3D Wavelet-based Fusion Techniques for Biomedical Imaging

Jose Luis Rubio Guivernau

Departamento de Ingeniería Electrónica
Escuela Técnica Superior de Ingenieros de Telecomunicación
Universidad Politécnica de Madrid

Ph.D. Thesis

3D Wavelet-based Fusion Techniques for Biomedical Imaging

Author

Jose Luis Rubio Guivernau

Advisors

María Jesús Ledesma Carbayo

Andrés Santos Lleó

Committee

D. Jose María Carazo García (*presidente*)	*Centro Nacional de Biotecnología - CSIC*
D. Juan Enrique Ortuño Fisac (*secretario*)	*Universidad Politécnica de Madrid*
D. Cristina Santa Marta Pastrana (*vocal*)	*Universidad Nacional de Educación a Distancia*
D. Arrate Muñoz Barrutia (*vocal*)	*Universidad de Navarra*
D. Dan Herzka (*vocal*)	*Johns Hopkins University*
D. Louise Duloquin (*suplente*)	*Institut de Neurobiologie Alfred Fessard - CNRS*
D. Esther Pérez David (*suplente*)	*Hospital General Universitario Gregorio Marañon*

Departamento de Ingeniería Electrónica

Escuela Técnica Superior de Ingenieros de Telecomunicación

Ph.D. Thesis

3D Wavelet-based Fusion Techniques for Biomedical Imaging

Author:

Jose Luis Rubio Guivernau

Telecommunication Engineer

Advisors:

María Jesús Ledesma Carbayo

Telecommunication Engineer, Ph.D.

Andrés Santos Lleó

Telecommunication Engineer, Ph.D.

September 2012

Resumen

Hoy en día las técnicas de adquisición de imágenes tridimensionales son comunes en diversas áreas, pero cabe destacar la relevancia que han adquirido en el ámbito de la imagen biomédica, dentro del cual encontramos una amplia gama de técnicas como la microscopía confocal, microscopía de dos fotones, microscopía de fluorescencia mediante lámina de luz, resonancia magnética nuclear, tomografía por emisión de positrones, tomografía de coherencia óptica, ecografía 3D y un largo etcétera. Un denominador común de todas esas aplicaciones es la constante necesidad por aumentar la resolución y la calidad de las imágenes adquiridas.

En algunas de dichas técnicas de imagen tridimensional se da una interesante situación: aunque que cada volumen adquirido no contiene información suficiente para representar el objeto bajo estudio dentro de los parámetros de calidad requeridos por algunas aplicaciones finales, el esquema de adquisición permite la obtención de varios volúmenes que representan diferentes vistas de dicho objeto, de tal forma que cada una de las vistas proporciona información complementaria acerca del mismo. En este tipo de situación es posible, mediante la combinación de varias de esas vistas, obtener una mejor comprensión del objeto que a partir de cada una de ellas por separado.

En el contexto de esta Tesis Doctoral se ha propuesto, desarrollado y validado una nueva metodología de proceso de imágenes basada en la transformada *wavelet* discreta para la combinación, o fusión, de varias vistas con información complementaria de un mismo objeto. El método de fusión propuesto aprovecha la capacidad de descomposición en escalas y orientaciones de la transformada *wavelet* discreta para integrar en un solo volumen toda la información distribuida entre el conjunto de vistas adquiridas.

El trabajo se centra en dos modalidades diferentes de imagen biomédica que permiten obtener tales adquisiciones multi-vista. La primera es una variante de la microscopía de fluorescencia, la microscopía de fluorescencia mediante lámina de luz, que se utiliza para el estudio del desarrollo temprano de embriones vivos en diferentes modelos animales, como el pez cebra o el erizo de mar. La segunda modalidad es la resonan-

cia magnética nuclear con realce tardío, que constituye una valiosa herramienta para evaluar la viabilidad del tejido miocárdico en pacientes con diversas miocardiopatías.

Como parte de este trabajo, el método propuesto ha sido aplicado y validado en ambas modalidades de imagen. En el caso de la aplicación a microscopía de fluorescencia, los resultados de la fusión muestran un mejor contraste y nivel de detalle en comparación con cualquiera de las vistas individuales y el método no requiere de conocimiento previo acerca la función de dispersión puntual del sistema de imagen. Además, los resultados se han comparado con otros métodos existentes. Con respecto a la aplicación a imagen de resonancia magnética con realce tardío, los volúmenes fusionados resultantes presentan una mejora cuantitativa en la nitidez de las estructuras relevantes y permiten una interpretación más sencilla y completa de la compleja estructura tridimensional del tejido miocárdico en pacientes con cardiopatía isquémica.

