly. However, it is difficult to know the extent to which creating artificial data is contributing meaningful statistical and anatomical information to the model. Another strategy consisted of combining statistical and physical models [16, 17], but these physical models are only fine-tuned approximations to a particular population and tissue, and cannot be extended trivially in a more general context. For the modeling of multi-part object ensembles, different strategies have been proposed such as through hierarchical representations [18, 19, 20, 21, 22], but they have not been applied to complex articulated systems such as the spine.

1.4 Objectives and Contributions of the Thesis

In this thesis, we aim to present novel solutions to the aforementioned problems, which will enable the construction of higher quality and richer SSMs for 3D medical image segmentation. This thesis is centered on these objectives:

1. To build SSMs that encode higher variability without new training samples (Addressed in Chapter 2).

2. To develop methods for modeling and segmenting geometrically complex structures (Addressed in Chapters 2 and 3).

3. To develop methods for modeling and segmenting multi-part articulated shapes (Addressed in Chapter 3).

4. To build SSMs that generalize well to the individual under scrutiny (Addressed in Chapters 2 and 4).

To achieve these objectives, this thesis proposes the following contributions:

Chapter 2 To increase the variability encoded by the model without additional training samples, by directly fusing pre-existing SSMs into a single one. A normalization technique is used to transform pre-existing SSMs into a common landmark representation. An eigenspace fusion technique is then employed to combine the mean and eigenvectors of the individual models without the original data. Such technique enables to merge models constructed in different research or clinical centers, with the ability to address multiple shape representations, image modalities, and anatomical populations. This increases the generalization ability of the resulting model while avoiding the very time-consuming annotation of new examples. This technique is validated with several anatomical structures including the human heart, brain and spine, and two imaging modalities, CT and MRI.

Chapter 3 To improve the statistical shape modeling of geometrically complex and articulated anatomical shapes, such as the spine. With this method, we decompose the individual vertebrae into statistically meaningful parts in order to improve the modeling of regions associated with fine details and geometrical complexity. The high-level inter-relationships between the vertebrae and the different regions are encoded using conditional models to ensure global consistency. The aim of such representation is to obtain a model that reduces the over-constraining problem of SSMs by replacing the global model with multi-part constraints that adapt better to the geometrical complexity of the articulated spine.

Chapter 4 To improve the specificity of the SSM for each individual by integrating individual-specific non-imaging information into the modeling and segmentation tasks. Specifically, we use demographic (e.g. age, sex), behavioral (e.g. smoking) and clinical indicators (e.g. weight, hypertension), which we call *metadata*, to constrain, and personalize the shape space of a specific subject based on conditional distributions relating the population's shape and metadata. In this manner, a SSM is created that is more representative of the subpopulation to which a specific subject belongs. We demonstrate these metadata-customized SSMs with cardiac segmentation in MRI.

In summary, we present in this thesis three complementary approaches to enhance the quality of the SSMs and to improve the accuracy of medical image segmentation in challenging applications, i.e. 1) by fusing pre-existing models, 2) by building multi-part models, 3) and

by customizing the models based on individual-specific metadata. We use different anatomical structures to demonstrate the value of these techniques, namely the human heart, brain and spine, based on two distinct imaging modalities, i.e, MRI and CT.

CHAPTER 2

A Framework for the Merging of Pre-Existing and Correspondenceless 3D Statistical Shape Models

2.1 Motivation

The construction of statistical shape models (SSMs) that are rich, i.e., that represent well the natural and complex variability of anatomical structures, is an important research topic in medical imaging. To this end, existing works have addressed the limited availability of training data by decomposing the shape vari-

ability hierarchically or by combining statistical and synthetic models built using artificially created modes of variation. In this paper, we present instead a method that merges multiple statistical models of 3D shapes into a single integrated model, effectively encoding extra variability anatomically meaningful, without the need for the original or new real datasets. The proposed framework has great flexibility due to its ability to merge multiple statistical models with unknown point correspondences. The approach is beneficial in order to re-use and complement pre-existing SSMs when the original raw data cannot be exchanged due to ethical, legal, or practical reasons. This paper describes two main stages, i.e., 1) statistical model normalization and 2) statistical model integration. The normalization algorithm uses surface-based registration to bring the input models into a common shape parameterization with point correspondence established across eigenspaces. This allows the model fusion algorithm to be applied in a coherent manner across models, with the aim to obtain a single unified statistical model of shape with improved generalization ability. The framework is validated with statistical models of the left and right cardiac ventricles, the L1 vertebra, and the caudate nucleus, constructed at distinct research centers based on different imaging modalities (CT and MRI) and point correspondences. The results demonstrate that the model integration is statistically and anatomically meaningful, with potential value for merging pre-existing multi-modality statistical models of 3D shapes.

2.2 Introduction

Constructing comprehensive statistical models of 3D shapes (SSMs) is a well-studied yet still challenging problem in medical image computing. One of the main difficulties remains the construction of shape models that are rich, i.e., that represent well the natural and complex variability of anatomical structures. However, this is often limited by the available delineated training shapes, which rarely exceed a few dozens. Consequently, such models often generalize poorly to new subjects due to over-fitting. Computational methods are required to enrich these SSMs and increase the quality of the representation of the variability between individuals and groups of individuals, without additional datasets. The latter requirement is important as the delineation of new training shapes is always laborious, particularly for building multi-modality SSMs from various imaging sources (e.g., MRI and CT).

Two main approaches have been proposed to improve the quality of SSMs without new datasets. First, several techniques reduce the gap between dimensionality and training size by considering object ensembles and/or object hierarchies and decomposing the shape models ([20, 22, 21, 18, 19, 23]). On a similar note, promoting sparsity in the statistical decomposition can provide shape models that are easier to manipulate as its modes of variation are local rather than global ([24, 25, 26, 27]). Perhaps these techniques do not explicitly enrich the models with extra anatomical variability but merely employs a shape decomposition that relaxes the shape constraints to give more

flexibility to the SSM during image interpretation.

To explicitly add shape variability to the models, a second popular approach consists of combining the statistical models with synthetic modes of variations ([15, 28, 29]). The deformation can be chosen to reflect expected variations, either through finite element analysis ([9, 17]) or heuristically ([30, 31]). However, despite the desired effect this approach appears to have, it remains a surrogate for having more artificial training data and its success depends on the parametrization of the method and its relative weight regarding the real data [20].

We present an alternative approach that considers the combination of different pre-existing shape models of the same shape class, without the assumption of point correspondence between the models. The training database is effectively increased, without the actual training data of either of the models used in the combination. When combining models from different clinical/research centers, it may be impossible to obtain the original training data for ethical, legal or practical reasons. However, the complementary richness of these models due to either the training population or the underlying imaging modality, may make it attractive to combine them.

The proposed workflow addresses two key technical issues. First, as these individual models are constructed at different centers with distinct delineation tools and protocols, the integration procedure must be able to handle differences in shape representations. While establishing point correspondence has been extensively investigated for training shapes ([32, 33, 34, 35, 36, 37, 38, 35, 39, 40, 41, 42, 43, 44,

45, 46]), this work addresses the problem of establishing correspondence across eigenspaces using a surface-based approach. Second, once the eigenspaces are normalized into a common representation, a fusion of the individual models into an integrated SSM is necessary so the relative importance of each individual model considered based on statistical criteria ([47, 48]).

The contributions of this work are: 1) From the point of view of the approach, to the best of our knowledge, this is the first work proposing to merge existing PDMs that have no point correspondence. This has a great impact in terms of re-usability of already constructed PDMs and allows to build enriched and more comprehensive models without new data. 2) From the technical point of view, we introduce a novel workflow using non-rigid registration of surfaces from [49], and PDM fusion method from [48]). We introduce a method to transform an eigenspace into a new shape representation by using the barycentric mapping in section 2.3.1, which allows us to obtain normalized PDMs that can be adequately merged. 3) From an application point of view, this paper is the first demonstration of the fusion of PDMs in the medical image analysis community. We show the performance of the technique with various experiments on cardiac, bone and brain data.

This chapter is organized as follows. In section 2.3.1 we introduce the algorithm used to establish point correspondence across the individual eigenspaces. Section 2.3.2 describes the method used to merge the transformed eigenspaces into a unique integrated SSM. In section 2.4 we evaluate the method with applications to cardiac statistical

shape modeling though the combination of existing MRI- and CT-based models, and to the brain's caudate nucleus, and vertebrae. We conclude in section 2.5 with a discussion of the SSM combination approach and conclusions.

2.3 Method

The proposed framework consists of two main stages: 1) the model normalization, and 2) the model merging. In section 2.3.1 we present the algorithm used to bring the input models into a unique shape parameterization with point correspondence established across mean shapes and eigenspaces. Then in section 2.3.2 we present the algorithm to merge the transformed input models into a single model that retains all the properties and characteristics of its constituent models.

2.3.1 Normalization algorithm

The first step in unifying a set of correspondenceless models is the creation of surface correspondence between the different PDMs. This is followed by the calculation of a barycentric mapping, which is used to transform the eigenspaces into the common shape space where they can be compared and unified into a single PDM.

Let $\Omega_i = (\bar{\mathbf{x}}_i, \Phi_i, \Lambda_i, K_i, V_i)$, for $i = 1, \ldots, N$ be the set of pre-built models of the same object with different landmark placement strategy. Here $\bar{\mathbf{x}}_i$ is a vector of concatenated vertex coordinates representing the mean shape, Φ_i and Λ_i are the eigenvector and eigenvalue matrices,

correspondingly, and K_i and V_i are the number of observations used to build the model, and their triangulation. Now, let $\mathcal{M}_i = (\bar{\mathbf{x}}_i, V_i)$ be the model mean surfaces defined by the vertices $\bar{\mathbf{x}}_i$ and triangulations V_i. In these steps we detail the process of finding point correspondence across the means of all models, and finding a suitable transformation for the normalization of all eigenvector spaces so we can perform fusion. Algorithm 1 summarizes the process described in the these sections.

Algorithm 1 Model normalization algorithm

Let: $\Omega_i = (\bar{\mathbf{x}}_i, \Phi_i, \Lambda_i, K_i, V_i)$ and $\mathcal{M}_i = (\bar{\mathbf{x}}_i, V_i)$
Let: $\mathcal{M}_{ref}$ be a reference mean surface from the models Ω_i ▷ See Sec. 2.3.1.1
Input: Ω_i, for $i = 1, ..., N$

for $i = 1 \rightarrow$ N
    $\mathcal{M}_i^{icp} = \text{icp}(\mathcal{M}_i, \mathcal{M}_{ref})$ ▷ Rigid registration. See Sec. 2.3.1.2
    $\mathcal{M}_i^{curr} = \text{currents}(\mathcal{M}_i^{icp}, \mathcal{M}_{ref})$ ▷ Non-rigid registration. See Sec. 2.3.1.2

    for $j = 1 \rightarrow$ numberOfVertices ▷ in reference mesh $\mathcal{M}_{ref}$
        $\mathbf{T}_{(i,j)} = \text{getBarycentricWeights}(\mathcal{M}_{(ref,j)}, \mathcal{M}_i^{curr})$
    end for

    $\bar{\mathbf{x}}_i' = \mathbf{T}_i \bar{\mathbf{x}}_i$ ▷ Transform the mean. See Sec. 2.3.1.3
    $\Phi_i' = \mathbf{T}_i \Phi_i$ ▷ Transform the eigenvectors. See Sec. 2.3.1.3
end for

Output: $\Omega_i' = (\bar{\mathbf{x}}_i', \Phi_i', \Lambda_i, K_i, V_{ref})$

2.3.1.1 Selecting a reference mean surface mesh

The first step of the normalization is the selection of a reference mesh from among the existing models. The most obvious criterion to choose the reference model is mesh resolution. Choosing the model with the highest mesh resolution will retain maximal information from all models. However, if retaining maximal information is not paramount, but rather preserving the same shape representation that is already being used in a given application, or to speed up a process, then choosing a more sparse shape representation may be more appropriate. The effects of mesh resolution of the reference mesh are detailed in section 2.4.5.

2.3.1.2 Computing surface correspondence

We wish to establish surface correspondence between the reference model and all other models in the set. To achieve, this we use a two-step registration approach. We perform an initial rigid registration using the iterative closest point (ICP) algorithm [50], and then perform diffeomorphic free-form registration on the shapes [49] to refine the matching. We estimate a combined transformation $\varphi_i : \mathscr{M}_i \rightarrow \mathscr{M}_{ref}$ that takes points from a surface $\mathscr{M}_i$ to $\mathscr{M}_{ref}$. This combined step provides meaningful morphological correspondence across the set of the different models' mean shapes.

We use ICP because, unlike methods that require point correspondence, such as Procrustes analysis, ICP iteratively associates points by the nearest neighbor criteria minimizing a cost function. Since we

want to register meshes with different numbers of landmarks, ICP is appropriate and computationally efficient. Applying ICP is an important step because although the mean shapes are translation and scale independent, different models may differ in pose. This initial alignment also saves considerable computational time for the non-rigid registration.

The registration [49] uses the definition of currents to represent surfaces. It defines a surface as the change in flux of a current, or integration of a varying set of vector fields $w \in W$, where the vector field space W is generated by a Gaussian kernel of the form $exp(-||\mathbf{x}-\mathbf{y}||^2/\lambda_W^2)$. The subset of linear functionals, or dual space of currents W^*, is the dense span of basis elements $\delta_{\mathbf{x}}^{\mathbf{a}}$, called Dirac delta currents, defined at the spatial position $\mathbf{x}$ and with direction $\mathbf{a}$. The method approximates the current of a triangulated surface $\mathcal{M}_i$, by the sum of the Dirac delta currents defined at the triangle barycenters $\mathbf{x}_i$, and oriented along the triangle normals $\mathbf{a}_k^i$,

$$\mathcal{M}_i = \sum_k \delta_{\mathbf{x}_k^i}^{\mathbf{a}_k^i}. \tag{2.1}$$

With this representation, and noting that the Gaussian kernel provides the space $\mathbf{W}^*$ with an inner product, the currents space has a norm, allowing for the definition of a distance measure between two surfaces. The distance between two surfaces being the norm of their difference,

$$d(\mathcal{M}_i, \mathcal{M}_j) = \|\mathcal{M}_i - \mathcal{M}_j\|_{\mathbf{W}^*} = \|\sum_k \delta_{\mathbf{x}_k^i}^{\mathbf{a}_k^i} - \sum_l \delta_{\mathbf{y}_l^j}^{\mathbf{b}_l^j}\|_{\mathbf{W}^*}. \tag{2.2}$$

Registration is performed using the Large Deformation Diffeomorphic Metric Mapping (LDDMM) method described in ([51]), where a deformation ϕ_t^v is estimated by integrating the time-varying transport equation $\partial\phi_t^v(\mathbf{x})/\partial t = v_t(\phi_t^v(\mathbf{x}))$ where geodesic flows v_t, $t \in [0,1]$, are families of elements of a Hilbert space V for which an inner product is defined. More concretely, ϕ_t^v is found by minimizing:

$$F_{\mathcal{M}_i,\mathcal{M}_j}(v) = \int_0^1 \|v_t\|_{\mathbf{V}}^2 dt + \|(\phi_1)\mathcal{M}_i - \mathcal{M}_j\|_{\mathbf{W}^*}^2, \tag{2.3}$$

where the first term is called the regularizing term and is associated with the rigidity of the diffeomorphism. The second term called the matching term contains ϕ, which is responsible for the geometrical transport of one surface to another. Within the currents space, the space V is produced by an isotropic Gaussian kernel with standard deviation λ_V which controls the extent of the diffeomorphism's spatial consistency. Finally, an alternate two-step process defines the template $\mathcal{M}_i$ and diffeomorphisms ϕ where the initial surface is approximated by the mean initial current, and the transformations are updated as the distance between currents is minimized.

Based on the above definition of the LDDMM framework, and within the currents definition of surfaces, there are two parameters we manipulated during the registration step of our method. The first parameter is the standard deviation λ_V of the velocity fields v. This parameter controls the rigidity of the deformation of the surface at any given barycenter. Large values of λ_V will propagate the transformation of the current barycenter onto a proportionally large radial

neighborhood following a symmetric Gaussian bell shape. Small values of lambda depending on the mesh resolution will tend to limit the influence of the transformation down to a single barycenter. The second parameter is the standard deviation of vector fields w, λ_W. This parameter defines a search radius for a matching surface and fine detail matching during registration. Large λ_W values will be more likely to find a matching structure, however at the cost of loosing high frequency accuracy of the final registration. λ_W defines a limit beyond which variations are discarded as noise.

For our purposes, we experimentally found that setting $\lambda_V = 10mm$ is appropriate for the typical size of the objects in this study. Finding a value for λ_W that can be applied across the board however is more difficult. Depending on the sample size from which the mean surface was extracted, the nature of the data (diseased population), or the acquisition method, relatively large geometrical differences can be found between mean surfaces. These potential differences make it difficult to ensure that a given λ_W will be large enough to find the correct matching surface, and small enough to ensure matching of the fine variations. For this reason, we follow a refining strategy of λ_W during matching. We found that geometrical differences between any two means did not exceed $20mm$. Based on this knowledge we followed an iterative approach setting the initial value of $\lambda_W = 20mm$, ensuring that all regions of the surface are within the radius of influence. Then we reduced λ_W proportionally to the decrease in error between the surfaces and stopped when either no significant improvement was

found or the distance between surfaces fell below $1mm$. Figure 2.1 illustrates the final result of the currents registration after the initial alignment obtained with the ICP matching.

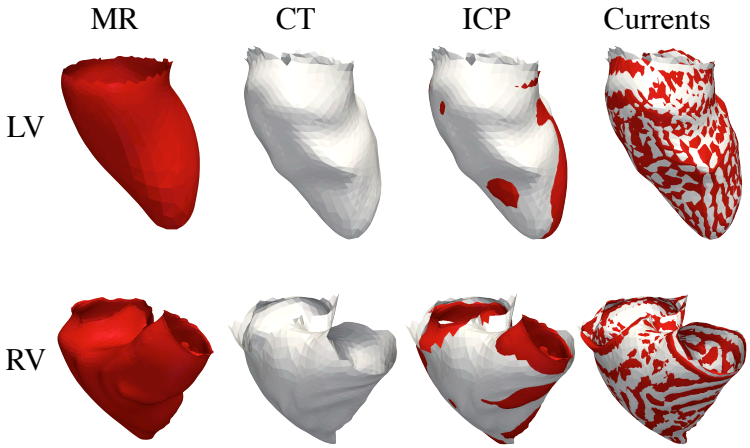

Figure 2.1: MR (Col. 1) and CT (Col. 2) mean shapes after ICP alignment. Their superposition before (Col. 3), and after (Col. 4) non-rigid currents registration. Top: LV; bottom: RV.

2.3.1.3 Computing inter-model landmark mapping

Once surface correspondence has been obtained, we find a suitable transformation that enables to achieve point correspondence between the input models with respect to the reference model. To achieve this, we introduce a barycentric mapping, which expresses every point on a reference mesh $\mathscr{M}_{ref} = (\bar{\mathbf{x}}_{ref}, V_{ref})$, as a linear combination of the vertices on the previously registered mean meshes $\mathscr{M}_i = (\varphi_i(\bar{\mathbf{x}}_i), V_i)$.

One key property of this mapping is that it preserves the linear nature of PCA-based point distribution models [9]. In other words, the merging of PDMs must produce a PDM.

A barycentric coordinate system describes a point as the center of mass of vertices of a simplex, which in our case are triangles (see Fig. 2.2). Barycentric transformations are often used in computer graphics to achieve linear interpolation of shapes [52, 53, 54].

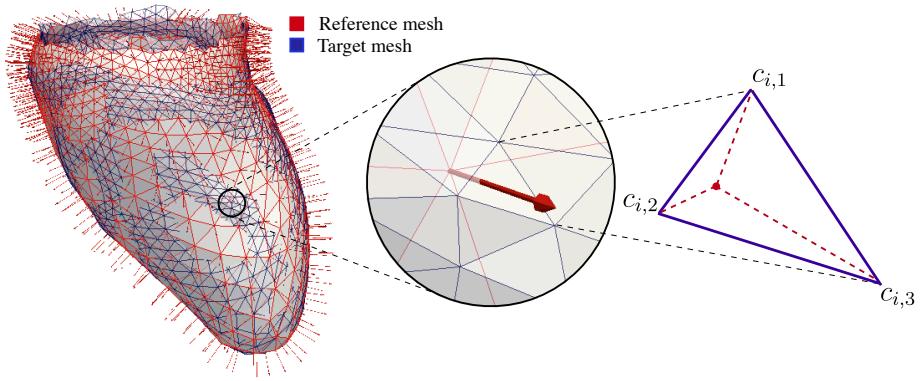

Figure 2.2: Barycentric mapping. Left: two registered meshes (reference and target). Middle: detail of the normal projection of a point in the reference mesh onto a target shape simplex. Right: The point at the plane-line intersection can then be expressed as the linear combination of coefficients $c_{i,1}, c_{i,2}, c_{i,3}$.

Let $\varphi_i(\bar{\mathbf{x}}_i)$ for $i = 1,\ldots,N$ be the previously aligned points (using currents registration) of the i^{th} mean shape on the reference surface $\mathscr{M}_{ref}$. Let $\mathbf{T}_i$ be a linear transformation matrix so $\bar{\mathbf{x}}_{ref} \approx \mathbf{T}_i\varphi_i(\bar{\mathbf{x}}_i)$. $\mathbf{T}_i\bar{\mathbf{x}}_i$ will change the parametrization of the surface $\mathscr{M}_i$ to match that of

$\mathscr{M}_{ref}$ and therefore it will be transformed into $\mathscr{M}'_i = (\mathbf{T}_i \bar{\mathbf{x}}_i, V_{ref})$ with established point correspondence.

A good candidate for $\mathbf{T}_i$ are the barycentric coordinates of the points $\bar{\mathbf{x}}_{ref}$ expressed in terms of points $\varphi_i(\bar{\mathbf{x}}_i)$ and triangulation V_i. We compute $\mathbf{T}_i$ as follows. For every point in the reference shape we compute the normal to its mesh surface, and the point of intersection with a plane described by a triangle on the target mesh. We then compute barycentric coefficients $c_{i,1}$, $c_{i,2}$, $c_{i,3}$ for each point in the triangle, and add them to a point transformation matrix $\mathbf{P}_j$ that transforms the j^{th} point

$$\mathbf{P}_j = \begin{bmatrix} c_{j,1} & 0 & 0 & 0 & 0 & c_{j,2} & 0 & 0 & 0 & 0 & c_{j,3} & 0 & 0 & \\ \cdots & 0 & c_{j,1} & 0 & 0 & \cdots & 0 & 0 & c_{j,2} & 0 & 0 & \cdots & 0 & 0 & c_{j,3} & 0 & \cdots \\ & 0 & 0 & c_{j,1} & 0 & & 0 & 0 & 0 & c_{j,2} & 0 & & 0 & 0 & 0 & c_{j,3} & \end{bmatrix}_{3\times 3m}.$$

The column-wise position of every 3×3 diagonal matrix $\mathbf{c}_j$, encodes the indexing of a simplex on the target mesh, thus preserving point correspondence between the reference and target meshes. m is the number of points in the target mesh, and the i^{th} mean shape vector is represented as $\bar{\mathbf{x}}_i = (x_{i,1}, y_{i,1}, z_{i,1}, x_{i,2}, y_{i,2}, z_{i,2}, \ldots, x_{i,m}, y_{i,m}, z_{i,m})^T$. Finally, the complete transformation matrix $\mathbf{T}_i$ is a row-wise concatenation of point transformation matrices $\mathbf{P}_j$,

$$\mathbf{T}_i = \begin{bmatrix} \mathbf{P}_{j,1}^T & \ldots & \mathbf{P}_{j,n}^T \end{bmatrix}_{3m\times 3n}^T,$$

where n is the number of points in the reference mesh. The transformations $\mathbf{T}_i$ are linear and therefore any shape $\mathbf{x}_i \approx \bar{\mathbf{x}}_i + \Phi_i \mathbf{b}_i$ can also

be re-parametrized as

$$\mathbf{T}_i\mathbf{x}_i \approx \mathbf{T}_i\bar{\mathbf{x}}_i + \mathbf{T}_i\Phi_i\mathbf{b}_i. \tag{2.4}$$

Having propagated the parametrization of $\mathscr{M}_{ref}$ to all the input models, we can also use its triangulation V_{ref} to obtain the resulting re-parametrized models as

$$\Omega_i' = (\mathbf{T}_i\bar{\mathbf{x}}_i, \mathbf{T}_i\Phi_i, \Lambda_i, K_i, V_{ref}), i = 1,\ldots,N, \tag{2.5}$$

which can now be merged using the fusion algorithm presented in the next section.

2.3.2 Model merging algorithm

Now that all eigenspaces have been transformed to having a common shape parametrization, we can detail the steps to merge them. To simplify notation, we denote the newly transformed mean shapes, and eigenvector matrices as $\bar{\mathbf{x}}_i' = \mathbf{T}_i\bar{\mathbf{x}}_i$, and $\Phi_i' = \mathbf{T}_i\Phi_i$ correspondingly.