Para ambas aplicaciones los resultados de esta tesis se encuentran actualmente en uso en los centros clínicos y de investigación con los que el autor ha colaborado durante este trabajo. Además se ha puesto a libre disposición de la comunidad científica la implementación del método de fusión propuesto. Por último, se ha tramitado también una solicitud de patente internacional que cubre el método de visualización desarrollado para la aplicación de Resonancia Magnética Nuclear.

Abstract

Nowadays three dimensional imaging techniques are common in several fields, but especially in biomedical imaging, where we can find a wide range of techniques including: Laser Scanning Confocal Microscopy, Laser Scanning Two Photon Microscopy, Light Sheet Fluorescence Microscopy, Magnetic Resonance Imaging, Positron Emission Tomography, Optical Coherence Tomography, 3D Ultrasound Imaging, etc. A common denominator of all those applications being the constant need for further increasing resolution and quality of the acquired images.

Interestingly, in some of the mentioned three-dimensional imaging techniques a remarkable situation arises: while a single volume does not contain enough information to represent the object being imaged within the quality parameters required by the final application, the acquisition scheme allows recording several volumes which represent different views of a given object, with each of the views providing complementary information. In this kind of situation one can get a better understanding of the object by combining several views instead of looking at each of them separately.

Within such context, in this PhD Thesis we propose, develop and test new image processing methodologies based on the discrete wavelet transform for the combination, or fusion, of several views containing complementary information of a given object. The proposed fusion method exploits the scale and orientation decomposition capabilities of the discrete wavelet transform to integrate in a single volume all the available information distributed among the set of acquired views.

The work focuses in two different biomedical imaging modalities which provide such multi-view datasets. The first one is a particular fluorescence microscopy technique, Light-Sheet Fluorescence Microscopy, used for imaging and gaining understanding of the early development of live embryos from different animal models (like zebrafish or sea urchin). The second is Delayed Enhancement Magnetic Resonance Imaging, which is a valuable tool for assessing the viability of myocardial tissue on patients suffering from different cardiomyopathies.

As part of this work, the proposed method was implemented and then validated on both imaging modalities. For the fluorescence microscopy application, the fusion results show improved contrast and detail discrimination when compared to any of the individual views and the method does not rely on prior knowledge of the system's point spread function (PSF). Moreover, the results have shown improved performance with respect to previous PSF independent methods. With respect to its application to Delayed Enhancement Magnetic Resonance Imaging, the resulting fused volumes show a quantitative sharpness improvement and enable an easier and more complete interpretation of complex three-dimensional scar and heterogeneous tissue information in ischemic cardiomyopathy patients.

In both applications, the results of this thesis are currently in use in the clinical and research centers with which the author collaborated during his work. An implementation of the fusion method has also been made freely available to the scientific community. Finally, an international patent application has been filed covering the visualization method developed for the Magnetic Resonance Imaging application.

A Leti.

Contents

Those are my principles. If you don't like them I have others.

Groucho Marx

Chapter 1

Motivation and Objectives

1.1 Motivation

Nowadays three dimensional (3D) imaging techniques are common in several fields. One example is fluorescence microscopy imaging techniques like: Laser Scanning Confocal Microscopy (LSCM), Laser Scanning Two Photon Microscopy (LSTPM), Light Sheet Fluorescence Microscopy (LSFM), etc. We can also find many 3D techniques in medical imaging, like: Magnetic Resonance Imaging (MRI), Positron Emission tomography (PET), Optical Coherence Tomography (OCT), 3D Ultrasound Imaging, etc ...

In all of those applications there is a constant need for further increasing resolution and quality of the acquired images. The main advances regarding this need usually come from new technological developments in the acquisition hardware, but given the state of the art at any time, new algorithmic methodologies are always available for getting the most of what current technology can obtain.

Interestingly, in some of the mentioned three-dimensional imaging techniques a remarkable situation often arises: a single volume acquired by the imaging equipment (MRI scanner, microscope, etc) does not contain enough information to represent the object being imaged within the quality parameters (sharpness, sampling frequency, ...) required by the final application, but at the same time the acquisition scheme allows to record several volumes representing different *views* (from different angles, for instance) of the object.

When this happens, it is usually the case that each of the *views* provides complementary information, and one can get a better understanding of the object by looking

at several *views* instead of looking at each of them separately. In figure 1.1a and figure 1.1b two examples are shown, one with MRI images of a human heart and another one with LSFM images of a live sea urchin embryo (*Paracentrotus lividus*), in which the situation described above occurs.