The goal of fusion is to compute such an eigenspace $\Omega = (\bar{\mathbf{x}}, \Phi, \Lambda, K)$, using the information only from $\Omega_i' = (\bar{\mathbf{x}}_i', \Phi_i', \Lambda_i, K_i)$, for $i = 1,...,N$, where $\bar{\mathbf{x}}_i'$ are the mean shapes, Φ_i' are the eigenvectors, Λ_i are the eigenvalues, and K_i is the number of shapes in the i^{th} training set. This equals building a model from the full set of original observations, to which it is assumed we have no access. Algorithm 2 outlines the steps, following [48].

Algorithm 2 Model merging algorithm

1: **Input:** $\Omega'_i = (\bar{\mathbf{x}}'_i, \Phi'_i, \Lambda_i, K_i)$, for $i = 1, ..., N$
2: and $\mathbf{D}_i = \Phi'_i \Lambda_i \Phi'^T_i$
3: $[\bar{\mathbf{x}}''_i, \mathbf{ROT}_i] = \text{procrustes}(\bar{\mathbf{x}}'_i)$ ▷ section 2.3.3
4: $\bar{\mathbf{x}} = \text{computeWeightedMean}(\bar{\mathbf{x}}''_i)$ ▷ section 2.3.3
5: $\Phi''_i = \text{rotateEigenvectors}(\Phi'_i, \mathbf{ROT}_i)$ ▷ section 2.3.4
6: $\Gamma = \text{buildOrthonormalAllSpanningBasis}(\Phi''_i)$ ▷ section 2.3.5
7: $\mathbf{D} = \text{computeWeightedCovarianceMatrix}(\mathbf{D}_i)$ ▷ section 2.3.5
8: $[\mathbf{R}, \Lambda] = \text{eig}(\Gamma \mathbf{D} \Gamma^T)$ ▷ section 2.3.5
9: $\Phi = \Gamma \mathbf{R}$
10: **Output:** $\Omega = (\bar{\mathbf{x}}, \Phi, \Lambda, K)$

2.3.3 Compute combined model mean

Since every PDM has a different mean shape, the first step is to align the means $\bar{\mathbf{x}}'_i$ of all PDMs using Procrustes Analysis [55]. After Procrustes Analysis, the aligned means $\bar{\mathbf{x}}'_i$ are used to estimate the combined weighted mean $\bar{\mathbf{x}}$:

$$\bar{\mathbf{x}} = \left(\sum_{i=1}^{N} w_i \right)^{-1} \cdot \sum_{i=1}^{N} w_i \bar{\mathbf{x}}'_i, \tag{2.6}$$

with

$$w_i = K_i \cdot \left(\sum_{j=1}^{N} K_j \right)^{-1}, \tag{2.7}$$

where w_i are the fusion weights and $\sum w_i = 1$.

2.3.4 Eigenvector rotation

During the alignment, the shapes are centered and rescaled to unit size and then a rotation matrix is estimated. Let $\mathbf{S}_i$ be the 3×3 rotation matrix from the transformation that aligns the shape $\bar{\mathbf{x}}'_i$ to the mean $\bar{\mathbf{x}}$. Let Ξ_i be a $3n \times 3n$ block-diagonal matrix with repeating $\mathbf{S}_i$ along its diagonal:

$$\Xi_i = \begin{bmatrix} \mathbf{S}_i & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \mathbf{S}_i \end{bmatrix}.$$

We then use transformations Ξ_i to reorient eigenvectors $\Xi_i \Phi_i$, so we work with eigenspaces $\Omega''_i = \left(\Xi_i \bar{\mathbf{x}}'_i, \Xi_i \Phi'_i, \Lambda_i, K_i\right)$, assuming the mean shapes are already translation, and scale-independent.

2.3.5 Combined eigenspace basis

We want to compute the eigenvalues and eigenvectors that satisfy $\mathbf{D} = \Phi \Lambda \Phi^T$, where $\mathbf{D}$ is a combined covariance matrix that can be obtained solely from information in models Ω'_i, and where each covariance matrix $\mathbf{D}_i = \Phi_i \Lambda_i \Phi_i^T$ is weighted according to the number of contributed shapes of each model. A complete mathematical derivation

of **D** is provided in Sec. 2.3.6. Other derivations can be found in [47] and [48]. The following is a concise description of the steps involved:

1. *Find an eigenspace that spans all models' eigenspaces:* Construct an orthonormal basis set Γ, that spans all the eigenspaces $\Xi_i \Phi'_i$ by creating a matrix **H** that concatenates all eigenvectors $\Xi_i \Phi'_i$, and all possible combinations of differences between the different model's mean vectors $(\Xi_i \bar{\mathbf{x}}'_i - \Xi_j \bar{\mathbf{x}}'_j)$, for $i, j = 1, ..., N$, and $j > i$:

   $$\mathbf{H} = [\Xi_1 \Phi'_1, ..., \Xi_N \Phi'_N, (\Xi_1 \bar{\mathbf{x}}'_1 - \Xi_2 \bar{\mathbf{x}}'_2), ..., (\Xi_{N-1} \bar{\mathbf{x}}'_{N-1} - \Xi_N \bar{\mathbf{x}}'_N)]$$

   and orthonormalize:

   $$\Gamma = Orth(\mathbf{H})$$

2. *Determine an intermediate eigenproblem:* Use Γ to derive an intermediate eigenproblem:

   $$\Gamma \mathbf{D} \Gamma^T = \mathbf{R} \Lambda \mathbf{R}^T$$

   whose solution provides the combined model eigenvalues Λ and eigenvectors **R** needed to correctly orient the all-spanning basis Γ.

3. *Compute the combined eigenvector matrix:* Finally, the eigenvectors of the combined model Φ are obtained by:

   $$\Phi = \Gamma \mathbf{R}$$

With $\bar{\mathbf{x}}$ from Sec. 2.3.3, Φ from Eq. (2.8), Λ from Eq. (2.8), and $K = \sum_{i=1}^{N} K_i$, we obtain the fused model

$$\Omega = (\bar{\mathbf{x}}, \Phi, \Lambda, K).$$

2.3.6 Covariance matrix derivation

We have shown procedurally how to obtain the merged mean, eigenvalues and eigenvectors of a model we claim to be equivalent to one constructed directly from a single set of shapes. In this section, we provide theoretical proof of their equivalence as follows.

Assume we have M models: $\Omega_j = (\bar{\mathbf{x}}_j, \mathbf{D}_j, K_j)$, for $j = 1,2,\ldots,K_j$ where $\mathbf{D}_j = \Phi_j \Lambda_j \Phi_j^T$ and $\sum K_j = K$. Then the fused mean is

$$\bar{\mathbf{x}} = \sum_{j=1}^{M} K_j p_j \left(\Xi_j \bar{\mathbf{x}}_j\right) \tag{2.8}$$

where Ξ_j is the transformation matrix from $\bar{\mathbf{x}}_j$ to $\bar{\mathbf{x}}$ (Procrustes alignment), and $p_j = w_j(\sum_{j=1}^{M} w_j K_j)^{-1}$ is the weight assigned to any shape in the j^{th} model.

Now let Ψ_{ij} be the transformation matrix from the unaligned original observation $\mathbf{x}_{ij}$, (from the set of observations j, $j = 1,2,\ldots,K_j$), to the corresponding mean $\bar{\mathbf{x}}_j$ defined as

$$\bar{\mathbf{x}} = \sum_{j=1}^{M} K_j p_j \left(\Xi_j \bar{\mathbf{x}}_j\right) = \sum_{j=1}^{M} p_j \left(\sum_{i=1}^{K_j} \Xi_j \Psi_{ij} \mathbf{x}_{ij} \right). \tag{2.9}$$

Let $\tilde{\mathbf{x}}$ be the mean of all the shapes $\mathbf{x}_{ij}$ estimated by Procrustes Analysis. Let Υ_{ij} be a transformation from $\mathbf{x}_{ij}$ to $\tilde{\mathbf{x}}$. If, $\Upsilon_{ij} = \Xi_j \Psi_{ij}$ then, $\tilde{\mathbf{x}} = \bar{\mathbf{x}}$. Assuming this is true, the covariance constructed from all the observations is

$$\mathbf{D} = \frac{1}{A} \sum_{j=1}^{M} \sum_{i=1}^{K_j} \left(\Upsilon_{ij} \mathbf{x}_{ij} - \tilde{\mathbf{x}}\right) \left(\Upsilon_{ij} \mathbf{x}_{ij} - \tilde{\mathbf{x}}\right)^T p_j, \tag{2.10}$$

where $A = 1 - \sum K_j p_j^2$ is the unbiased normalizing factor of the variance. Then, using $\Upsilon_{ij} = \Xi_j \Psi_{ij}$ we obtain

$$\mathbf{D} = \frac{1}{A} \sum_{j=1}^{M} \sum_{i=1}^{K_j} \left(\Xi_j \Psi_{ij} \mathbf{x}_{ij} - \Xi_j \bar{\mathbf{x}}_j \right) \left(\Xi_j \Psi_{ij} \mathbf{x}_{ij} - \Xi_j \bar{\mathbf{x}}_j \right)^T p_j. \tag{2.11}$$

Factoring out Ξ_j and Ξ_j^T we obtain:

$$\mathbf{D} = \frac{1}{A} \sum_{j=1}^{M} \Xi_j \left(\sum_{i=1}^{K_j} \left(\Psi_{ij} \mathbf{x}_{ij} - \bar{\mathbf{x}}_j \right) \left(\Psi_{ij} \mathbf{x}_{ij} - \bar{\mathbf{x}}_j \right)^T \right) \Xi_j^T p_j, \tag{2.12}$$

where the expression in parenthesis is in fact the j^{th} unnormalized covariance matrix $(K_j - 1)\mathbf{D}_j$.

$$\mathbf{D} = \frac{1}{A} \sum_{j=1}^{M} \left(K_j - 1 \right) \Xi_j \mathbf{D}_j \Xi_j^T p_j, \tag{2.13}$$

$$\mathbf{D} = \frac{1}{A} \sum_{j=1}^{M} \left(K_j - 1 \right) \left(\Xi_j \Phi_j \right) \Lambda_j \left(\Xi_j \Phi_j \right)^T p_j \tag{2.14}$$

Equation (2.14) is the expression for the covariance matrix of the fused model estimated through the fusion algorithm. It was derived from the set of all unaligned observations. By demonstrating that it is exactly the linear combination of the transformed component covariance matrices $\mathbf{D}_j$ we demonstrated that the fusion algorithm gives exactly the same result as a batch construction of one single model from all the shapes.

2.4 Experiments and Results

2.4.1 Data

The proposed technique is validated using two sets of data for both the left ventricle (LV) and the right ventricle (RV), obtained from distinct clinical centers and using different imaging modalities. The first sample consists of 42 MRI datasets from the Sunnybrook cardiac MRI database [56]. These were acquired using a Steady State Free Precession (SSFP) sequence on a GE Signa CVi-HDx 1.5T scanner (General Electric, Milwaukee, USA). The resulting images consisted of short-axis slices with in-slice resolution = $1.56 \times 1.56mm$ and thickness = $8mm$.

The second sample consists of 134 CT datasets acquired using a Toshiba Aquilion 64 multi-slice CT scanner (Toshiba Medical Systems, Tochigi, Japan). The in-slice resolution in these images was $0.4 \times 0.4mm$ with slice spacing of $2.0mm$. The image volumes were segmented using an atlas-based approach [46].

The landmark distributions differ between the two sets; the LV is described by 880 landmarks in the MR dataset, and by 1347 landmarks in the CT dataset. For the RV these numbers are 1200 and 1748, respectively.

The vertebra spine images consist of 60 cross-sectional CT datasets of the lumbar spine collected at the CETIR Medical Center (Barcelona, Spain). The CT scans were performed using the Philips Gemini GXL 16 system (Philips Healthcare, Best, The Netherlands). The images

have a pixel resolution ranging between 0.47 and 1.04mm and a slice thickness of 0.5mm.

The caudate data were obtained from a total of 66 brain MR images through the Caudate Segmentation Evaluation MICCAI Challenge website[1] [57].

2.4.2 Evaluation Metrics

To quantitatively assess the performance of the proposed technique, we compare the quality of the original and the fused models using the well established measures of generalization ability, specificity, and compactness ([34]).

The generalization ability $G(M)$ is defined as

$$G(M) = \frac{1}{n_s} \sum_{i=1}^{n_s} \mathscr{D}(\mathbf{x}_i'(M), \mathbf{x}_i) \tag{2.15}$$

where n_s is the number of shapes in the dataset. $\mathbf{x}_i'(M)$ is the reconstruction of shape $\mathbf{x}_i$ using a model that excludes the instance being evaluated. M is the number of modes of variation used in the reconstruction, and $\mathscr{D}(\cdot,\cdot)$ indicates the distance computed between the shapes. In our case, we use the Root Mean Squared Error measure (RMSE) throughout the paper. Note that the generalization ability for a given number of modes of variation M is equivalent to the shape reconstruction error.

[1] `http://cause07.comicframework.org`

The specificity $S(M)$ measures the distance between a model-based synthetic population and the training set, and is defined as

$$S(M) = \frac{1}{N} \sum_{j=1}^{N} \mathscr{D}(\mathbf{x}_j(M), \mathbf{x}_j) \tag{2.16}$$

where N is the size of the synthetic population, set here to $N = 10,000$. $\mathbf{x}_j(M)$ are each of the synthetic instances created, and $\mathbf{x}_j$ is the closest element from the training set. The random instances were generated within 3σ from the mean. The compactness $C(M)$ is defined as

$$C(M) = \sum_{m=1}^{M} \lambda^m, \tag{2.17}$$

the summation of the first M eigenvalues in the model.

2.4.3 Merging two models

We evaluate numerically whether the proposed normalization and fusion framework translates in any loss in the generalization ability, specificity or compactness of the fused model. We perform a set of experiments through which we subdivide the available shapes for all anatomical structures into two randomly selected and non-intersecting subsets. The shape representation of one of the two sub-models is modified, before the application of the proposed fusion framework and comparison to the model obtained directly using all the training shapes with point correspondence.

Here we define Ω_S as the PDM constructed from all the original shapes S with point correspondence. The two sub-models are Ω_{S_A}, and

$\Omega_{S'_B}$, where $S_A \cup S_B = S$ and $S_A \cap S_B = \varnothing$. The prime in $\Omega_{S'_B}$ indicates a change in mesh connectivity and point density of shapes S_B.
The experiment comprises these steps:

1. Define a ground truth PDM Ω_{OM} (original model) constructed from S_{all}.

2. Partition S_{OM} into randomly selected subsets S_A, S_B so $S_A \cap S_B = \emptyset$, $S_A \cup S_B = S_{OM}$ (repeated 100 times).

3. Without loss of generality, choose subset S_B and re-sample all shapes to obtain S'_B so point correspondence between sets S_A and S_B is lost.

4. Construct the models Ω_{S_A} and $\Omega_{S'_B}$ from S_A and S'_B, respectively.

5. Apply the normalization and merging algorithm to the PDMs Ω_{S_A} and $\Omega_{S'_B}$ to obtain normalized and merged model Ω_{NM}.

6. Compare the generalization, specificity, and compactness of Ω_{NM} versus Ω_{OM}.

All models were built preserving 98% of total variance.

Elimination of point correspondence in the experiment (step 3) was achieved by using two iterations of the Loop subdivision algorithm [58] to the original meshes. Then, we decimated the meshes using a centroidal Voronoi diagram decimation algorithm as described in [59]. Since this algorithm computes vertices as the centroid of Voronoi-like regions comprising existing faces, the new vertices bear no linear

relationship to the original vertices. Additionally, this approach gives us full control over the number of vertices, which in this experiment was set to match the original number of landmarks.

Figs. 2.3, 2.4, and 2.5 plot the $G(M)$, $S(M)$ and $C(M)$ metrics for all anatomical structures and modes of variation. It can be seen that the curves are almost identical for all measures and modes, which indicates that the obtained models after normalization and merging are equivalent independently of the complexity of the shapes.

2.4.4 Convergence of fusing single-shape models

To further demonstrate the convergence properties of the merging algorithm, we evaluate in this section the limit case in which every individual shape is treated as a single model. In this case each PDM has a mean shape equal to the individual shape and a covariance matrix equal to zero. The proposed framework is then applied to this limit case and the result is compared to the model obtained directly by principal component analysis of all the training shapes. Table 2.1 summarizes the differences in generalization ability, specificity and compactness between the fused model of all the single-shape PDMs, and the ground truth model. It is evident from the results that the merging framework converges well to the original model even in the case of multiple input PDMs. We also see that despite having the covariance associated to each individual model be equal to zero, the method is able to accurately compute the covariance of the set.

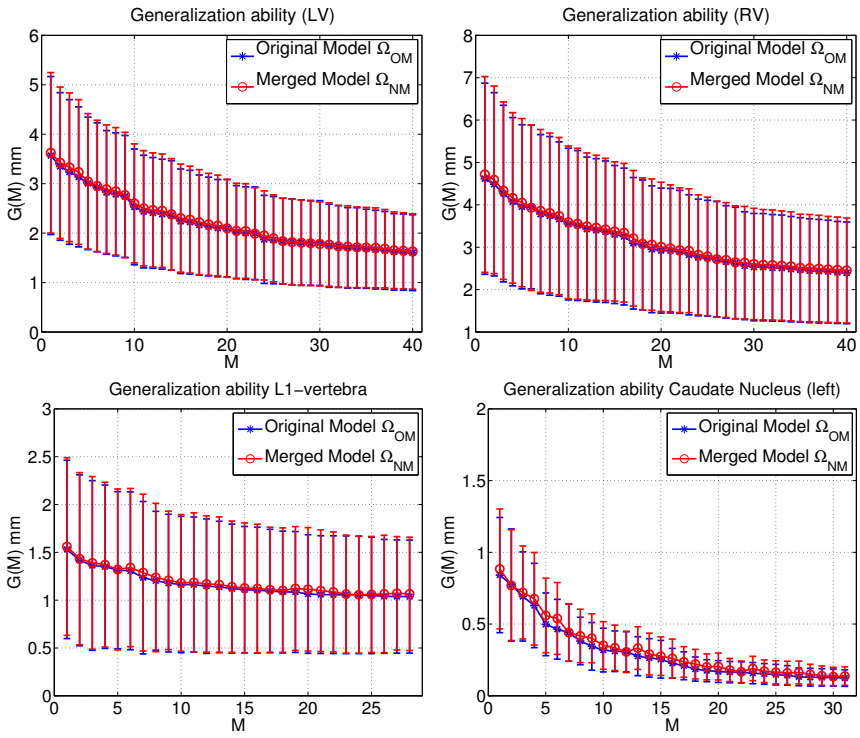

Figure 2.3: Generalization ability and standard deviation of the merged model (red) vs. the original shapes' model (blue) for the first M modes of variation.

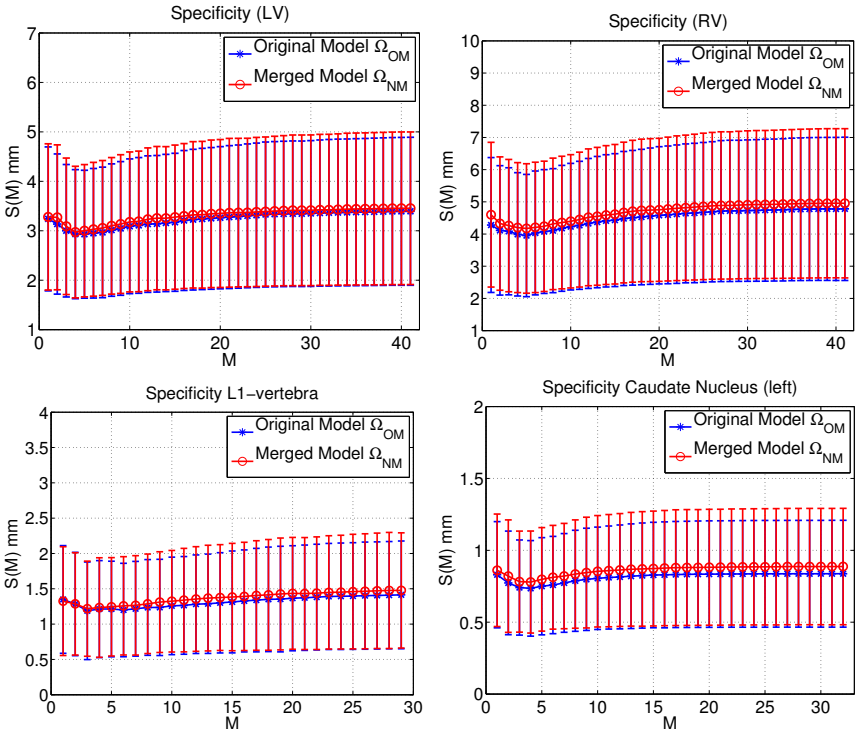

Figure 2.4: Specificity and standard deviation of the merged model (red) vs. the original shapes' model (blue) for the first M modes of variation.

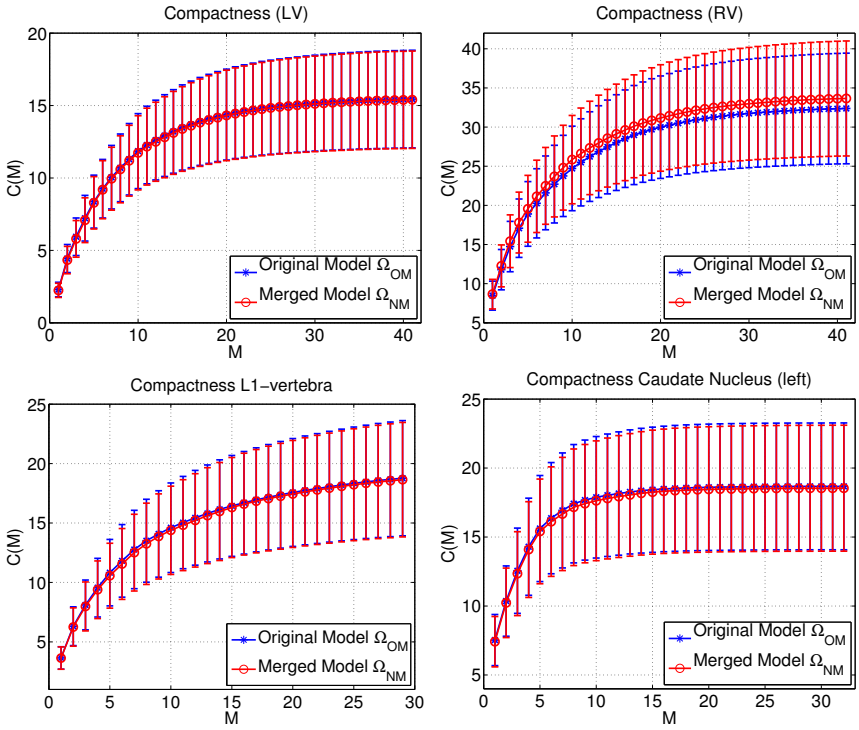

Figure 2.5: Compactness and standard deviation of the merged model (red) vs. the original shapes' model (blue) for the first M modes of variation.

Table 2.1: Unsigned difference between the Ω_{OM} and Ω_{NM} models' evaluation metrics preserving 98% of total variation. Compactness is expressed as a percentage of the total variance obtained for the Ω_{OM}.

Structure	$\lvert\Delta G\rvert\ (mm)$	$\lvert\Delta S\rvert\ (mm)$	$\frac{\lvert\Delta C\rvert}{C_{OM}}\ 100\%$
Left ventricle	0.002	0.448	2.281
Right ventricle	0.001	0.324	1.025
L-1 vertebra	0.000	0.040	0.407
Caudate nucleus	0.018	0.081	0.601

2.4.5 Choice of the reference model

We evaluate the effect of the choice of the reference model, both with respect to its mesh resolution and the number of training samples it is derived from. First, in order to test the effect of mesh resolution, we apply the proposed framework to two randomly generated subsets, while fixing the mesh resolution of the second PDM and varying the shape representation of the reference model. We re-sample the reference mesh for all structures to 9 different resolutions well below and above the fixed resolution of the second model.

Figure 2.6 displays the difference in shape reconstruction errors between the fused and original models. It can be seen that the differences are negligible for all four structures, with the largest error obtained for the RV of about $0.25mm$.

Second, we evaluate the effect of choosing the reference model as the one with the highest or lowest number of training samples. We

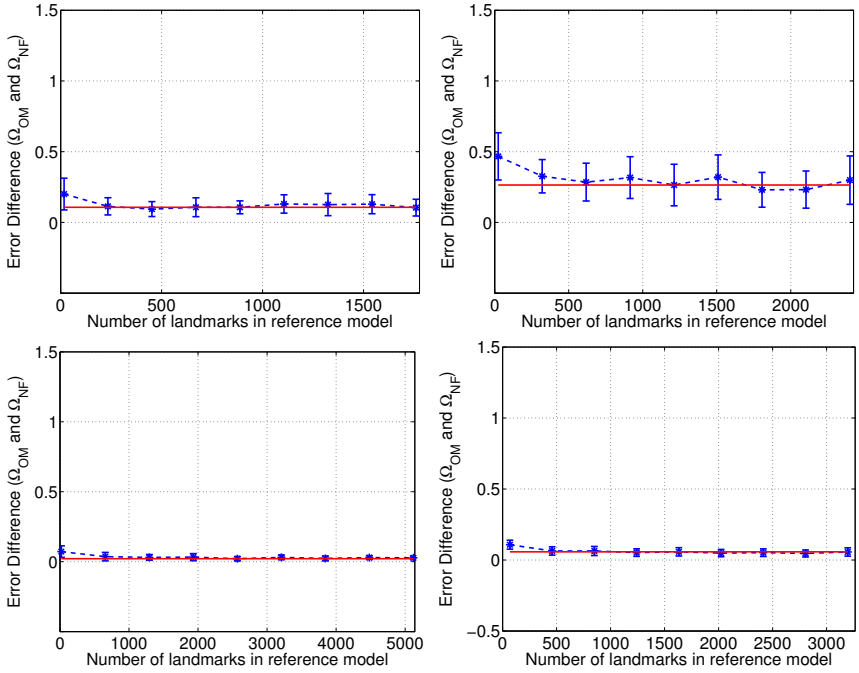

Figure 2.6: Sensitivity to mesh resolution of the reference shape. Showing the distance between reconstructions with the Ω_{OM}, and the Ω_{NM} resampling the reference mesh to nine resolutions. The horizontal continuous line indicates the mean distance obtained for equal reference-target mesh resolution.

varied the proportion of shapes in the reference model from 10 to 90 percent with respect to the second model. The flat curves in Figure 2.7 demonstrate that the proposed technique has the same performance independently of the relative number of training shapes in the input PDMs.

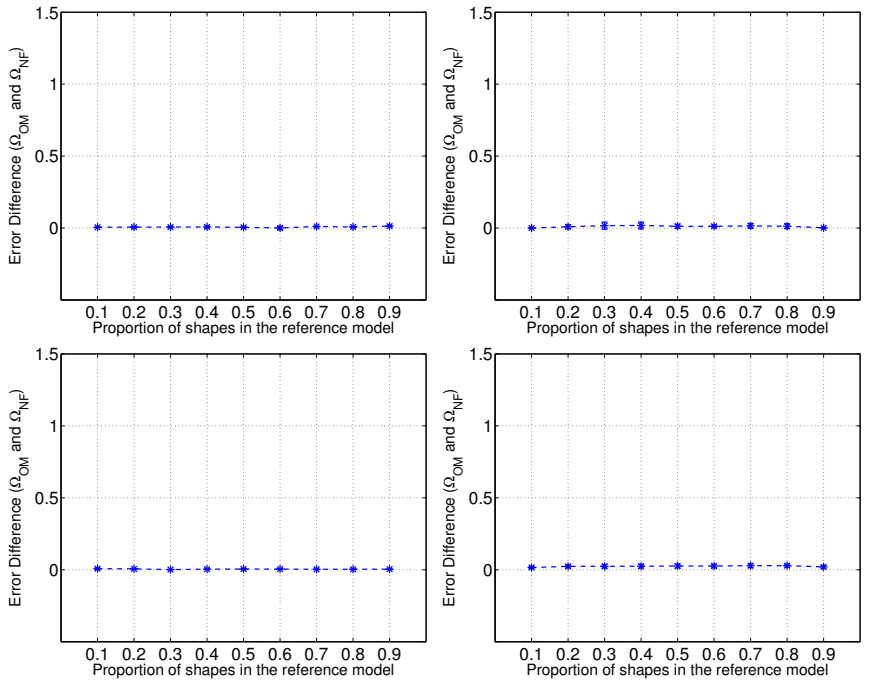

Figure 2.7: Sensitivity to proportion of data samples in the reference model. Showing the distance between reconstructions with the Ω_{OM}, and the Ω_{NM}.

2.4.6 Examples

Examples of the model fusion for all structures are displayed in Figs. 2.8, 2.9, 2.10, 2.11 for the LV, RV, L1-vertebra and caudate nucleus, respectively. Rows 1 and 3 show the first mode of variation ($\pm 1.5\sigma$, and 3.0σ) for two distinct models obtained through random partition of the original dataset. Subsequently, the normalization procedure is applied and the resulting first mode for model 1 can be seen (row 2) to be very similar to the original variability (row 1). Finally, the two transformed PDMs are merged using the proposed algorithm, resulting in a unified model (row 4), which is virtually identical to the PDM obtained with the original datasets (row 5). This is also evident from the sets of eigenvalues of the original and merged models as shown below corresponding Figs. 2.8, 2.9, 2.10, and 2.11.

2.4.7 Segmentation using a merged MR-CT model

To illustrate the benefits and the applicability of PDM fusion in medical image analysis, we apply the proposed technique in the context of image segmentation, which is an important yet challenging application for statistical models of anatomical shape. We would like to automatically segment the 42 cardiac MRI images used in the previous sections by using two distinct PDMs. The first PDM is built directly from the MRI data and is referred to as the MR model. The second model is constructed using the proposed framework by fusing the MR model with a publicly available PDM constructed from a population of CT

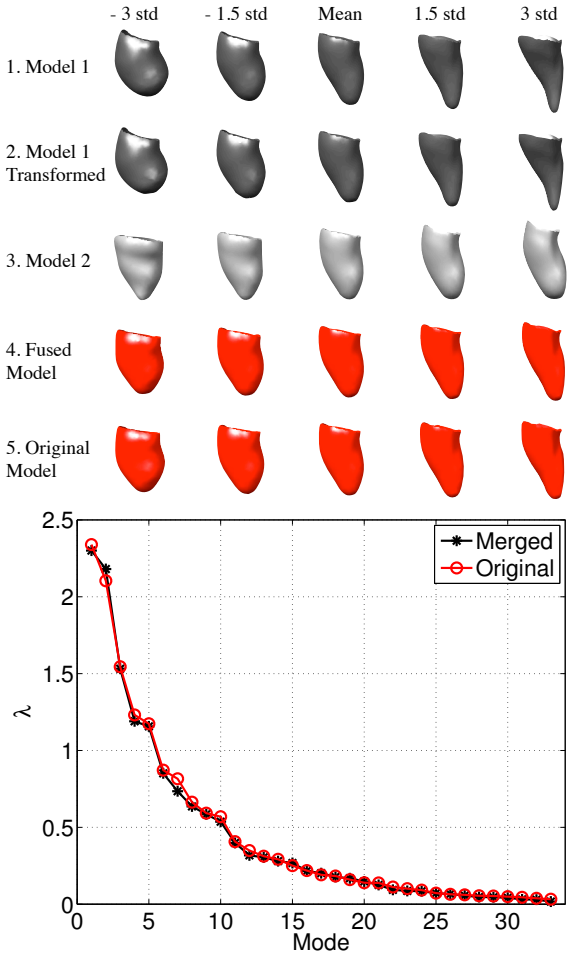

Figure 2.8: (a) shows the first mode of variation of two differently sampled sub-models of the LV. Rows 1 and 2 show model 1 before and after normalization. Row 3 shows model 2, and rows 4 and 5 their integration and the original shapes model. (b) shows the eigenvalues of the merged (Ω_{NM}) and original models (Ω_{OM}).

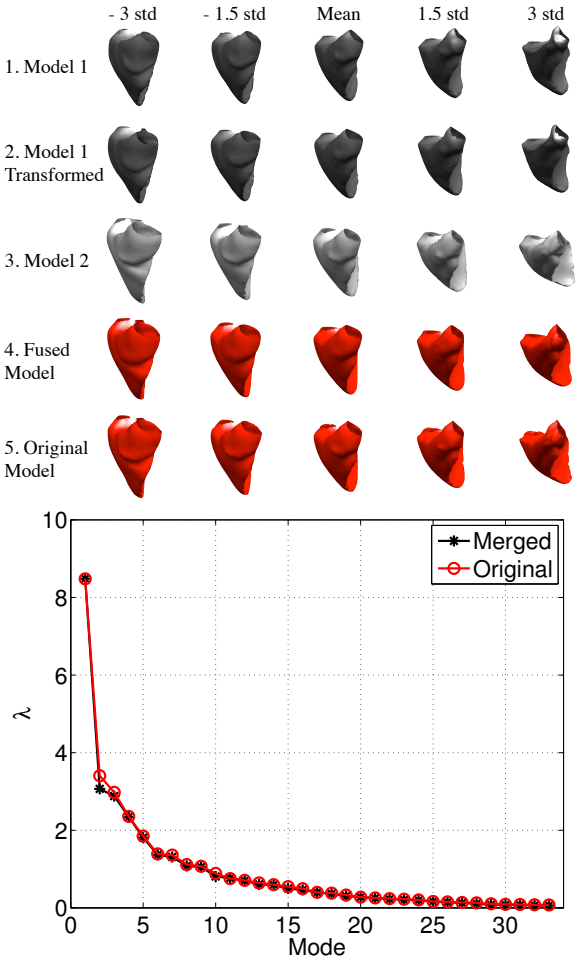

Figure 2.9: (top) shows the first mode of variation of two differently sampled sub-models of the RV. Rows 1 and 2 show model 1 before and after normalization. Row 3 shows model 2, and rows 4 and 5 their integration and the original shapes model. (bottom) shows the eigenvalues of the merged (Ω_{NM}) and the original models (Ω_{OM}).

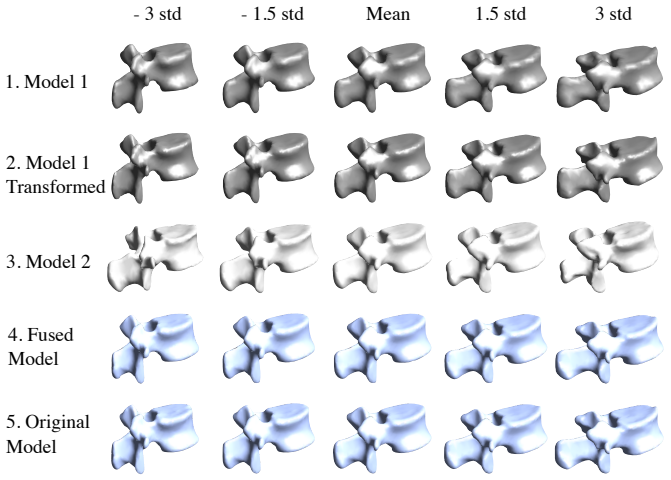

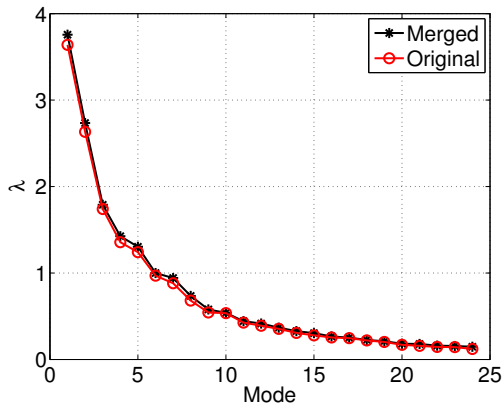

Figure 2.10: (top) shows the first mode of variation of two differently sampled sub-models of the L1-vertebra. Rows 1 and 2 show model 1 before and after normalization. Row 3 shows model 2, and rows 4 and 5 their integration and the original shapes model. (bottom) shows the eigenvalues of the merged (Ω_{NM}) and original models (Ω_{OM}).

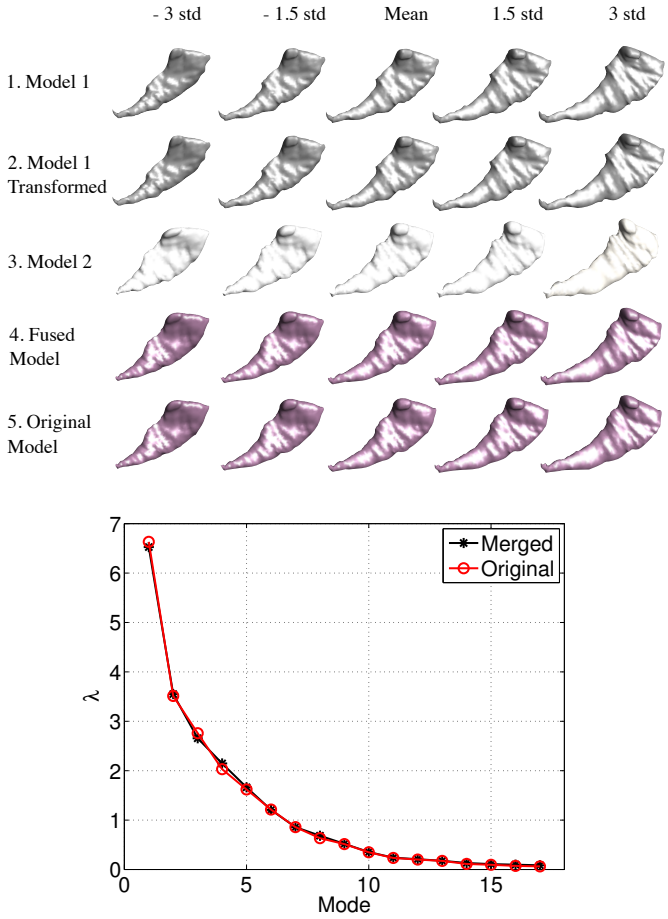

Figure 2.11: (top) shows the first mode of variation of two differently sampled sub-models of the caudate nucleus. Rows 1 and 2 show model 1 before and after normalization. Row 3 shows model 2, and rows 4 and 5 their integration and the original shapes model. (bottom) shows the eigenvalues of the merged (Ω_{NM}) and original models (Ω_{OM}).

datasets[2] [46]. We refer to the fused model as the MR-CT model.

First, we illustrate visually the benefits of the fusion approach in encoding extra variability in Fig. 2.12. The first mode of variation of the MR and CT models is shown in rows 1 and 3 respectively, for the LV (top), and the RV (bottom). It can be seen that the variability captured by the MR and CT models differ since they were constructed from different populations. In particular, more localized variation is encoded by the CT models for both the LV and the RV.

The result of the integration scheme is displayed in Fig. 2.12 row 4 for the LV (top), and the RV (bottom). Visually, it is evident that features of variation found in both the MR and CT models are incorporated into the unified model. In particular, for both the LV and RV, the first mode of variation for the merged model incorporates a pattern describing a transition from the high resolution (higher detail) CT model to the smoother representation found in MRI. The unified model integrates detailed, as well as smoother surface representations, and inter-model differences in resolution can be encoded as patterns of shape variability into the unified model. In the case of the RV, the regional variation at the valvular level is mostly related to the MR model, while the CT model contributes with the global variation in morphology.

We then performed image segmentation on the 42 MRI datasets. We segmented the left and right ventricles (LV-endo, LV-epi and RV) using the Sparse-ASM (SPASM) algorithm in ([60]). The segmentation errors

[2] `https://sites.google.com/site/cornehoogendoorn2013/home/publications/downloads`

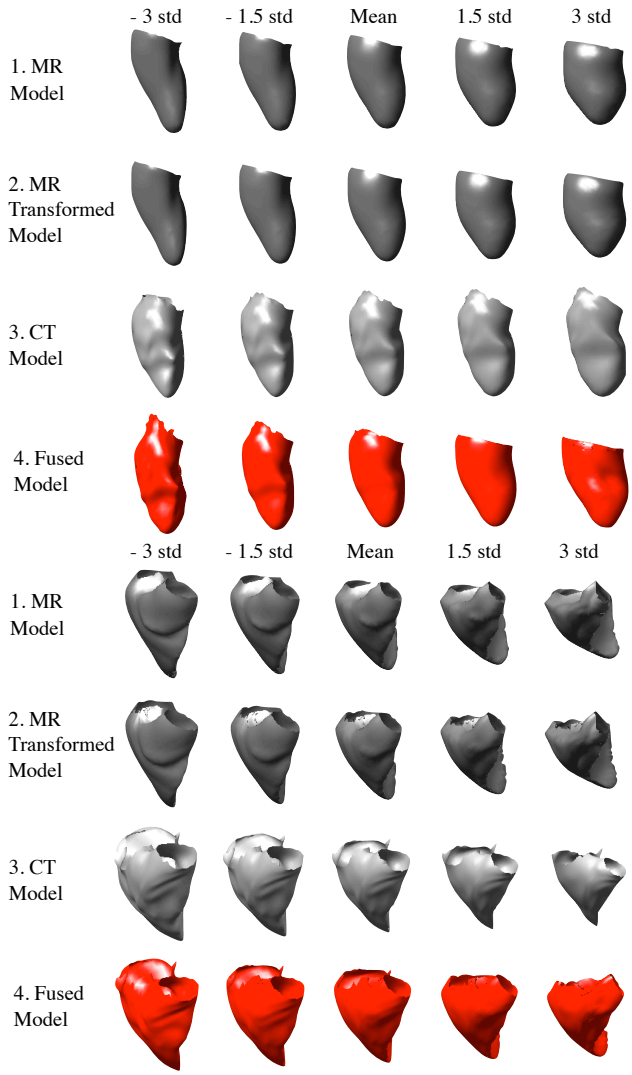

Figure 2.12: First mode of variation for the LV (top) and the RV (bottom), of two models computed from two different image modalities (MR, CT), (rows 2, 3). Their integration (row 4), and the pre-transformed MR model (row 1) i.e. before establishing correspondence.

for all cases are displayed in Fig. 2.13 and summarized in Table 2.2. It is evident from the results that the fused model improves significantly the segmentation results, from an average error of 3.4 mm for the MR-model to 2.3 mm with the merged model, which is equivalent to an improvement of 32%. Such improvement can be obtained simply by downloading an online statistical model such as the one used in this section and without the need for additional raw data that would require costly and tedious collection and processing.

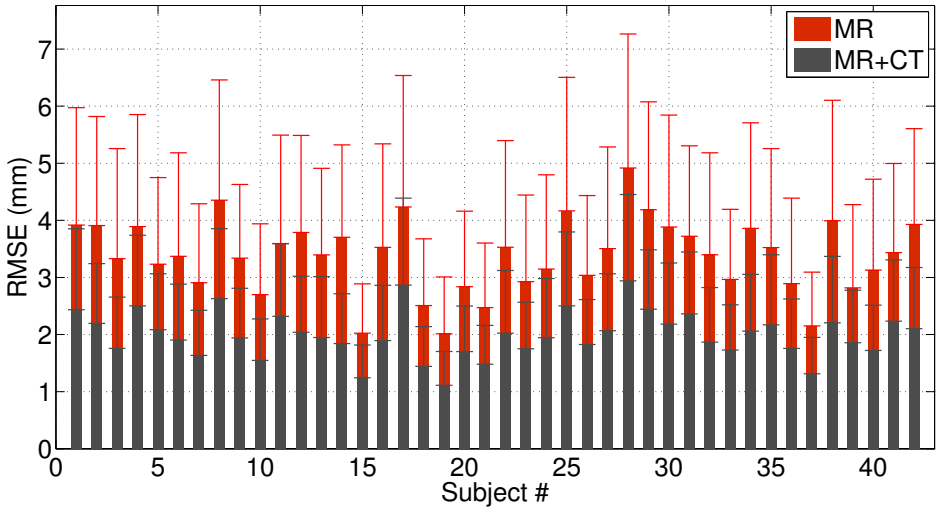

Figure 2.13: Segmentation error and standard deviation, on a per subject basis, of the MR image dataset of a MR-only model (red), and a MR-CT model (dark gray), for the LV epi/endo, and the RV.

Figures 2.14 and 2.15 show two examples of the improvement in generalization ability that results from merging the variability of the

Table 2.2: Segmentation errors (mm) for the MR dataset (LV-endo, LV-epi, RV) comparing the performance of a MR-only model Ω_{MR}, and a Ω_{NM} (MR-CT) model.

model	*mean ± stdev (mm)*	*max (mm)*
MR model	3.4 ± 1.7	4.9
MR-CT model	2.3 ± 1.2	3.4

MR and CT models using our method. Figure 2.14 (left) shows the inability of the MR model to accurately represent the right ventricular wall at the regions indicated by the arrows. In contrast, the MR-CT model is able to correctly represent the contour. Figure 2.15 shows a 3D view example of the segmentation obtained with the MR (left) and MR-CT model (right). Again, the MR model fails to properly adapt to the image, showing the largest errors on the right ventricular surface. For the same areas, the MR-CT model is able to better adapt to the surface.

2.5 Discussion and conclusions

Statistical models of anatomical shapes are routinely built in clinical/research centers in medical imaging. Such models generally encode various types of anatomical variability due to the differences in the imaging modalities and/or clinical populations (e.g., healthy or abnormal subjects). In this chapter we presented a workflow that reuses pre-existing SSMs in a complementary fashion, with the aim to de-

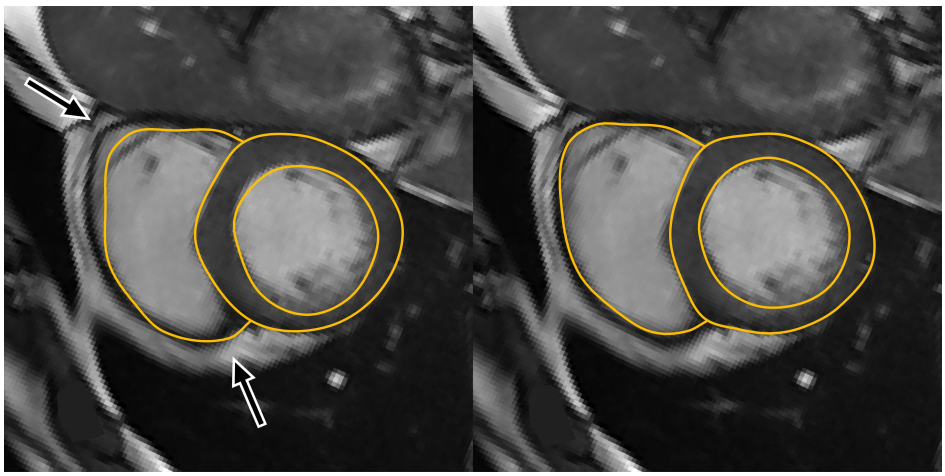

Figure 2.14: Example of the improvement in segmentation accuracy obtained with the combined model. The left image shows the MR-model segmentation, and the arrows indicate the most prominent areas of error. The right image is the MR-CT model.

rive a single integrated model with better coverage of the anatomical variability. This is achieved without the original raw data or for an additional labor-intensive data collection. The framework provides great flexibility due to its ability to handle models with distinct shape representations and resolutions.

The proposed framework can also play an important role in collaborative research. Often, research centers will share or make public their statistical models or atlases but not the original raw data (due to various ethical, legal or practical reasons). With the increasing number of integrated projects and initiatives where information, knowledge,

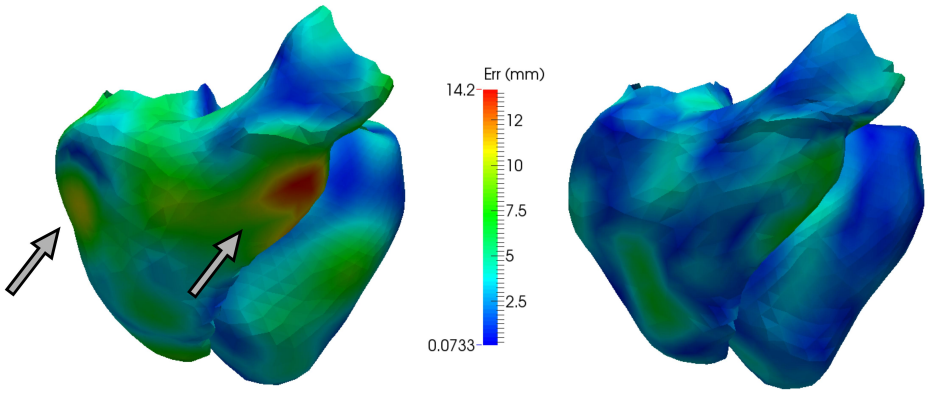

Figure 2.15: Color-coded example of the over-constrained generalization ability of the MR model (left), compared to the combined MR-CT model (right). The arrows indicate two of the areas with highest error.

and models must be merged into unified representations [5], the proposed framework can contribute in promoting such exchanges and build more comprehensive models of anatomy.

As demonstrated in the integrated MRI/CT cardiac shape modeling example, the proposed workflow can be useful to build multi-modality and multi-population SSMs. The generalization ability of the integrated model is effectively increased by adding anatomically plausible variability, without the original or any new training data. As a future work, one can also extend the method to merge spatially resolved models (e.g. from MRI or CT) with temporally resolved models (e.g. 3D US). The proposed work is only the first step in this direction but one that breaks the need for correspondences across eigenspaces.

CHAPTER 3

ACCURATE SEGMENTATION OF VERTEBRAL BODIES AND PROCESSES USING STATISTICAL SHAPE DECOMPOSITION AND CONDITIONAL MODELS

3.1 Motivation

Detailed segmentation of the vertebrae is an important prerequisite in various applications of image-based spine assessment, surgery and biomechanical modeling. Accurate segmentation of the processes is required for image-guided interventions, for example for optimal placement of bone grafts between the transverse

processes. Furthermore, the geometry of the processes is now required in musculoskeletal models due to their interaction with the muscles and ligaments. In this chapter, we present a new method for detailed segmentation of both the vertebral bodies and processes based on statistical shape decomposition and conditional models. The proposed technique is specifically developed with the aim to handle the complex geometry of the processes and the large variability between individuals. The key technical novelty in this work is the introduction of a part-based statistical decomposition of the vertebrae, so the complexity of the subparts is effectively reduced, and model specificity is increased. Then, to maintain the statistical and anatomic coherence of the ensemble, conditional models are used to model the statistical inter-relationships between the different subparts. For shape reconstruction and segmentation, a robust model fitting procedure is used to exclude improbable inter-part relationships in the estimation of the shape parameters. Segmentation results based on a dataset of 30 healthy CT scans and a dataset of 10 pathological scans show a point-to-surface error improvement of 20% and 17% respectively, and the potential of the proposed technique for detailed vertebral modeling.