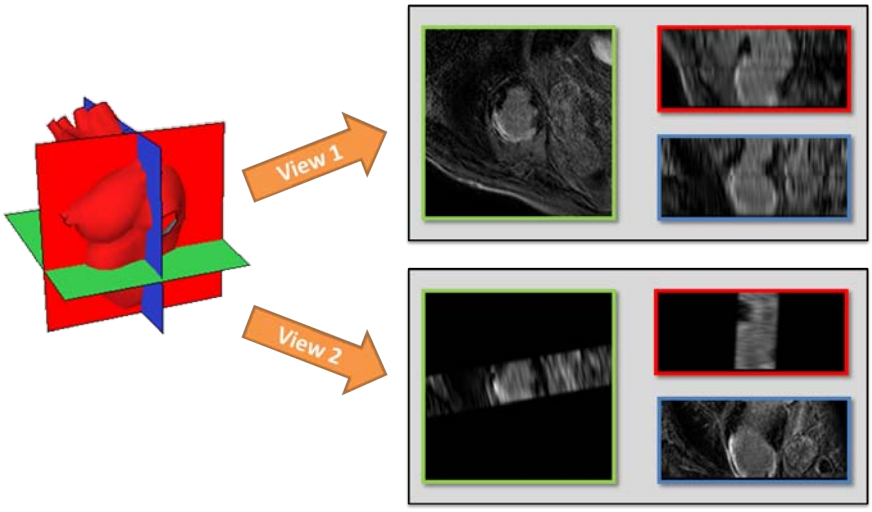

(a) Multi-View MRI imaging of a human heart

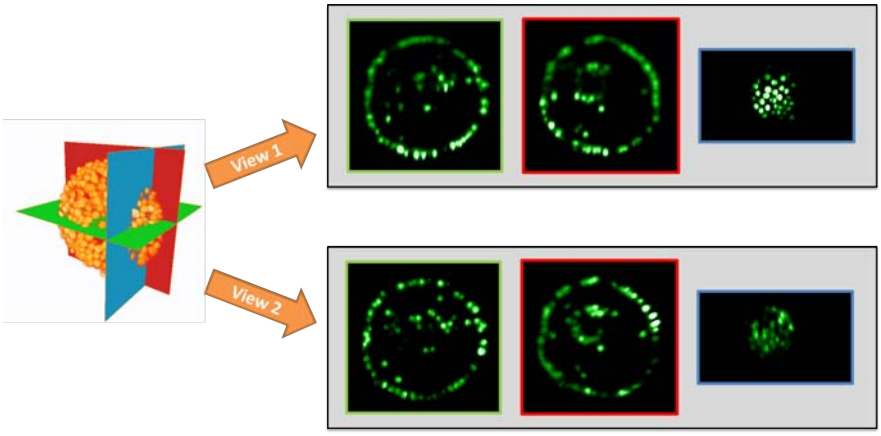

(b) Multi-View LSFM imaging of a sea urchin embryo

Figure 1.1: Two examples of multi-view three-dimensional imaging techniques: Cardiac MRI and Light-Sheet Fluorescence Microscopy. For each modality, two different volumes or *views* of the same object are shown. Each volume or *view* is depicted by three orthogonal slices, and it can be appreciated how on both cases the information provided by the views is complementary.

It is apparent from figure 1.1 that in such multi-view imaging schemes the information acquired from the object under study is spread among several volumes or *views*, which hinders its use in many applications. In order to make the most of all the available

information, it is desirable to have it all merged or *fused* into one single volume with a richer information content. This will facilitate any subsequent use of the acquired information (further processing steps, direct interpretation, etc).

1.2 Problem Outline

The general problem we are approaching in this PhD thesis can be generally named "Three-Dimensional Multi-View Imaging", and can certainly appear in many imaging fields. However the focus of this work will be on biomedical images given the special relevance of three dimensional imaging in this field.

For a given imaging scheme to fall into this category, the following combination of factors must arise:

- We are able to measure three-dimensional distributions of some physical property from a given object or scene. We will call $f(x, y, z)$ the spatial distribution of said physical property.
- The imaging scheme allows us to acquire several three-dimensional *views*, which we will call $\{X_1, X_2, \ldots, X_V\}$ (with V the number of *views*).
- The spatial distribution of the imaged physical property does not change significantly during the time it takes to acquire the set of V *views*, so that we can assume that all of them represent the same object or scene.
- Due to the imaging process, none of the *views* alone contains a satisfactory representation of the object, according to the quality parameters (sharpness, sampling frequency, ...) required by the final application.