3.2 Introduction

Segmentation of the vertebrae is an important pre-requisite for several clinical applications, ranging from the assessment of spinal disorders and image-guided surgery [61], [62] to biomechanical modeling for

patient-specific planning of interventions [63], [64]. For such applications, in addition to the segmentation of the vertebral bodies, accurate and detailed knowledge of the vertebral processes is necessary (Fig. 3.1). For spinal fusion surgery, for example, precise delineation of the processes can lead to an improved placement of the bone graft between the transverse processes of the affected vertebrae [65]. In biomechanical modeling of the spine, accurate definition of the spinous process is also critical due its interaction with the ligaments and the muscles [66].

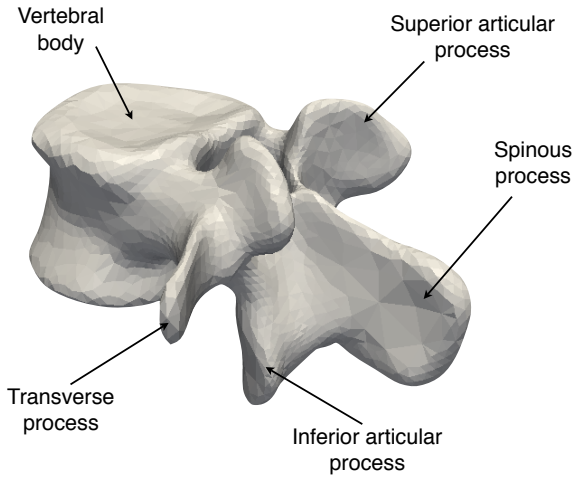

Figure 3.1: A lumbar vertebra and its main regions.

However, automatic and detailed segmentation of vertebrae, and in particular of its processes and pathological cases, has proven to be a difficult task due to the complexity of the shapes and the high variability between individuals. As shown in Figs. 3.1 and 3.2, the

vertebral processes consist of various areas of distinct geometrical but equally complex characteristics, with several convex/concave structures, and thin lobe-like elongated regions. As shown in Fig. 3.2 (bottom), trauma patients with fractured vertebrae present statistically anomalous shapes that present a challenge for straightforward shape modeling. The precise modeling and segmentation of the vertebral processes remains a challenging research topic within spine imaging.

There exists a wide range of approaches in the existing literature for spine and vertebral segmentation [67], [68], [69], [70], [71], [72], [73], [74], [75], [76], [77], [78], [79], [80]. Table 3.1 lists the main papers and reported segmentation results on spine segmentation in the literature. Of these, methods based on the SSM paradigm [67], [68], [70], [79], [78] provide the image segmentation a shape prior to increase the robustness to image inhomogeneities. For example, Kadoury et al. [79] recently developed a method combining global shape appearance and local statistical shape models for each vertebra. Rasoulian et al. [78] combined all the vertebrae into a single shape model, with a statistical pose model. In both cases, these methods consider at best a whole vertebra as the smallest unit to construct the SSM. Due to the large variability of the vertebrae in particular for pathological instances and at the processes, and the generally small number of samples available for training, the obtained models are typically too constraining and not flexible enough to localize the fine details and areas of high curvatures and complexity within the vertebral body and processes, as illustrated in Fig. 3.2.

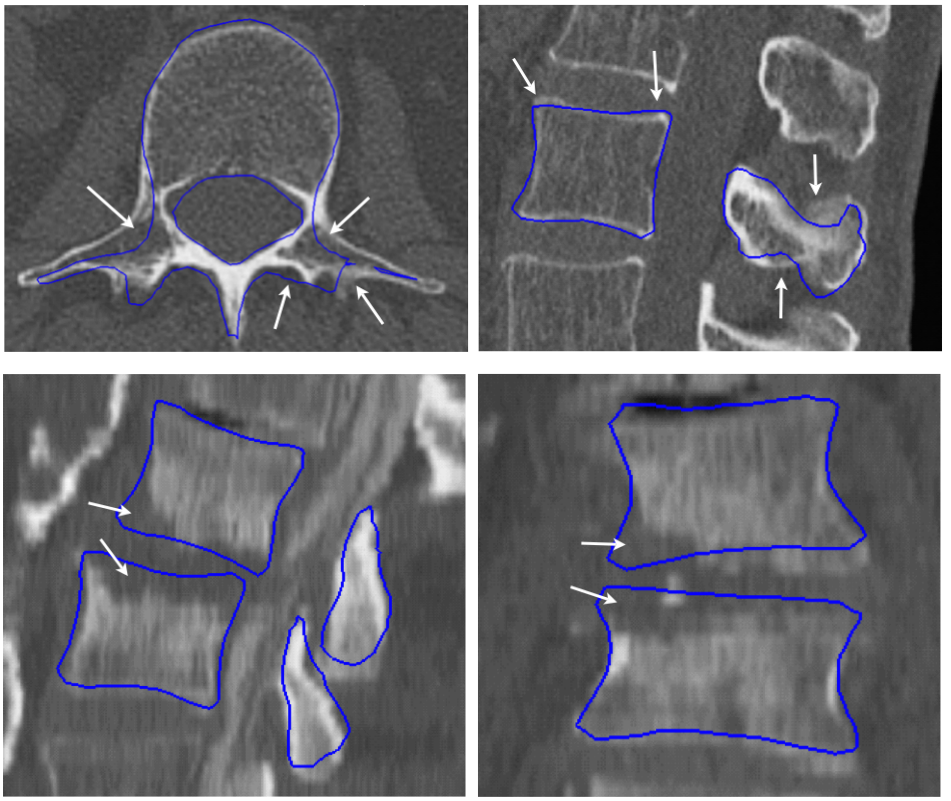

Figure 3.2: Examples of suboptimal segmentations in areas of complex geometry and high curvature (top), and fractured vertebrae (bottom). These segmentations (in blue) were obtained using the image search method [81] described in section 3.3.3.1 and constrained with a whole-vertebra PDM.

We present a new method for detailed modeling and segmentation of the vertebrae based on statistical shape decomposition. Multi-part shape models have been proposed in the literature with the aim to relax the shape constraints in the presence of a few training samples, or extract additional information in the relationship between objects to aid the segmentation process [20], [22], [21], [82]. Other methods have proposed subdivision of the parametric shape-space rather than the shapes themselves in order to better approximate the actual shape distribution of the object class [83], [84]. But such methods have not been applied to the vertebrae as the subdivision of such a complex shape is a non-trivial problem.

We propose an algorithm for statistical decomposition of the vertebra, and for modeling the relationship between the parts. The proposed shape decomposition effectively reduces the complexity of each constituent model, while increasing their specificity. The proposed approach is useful to model difficult regions of the vertebra such as the processes and pathological cases such as fractured vertebrae improving segmentation accuracy. Then to maintain the coherence of the ensemble, conditional models are used to model the statistical inter-relationships between the different subparts. For spine image segmentation, a robust model fitting procedure is then introduced to exclude inconsistent inter-part relationships during the estimation of the shape parameters.

The segmentation accuracy of the proposed technique was tested on two CT scan datasets. One dataset of 30 healthy, and a dataset of 10

pathological cases. Training was performed on the healthy population following a leave-one-out scheme (29 training, 1 testing) to test the healthy cases, and the complete healthy dataset was used to test the pathological cases.

This work is based on a conference paper [85], which we extend by developing a new statistical decomposition of the vertebrae, with more detailed evaluation of the properties of the algorithm and testing segmentation results on both healthy and pathological patients.

Table 3.1: Spine segmentation: State-of-the-Art.

Study	**Region** (C)ervical (T)horacic (L)umbar	**Modality**	**No. data**	**Dims**	**Metric**	**Accuracy**
[67]	T, L	X-ray	57	2, 3	P2S[1]	Lumbar $0.71 \pm 0.06mm$
[68]	T, L	DXA	78	2	P2C[2]	$0.88mm$
[70]	L	X-ray	91	2	P2C[2]	$1.4mm$
[72]	C, T, L	MRI	5	3	CD [3]	See Study.
[71]	L	CT	41	3	DC [4]	0.98
[73]	C, T, L	MRI	22	3	DC [4]	0.96
[74]	L	CT	50	3	FC [5]	See Study.
[75]	C, T, L	CT	64	3	P2S[1]	$1.12 \pm 1.04mm.$
[76]	T	CT	40	3	P2S[1]	$0.95 \pm 0.91mm.$
[79]	T, L	CT, MRI	29	3	P2S[1]	$1.6mm$
[78]	L	CT	32	3	P2P[6]	$< 2mm$

1 = Point-to-Surface, 2 = Point-to-Contour, 3 = Corner detection (%),
4 = Dice Coefficient, 5 = Fence Correctness, 6 = Point-to-Point

3.3 Method

The proposed framework consists of three main stages. First, in section 3.3.1, a subdivision of each vertebra into several subparts is proposed based on a statistically driven region decomposition. Then the conditional models describing the statistical inter relationships between the subparts are presented in section 3.3.2. Finally, a model fitting approach based on all pairwise conditional models is introduced in section 3.3.3.2, with the aim to estimate the shape parameters for each subpart robustly during image segmentation. Figure 3.3 shows a schematic of the method's workflow.

3.3.1 Statistical Vertebral Decomposition

A common problem in the representation of complex and highly variable anatomical objects using SSMs is that usually there are too few available examples from which to obtain a sufficiently detailed representation of the population and its natural variability. With the vertebra, few samples will obtain a gross approximation of the global shape distribution of a given population, but this is often not enough to encode the finer details of the vertebral processes (Fig. 3.2) or to represent instances that deviate far from the mean of the population. In this chapter, we address these issues by developing a statistical part-based decomposition to better encode the statistical variability of each region of the vertebra, and to better generalize to instances not present in the training set. However, such subdivision is not trivial as

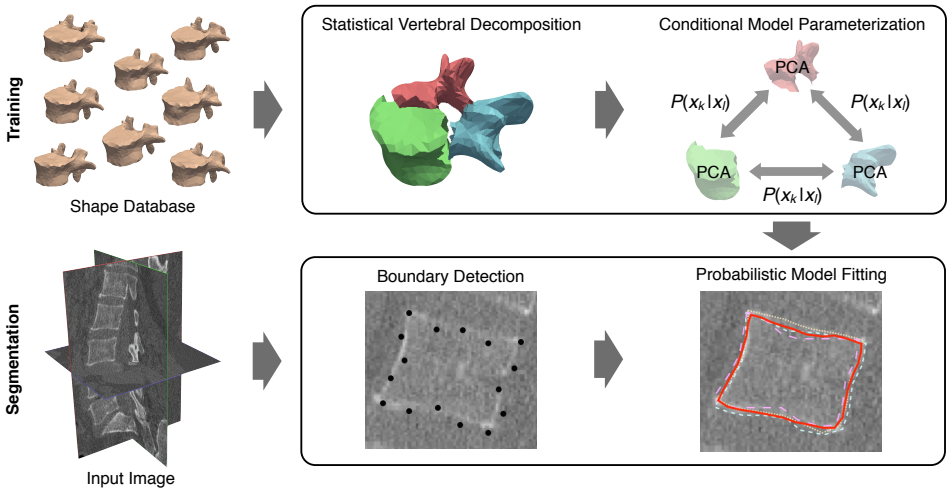

Figure 3.3: The proposed method consists of two training steps (top): 1) vertebral shape subdivision (See Sec. 3.3.1), and 2) construction of conditional models (See Sec. 3.3.2). The segmentation process also has two steps (bottom): 1) initial boundary detection (See Sec. 3.3.3.1), and 2) Multiple model fittings given all conditioning subshapes (See Sec. 3.3.3.2). These model fittings are shown in dotted lines (bottom right). The final segmentation is computed as the median estimation of all model fittings (continuous red line).

the statistical variability and geometrical complexity in the vertebra is uneven.

In this work, the proposed shape subdivision considers the statistical properties of the parts, and therefore provides a statistically coherent subdivision that minimizes bias towards any of its parts. The proposed method ensures that the variability of the whole shape is equitably distributed into a specified number of regions so Point

Distribution Models (PDMs) constructed from these regions encode similar variability.

The algorithm has three parts: 1) seed placement, 2) initial region labeling, 3) statistical region optimization.

3.3.1.1 Seed Placement

To subdivide the shape into the desired number of regions, the user must specify the number of regions K. Let us denote $\mathbf{x}_i = (\mathbf{x}_1 \dots \mathbf{x}_r)^T$, $i = 1 \dots N$ the landmark-based shape representation of each vertebra, where r is the number of landmarks, N is the number of shapes, and $\bar{\mathbf{x}}$ is the mean shape computed from all shapes $\mathbf{x}_i$. The aim is to subdivide $\bar{\mathbf{x}}$ into K sub-parts $\bar{\mathbf{x}}_k$. We use the mean shape $\bar{\mathbf{x}}$ so that the result is not biased toward any one sample. After K is specified an initial seed point is randomly selected from the vector $\bar{\mathbf{x}}$. Then, the remaining $K-1$ points are selected so the Euclidean distance between the k^{th} point and all previously selected seed points is maximized. This strategy ensures that the initial region seeds are uniformly distributed throughout the shape, and helps to minimize computational time during the statistical region optimization (step 3) of the algorithm. See Algorithm 3.

3.3.1.2 Initial Region Labeling

Based on the K seed points, a partition into K regions R_k is obtained. Initially, each region R_k contains only its corresponding seed point, and all points not assigned to any region are said to belong to the $null$ region R_0. We then iterate through the newly initialized regions,

Algorithm 3 Seed placement

Input: Number of regions: K
Input: Mean shape: $\bar{\mathbf{x}}$

▷ Randomly determine initial region seed.
[**point**, $pointIndex$] = getRandomPoint($\bar{\mathbf{x}}$)
seeds(1) = **point**
$\bar{\mathbf{x}}(pointIndex)$ = [] ▷ Delete seed point from $\bar{\mathbf{x}}$.

for $k = 2 \rightarrow K$ ▷ For all regions.
  $maxDist = 0$
  $nPoints$ =getNumberOfPoints($\bar{\mathbf{x}}$)
  $nSeeds$ =getNumberOfPoints(**seeds**)

  for $p = 1 \rightarrow nPoints$ ▷ For all points in the shape.

    for $s = 1 \rightarrow nSeeds$ ▷ For all seed points.
    Comment: Get distance from point p to all seeds.
      dists(s) = getDist($\bar{\mathbf{x}}(p)$, **seeds**(s))
    end for

    $distsSum$ = sum(**dists**) ▷ Sum dists to all seeds.

      ▷ If current point is further from all seeds than previous point, store index.
      if $distsSum > maxDist$
        $maxDist = distsSum$
        $nextSeed = p$
      end if

  end for

  seeds(k) = $\bar{\mathbf{x}}(nextSeed)$ ▷ Store next seed.
  $\bar{\mathbf{x}}(nextSeed)$ = [] ▷ Delete seed point from **x**.
end for

Output: seeds

and for each region we find all points at its boundary using its mesh triangulation. If these boundary points belong to the $null$ region R_0, they are removed from R_0 and assigned to the region R_k. We repeat this process until the $null$ region is empty, and all points in the mesh have been added to some region R_k.

Based on the obtained regions R_k we can now subdivide shapes $\mathbf{x}_i$ into subparts $\mathbf{x}_{i,k}$. From the initialization and initial region label assignment these regions have approximately the same number of points, however, they may vary greatly in terms of their variance across the population, particularly regions at the processes will contain a variance higher than regions on the vertebral body. An optimization is required in order to equalize the variability of all regions so that modeling of the vertebra is not biased by any region.

3.3.1.3 Statistical Region Optimization

The aim of the optimization is to modify the obtained region subdivision so all regions have approximately the same variance across all samples. Note that all computation in the previous two steps was performed on a single shape (the mean shape $\bar{\mathbf{x}}$), however, now we consider the variation across the population. For this purpose we first align all shapes by performing Procrustes analysis [86]. We then define the global covariance matrix

$$\mathbf{C} = \frac{1}{N-1}\sum_{i=1}^{N}(\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T, \tag{3.1}$$

and the total population variance as

$$\mathrm{Var}_{total} = \mathrm{tr}(\mathbf{C}). \tag{3.2}$$

Similarly we define the regional covariance matrices as

$$\mathbf{C}_k = \frac{1}{N-1}\sum_{i=1}^{N}(\mathbf{x}_{i,k} - \bar{\mathbf{x}}_k)(\mathbf{x}_{i,k} - \bar{\mathbf{x}}_k)^T, \tag{3.3}$$

and the regional variances as

$$\mathrm{Var}_k = \mathrm{tr}(\mathbf{C}_k). \tag{3.4}$$

The algorithm first determines the target variance for each region as

$$\mathrm{Var}_{target} = \frac{\mathrm{Var}_{total}}{K}. \tag{3.5}$$

Then, the algorithm iterates through the regions and computes the variance of the current region. If the variance of the current region Var_k is less than the target variance Var_{target}, we reassign all adjacent points to the perimeter of the current region R_k from adjacent regions R_l, $k \neq l$. In Algorithm 4, line 14, we denote this operation `expandPerimeter`. Similarly, If the variance of the current region Var_k is greater than the target variance Var_{target} we reassign all points at the perimeter of the current region to those regions adjacent to it. In Algorithm 4, line 20, we denote this point reassignment with the function name `shrinkPerimeter`.

The algorithm iterates until the standard deviation of all region variances falls below 5% of the target variance, i.e., $\mathrm{std}(\mathrm{Var}_k) < 0.05 \cdot$

Var_{target} (see Algorithm 4), or no further changes in region variances occur. Once the optimization converges the vertebral shape is effectively subdivided into regions of similar variability. For all experimental results reported in the results section (Sec. 3.4) of this chapter, the shape $\mathbf{x}_i$ being tested was removed from the training set, i.e, leave-one-out scheme.

Figure 3.4 shows examples of the convergence of the algorithm for three different shape subdivisions (2, 3 and 4 regions). The figure shows how region variances at iteration 1 are unevenly distributed across the shape and how they converge for an even distribution of the variance. Figure 3.5 shows two examples of the final vertebral decomposition for 2 and 5 regions. Note that the region/s describing the vertebral processes comprise fewer points indicating higher variability, whereas the regions at the vertebral body are larger. Also in the 5-region case more regions are necessary to represent the variability at the processes, whereas only 2 larger regions represent the vertebral body.

Algorithm 4 Statistical region optimization

Input: Region point matrices: $\mathbf{X}(1)\ldots\mathbf{X}(K)$
Input: Region point connectivity: $\mathbf{V}(1)\ldots\mathbf{X}(K)$
Input: Number of regions: K

for $k = 1 \rightarrow K$ ▷ For all regions.
  $\mathbf{C}_k = \text{cov}(\mathbf{X}(k))$ ▷ Covariance of region k.
  $regVar(i) = \text{trace}(\mathbf{C}_k)$ ▷ Get region variance.
end for

$varTotal = \text{sum}(regVar)$ ▷ Sum all region variances.
$varTarget = varTotal/K$ ▷ Get target variance.
$allowedDispersion = 0.05 * varTarget$ ▷ 5% of target variance
$dispersion = \text{std}(regVar)$ ▷ Get region variances' dispersion.

while $dispersion > allowedDispersion$

  for $i = 1 \rightarrow K$ ▷ For all regions.

    if $regVar(i) < varTarget$ ▷ If Var is below target.
      $\mathbf{X}(i) = \text{expandPerimeter}(\mathbf{X}(\text{i}), \mathbf{V}(\text{i}))$ ▷ Enlarge region.

      for $j = 1 \rightarrow K$ ▷ Update region variances.
        $\mathbf{C}_k = \text{cov}(\mathbf{X}(j))$
        $regVar(j) = \text{trace}(\mathbf{C}_k)$
      end for

    else ▷ If region variance is above target.
      $\mathbf{X}(i) = \text{shrinkPerimeter}(\mathbf{X}(\text{i}), \mathbf{V}(\text{i}))$ ▷ Contract region.

      for $j = 1 \rightarrow K$ ▷ Update region variances.
        $\mathbf{C}_k = \text{cov}(\mathbf{X}(j))$
        $regVar(j) = \text{trace}(\mathbf{C}_k)$
      end for

    end if

  end for

  $dispersion = \text{std}(regVar)$ ▷ Update variance dispersion.
end while

Output: Region matrices $\mathbf{X}(1)\ldots\mathbf{X}(K)$ ▷ Modified regions.

Figure 3.4: Illustration of the region variance evolution produced by the statistical vertebral decomposition algorithm. It can be seen that the variability of the subparts converges to a nearly equal value.

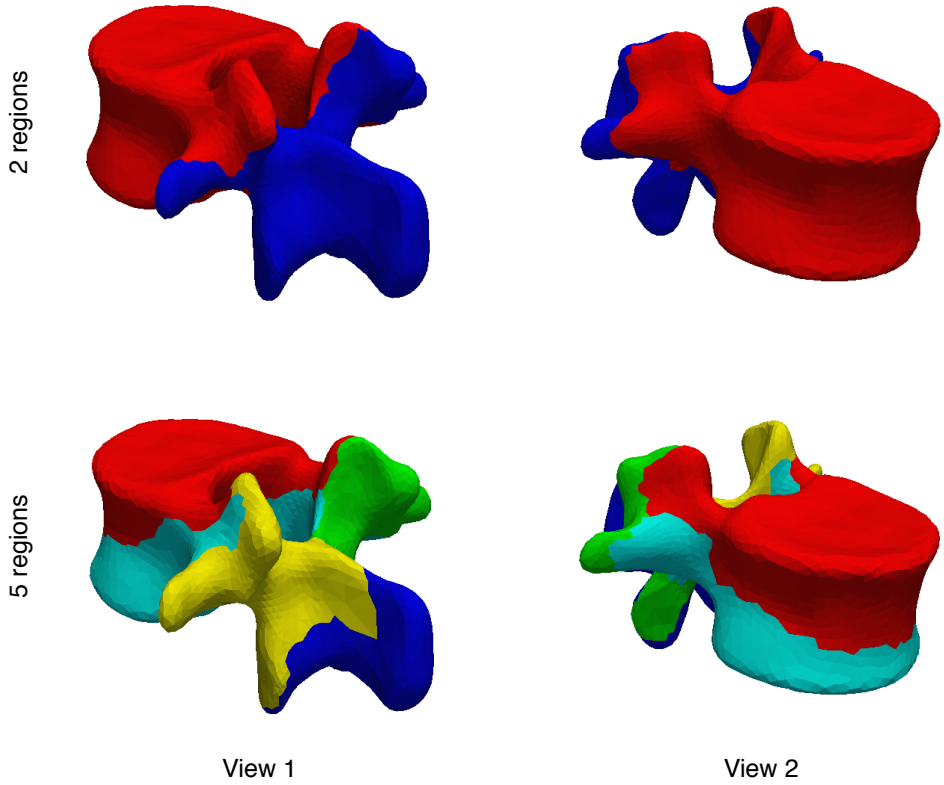

Figure 3.5: Examples of the obtained statistical decomposition with 2 regions (top), and 5 regions (bottom). The left and right columns show a latero-posterior and latero-anterior views. In the 5-region case more regions are necessary to represent the processes due to their high variability.

3.3.2 Conditional Model Parametrization

In the previous section we subdivided the shape of all vertebrae into K subparts $\mathbf{x}_{i,k}$, $k = 1,\ldots,K$. The aim of this section is to describe the statistical modeling of the inter-part probability distributions, i.e. $P(\mathbf{x}_{i,k}|\mathbf{x}_{i,l})$, where $k,l = 1,\ldots,K$ and $k \neq l$. We would like to compute a PDM for each part $\mathbf{x}_{i,k}$ based on its conditional relationship with $\mathbf{x}_{i,l}$, a mean $\bar{\mathbf{x}}_{k|l}$ and covariance matrix $\Sigma_{k|l}$. In this work, we model $P(\mathbf{x}_{i,k}|\mathbf{x}_{i,l})$ using a normal probability distribution. The conditional mean and covariance estimates that relate shapes $\mathbf{x}_{i,k}$, and $\mathbf{x}_{i,l}$ are calculated as

$$\bar{\mathbf{x}}_{k|l} = \bar{\mathbf{x}}_k + \Sigma_{kl}\Sigma_{ll}^{-1}(\mathbf{x}_{i,l} - \bar{\mathbf{x}}_l) \tag{3.6}$$

$$\Sigma_{k|l} = \Sigma_{kk} - \Sigma_{kl}\Sigma_{ll}^{-1}\Sigma_{lk}, \tag{3.7}$$

where the covariance matrices in Eqs. 3.6 and 3.7 are obtained from a block covariance matrix

$$\Sigma = \begin{bmatrix} \Sigma_{kk} & \Sigma_{kl} \\ \Sigma_{kl}^T & \Sigma_{ll} \end{bmatrix}. \tag{3.8}$$

Through eigendecomposition of Eq. 3.7 we obtain Eq. 3.9, which provides the eigenvalues $\Lambda_{k|l}$, and eigenvectors $\Phi_{k|l}$ that represent the conditional variability between shapes $\mathbf{x}_{i,k}$ and $\mathbf{x}_{i,l}$.

$$\Sigma_{k|l} = \Phi_{k|l}\Lambda_{k|l}\Phi_{k|l}^T, \tag{3.9}$$

Putting Eqs. 3.6 and 3.9 together we obtain the new conditional models

$$\Omega_{k|l} = (\bar{\mathbf{x}}_{k|l}, \Phi_{k|l}, \Lambda_{k|l}) \tag{3.10}$$

that we will use to constrain the image segmentation process.

To compute the conditional mean $\bar{\mathbf{x}}_{k|l}$ and covariance matrix $\Sigma_{k|l}$ we must compute the inverse of the covariance matrix of the predictor shape Σ_{ll}^{-1}. However, as the dimensionality of the shapes is much larger than the number of training samples available, the sample covariance matrix becomes singular, and cannot be inverted. A solution to this is using ridge regression [87], where a small constant is added to the diagonal of the covariance matrix $\hat{\Sigma}_{ll} = \Sigma_{ll} + \epsilon\mathbf{I}$, where $\mathbf{I}$ is the identity matrix.

The computational burden of inverting matrices representing several thousands of points can be considerable, especially given that we need to compute a conditional PDM for each pair-wise relationship between the shape subdivisions, and so the number of matrix inversions to compute Eqs. 3.6 and 3.7 grows linearly with K. We address this issue by reducing the dimensionality of the problem using PCA on the sub-shapes $\mathbf{x}_{i,k}$ prior to computation of the mean and covariance matrix as follows [88]. Given subshapes $\mathbf{x}_{i,k}$ represented as,

$$\mathbf{x}_{i,k} = \overline{\mathbf{x}}_k + \Phi_k \mathbf{b}_{i,k}, \tag{3.11}$$

then their parametric representation is

$$\mathbf{b}_{i,k} = \Phi_k^T(\mathbf{x}_{i,k} - \bar{\mathbf{x}}_k). \tag{3.12}$$

Now let us denote $\mathbf{B}_k$ the column-wise concatenation of parametric shape vectors $\mathbf{b}_{i,k}$ from equation 3.12. With this new representation of the shapes, the cross-covariance matrix Σ_{kl} on Eqs. 3.6, 3.7 and 3.8

becomes

$$\Sigma_{kl}^{(b)} = \frac{1}{N-1}\mathbf{B}_k\mathbf{B}_l^T, \tag{3.13}$$

where the superscript (b) in Eq. 3.13 indicates that the covariance matrix is computed from parametric shape vectors $\mathbf{b}_{i,k}$. The block covariance matrix of Eq. 3.8 can now be replaced by

$$\Sigma^{(b)} = \begin{bmatrix} \Lambda_k & \Sigma_{kl}^{(b)} \\ \Sigma_{kl}^{(b)^T} & \Lambda_l \end{bmatrix}, \tag{3.14}$$

where Λ_k and Λ_l are diagonal eigenvalue matrices obtained through eigendecomposition of the original sub-shapes $\mathbf{x}_{i,k}$ and $\mathbf{x}_{i,l}$.

With this new representation Eqs. 3.6 and 3.7 can be rewritten as

$$\bar{\mathbf{x}}_{k|l} = \bar{\mathbf{x}}_k + \Phi_k(\Sigma_{kl}^{(b)}\Lambda_l^{-1}\mathbf{b}_{i,l}) \tag{3.15}$$

$$\Sigma_{k|l} = \Phi_k(\Lambda_k - \Sigma_{kl}^{(b)}\Lambda_l^{-1}\Sigma_{kl}^{(b)^T})\Phi_k^T, \tag{3.16}$$

where the expression in parenthesis on Eq. 3.15 is the regressed parametric shape estimate $\bar{\mathbf{b}}_{k|l}$, i.e.,

$$\bar{\mathbf{b}}_{k|l} = \Sigma_{kl}^{(b)}\Lambda_l^{-1}\mathbf{b}_{i,l}, \tag{3.17}$$

and the expression in parenthesis on Eq. 3.16 is the conditional model variance

$$\Lambda_{k|l} = \Lambda_k - \Sigma_{kl}^{(b)}\Lambda_l^{-1}\Sigma_{kl}^{(b)^T} \tag{3.18}$$

required to obtain the desired model $\Omega_{k|l} = (\bar{\mathbf{x}}_{k|l}, \Phi_k, \Lambda_{k|l})$.

The proposed shape subdivision and conditional model parameterization have two important goals. First, it decreases the over-constraining nature of the global model caused by the dimensionality disparity between the available samples and dimensionality of the shapes. Second, the inter-part conditioning is a mechanism to find the optimal domain of valid subregions and exclude incorrect localized segmentations due to insufficient image information.

3.3.3 Image Segmentation

3.3.3.1 Boundary Detection

To detect the vertebral boundary in the image we followed the feature training method introduced in [81]. Training was performed on the database of 30 healthy patients detailed in section 3.4.1 leaving the test instance out at each trial. The features tested were:

1. Directional derivative along the normal profile pointing outwards.
2. Directional derivative along the normal profile pointing inwards.
3. Maximum intensity profile value.
4. Minimum intensity profile value.
5. Distance to the previous point location at each iteration.
6. Mahalanobis distance to the mean intensity profile.

Each of the previous features was optimized independently within a standard ASM segmentation framework. A greedy optimization heuristic was used where at each iteration a set of uniformly distributed weights $w = [0,1]$ was used for segmentation and the best performing weight was selected based on segmentation accuracy. At the next iteration a new set of weights was uniformly chosen within a neighborhood of the best weight from the previous iteration. This procedure was repeated until no significant improvements in segmentation accuracy were obtained. For more details on the optimization procedure please refer to [81].

After individual optimization of the features, the best-performing three features were chosen and normalized. The selected features were: 1) the directional derivative along the outward-pointing normal, 2) the image intensity, and 3) the distance between the current point candidate and its previous location. The corresponding feature weights were: $w_1 = 0.55$, $w_2 = 0.25$, and $w_3 = 0.20$.

$$E = w_1(\hat{\mathbf{n}}(\mathbf{p}_r) \cdot \nabla \mathbf{I}(\mathbf{p}'_r)) + w_2 \mathbf{I}(\mathbf{p}'_r) + w_3 \|\mathbf{p}'_r - \mathbf{p}_r\|_2. \tag{3.19}$$

Eq. 3.19 shows the energy function maximized during image search, where $\mathbf{I}$ is the intensity image, $\mathbf{p}_r$ is the current position of the landmark r, $\mathbf{p}'_r$ is a vector of candidate landmark positions normally and outwardly oriented regarding the surface mesh at $\mathbf{p}_r$, and $\hat{\mathbf{n}}$ is the normal direction at $\mathbf{p}_r$.

The segmentation process was initialized by rigidly aligning the mean vertebral shape, for the structure (L1-L5), with a manually selected point placed roughly at the center of mass of the vertebral body

on a sagittal view of the image. Then the algorithm determined the optimal placement of each landmark in the mesh based on minimization of Eq. 3.19.

Let us denote the resulting feature points $\mathbf{x}'$. At this point $\mathbf{x}'$ can be subdivided into K statistically optimized regions $\mathbf{x}'_k$, $k = 1 \ldots K$, as described in Sec. 3.3.1. The following section describes the process through which information from all other $K-1$ subregions inform the optimal shape parameters $\mathbf{b}_k$ of $\mathbf{x}'_k$ to maintain anatomical coherence of the ensemble of parts.

3.3.3.2 Probabilistic Model Fitting

To preserve the anatomic validity of the segmentation despite the shape decomposition, the estimation of the shape parameters must be carried out by considering all pairwise conditional probabilities $P(\mathbf{x}_{i,k}|\mathbf{x}_{i,l})$. We first calculate the initial shape parameters $\mathbf{b}_{i,k}$ by projecting the boundary feature points (obtained during image search) onto the standard PDM of $\mathbf{x}_{i,k}$. Then we calculate $K-1$ shape parameters $\mathbf{b}_{i,k|l}$ by considering the $K-1$ shape constraints formed by the conditional mean parameter $\bar{\mathbf{b}}_{k|l}$ and its corresponding bounds $\lambda_{k|l}$ obtained from the diagonal of matrix $\Lambda_{k|l}$, i.e.,

$$\mathbf{b}_{i,k|l} = \begin{cases} \mathbf{b}^0_{i,k} & \text{if } |\mathbf{b}^0_{i,k} - \bar{\mathbf{b}}_{k|l}| \leq 3\sqrt{\lambda_{k|l}} \\ \bar{\mathbf{b}}_{k|l} + 3\sqrt{\lambda_{k|l}} & \text{if } \mathbf{b}^0_{i,k} > \bar{\mathbf{b}}_{k|l} + 3\sqrt{\lambda_{k|l}} \\ \bar{\mathbf{b}}_{k|l} - 3\sqrt{\lambda_{k|l}} & \text{if } \mathbf{b}^0_{i,k} < \bar{\mathbf{b}}_{k|l} - 3\sqrt{\lambda_{k|l}} \end{cases} \tag{3.20}$$

Equation 3.20 is applied to each shape pair $\{\mathbf{x}_{i,k}, \mathbf{x}_{i,l}\}$:

1. For each subpart $\mathbf{x}_{i,k}$ we have a prediction subregion defined by all the points $\mathbf{b}_{i,k}$ defined at less than 3 standard deviations from the conditional mean $\bar{\mathbf{b}}_{k|l}$. Let us denote this subregion the conditional prediction interval $T_{k|l}$ (Algorithm 5, line 15).

2. If $\mathbf{b}^0_{i,k}$ is inside the conditional prediction interval, $\mathbf{b}^0_{i,k} \in T_{k|l}$, we consider the conditional prediction $\mathbf{b}_{i,k|l}$ as the same as $\mathbf{b}^0_{i,k}$. No extra information is provided by the conditional model.

3. In case $\mathbf{b}^0_{i,k}$ is outside $T_{k|l}$, then it is projected to obtain the closest point inside $T_{i,k|l}$. This point is then considered the prediction $\mathbf{b}_{i,k|l}$.

The difficulty with this approach is that at the segmentation stage, all subparts are being segmented and therefore uncertainty surrounds the correctness of the different conditioning shapes $\mathbf{x}_{i,l}$ in $P(\mathbf{x}_{i,k}|\mathbf{x}_{i,l})$. This can lead to inaccurate constraining and parameter estimation of $\mathbf{x}_{i,k|l}$ if some of the $\mathbf{x}_{i,l}$, $l = 1 \ldots K$, $k \neq l$ are erroneous during the segmentation procedure. To exclude these values and obtain a consensual and robust estimation of the shape parameters, we use the marginal median (component-wise median) as the final estimation of $\mathbf{b}_{i,k|l}$, i.e.,

$$\mathbf{b}^{final}_{i,k|l} = \text{median}(\mathbf{b}_{i,k|l}). \tag{3.21}$$

Algorithm 5 presents the step by step model fitting procedure described in this section.

Algorithm 5 Model fitting procedure

Input: Image: $\mathbf{I}$
Input: Mean shape manually initialized: $\bar{\mathbf{x}}$
Input: $\Omega_{k|l} = (\bar{\mathbf{x}}_{k|l}, \Phi_k, \Lambda_{k|l})$
Input: Region indexes: **regIds**

for $r = 1 \rightarrow nLandmarks$
  $\mathbf{x}'_r$ = displaceLandmark($\mathbf{I}, \bar{\mathbf{x}}_r$) ▷ Minimize Eq. 3.19
end for

▷ Subdivide feature-point shape into K regions.
$\mathbf{x}'_k$ = subdivideShapeByRegions($\mathbf{x}'$, **regIds**)

for $k = 1 \rightarrow K$ ▷ For each region being predicted k

▷ Project onto self-PDM.
  $\mathbf{b}'^0_k$ = constrain($\mathbf{x}'_k, \bar{\mathbf{x}}_{k|l}, \Lambda_{k|l}$) ▷ $k = l$. See Eq. 3.12.

  for $l = 1 \rightarrow K$ ▷ For each predictor region l.

    $\mathbf{b}'_l = \Phi_l^T(\mathbf{x}'_l - \bar{\mathbf{x}}_l)$ ▷ Parametric shape estimate $\mathbf{b}'_l$.

    $\bar{\mathbf{b}}_{k|l} = \Sigma_{kl}^{(b)} \Lambda_l^{-1} \mathbf{b}'_l$ ▷ Conditional mean. See Eq. 3.17.

    $T_{k|l}$ = ComputeValidShapeInterval($\bar{\mathbf{b}}_{k|l}, \Lambda_{k|l}$)

      if $\mathbf{b}'^0_k$ is within interval $T_{k|l}$

        $\mathbf{b}_{k|l}(l) = \mathbf{b}'^0_k$ ▷ No conditional information.

      else

        $\mathbf{b}'_k$ = constrain($\mathbf{b}'^0_k, \bar{\mathbf{b}}_{k|l}, \Lambda_{k|l}$) ▷ $k \neq l$. See Eq. 3.12.

        $\mathbf{b}_{k|l}(l) = \mathbf{b}'_k$
      end if

  end for

  $\mathbf{b}^{final}_{k|l}(k) = \text{median}(\mathbf{b}_{k|l})$

  $\mathbf{x}^{final}_{k|l}(k) = \bar{\mathbf{x}}_k + \Phi_k \mathbf{b}^{final}_{k|l}$

end for

Output: $\mathbf{x}^{final}_{k|l}$

3.4 Results

3.4.1 Data

We first trained and validated our method using a database of lumbar spine (L1-L5) CT images of 30 healthy patients reporting lower back pain. The images were collected at the National Center for Spinal Disorders (Budapest, Hungary). The data were acquired with a Hitachi Presto CT scanner. No contrast agent was administered to the patients. The volumes have an in-plane resolution of $0.608 \times 0.608mm$ and slice spacing of $0.62mm$. Patients were 13 males and 17 females with a mean age of 40 (age interval: 27-62 years). Those patients were selected for participating in the European Commission funded MySpine project (FP7-ICT-2009-6-269909) [89].

To assess the strength of the proposed technique in the presence of abnormalities, a second set of 10 scans were obtained to evaluate segmentation. The images were obtained from a publicly available database [90] of CT scans of adult patients with varying types of pathologies including pathological curvature (scoliotic and kyphotic), and fractured vertebrae. The data were acquired at the Department of Radiology, University of Washington, Seattle, USA. The images were acquired with General Electric multidetector CT scanners and a standard bone algorithm. For our purposes, 10 image volumes containing full lumbar spines (L1-L5) were randomly selected and manually segmented. The images have varying in-plane resolution between $0.31mm$ and $0.41mm$, and a slice spacing of $2.5mm$. Patients were 5 males and

5 females with a mean age of 41 (age interval: 16-61 years).

All computer processes were run on a 2.8 GHz Intel i7 processor on 6 GB DDR memory running a single-threaded CPU bound process. All PDMs were trained on the healthy patient database. For the segmentation of healthy patients we followed a leave-one-out scheme. To segment pathological patients we used all 30 healthy patients for training. All segmentations were performed by preserving 98% of the model's total variance, and allowing ±3 standard deviations from the mean. The volumes were manually segmented by an image expert using open source software (ITK-SNAP). Accuracy was measured as the RMS point-to-surface distance between manual segmentations and reconstructions.

3.4.2 Optimal Number of Subparts

The choice of the number of shape subdivisions using the proposed statistical decomposition is important to obtain the best possible segmentations of the spine. A few of subparts might not allow to decompose sufficiently the shape constraints and to adapt to all the regions of the vertebrae. The model fitting stage as introduced in section 3.3.3.2 requires enough subparts to allow a suitable probabilistic weighting of the multiple conditional models and to eliminate potentially incorrect local segmentations.

Many subparts (the extreme case being the modeling of each single landmark as one subpart) might lead to constraints too weak to adequately guide the image segmentation process, i.e., so it achieves

robustness to image inhomogeneities.

In this section, we perform a sensitivity experiment on the healthy datasets, through which we apply the proposed statistical decomposition with a varying number of subparts (from 2 to 20). We then apply the segmentation technique based on the derived conditional models and we estimate the segmentation accuracy for each subdivision. The obtained results in Fig. 3.6 show that the segmentation errors decrease after two subparts, then stabilize between $k = 5$ through 17 subparts, and then rise again after 17 subparts, indicating that the number of subparts becomes too high to allow adequate constraining of the segmentation procedure.

The optimal results are obtained for $k = 15$ and we use this decomposition for the remainder of the validation.

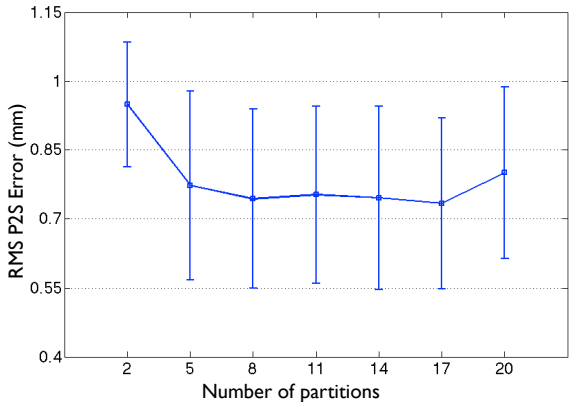

Figure 3.6: point-to-surface segmentation error as a function of number of regions in the decomposition.

3.4.3 Segmentation Accuracy - Healthy Population

We evaluated the performance of our algorithm on the healthy datasets described in section 3.4.1. We performed segmentation on the 30 subjects leaving-one-out both using a whole-vertebra PDM, and our technique.

Fig. 3.7 shows the segmentation errors for all 30 scans using both ASM methods. The proposed technique outperforms the whole-vertebra model ASMs in all cases. The median improvement is of 20% and sometimes the improvement is over 30% due to the ability of the proposed technique to better encode the fine details of the vertebrae.

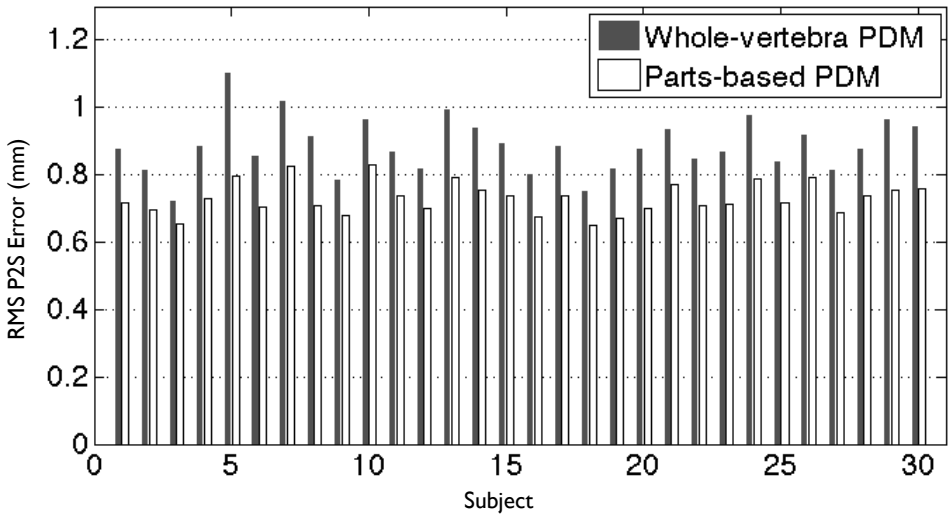

Figure 3.7: Normal population: point-to-surface segmentation error comparison between the proposed method and the whole-vertebra ASM.

Table 3.2 summarizes the segmentation results for the whole-vertebra ASM, and the proposed technique for the different lumbar vertebrae (L1 to L5). It can be seen that performing the proposed technique is consistently better for the entire lumbar spine. Particularly, the right-most column of the table shows an average improvement of 47% in the dispersion of the error.

Table 3.2: Normal population: Image segmentation errors (mm) comparing the performance of our technique against a whole-vertebra PDM segmentation. Errors are shown individually for each lumbar vertebra.

	Whole-vertebra PDM (mm)				Proposed technique (mm)				Improvement (%)			
	Mean	± Std	Max	Min	Mean	± Std	Max	Min	Mean	± Std	Max	Min
L1	0.83	0.09	1.12	0.66	0.67	0.06	0.79	0.61	19%	33%	29%	8%
L2	0.84	0.07	1.12	0.64	0.68	0.06	0.83	0.60	19%	14%	26%	6%
L3	0.88	0.09	1.17	0.68	0.71	0.08	0.95	0.60	20%	11%	19%	11%
L4	0.90	0.08	1.15	0.66	0.73	0.06	0.83	0.64	18%	25%	28%	3%
L5	1.02	0.11	1.32	0.82	0.80	0.06	0.89	0.72	22%	45%	33%	12%
All	0.89	0.09	1.18	0.70	0.72	0.07	0.86	0.63	20%	26%	27%	8%

Figure 3.8 compares the average segmentation error distribution for the whole-vertebra PDM segmentation, and the proposed technique. It can be seen that the errors introduced locally by using a whole-vertebra model are corrected by the proposed approach. In both views our errors are consistently low in all regions of the vertebra. Also the major improvements in accuracy stem from improved fitting of high curvature regions at all processes i.e., spinous, transverse and articular.

Figure 3.8: Normal population: Comparison of the mean error distribution between the whole-vertebra PDM segmentation and our part-based PDM segmentation. Showing lateral and posterior views from left to right, respectively.

Some examples are given in Fig. 3.9 to visually illustrate the strength of the proposed technique. The axial views in columns (a) and (b) show how the segmentation using a whole-vertebra PDM (column (a)) can be typically affected in various areas due to the geometrical complexity and high variability involved, as shown by the arrows. The proposed technique (column (b)), due to its use of decomposed statistical constraints adapts better to the areas of high geometrical complexity.

Columns (c) and (d) show a sagittal view of the segmentations. It can be seen that both the whole-vertebra PDM and the proposed technique have a similar performance for the main body of the vertebrae, as it is geometrically less complex. The whole-vertebra PDM introduces significant errors in the regions of the processes due their more complex nature as indicated by the arrows. These errors are corrected by the proposed technique.

3.4.4 Segmentation Accuracy - Pathological Population

In this section we test whether the technique also improves the segmentation of pathological cases. We train the statistical models on the population of healthy patients as it is the larger sample providing more class specific variability. Fig. 3.10 shows the segmentation errors for the 10 pathological scans using both the whole-vertebra ASM and our method. The plot shows that the proposed technique outperforms the whole-vertebra model ASM for all cases. The average improvement

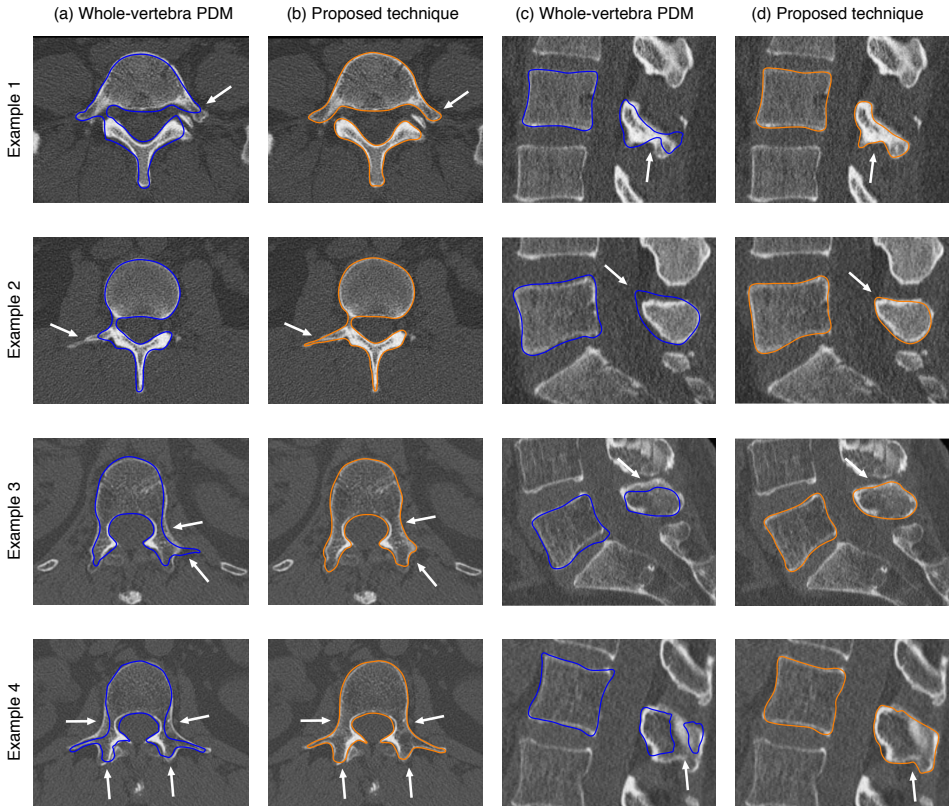

Figure 3.9: Normal population: Four examples (rows) of the improvement in segmentation accuracy obtained with the proposed technique (columns (b) and (d) in yellow), compared against the results obtained with a whole-vertebra PDM (columns (a) and (c) in blue). Showing axial and sagittal views from left to right of vertebrae L3, L5, L5 and L3 (from top to bottom) of 4 patients.

is of 17% with the largest improvement at 32% for patient 2, and least improvement at 8% for patient 4.

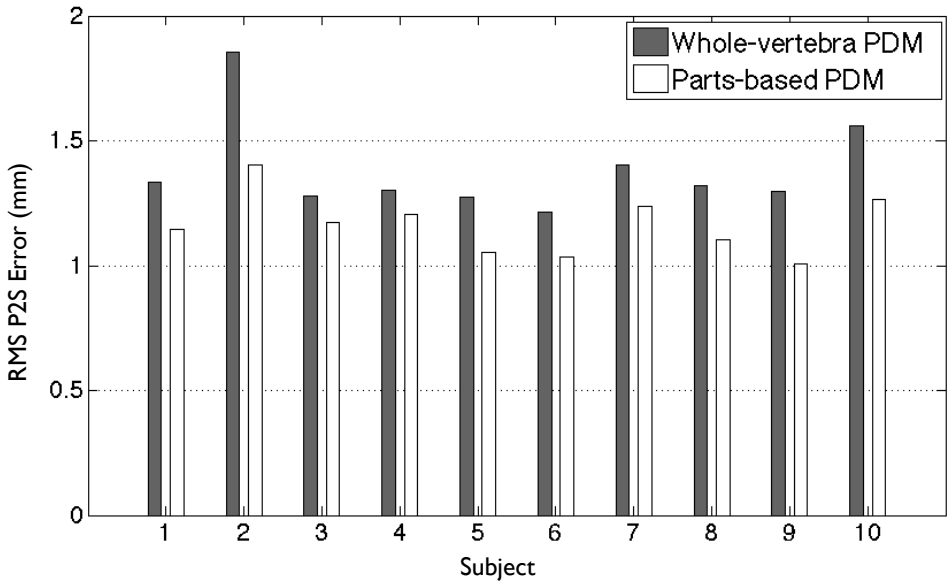

Figure 3.10: Abnormal population: point-to-surface segmentation error comparison between the proposed method and the whole-vertebra ASM.

It is apparent from the magnitude of the errors that both the whole-vertebra ASM and our technique perform worse on pathological cases as compared to the healthy patients. The error is increased by an average of 36% compared to the healthy patient population. But it can be seen that the relative improvement of the proposed algorithm remains constant. All of the pathological images have a lower resolution along the z-axis which may account for the increased

errors, nonetheless, these images also have higher in-plane resolution compared to the healthy images.

Table 3.3 summarizes the segmentation results for the whole-vertebra ASM, and the proposed technique for all lumbar vertebrae (L1 to L5). It can be seen that performing the proposed technique is consistently better for the entire lumbar spine and particularly for the L3 at 19%. On average, the maximum error was also reduced by 19%, with the highest error reduction for the L2 and L3 at 20%.

Table 3.3: Abnormal population: Image segmentation errors (mm) comparing the performance of our technique against a whole-vertebra PDM segmentation. Errors are shown individually for each lumbar vertebra.

	Whole-vertebra PDM (mm)				Proposed technique (mm)				Improvement (%)			
	Mean	± Std	Max	Min	Mean	± Std	Max	Min	Mean	± Std	Max	Min
L1	1.29	0.09	1.41	1.14	1.09	0.05	1.15	0.95	16%	45%	19%	17%
L2	1.37	0.24	2.01	1.17	1.14	0.17	1.62	1.01	17%	30%	20%	14%
L3	1.42	0.13	1.56	1.18	1.16	0.06	1.25	0.99	19%	54%	20%	17%
L4	1.37	0.19	1.81	1.20	1.16	0.13	1.49	1.01	16%	32%	18%	16%
L5	1.47	0.15	1.68	1.27	1.22	0.09	1.41	1.09	18%	40%	17%	15%
All	1.38	0.16	1.69	1.19	1.15	0.10	1.38	1.01	17%	40%	19%	16%

Two examples are shown in Fig. 3.11 comparing the performance of the two techniques. Both examples are of patients with at least one fractured vertebra, and both display a coronal view (columns (a) and (b)), and a sagittal view (columns (c) and (d)) of the same patient. In Example 1 (top), two vertebrae (L2 and L3) are fractured (see white arrows), and it can be seen that the proposed algorithm (columns (b) and (d)) has the flexibility to represent the fracture for both L2 and L3, whereas the whole-vertebra PDM fails to adapt to the contour.

Similarly, Fig. 3.11 Example 2 shows a fractured L3 vertebra (see white arrows). Columns (a) and (c) show how the whole-vertebra PDM does not have the flexibility to adapt to the pathological contour while our method (columns (b) and (d)) is able to represent the pathology.

Figure 3.12 compares the average segmentation error distribution for the whole-vertebra PDM segmentation, and the proposed technique. It can be seen that local errors introduced by the whole-vertebra model on high curvature regions of the transverse and spinous processes are corrected by the proposed approach. In both sagittal and coronal views errors are consistently lower using our method. Major improvements are observed on the L2 and L3 vertebrae where some pathological cases present fractures. In this cases, our method is able to represent the abnormal anatomies.

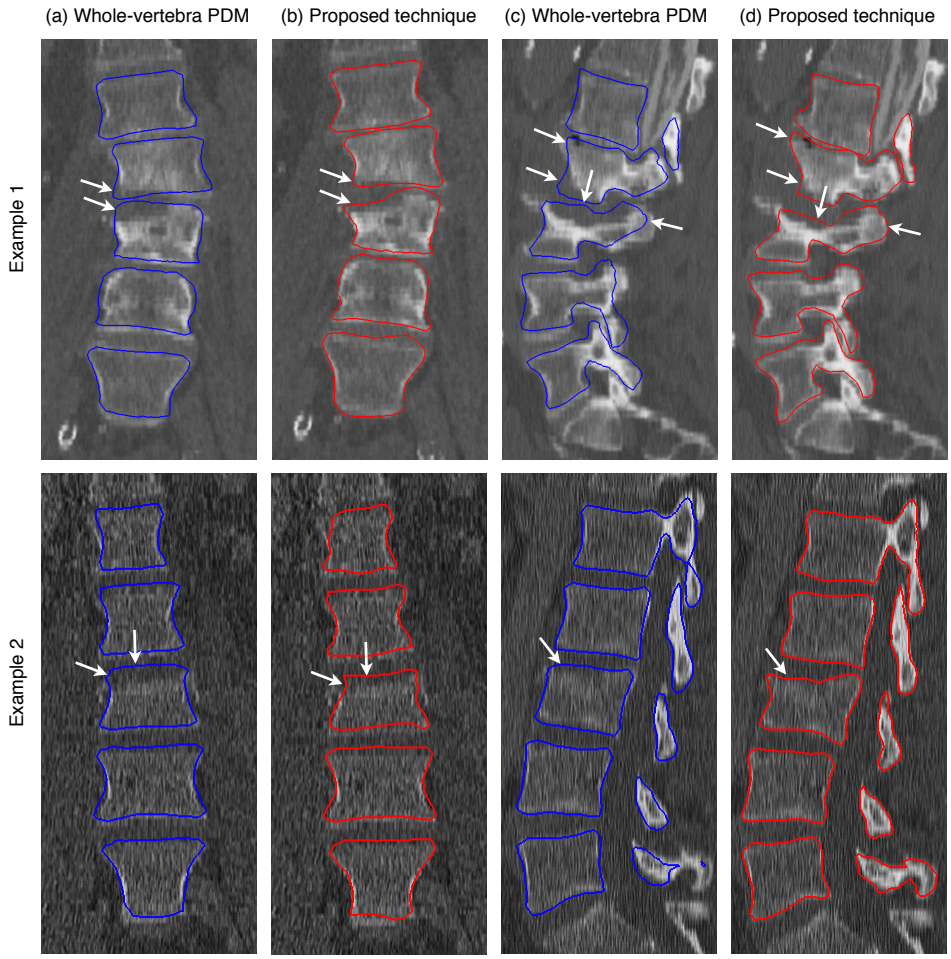

Figure 3.11: Abnormal population: Two examples (top and bottom) of the improvement in segmentation accuracy obtained with the proposed technique (columns (b) and (d) in red), compared against the results obtained with a whole-vertebra PDM (columns (a) and (c) in blue). Showing coronal and sagittal views from left to right.

Figure 3.12: Abnormal population: Comparison of the mean error distribution between the whole-vertebra PDM reconstruction and our part-based PDM reconstruction. Showing lateral and anterior views from left to right, respectively.

3.5 Conclusion

Detailed segmentation of the spine is challenged by the geometrical complexity and the high variability of the vertebrae particularly at the processes and in pathological cases. To address this issue, in this chapter we presented a novel solution based on a statistical part-based decomposition of the vertebral shape, so the total variance of the sample population is evenly distributed among the different subregions. Conditional inter-part models are then constructed to maintain the statistical coherence of the ensemble of shapes during the model fitting procedure at the time of image segmentation. In addition, a probabilistic model fitting approach is introduced to robustly select the most likely shape parameters of each subregion.

The obtained segmentation results indicate that our approach can provide highly accurate and consistent segmentations throughout different individuals and regions of the vertebra. The segmentation adapts well to areas geometrically complex or highly curved such as the vertebral processes, and to abnormalities such as fractured vertebrae. We also show that the proposed method outperforms the segmentation accuracy obtained with a whole-vertebra PDM.

Finally, the proposed approach, while validated with vertebral segmentation, is generic and can be applied for image segmentation of other complex or multi-part structures, such as, multi-chamber heart, simultaneous segmentation of multiple brain structures, and complex vascular structures among others.

3.5.1 Limitations

One limitation of the proposed technique is related to the computational costs it requires. At training, the statistical decomposition due to its iterative nature is computationally demanding. For 15 subdivisions, for example, the algorithm converges after 6.8 minutes. However, this stage is performed once offline and the conditional models are then saved for image segmentation. At the segmentation stage, the proposed technique is more expensive than the whole-vertebra model segmentation, taking 32 seconds on average per patient (L1 through L5), due to using multiple conditional shape models, while the whole-vertebra model segmentation takes 18 seconds to complete the same patient.

A potential challenge of the decomposition approach is that it can theoretically produce discontinuities between two adjacent subparts of the vertebrae. This is because we only use implicit constraints based on the conditional models to obtain a statistical coherence of the ensemble, but this does not guarantee explicitly smoothness between adjacent subregions. In our experimental results, however, we found these implicit constraints, with the probabilistic model fitting approach, have a good performance at maintaining the smoothness of the vertebrae, as illustrated in the examples of Figs. 3.9, and 3.11. One could consider adding an extra step after the segmentation, where discontinuous transitions between subregions are identified and corrected using smoothing.

3.5.2 Future Work

One interesting avenue for future work would be the simultaneous segmentation of the complete lumbar structure, or larger vertebral groups that include thoracic and/or cervical vertebrae. As it is presented, our algorithm is sequential i.e., it is applied to one vertebra at a time, and all the statistical constraints are within-vertebra constraints. However, we would like to extend the probabilistic framework to include interaction between adjacent vertebral structures assuming shape correlations between them. Considering inter-vertebral relationships would not only account for a more holistic model, but it would also aid in the automatic initialization of most vertebrae reducing the human interaction for segmentation.

A natural consequence of considering interactions between neighboring structures is then the development of a system of weights to balance the influence of the inter vs. the intra probabilistic relationships between the different objects and their constituent subparts. These weights can be shape correlation-driven or image-driven. We expect to continue work in this direction.

CHAPTER 4

PATIENT METADATA-CONSTRAINED SHAPE MODELS FOR CARDIAC IMAGE SEGMENTATION

4.1 Motivation

Patient metadata such as demographic information and cardiovascular disease (CVD) risk factors, genetics and other indicators are valuable data readily available about subjects participating in population imaging studies. This information can inform the construction of customized SSMs fitting the patient's unique characteristics. However, to the best of our knowledge, no studies have reported

using these types of metadata in constructing shape models for image segmentation. In this thesis, we propose the use of a conditional model framework to include these patient metadata in constructing a personalized shape model and evaluate its effect on image segmentation. Our validation on a dataset of 250 asymptomatic cardiac MR images shows an average segmentation improvement of 7% and sometimes up to 30% over a conventional PCA-based framework. These results show the potential of our technique for improved shape analysis.

4.2 Introduction

Cardiac segmentation is a prerequisite for several important clinical applications ranging from the relatively simple computation of ejection fraction (EF), to understanding disease progression through shape analysis in longitudinal studies, to more complex tasks, such as, simulating cardiovascular function and electrophysiology for treatment planning and intervention. These tasks require accurate segmentation of the cardiac structure in order to obtain reliable and meaningful outputs.

Segmentation of cardiac structures, however, remains a challenging task due to the high geometric complexity of the organ, and image inhomogeneities. To overcome these problems, statistical shape model-based methods have been widely adopted due to their ability to simplify shape complexity and overcome noisy image information through the model fitting process. But these model-based techniques have tradi-

tionally only focused on the use of shape information for constructing models and do not include other potentially useful non-image derived information typically found on a patient's clinical chart.

In our review of the literature we found a study by Wolz et al. [91] that reports using both patient and image metadata to enhance their manifold learning technique for brain image classification. In another study, Blanc et al. [92] use what they call surrogate variables of the femur bone i.e., different anatomical lengths and angles within the bone and patient information, to further constrain the shape space of a PCA derived model and study their influence. Grbić et al. [93] use landmark derived features to build a patient specific model with a reduced set of shapes for image segmentation. Finally, Medrano-Gracia et al. [94] quantify and show significant morphological differences in asymptomatic cardiac shape between different demographic and risk factor sub-groups. They do this by performing PCA on the sub-groups, and evaluating the statistical significance of differences between their principal modes of variation. Although these works address including other non-image data in the analysis of shape, none use patient-metadata for cardiac segmentation, which is the focus of this chapter.

In this chapter, we propose the use of a conditional shape model framework [95] for the definition of a shape distribution constrained by demographic data and CVD indicators and apply it to cardiac image segmentation. We validate our framework by comparing the segmentation accuracy of our method with that of a standard PCA-based model.

The remainder of this chapter is organized as follows. In section 4.3, we describe the details regarding model construction and segmentation. In section 4.4, we present the obtained segmentation results and comparison to the standard PCA-based model. Finally, conclusions are drawn in section 4.5.

4.3 Method

This section describes the statistical modeling of the probability distribution of shape conditioned on metadata, i.e. $P(\mathbf{x}|\mathbf{m}_j)$. We would like to compute a PDM for shapes $\mathbf{x}_i$ based on their conditional relationship with each metadata field $\mathbf{m}_j$, a mean $\bar{\mathbf{x}}_{\mathbf{m}_j}$, and covariance matrix $\Sigma_{\mathbf{x}|\mathbf{m}_j}$. Then we combine these conditional PDMs into one unified PDM that accounts for all shape-metadata relationships. Last, we describe the image feature search and model fitting process.

Our method consists of four main steps:

1. **Metadata-constrained allowable domain**: Given a dataset of shape vectors $\mathbf{x}_i$, $i = 1,...,N$ obtained from segmented images, and their corresponding metadata fields $\mathbf{m}_{i,j}$, $j = 1,...,M$, we compute a set of eigenspaces $E_j = \{\Phi, \Lambda_j\}$ that describe the allowable shape space for shapes $\mathbf{x}_i$ conditioned on metadata $\mathbf{m}_{i,j}$. To train these models, we use the definitions of multivariate conditional distributions described in section 4.3.1.

2. **Metadata-constrained mean shape**: Based on a new patient's metadata $\mathbf{m}_j^{new}$, we compute M new mean shape esti-

mates $\bar{\mathbf{x}}_{\mathbf{m}_j}$, and combine them to obtain the final model estimate $\bar{\mathbf{x}}_c$. (see sec. 4.3.2).

3. **Metadata constraints combination**: The final allowable domain Λ_c is computed as the intersecting space from the different metadata-constrained variances as estimated in step 1, centered on their corresponding mean estimates $\bar{\mathbf{x}}_{\mathbf{m}_j}$ computed on step 2. (see sec. 4.3.3).

4. **Application to image search**: We use the obtained metadata constrained PDM to guide the image segmentation of a new patient's image volume. We use the Sparse Active Shape Models (SPASM) framework [96] to perform the cardiac image segmentations reported in this chapter. (see sec. 4.3.4).

4.3.1 Metadata-Constrained Allowable Domain

Let us define an augmented shape data matrix $\mathbf{D}$ by appending metadata values $\mathbf{m}_{i,j}$ to the last index of shape vectors $\mathbf{x}_i$ (Eq. 4.1).

$$\mathbf{D} = \left[\left\{ \begin{array}{c} \mathbf{x}_1 \\ \mathbf{m}_{1,j} \end{array} \right\}, \ldots, \left\{ \begin{array}{c} \mathbf{x}_N \\ \mathbf{m}_{N,j} \end{array} \right\} \right] \tag{4.1}$$

Using $\mathbf{D}$ we compute the block covariance matrix $\Sigma_{\mathbf{DD}}$ as shown on Eq 4.2.

$$\Sigma_{\mathbf{DD}} = \left[\begin{array}{cc} \Sigma_{\mathbf{xx}} & \Sigma_{\mathbf{xm}_j} \\ \Sigma^T_{\mathbf{xm}_j} & \Sigma_{\mathbf{m}_j\mathbf{m}_j} \end{array} \right]. \tag{4.2}$$

The conditional covariance estimates that relate shape in $\mathbf{x}_i$, and metadata $\mathbf{m}_j$ are calculated using Eq. 4.3.

$$\Sigma_{\mathbf{x}|\mathbf{m}_j} = \Sigma_{\mathbf{xx}} - \Sigma_{\mathbf{xm}_j}\Sigma^{-1}_{\mathbf{m}_j\mathbf{m}_j}\Sigma^{T}_{\mathbf{xm}_j}. \tag{4.3}$$

The covariance matrices in Eq. 4.3 are obtained from the block covariance matrix in Eq. 4.2. $\Sigma_{\mathbf{xm}_j}\Sigma^{-1}_{\mathbf{m}_j\mathbf{m}_j}$ are the regression coefficients that model the relationship between shape and the metadata fields. These coefficients are stored and used to compute the final mean shape estimate of the model in the next section.

Since for N models perform N eigendecompositions on relatively large covariance matrices $\Sigma_{\mathbf{x}|\mathbf{m}_j}$, we reduce the dimensionality of the problem by working with parametric shape vectors $\mathbf{b}_i = \Phi^T(\mathbf{x}_i - \bar{\mathbf{x}})$ through eigendecomposition of $\Sigma_{\mathbf{xx}}$, rather than 3D shape vectors $\mathbf{x}_i$. We retained 98% of variance after eigendecomposition, which, depending on the size of the training set is accounted for by 20 to 30 modes of variation. With this new representation, Eq. 4.2 is replaced by Eq. 4.4,

$$\Sigma^{(\mathbf{b})}_{\mathbf{DD}} = \begin{bmatrix} \Lambda_{\mathbf{bb}} & \Sigma_{\mathbf{bm}_j} \\ \Sigma^{T}_{\mathbf{bm}_j} & \Sigma_{\mathbf{m}_j\mathbf{m}_j} \end{bmatrix} \tag{4.4}$$

where $\Lambda_{\mathbf{bb}}$ are the eigenvalues of $\Sigma_{\mathbf{xx}}$, and subindex $\mathbf{b}$ indicates parametric shape.

Similarly, equation 4.3 is replaced by equation 4.5,

$$\Lambda_{\mathbf{m}_j} = \Lambda_{\mathbf{bb}} - \Sigma_{\mathbf{bm}_j}\Sigma^{-1}_{\mathbf{m}_j\mathbf{m}_j}\Sigma^{T}_{\mathbf{bm}_j}, \tag{4.5}$$

where $\Lambda_{\mathbf{m}_j}$ is a subset of the variance in $\Lambda_{\mathbf{bb}}$, and has the same coordinate system Φ that represents parametric shapes $\mathbf{b}_i$.

Using equations 4.4 and 4.5 we obtain M eigenvalue matrices $\Lambda_{\mathbf{m}_j}$ that represent the conditional shape variability of $\mathbf{x}_i$ regarding each metadata field $\mathbf{m}_j$.

We have a coordinate system Φ and M conditional variance matrices $\Lambda_{\mathbf{b}|\mathbf{m}_j}$. However, we still must compute M corresponding conditional mean vectors $\bar{\mathbf{b}}_{\mathbf{m}_j}$, to which $\Lambda_{\mathbf{b}|\mathbf{m}_j}$ are centered, and combine these models into a unique PDM that encodes all conditional relationships between shape and metadata.

4.3.2 Metadata-Constrained Mean Shape

To compute the final model's mean shape, we first must compute M mean estimates $\bar{\mathbf{b}}_{\mathbf{m}_j}$ using the regression coefficients $\Sigma_{\mathbf{b}\mathbf{m}_j}\Sigma^{-1}_{\mathbf{m}_j\mathbf{m}_j}$ obtained in section 4.3.1. Using these coefficients, and the definitions of multivariate conditional normal distributions, we obtain M conditional mean estimates based on each metadata field (Equation 4.6).

$$\bar{\mathbf{b}}_{\mathbf{m}_j} = \Sigma_{\mathbf{b}\mathbf{m}_j}\Sigma^{-1}_{\mathbf{m}_j\mathbf{m}_j}(\mathbf{m}_j^{new} - \bar{\mathbf{m}}_j) \tag{4.6}$$

Notice that in Eq. 4.6, $\mathbf{m}_j^{new}$ is the j^{th} metadata field of a new subject for whom shape is unknown.

The question then becomes how to combine these mean shape estimates into the final mean. Given that not all metadata are good predictors of shape, we weigh each mean estimate proportionally to the correlation between shape and the different metadata fields (Eq. 4.7),

$$w_j = \frac{\rho_{\mathbf{b}_i\mathbf{m}_j}}{\sum \rho_{\mathbf{b}_i\mathbf{m}_j}} \tag{4.7}$$

where $\rho_{\mathbf{x}_i \mathbf{m}_j}$ is the correlation coefficient between shape and each metadata field.

The final parametric mean estimate $\bar{\mathbf{b}}$ is a weighted average of the obtained mean estimates $\bar{\mathbf{b}}_{\mathbf{m}_j}$ (Eq. 4.8). The final mean shape in 3D space is recovered using Eq. 4.9. We denote the final conditional model parameters with subindex c.

$$\bar{\mathbf{b}} = \sum_{j=1}^{M} \bar{\mathbf{b}}_{\mathbf{m}_j} w_j \tag{4.8}$$

$$\bar{\mathbf{x}}_c = \bar{\mathbf{x}} + \Phi \bar{\mathbf{b}} \tag{4.9}$$

4.3.3 Metadata Constraints Combination

The variance associated to the final mean, $\bar{\mathbf{x}}_c$, is computed by centering the obtained eigenvalue matrices $\Lambda_{\mathbf{m}_j}$ at their corresponding mean estimates $\bar{\mathbf{b}}_{\mathbf{m}_j}$, and calculating the overlapping space between all eigenspaces $E_j = \{\Phi, \Lambda_{\mathbf{m}_j}\}$. Eq. 4.10 describes the intersecting space, where $\lambda_{\mathbf{m}_j}$ are the diagonal entries of eigenvalue matices $\Lambda_{\mathbf{m}_j}$.

$$\lambda_c = \frac{1}{2} \left\{ \min_j (\bar{\mathbf{b}}_{\mathbf{m}_j} + \lambda_{\mathbf{m}_j}) + |\max_j (\bar{\mathbf{b}}_{\mathbf{m}_j} - \lambda_{\mathbf{m}_j})| \right\} \tag{4.10}$$

By assembling a diagonal matrix with the entries of λ_c, we obtain the final eigenvalue matrix Λ_c describing the conditional variability of shapes $\mathbf{x}_i$ given metadata $\mathbf{m}_j$. The final PDM described in Eq. 4.11 will constrain image segmentation in the next section.

$$\Omega_{\mathbf{x}|\mathbf{m}} = (\bar{\mathbf{x}}_c, \Phi, \Lambda_c) \tag{4.11}$$

4.3.4 Application to Image Search

4.3.4.1 Boundary Detection

Image boundaries were detected by minimizing the Mahalanobis distance between a profile of grey levels $\mathbf{g}$ sampled from each landmark of the current shape estimate, and the mean profile from the training set $\bar{\mathbf{g}}$. The appearance model for each landmark is trained by computing the mean $\bar{\mathbf{g}}$, and covariance matrix $\Sigma_{\mathbf{gg}}$, for the profiles of corresponding landmarks across the training set. Landmark profiles $\mathbf{g}$ are sampled by projecting a normal vector to the shape's surface onto the image slice closest to that landmark. Then the Mahalanobis distance between every point in the model and the profile is computed. Finally, the landmark is displaced to the location that minimizes the distance on Eq. 4.12.

$$D(\mathbf{g}) = \sqrt{(\mathbf{g}-\bar{\mathbf{g}})^T \Sigma_{\mathbf{gg}}^{-1} (\mathbf{g}-\bar{\mathbf{g}})} \tag{4.12}$$

4.3.4.2 Shape Model Fitting

After the landmark displacement procedure described in the previous section, the new feature points must be constrained to ensure they constitute a valid shape (described by our PDM). We do this following a standard procedure [97] to find the pose and shape parameters that best fit the new feature points. Let us assume we wish to fit a new model instance $\tilde{\mathbf{x}}$ to feature points $\mathbf{x}'$.

1. Initialize parameter vector $\mathbf{b}$ to zero.

2. Generate new conditional model instance $\tilde{\mathbf{x}} = \bar{\mathbf{x}}_c + \Phi\mathbf{b}$.

3. Align feature points $\mathbf{x}'$ to model $\tilde{\mathbf{x}}$.

4. Update model parameters to match the feature points, $\mathbf{b} = \Phi^T(\mathbf{x}' - \bar{\mathbf{x}}_c)$.

5. Constrain $\mathbf{b}$ to be within $\pm 3\sqrt{\lambda_c}$.

6. Recover current 3D shape $\tilde{\mathbf{x}} = \bar{\mathbf{x}}_c + \Phi\mathbf{b}$.

7. Using the current shape perform landmark displacement and iterate from Step 1.

Alignment of points is performed using Procrustes analysis [98] to eliminate rotation and translation effects. However, scale is preserved as it is an important feature in cardiac image segmentation.

4.4 Results

The goal of this chapter is to show the potential in using commonly available patient information to construct personalized shape models that improve the accuracy of image segmentation. To test the extent of this improvement, we compare the segmentation accuracy obtained with a standard shape-only PDM to the accuracy obtained with the proposed metadata-constrained PDM. Both PDMs were trained using 150 datasets, and tested on the remaining 100 cases.

4.4.1 Data

4.4.1.1 Image

We used 250 cardiac magnetic resonance imaging (CMR) datasets obtained from the Cardiac Atlas Project (CAP). CAP is a web-accessible resource[1], comprising a population atlas of asymptomatic and pathological hearts [99]. For this study we use 250 asymptomatic cases from the Multi Ethnic Study of Atherosclerosis (MESA) study [100]. The MESA protocol used fast gradient-recalled echo (GRE) imaging with 10-12 short axis slices with typical parameters 6 mm thickness, 4 mm gap, field of view 360-400mm, 256x160 matrix, and pixel size from 1.4-2.5$mm/pixel$ depending on patient size.

4.4.1.2 Shape

Contours were manually drawn as a series of points by the MESA CMR core lab on short-axis slices for all cases at end-diastole (ED). These contours were fitted by a finite element model by linear least squares as described in [94]. The resulting models are 250 triangular meshes with point correspondence. The meshes comprise 1570 points, of which 785 describe the endocardial surface at ED. In this chapter we focus on the shape analysis of the LV endocardial surface at ED.

[1] `www.cardiacatlas.org`

4.4.1.3 Metadata

In this study we include these metadata fields provided with the image data: age, gender, race, height, weight, systolic/diastolic blood pressure, heart rate, hypertension, smoker, alcohol. Some fields are continuous variables and others are categorical. All categorical variables were replaced by binary codes of length n, where n is the number of categories for the variables.

4.4.2 Standard vs. Metadata-Constrained Segmentation

All image volumes were initialized by artificially aligning an instance of the PDM's mean shape with the ground truth. The dimension of image profiles for feature extraction was set to 7x1 voxels (3 on each side) projected inwardly and outwardly from the landmark point. The stopping criterion for the ASM was set at 30 iterations.

Table 4.1 shows the summary statistics for the segmentation. The RMS point-to-surface segmentation error across all cases for the standard PDM was $2.18mm$, whereas the error for the proposed technique was $2.03mm$. This represents an accuracy improvement of 7% over the standard PDM. Also, the standard deviation of the error was improved by 13% over the standard PDM.

Figure 4.1 shows the percentage of improvement obtained for each individual subject plotted from lowest to highest improvement. From the figure we can see that using metadata improves the segmentation

Table 4.1: Summary statistics comparing point-to-surface RMS segmentation errors of the standard PDM vs. our proposed metadata-constrained PDM.

	Mean $\pm$ σ (mm)	Min.	Max.
Standard PDM	2.18 $\pm$ 0.62	1.41	4.33
Conditional PDM	2.03 $\pm$ 0.54	1.38	4.16
Improvement (%)	7% $\pm$ 13%	-7.5%	30.4%

of nearly 80% of the subjects. And in more than 20% of subjects the segmentation accuracy was improved by a significant 15% to 30%.

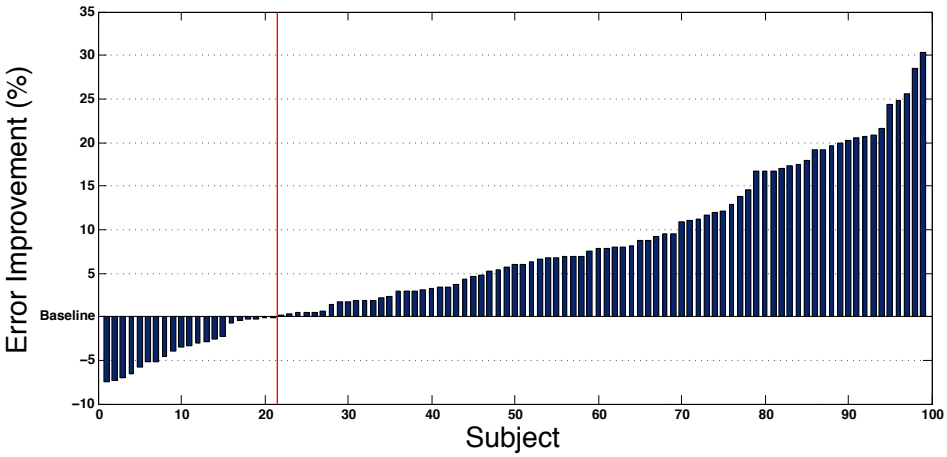

Figure 4.1: Percentage of improvement provided by our metadata-constrained models. Improvements are shown in ascending order for all subjects.

Table 4.2 shows the distribution of the errors in the test sample for both the conditional- and standard-PDM segmentations. It can be seen how the use of the metadata reduces the number of cases with large errors. For example, by using the standard shape models, 59% of the datasets are segmented with over $2mm$ errors. This number is reduced to 39% when adding the metadata during segmentation. Similarly, with the standard shape models, only 3% are segmented with less than $1.5mm$ errors, while this number increases to 14% by using the proposed metadata constrained segmentation framework. Table 4.2 illustrates the positive effect of considering the patient metadata, besides shape and image information, during cardiac segmentation.

Table 4.2: Distribution of cases by error range between the standard PDM and our proposed metadata-constrained PDM.

	Error $< x$ mm (% of Cases)			
	$< 1.5mm$	$< 2mm$	$< 2.5mm$	$< 3mm$
Standard PDM	3%	41%	76%	94%
Conditional PDM	14%	61%	88%	97%

4.4.3 Examples

Figure 4.2 shows two typical examples where our technique outperforms the standard PDM. Both panels, left and right, show the segmentation obtained with the standard PDM (left), and our technique

(right) superimposed on the ground truth shape. On Example 1 (left panel), the standard PDM (left) fails to correctly match lateral wall of the LV, whereas our method (right) better approximates the region. Some arrows are placed on the figure to show the regions where these differences are most significant. Similarly, example 2 (right panel) shows large errors, on the basal, apical and lateral wall regions for the standard PDM (left), while our method (right) can match those same regions with smaller errors.

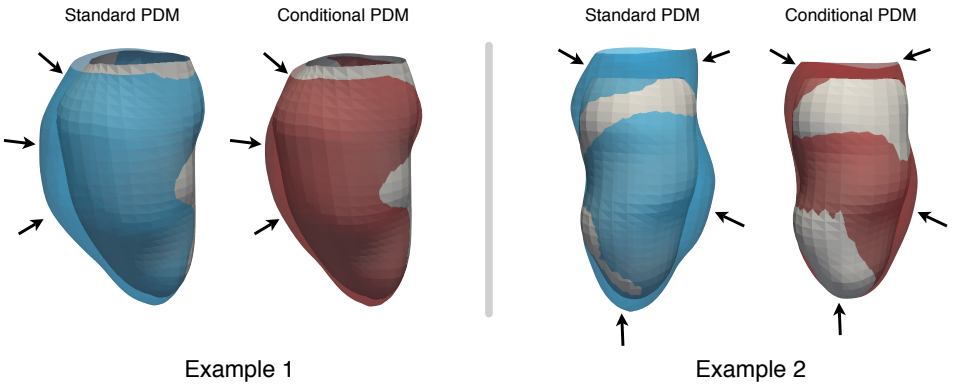

Figure 4.2: Segmentation examples. The blue and red shapes (left and right on each example) are segmentations obtained with the standard PDM, and our conditional PDM respectively. The ground truth shape is overlaid in gray color for comparison.

4.5 Conclusion

We presented a method to construct statistical shape models that incorporate non-image information from the patient. The proposed models reduce the shape domain to custom fit the patient's unique characteristics. We achieved this by using multivariate conditional distributions to regress a model that represents the most likely shape variation given the patient's metadata. We validated our method by comparing the segmentation accuracy obtained with a standard PCA-based method with the accuracy obtained with our technique. Results showed a 7% average segmentation improvement and often improvements of up to 30% over the standard PCA model. As future work, we would like to identify the metadata that best predicts shape and a methodology to optimally combine them so segmentation accuracy is maximally improved.

CHAPTER 5

GENERAL CONCLUSIONS

5.1 Overview

The aim of this thesis has been the development of new methods capable to build enhanced statistical models of anatomical shape from medical image data. We have presented three independent, yet, complementary methods, to independently address the shortcomings of traditional SSMs, in particular their limited quality when dealing highly variable and complex anatomies. We have shown on multiple applications including the analysis of cardiac, spine, and brain structures that the methods presented can improve the generalization and specificity of the representation of these structures, and as

a consequence, the accuracy and reliability of medical image segmentation. In these paragraphs, we summarize the specific contributions of this dissertation.

First, we presented a method for integrating multiple SSMs by solving the correspondence problem across two or more models with distinct shape representation, and applying existing fusion techniques for the aggregation of the individual model's statistical properties. Having the ability to merge multiple SSMs allows the construction of models statistically richer, without the need to generate or gain access to the original observations. Such utility can facilitate research collaboration, and encourage reusability of statistical data. The ability to merge SSMs also allows for integrating shape statistics from multiple image modalities, such as CT and MRI, and the integration of different statistical populations, such as, normal and pathological samples. A natural application of our technique would be the integration of statistical models emerging from multiple large scale population imaging studies being undertaken all over the world including, the MESA project with $N > 6,500$ subjects, the DETERMINE project with $N > 10,000$, the UK Biobank aiming to image over 100,000 subjects, and at least 30 other European projects that have emerged in the last 10 years [101].

Second, we developed a framework for the modeling and segmentation of highly complex multi-part anatomies. Our method proposes, 1) an iterative approach to finding an optimal subdivision of the anatomy of the training samples, based on the equalization of the variance

within each part of the complete anatomy, 2) the independent modeling of each of the resulting anatomies to capture more local and subtle geometrical details, 3) the use of conditional probabilities to model the statistical relationships between all constituent sub-parts so their combination yield geometrically valid shapes, and 4) a segmentation mechanism that finds the boundary of each sub-part by considering its conditional relationships to all other parts, and updates said probabilities online as it iterates through the target shape. Our approach is ideal because it allows to maintain the simplicity of linear models, while at the same time handling the complexity of highly non-linear structures. Our technique is well suited for anatomies comprising multiple components. As future work, this technique could model the complete spine including thoracic, and cervical vertebrae, and intervertebral discs. Our method would also be well suited for the modeling and segmentation of the 4-chamber heart, as the complex interactions between the different structures can be easily decomposed and modeled using our approach.

Third, we proposed a method that extends and fine-tunes the SSMs beyond considering the shape statistics that can be found on image data, by including other relevant non-imaging information about the subject into constructing shape models. We suggested the use of multivariate Gaussian conditional probabilities to model the relationship between shape data extracted from the image, and demographic and/or clinical indicators that can typically be found on a patient's clinical chart. By modeling the relationship between shape and metadata, we

can narrow down the uncertainty surrounding a given subject's shape during the image segmentation process. By modeling the conditional probabilities between shape and the metadata independently, we can obtain multiple independent predictions from each of the metadata and combine them into a more robust majority agreement segmentation. In line with this approach, our future objective is to study the impact that a more extensive set of metadata can have on increasing the specificity of SSMs and to what degree this information can lead to personalized models that deliver on the promise of precision imaging and medicine.

5.2 Future Perspectives

We are in the early years of personalized medicine. The term personalized, or precision medicine, refers to the identification of subpopulations, or groups of patients, to whom specifically tailored treatments and practices are provided. The personalized medicine revolution began with developing genomics, and the possibility of stratifying the population according to genetic biomarkers. Today the tenets of personalized medicine can, and are being applied to all areas of medicine and healthcare, including medical imaging and radiology. Increasingly, clinicians and researchers have access to larger and more complex population studies, which opens the door for statistically driven patient stratification. One example of this is the UK Biobank imaging study [102], which aims to conduct detailed MRI imaging scans of the vital

organs of over 100,000 participants, making it the largest study of its kind in the world. The quantitative assessment of large populations is essential to determine baseline reference values against which to compare new observations, which supports true population stratification. In this context, enhanced SSMs such as those described in this thesis, i.e., models with increased generalization ability, accuracy and specificity, will play an important role in many areas as detailed in these sections.

5.2.1 Enhanced SSMs and Large-Scale Image Analysis

To obtain clinically relevant results from large cohorts such as the UK Biobank, robust and scalable full data analysis pipelines are needed that can manage every step of data processing, from data scrubbing, to image segmentation, to the computation of morphological and functional cardiovascular parameters. The large variability in studies of this magnitude requires the use of model-based segmentation techniques that can handle large inter-subject and intra-subject variability. It is our belief that a method that can combine the generalization ability of rich models [103] (Chapter 2), the potential to integrate behavioural, phenotypic, and demographic factors into a personalized model [104] (Chapter 4), and can simultaneously segment multiple structures such as adjacent organs or multi-part complex objects, such as the spine and heart [105] (Chapter 3), is one important step towards providing highly reliable and accurate segmentations. Such models

would be fundamental for the extraction of accurate and relevant biomarkers for biomedical research and clinical practice.

5.2.2 Enhanced SSMs and Content-Based Image Retrieval

SSMs have the potential to play a critical role in content-based medical image retrieval (CBIR) systems. CBIR is an image search technique that complements the conventional text-based retrieval of images by using visual features, such as texture, and shape, as search criteria. Large image repositories such as those in hospitals, and research repositories such as the UK Biobank, offer the opportunity to develop more sophisticated tools for clinicians that allow querying large databases for images that match the shape characteristics of a particular case of interest [106]. In the medical context, shape information is one of the strongest factors in detecting certain diseases/lesions, and in understanding their evolution. Shape-based descriptors are likely to be critical to characterize the detail required for medical image retrieval [107]. Today, however, most of the medical CBIR systems do not exploit the full potential of the shape information as they either use indirect correlates of the shape such as texture measurements or employ very global and simple shape description schemes incapable of capturing the required classification granularity. Accurate and reliable model-based segmentation techniques such as those presented in this thesis become essential to capture the required level of detail for successful CBIR systems.

5.2.3 Enhanced SSMs and Computer-Aided Diagnosis

A clinical application for which enhanced SSMs can play an important role is computer-aided diagnosis, in particular for the automatic localization and classification of abnormalities. For example, vertebral fractures generally occur earlier in life than hip fractures and can be important early indicators of poor skeletal health [108]. They are predictors of future spine and hip fractures, so accurate diagnosis and clear reporting of these fractures is essential, yet only about a third of vertebral fractures come to clinical attention [109]. There is considerable evidence that vertebral fractures are under-reported, and when reported, appropriate intervention is often not initiated. Studies show that vertebral fractures are often not diagnosed so, only about 30% of vertebral fractures come to medical attention [110], as they are hard to detect despite training; they can easily be misdiagnosed; they are reported in ambiguous language, or they are not mentioned in radiology reports. This under-reporting problem could be reduced by statistical model-based approaches. A pipeline that includes the training of high-quality SSMs to detect abnormally shaped vertebrae and points out the location and abnormality to the physician could be an invaluable tool in routine clinical practice. Such approach to fracture detection has the potential of increasing detection rates, reducing medical care costs, and most important increasing the patient's quality of life.

In summary, fueled by unprecedented access to large image datasets

and technological innovations, a new era of personalized medicine is expected to revolutionize health care, and with it, new challenges and opportunities to improve biomedical research and patient care are emerging. We will see the rapid, and widespread implementation of big data analytics across health-care systems, and with them, the ability to provide focalized treatments, and prevention strategies for specific patients based on large population statistics. The development of tools for the proper analysis of image data, such as those presented in this thesis, is a crucial step toward realizing the promise of personalized medicine.

BIBLIOGRAPHY

[1] N. J. Schork, "Personalized medicine: time for one-person trials," *Nature*, vol. 520, no. 7549, pp. 609–611, 2015.

[2] "European perspectives in personalised medicine," tech. rep., May 2011.
`https://ec.europa.eu/research/health/pdf/personalised-medicine-conference-report_en.pdf`.

[3] "Medical imaging in personalised medicine: a white paper of the research committee of the European Society of Radiology (ESR)," *Insights Imaging*, vol. 6, no. 2, pp. 141–155, 2015.

[4] European Society of Radiology (ESR), "White paper on imaging biomarkers.," *Insights Imaging*, vol. 1, pp. 42–45, may 2010.

[5] M. Viceconti, G. Clapworthy, and S. V. S. Jan, "The Virtual Physiological Human - A European Initiative for in silico Human Modelling -," *J. Physiol. Sci.*, vol. 58, no. 7, pp. 441–446, 2008.

[6] NHS, "NHS imaging and radiodiagnostic activity in england - 2012/13 release," Apr. 2012.

https://www.england.nhs.uk/statistics/wp-content/uploads/sites/2/2013/04/KH12-release-2012-13.pdf.

[7] T. Heimann and H.-P. Meinzer, "Statistical shape models for 3D medical image segmentation: A review," *Med. Image Anal.*, vol. 13, pp. 543–563, Aug. 2009.

[8] T. Cootes, D. Cooper, C. Taylor, and J. Graham, "Trainable method of parametric shape description," *Image Vis. Comput.*, vol. 10, no. 5, pp. 289 – 294, 1992.

[9] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models—their training and application," *Comput. Vis. Image Underst.*, vol. 61, no. 1, pp. 38–59, 1995.

[10] I. Jolliffe, "Principal component analysis," 2002.

[11] B. de Vos, J. Wolterink, P. de Jong, T. Leiner, M. Viergever, and I. Isgum, "ConvNet-Based Localization of Anatomical Structures in 3D Medical Images," *IEEE Trans. Med. Imaging*, pp. 1–1, 2017.

[12] X. Albá, K. Lekadir, M. Pereañez, P. Medrano-Gracia, A. Young, and A. Frangi, "Fully Automatic Image Parsing and Quality Control for Population Cardiac MRI," *Med. Image Anal.*, vol. -, no. -, pp. –, under review.

[13] I. Castro-Mateos, J. M. Pozo, M. Pereanez, K. Lekadir, A. Lazary, and A. F. Frangi, "Statistical Interspace Models (SIMs): Application to Robust 3D Spine Segmentation," *IEEE Trans. Med. Imaging*, vol. 34, pp. 1663–1675, aug 2015.

[14] T. Toumanidou and J. Noailly, "Musculoskeletal Modeling of the Lumbar Spine to Explore Functional Interactions between Back Muscle Loads and Intervertebral Disk Multiphysics," *Front. Bioeng. Biotechnol.*, vol. 3, no. August, pp. 1–13, 2015.

[15] J. Lötjönen, K. Antila, E. Lamminmäki, J. Koikkalainen, M. Lilja, and T. Cootes, "Artificial enlargement of a training set for statistical shape models: application to cardiac images," in *Int. Workshop on Functional Imaging and Modeling of the Heart (FIMH)*, vol. 3504 of *LNCS*, pp. 92–101, 2005.

[16] T. Cootes and C. Taylor, "Combining point distribution models with shape models based on finite element analysis," *Image Vis. Comput.*, vol. 13, no. 5, pp. 403 – 409, 1995.

[17] Y. Wang and L. Staib, "Boundary finding with prior shape and smoothness models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 7, pp. 738 –743, 2000.

[18] O. Ecabert, J. Peters, H. Schramm, C. Lorenz, J. von Berg, M. J. Walker, M. Vembar, M. E. Olszewski, K. Subramanyan, G. Lavi, and J. Weese, "Automatic model-based segmentation of the heart in CT images," *IEEE Trans. Med. Imaging*, vol. 27, pp. 1189–1201, September 2008.

[19] Y. Zheng, B. Georgescu, A. Barbu, M. Scheuering, and D. Comaniciu, "Four-chamber heart modeling and automatic segmentation for 3D cardiac CT volumes," *IEEE Trans. Med. Imaging*, vol. 27, pp. 1668–1681, November 2008.

[20] C. Davatzikos, X. Tao, and D. Shen, "Hierarchical active shape models, using the wavelet transform," *IEEE Trans. Med. Imaging*, vol. 22, pp. 414–423, March 2003.

[21] T. Okada, K. Yokota, M. Hori, M. Nakamoto, H. Nakamura, and Y. Sato, "Construction of hierarchical multi-organ statistical atlases and their application to multi-organ segmentation from CT images," in *Med. Image Comput. Comput. Assist. Interv.* (D. M. et al., ed.), vol. 5241, pp. 502–509, 2008.

[22] T. Okada, R. Shimada, Y. Sato, M. Hori, K. Yokota, M. Nakamoto, Y.-W. Chen, H. Nakamura, and S. Tamura, "Automated segmentation of the liver from 3D CT images using probabilistic atlas and multi-level statistical shape model," in *Med. Image Comput. Comput. Assist. Interv.* (N. A. et al., ed.), vol. 4792, pp. 86–93, 2007.

[23] K. Lekadir, N. Keenan, D. Pennell, and G. Yang, "An inter-landmark approach to 4-d shape extraction and interpretation: application to myocardial motion assessment in mri.," *IEEE Trans. Med. Imaging*, vol. 30, no. 1, pp. 52–68, 2011.

[24] M. Üzümcü, A. Frangi, M. Sonka, J. Reiber, and B. Lelieveldt, "ICA vs. PCA active appearance models: Application to cardiac MR segmentation,"

[25] M. B. Stegmann, K. Sjöstrand, and R. Larsen, "Sparse modeling of landmark and texture variability using the orthomax criterion," in *Proc. SPIE Medical Imaging*, vol. 6144, p. 61441G, 2006.

[26] M. Loog, "Localized maximum entropy shape modelling," in *Inf. Process Med. Imaging*, vol. 4584 of *LNCS*, pp. 619–629, 2007.

[27] K. Sjöstrand, E. Rostrup, C. Ryberg, R. Larsen, C. Studholme, H. Baezner, J. Ferro, F. Fazekas, L. Pantoni, D. Inzitari, and G. Waldemar, "Sparse decomposition and modeling of anatomical shape variation," *IEEE Trans. Med. Imaging*, vol. 26, no. 12, pp. 1625–1635, 2007.

[28] T. Tölli, J. Koikkalainen, K. Lauerma, and J. Lötjönen, "Artificially enlarged training set in image segmentation," in *Med. Image Comput. Comput. Assist. Interv.*, pp. 75–82, Springer, 2006.

[29] J. Koikkalainen, T. Tölli, K. Lauerma, K. Antila, E. Mattila, M. Lilja, and J. Lötjönen, "Methods of artificial enlargement of the training set for statistical shape models," *IEEE Trans. Med. Imaging*, vol. 27, pp. 1643–1654, Nov. 2008.

[30] M. de Bruijne, B. van Ginneken, M. A. Viergever, and W. J. Niessen, "Adapting active shape models for 3D segmentation of tubular structures in medical images," in *Inf. Process Med. Imaging*, vol. 2732 of *LNCS*, pp. 136–147, 2003.

[31] S. Sclaroff and J. Isidoro, "Active blobs: region-based, deformable appearance models," *Comput. Vis. Image Underst.*, vol. 89, no. 2, pp. 197–225, 2003.

[32] A. C. Kotcheff and C. J. Taylor, "Automatic construction of eigenshape models by direct optimization," *Med. Image Anal.*, vol. 2, no. 4, pp. 303–314, 1998.

[33] D. Meier and E. Fisher, "Parameter space warping: Shape-based correspondence between morphologically different objects," *IEEE Trans. Med. Imaging*, vol. 21, pp. 31–47, Jan. 2002.

[34] R. H. Davies, *Learning Shape: Optimal Models for Analysing Natural Variability*.
PhD thesis, Univ. of Manchester, UK, 2002.

[35] T. Heimann, I. Wolf, and H.-P. Meinzer, "Automatic generation of 3D statistical shape models with optimal landmark distributions," *Methods Inf. Med.*, vol. 46, no. 3, pp. 275–281, 2007.

[36] R. H. Davies, C. J. Twining, T. F. Cootes, and C. J. Taylor, "Building 3D statistical shape models by direct optimisation," *IEEE Trans. Med. Imaging*, vol. 29, pp. 961–981, Apr. 2010.

[37] A. D. Brett and C. J. Taylor, "A method of automated landmark generation for automated 3D PDM construction," *Image Vis. Comput.*, vol. 18, no. 9, pp. 739–748, 2000.

[38] A. Hill, C. J. Taylor, and A. D. Brett, "A framework for automatic landmark identification using a new method of nonrigid correspondence," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 3, pp. 241–251, 2000.

[39] K. B. Hilger, R. R. Paulsen, and R. Larsen, "Markov random field restoration of point correspondences for active shape modeling," in *Proc. SPIE Medical Imaging*, pp. 1862–1869, 2004.

[40] I. Oguz, J. Cates, T. Fletcher, R. Whitaker, D. Cool, S. Aylward, and M. Styner, "Cortical correspondence using entropy-based

particle systems and local features," in *Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI)*, pp. 1637–1640, 2008.

[41] M. Datar, J. Cates, P. T. Fletcher, S. Gouttard, G. Gerig, and R. Whitaker, "Particle based shape regression of open surfaces with applications to developmental neuroimaging," in *Med. Image Comput. Comput. Assist. Interv.*, vol. 5762 of *LNCS*, pp. 167–174, 2009.

[42] M. Datar, Y. Gur, B. Paniagua, M. Styner, and R. Whitaker, "Geometric correspondence for ensembles of nonregular shapes," in *Med. Image Comput. Comput. Assist. Interv.*, vol. 6892 of *LNCS*, pp. 368–375, 2011.

[43] M. R. Kaus, V. Pekar, C. Lorenz, R. Truyen, S. Lobregt, and J. Weese, "Automated 3-D PDM construction from segmented images using deformable models," *IEEE Trans. Med. Imaging*, vol. 22, no. 8, pp. 1005–1013, 2003.

[44] Z. Zhao and E. K. Teoh, "A new scheme for automated 3D PDM construction using deformable models," *Image Vis. Comput.*, vol. 26, pp. 275–288, Feb. 2008.

[45] A. F. Frangi, D. Rueckert, J. A. Schnabel, and W. J. Niessen, "Automatic construction of multiple-object three-dimensional statistical shape models: Application to cardiac modeling," *IEEE Trans. Med. Imaging*, vol. 21, no. 9, pp. 1151–1166, 2002.

[46] C. Hoogendoorn, N. Duchateau, D. Sánchez-Quintana, T. Whitmarsh, F. Sukno, M. De Craene, K. Lekadir, A. Frangi, *et al.*, "A

high-resolution atlas and statistical model of the human heart from multislice CT," *IEEE Trans. Med. Imaging*, vol. 32, no. 1, pp. 28–44, 2013.

[47] P. Hall, D. Marshall, and R. Martin, "Merging and splitting eigenspace models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 9, pp. 1042–1049, 2000.

[48] C. Butakoff and A. F. Frangi, "A framework for weighted fusion of multiple statistical models of shape and appearance," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 11, pp. 1847–1857, 2006.

[49] S. Durrleman, X. Pennec, A. Trouvé, and N. Ayache, "Statistical models of sets of curves and surfaces based on currents," *Med. Image Anal.*, vol. 13, no. 5, pp. 793–808, 2009.

[50] P. J. Besl and N. D. McKay, "Method for registration of 3D shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 239–256, 1992.

[51] M. Vaillant and J. Glaunès, "Surface matching via currents," in *Inf. Process Med. Imaging*, vol. 3565 of *LNCS*, pp. 1–5, 2005.

[52] H. Coxeter, "Barycentric coordinates," *Introduction to Geometry*, vol. 13, 1969.

[53] M. S. Floater, "Parametrization and smooth approximation of surface triangulations," *Computer aided geometric design*, vol. 14, no. 3, pp. 231–250, 1997.

[54] N. Sukumar and E. Malsch, "Recent advances in the construction of polygonal finite element interpolants," *Archives of Computational Methods in Engineering*, vol. 13, no. 1, pp. 129–163, 2006.

[55] B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," *J. Opt. Soc. Am. A Opt. Image Sci. Vis.*, vol. 4, no. 4, pp. 629–642, 1987.

[56] P. Radau, Y. Lu, K. Connelly, G. Paul, A. J. Dick, and G. A. Wright, "Evaluation framework for algorithms segmenting short axis cardiac MRI," *The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge*, 2009.

[57] B. Van Ginneken, T. Heimann, and M. Styner, "3D segmentation in the clinic: A grand challenge," *MICCAI Workshop on 3D segmentation in the clinic: a grand challenge*, pp. 7–15, 2007.

[58] C. Loop, "Smooth subdivision surfaces based on triangles," Master's thesis, Department of mathematics, Univ. of Utah, Utah, USA, 1987.

[59] S. Valette and J.-M. Chassery, "Approximated centroidal voronoi diagrams for uniform polygonal mesh coarsening," in *Comput. Graph. Forum*, vol. 23, pp. 381–389, Wiley Online Library, 2004.

[60] H. C. van Assen, M. G. Danilouchkine, A. F. Frangi, S. Ordás, J. J. M. Westenberg, J. H. C. Reiber, and B. P. F. Lelieveldt, "SPASM: A 3D-ASM for segmentation of sparse and arbitrarily

oriented cardiac MRI data," *Med. Image Anal.*, vol. 10, no. 2, pp. 286–303, 2006.

[61] H.-M. Mayer, *Minimally invasive spine surgery*. Springer, 2006.

[62] D. Schlenzka, T. Laine, and T. Lund, "Computer-assisted spine surgery," *European Spine Journal*, vol. 9, no. 1, pp. S057–S064, 2000.

[63] A. C. Jones and R. K. Wilcox, "Finite element analysis of the spine: towards a framework of verification, validation and sensitivity analysis," *Medical engineering & physics*, vol. 30, no. 10, pp. 1287–1304, 2008.

[64] A. Middleditch and J. Oliver, *Functional anatomy of the spine*. Elsevier Health Sciences, 2005.

[65] H. A. Yuan, S. R. Garfin, C. A. Dickman, and S. M. Mardjetko, "A historical cohort study of pedicle screw fixation in thoracic, lumbar, and sacral spinal fusion," *Spine*, vol. 19, no. 20, pp. 2279S–2296S, 1994.

[66] L. Hansen, M. De Zee, J. Rasmussen, T. B. Andersen, C. Wong, and E. B. Simonsen, "Anatomy and biomechanics of the back muscles in the lumbar spine with reference to biomechanical modeling," *Spine*, vol. 31, no. 17, pp. 1888–1899, 2006.

[67] S. Benameur, M. Mignotte, S. Parent, H. Labelle, W. Skalli, and J. de Guise, "3D/2D registration and segmentation of scoli-

otic vertebrae using statistical models," *Computerized Medical Imaging and Graphics*, vol. 27, no. 5, pp. 321–337, 2003.

[68] M. G. Roberts, T. F. Cootes, and J. E. Adams, "Linking sequences of active appearance sub-models via constraints: An application in automated vertebral morphometry.," in *British Machine Vision Conference*, pp. 1–10, Citeseer, 2003.

[69] J. Kaminsky, P. Klinge, T. Rodt, M. Bokemeyer, W. Luedemann, and M. Samii, "Specially adapted interactive tools for an improved 3D-segmentation of the spine," *Computerized Medical Imaging and Graphics*, vol. 28, no. 3, pp. 119–127, 2004.

[70] M. de Bruijne and M. Nielsen, "Image segmentation by shape particle filtering," in *Proc. Int. Conf. on Pattern Recognition*, pp. 722–725, 2004.

[71] A. Mastmeyer, K. Engelke, C. Fuchs, and W. A. Kalender, "A hierarchical 3D segmentation method and the definition of vertebral body coordinate systems for QCT of the lumbar spine," *Med. Image Anal.*, vol. 10, no. 4, pp. 560–577, 2006.

[72] Z. Peng, J. Zhong, W. Wee, and J.-h. Lee, "Automated vertebra detection and segmentation from the whole spine MR images," in *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, pp. 2527–2530, IEEE, 2006.

[73] S.-H. Huang, Y.-H. Chu, S.-H. Lai, and C. L. Novak, "Learning-based vertebra detection and iterative normalized-cut segmen-

tation for spinal MRI," *IEEE Trans. Med. Imaging*, vol. 28, no. 10, pp. 1595–1605, 2009.

[74] Y. Kim and D. Kim, "A fully automatic vertebra segmentation method using 3D deformable fences," *Computerized Medical Imaging and Graphics*, vol. 33, no. 5, pp. 343–352, 2009.

[75] T. Klinder, J. Ostermann, M. Ehm, A. Franz, R. Kneser, and C. Lorenz, "Automated model-based vertebra detection, identification, and segmentation in CT images," *Med. Image Anal.*, vol. 13, no. 3, pp. 471–482, 2009.

[76] J. Ma, L. Lu, Y. Zhan, X. Zhou, M. Salganicoff, and A. Krishnan, "Hierarchical segmentation and identification of thoracic vertebra using learning-based edge detection and coarse-to-fine deformable model," in *Med. Image Comput. Comput. Assist. Interv.*, pp. 19–27, Springer, 2010.

[77] D. Štern, B. Likar, F. Pernuš, and T. Vrtovec, "Parametric modelling and segmentation of vertebral bodies in 3D CT and MR spine images," *Physics in medicine and biology*, vol. 56, no. 23, p. 7505, 2011.

[78] A. Rasoulian, R. Rohling, and P. Abolmaesumi, "Lumbar spine segmentation using a statistical multi-vertebrae anatomical shape+pose model," *IEEE Trans. Med. Imaging*, vol. 32, pp. 1890–1900, Oct 2013.

[79] S. Kadoury, H. Labelle, and N. Paragios, "Spine segmentation in medical images using manifold embeddings and higher-order

MRFs," *IEEE Trans. Med. Imaging*, vol. 32, no. 7, pp. 1227–1238, 2013.

[80] J. Huang, F. Jian, H. Wu, and H. Li, "An improved level set method for vertebra CT image segmentation," *Biomedical engineering online*, vol. 12, no. 1, p. 48, 2013.

[81] I. Castro-Mateos, J. M. Pozo, P. E. Eltes, L. Del Rio, A. Lazary, and A. F. Frangi, "3D segmentation of annulus fibrosus and nucleus pulposus from T2-weighted magnetic resonance images," *Physics in Medicine and Biology*, vol. 59, no. 24, p. 7847, 2014.

[82] J. Yang, L. H. Staib, and J. S. Duncan, "Neighbor-constrained segmentation with level set based 3-D deformable models," *IEEE Trans. Med. Imaging*, vol. 23, pp. 940–948, August 2004.

[83] T. Heap and D. Hogg, "Improving specificity in PDMs using a hierarchical approach.," in *British Machine Vision Conference*, 1997.

[84] T. F. Cootes and C. J. Taylor, "A mixture model for representing shape variation," *Image and Vision Computing*, vol. 17, no. 8, pp. 567–573, 1999.

[85] M. Pereañez, K. Lekadir, C. Hoogendoorn, I. Castro-Mateos, and A. Frangi, "Detailed vertebral segmentation using part-based decomposition and conditional shape models," in *Recent Advances in Computational Methods and Clinical Applications for Spine Imaging* (J. Yao, B. Glocker, T. Klinder, and S. Li, eds.), vol. 20 of *LNCVB*, Springer, 2015.

[86] C. Goodall, "Procrustes methods in the statistical analysis of shape," *J. R. Stat. Soc. Series B Stat. Methodol.*, vol. 53, no. 2, pp. 285–339, 1991.

[87] A. E. Hoerl and R. W. Kennard, "Ridge regression: Biased estimation for nonorthogonal problems," *Technometrics*, vol. 12, no. 1, pp. 55–67, 1970.

[88] C. Metz, N. Baka, H. Kirisli, M. Schaap, T. van Walsum, S. Klein, L. Neefjes, N. Mollet, B. Lelieveldt, M. de Bruijne, *et al.*, "Conditional shape models for cardiac motion estimation," in *Med. Image Comput. Comput. Assist. Interv.*, pp. 452–459, Springer, 2010.

[89] CORDIS, "Functional prognosis simulation of patient-specific spinal treatment for clinical use," Apr. 2017.
`http://cordis.europa.eu/project/rcn/97394_en.html`.

[90] B. Glocker, D. Zikic, E. Konukoglu, D. R. Haynor, and A. Criminisi, "Vertebrae localization in pathological spine CT via dense classification from sparse annotations," in *Med. Image Comput. Comput. Assist. Interv.*, pp. 262–270, Springer, 2013.

[91] R. Wolz, P. Aljabar, J. V. Hajnal, J. Lötjönen, D. Rueckert, A. D. N. Initiative, *et al.*, "Nonlinear dimensionality reduction combining MR imaging with non-imaging information," *Med. Image Anal.*, vol. 16, no. 4, pp. 819–830, 2012.

[92] R. Blanc, M. Reyes, C. Seiler, and G. Székely, "Conditional variability of statistical shape models based on surrogate variables,"

Med. Image Comput. Comput. Assist. Interv., pp. 84–91, 2009.

[93] S. Grbić, J. K. Swee, and R. Ionasec, "Shapeforest: Building constrained statistical shape models with decision trees," in *Computer Vision–ECCV 2014*, pp. 597–612, Springer, 2014.

[94] P. Medrano-Gracia, B. R. Cowan, B. Ambale-Venkatesh, D. A. Bluemke, J. Eng, J. P. Finn, C. G. Fonseca, J. A. Lima, A. Suinesiaputra, and A. A. Young, "Left ventricular shape variation in asymptomatic populations: the multi-ethnic study of atherosclerosis," *Cardiovasc. Res.*, vol. 16, no. 1, p. 56, 2014.

[95] M. Pereanez, K. Lekadir, I. Castro-Mateos, J. Pozo, A. Lazary, and A. Frangi, "Accurate segmentation of vertebral bodies and processes using statistical shape decomposition and conditional models," *IEEE Trans. Med. Imaging*, vol. 34, pp. 1627–1639, Aug 2015.

[96] H. C. Van Assen, M. G. Danilouchkine, A. F. Frangi, S. Ordás, J. J. Westenberg, J. H. Reiber, and B. P. Lelieveldt, "SPASM: a 3D-ASM for segmentation of sparse and arbitrarily oriented cardiac MRI data," *Med. Image Anal.*, vol. 10, no. 2, pp. 286–303, 2006.

[97] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models-their training and application," *Comput. Vis. Image Understand.*, vol. 61, no. 1, pp. 38–59, 1995.

[98] C. Goodall, "Procrustes methods in the statistical analysis of shape," *J. Roy. Stat. Soc. B Stat. Meth.*, pp. 285–339, 1991.

[99] C. G. Fonseca, M. Backhaus, D. A. Bluemke, R. D. Britten, J. Do Chung, B. R. Cowan, I. D. Dinov, J. P. Finn, P. J. Hunter, A. H. Kadish, *et al.*, "The cardiac atlas project-an imaging database for computational modeling and statistical atlases of the heart," *Bioinformatics*, vol. 27, no. 16, pp. 2288–2295, 2011.

[100] D. E. Bild, D. A. Bluemke, G. L. Burke, R. Detrano, A. V. D. Roux, A. R. Folsom, P. Greenland, D. R. JacobsJr, R. Kronmal, K. Liu, *et al.*, "Multi-ethnic study of atherosclerosis: objectives and design," *Am. J. epidemiol*, vol. 156, no. 9, pp. 871–881, 2002.

[101] European Population Imaging Infrastructure, "Population Imaging - European Cohorts," Apr. 2017.
`http://populationimaging.eu/europeancohorts/`.

[102] S. E. Petersen, P. M. Matthews, F. Bamberg, D. A. Bluemke, J. M. Francis, M. G. Friedrich, P. Leeson, E. Nagel, S. Plein, F. E. Rademakers, A. A. Young, S. Garratt, T. Peakman, J. Sellors, R. Collins, and S. Neubauer, "Imaging in population science: cardiovascular magnetic resonance in 100,000 participants of UK Biobank - rationale, challenges and approaches," *J. Cardiovasc. Magn. Reson.*, vol. 15, no. 1, p. 46, 2013.

[103] M. Pereañez, K. Lekadir, C. Butakoff, C. Hoogendoorn, and A. F. Frangi, "A framework for the merging of pre-existing and correspondenceless 3D statistical shape models," *Med. Image Anal.*, vol. 18, no. 7, pp. 1044–1058, 2014.

[104] M. Pereañez, K. Lekadir, I. Castro-Mateos, J. M. Pozo, Á. Lazáry,

and A. F. Frangi, "Accurate segmentation of vertebral bodies and processes using statistical shape decomposition and conditional models," *IEEE Trans. Med. Imaging*, vol. 34, no. 8, pp. 1627–1639, 2015.

[105] M. Pereañez, K. Lekadir, X. Albà, P. Medrano-Gracia, A. A. Young, and A. Frangi, "Patient metadata-constrained shape models for cardiac image segmentation," in *International Workshop on Statistical Atlases and Computational Models of the Heart*, pp. 98–107, Springer International Publishing, 2015.

[106] A. Kumar, J. Kim, W. Cai, M. Fulham, and D. Feng, "Content-based medical image retrieval: a survey of applications to multidimensional and multimodality data.," *J. Digit. Imaging*, vol. 26, pp. 1025–39, dec 2013.

[107] C. B. Akgül, D. L. Rubin, S. Napel, C. F. Beaulieu, H. Greenspan, and B. Acar, "Content-based image retrieval in radiology: current status and future directions.," *J. Digit. Imaging*, vol. 24, pp. 208–22, apr 2011.

[108] C. Holroyd, C. Cooper, and E. Dennison, "Epidemiology of osteoporosis," *Best Practice & Research Clinical Endocrinology & Metabolism*, vol. 22, no. 5, pp. 671–685, 2008.

[109] H. A. Fink, D. L. Milavetz, L. Palermo, M. C. Nevitt, J. A. Cauley, H. K. Genant, D. M. Black, and K. E. Ensrud, "What Proportion of Incident Radiographic Vertebral Deformities Is Clini-

cally Diagnosed and Vice Versa?," *J. Bone Miner. Res.*, vol. 20, pp. 1216–1222, mar 2005.

[110] P. D. Delmas, L. van de Langerijt, N. B. Watts, R. Eastell, H. Genant, A. Grauer, D. L. Cahall, and IMPACT Study Group, "Underdiagnosis of Vertebral Fractures Is a Worldwide Problem: The IMPACT Study," *J. Bone Miner. Res.*, vol. 20, pp. 557–563, dec 2004.

AUTHOR PUBLICATIONS: JOURNALS

[1] A. Suinesiaputra, P. Ablin, X. Alba, M. Alessandrini, J. Allen, W. Bai, S. Cimen, P. Claes, B. R. Cowan, J. D'hooge, N. Duchateau, J. Ehrhardt, A. F. Frangi, A. Gooya, V. Grau, K. Lekadir, A. Lu, A. Mukhopadhyay, I. Oksuz, N. Parajali, X. Pennec, **M. Pereañez**, C. Pinto, P. Piras, M. M. Rohe, D. Rueckert, D. Saring, M. Sermesant, K. Siddiqi, M. Tabassian, L. Teresi, S. A. Tsaftaris, M. Wilms, A. A. Young, X. Zhang, and P. Medrano-Gracia, "Statistical shape modeling of the left ventricle: myocardial infarct classification challenge," *IEEE Journal of Biomedical and Health Informatics*, vol. PP, no. 99, pp. 1–1, 2017.

[2] X. Albá, **M. Pereañez**, C. Hoogendoorn, A. J. Swift, J. M. Wild, A. F. Frangi, and K. Lekadir, "An algorithm for the segmentation of highly abnormal hearts using a generic statistical shape model," *IEEE Transactions on Medical Imaging*, vol. 35, pp. 845–859, March 2016.

[3] **M. Pereañez**, K. Lekadir, I. Castro-Mateos, J. M. Pozo, Á. Lazáry, and A. F. Frangi, "Accurate segmentation of verte-

bral bodies and processes using statistical shape decomposition and conditional models," *IEEE transactions on medical imaging*, vol. 34, no. 8, pp. 1627–1639, 2015.

[4] I. Castro-Mateos, J. M. Pozo, **M. Pereañez**, K. Lekadir, A. Lazary, and A. F. Frangi, "Statistical interspace models (sims): application to robust 3d spine segmentation," *IEEE transactions on medical imaging*, vol. 34, no. 8, pp. 1663–1675, 2015.

[5] K. Lekadir, A. Pashaei, C. Hoogendoorn, **M. Pereañez**, X. Albà, and A. F. Frangi, "Effect of statistically derived fiber models on the estimation of cardiac electrical activation," *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 11, pp. 2740–2748, 2014.

[6] **M. Pereañez**, K. Lekadir, C. Butakoff, C. Hoogendoorn, and A. F. Frangi, "A framework for the merging of pre-existing and correspondenceless 3d statistical shape models," *Medical image analysis*, vol. 18, no. 7, pp. 1044–1058, 2014.

[7] K. Lekadir, C. Hoogendoorn, **M. Pereañez**, X. Albà, A. Pashaei, and A. F. Frangi, "Statistical personalization of ventricular fiber orientation using shape predictors," *IEEE transactions on medical imaging*, vol. 33, no. 4, pp. 882–890, 2014.

AUTHOR PUBLICATIONS: CONFERENCES

[1] **M. Pereañez**, K. Lekadir, X. Albà, P. Medrano-Gracia, A. A. Young, and A. Frangi, "Patient metadata-constrained shape models for cardiac image segmentation," in *International Workshop on Statistical Atlases and Computational Models of the Heart*, pp. 98–107, Springer International Publishing, 2015.

[2] K. Lekadir, X. Albà, **M. Pereañez**, and A. F. Frangi, "Statistical shape modeling using partial least squares: Application to the assessment of myocardial infarction," in *International Workshop on Statistical Atlases and Computational Models of the Heart*, pp. 130–139, Springer International Publishing, 2015.

[3] **M. Pereañez**, K. Lekadir, C. Hoogendoorn, I. Castro-Mateos, and A. Frangi, "Detailed vertebral segmentation using part-based decomposition and conditional shape models," in *Recent Advances in Computational Methods and Clinical Applications for Spine Imaging*, pp. 95–103, Springer International Publishing, 2015.

[4] X. Albà, K. Lekadir, C. Hoogendoorn, **M. Pereañez**, A. J. Swift,

J. M. Wild, and A. F. Frangi, "Reusability of statistical shape models for the segmentation of severely abnormal hearts," in *International Workshop on Statistical Atlases and Computational Models of the Heart*, pp. 257–264, Springer International Publishing, 2014.

[5] K. Lekadir, A. Pashaei, C. Hoogendoorn, **M. Pereañez**, X. Albà, and A. F. Frangi, "Personalized modeling of cardiac electrophysiology using shape-based prediction of fiber orientation," in *International Workshop on Statistical Atlases and Computational Models of the Heart*, pp. 196–203, Springer Berlin Heidelberg, 2013.

[6] **M. Pereañez**, K. Lekadir, C. Butakoff, C. Hoogendoorn, and A. Frangi, "Fusing correspondenceless 3d point distribution models," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 251–258, Springer Berlin Heidelberg, 2013.

BIOSKETCH

Marco Pereañez was born on September 3, 1981, in Medellin, Colombia. After his high-school studies, he moved to New York and obtained his BA in Mathematics and Computer Science from the State University of New York (SUNY) at Purchase in 2006. After a two year hiatus working at the Whitney Museum of American Art, he earns an Erasmus Mundus scholarship and moves to Europe to obtain his MSc in Computer Vision and Robotics (ViBot) from the Universities of Burgundy (France), Heriot-Watt (UK), and Girona (Spain) in 2010. After graduation, he joined the Center for Computational Imaging and Simulation Technologies in Biomedicine (CISTIB) at the Universitat Pompeu Fabra (UPF) in Barcelona. In 2011 he enrolled in the PhD program of the Department of Information and Communication Technologies (DTIC) at the UPF. During his studies, Marco taught Probability and Statistics to second year students of the Informatics and Multimedia Program at the UPF, and

since 2015 he has reviewed for the journal of Medical Image Analysis (MedIA). He works as Research Assistant in the Department of Electronic and Electrical Engineering at the University of Sheffield in the UK. Marco's research emphasis during his PhD has been on statistical shape modeling, and its application to personalized medical image analysis.

Dedication and acknowledgements

On the face of it, this thesis is a reflection of the work done in the past few years. On closer inspection however, it is evidence of the support of those who have been next to me along the way.

I would like to thank Prof. Frangi for enabling my research and always striving to show me the bigger picture. I am thankful to Dr. Karim Lekadir, my thesis supervisor, who has taught me a bit about tenacity and integrity.

I also want to say thanks to Themis, one of the micest people I know, for her constant support, caring and countless laughs. A special thanks to my friend Flokas, for playing with me constantly.

Gracias papa por todo el amor y mostrarme la belleza en tantas cosas. Lo extraño todos los días.

Finally, I want to thank the most important person in my life, and without whom I would not be here today, my Mother Carmen Vargas. Gracias mama por todo el amor, el apoyo, la paciencia, y por ser ejemplo e inpiración para seguir adelante. La quiero.

TESI DOCTORAL UPF / 2017