In figure 1.2 we show a simplified representation of the three-dimensional multi-view situation we just described. In such a multi-view imaging scheme, the information acquired from the object under study might be spread among the set of *views*. This is the case when there is complementarity among the information contained in different views. There are several ways in which this might happen, but in this work we will be focusing on applications in which complementarity is given by a combination of the following two situations (illustrated in figure 1.3):

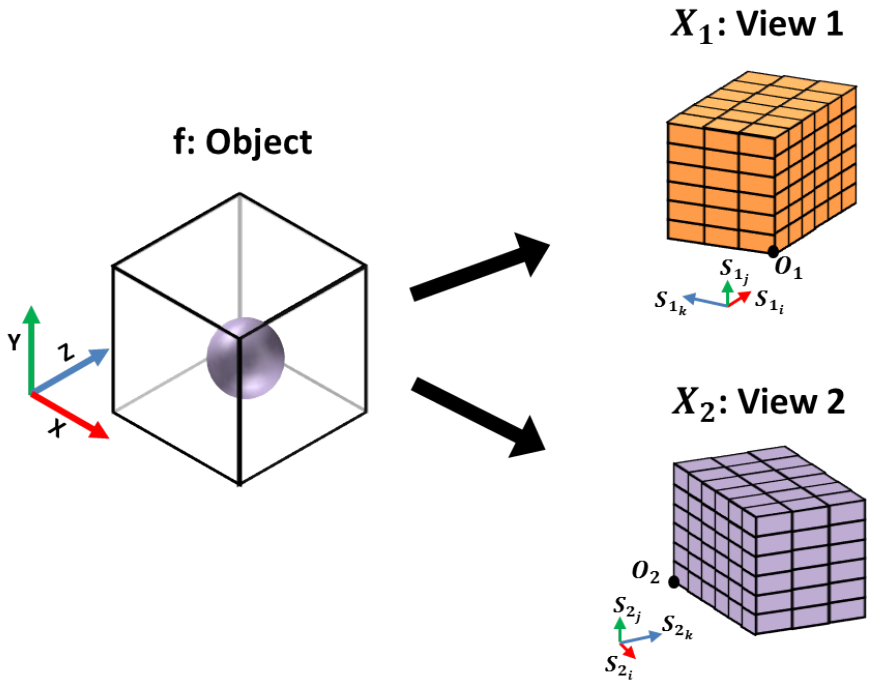

Figure 1.2: Representation of an object by different views with complementary information.

- Each of the views provides sharp information only on some regions of the total volume. This could happen for instance if there is some kind of spatially-variant blur in the acquisition process, which changes between views (Figure 1.3a).

- The sampling pattern is different for each view. This could be the case when "voxels" are anisotropic in each view, but the orientations with higher and lower spatial detail exchange roles on each of the views (Figure 1.3b).

It is important to highlight that there is an additional pre-requisite for a given problem to be suitable for the proposed fusion method. Before the fusion process can take place, all the *views* must be registered, i.e., they must be brought to a common spatial frame of reference. Therefore, although it is not the purpose of this PhD Thesis to discuss or develop registration methods, we will be using existing ones when we apply the fusion method to specific problems.

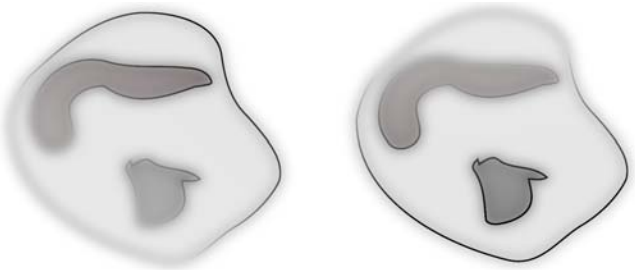

(a) Spatially-Variant Blur

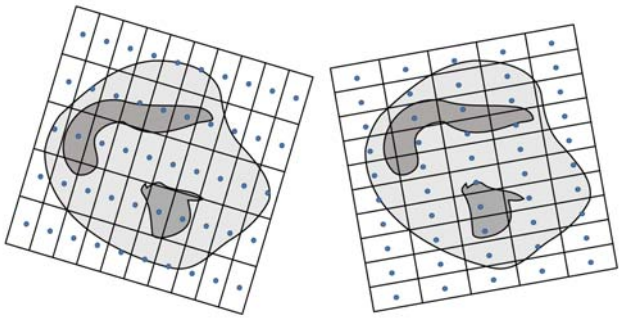

(b) Different Sampling Patterns

Figure 1.3: Complementary information from different *views*

1.3 Objectives

Within this context, the objective of this work will be **to propose, develop and test new image processing methodologies based on the discrete wavelet transform for the combination (*fusion*) of several *views* containing complementary information of a given object**. We can divide this general objective into several secondary objectives, namely:

- Propose and develop a new fusion methodology for multi-view three-dimensional datasets, based on the discrete wavelet transform and exploiting its scale and orientation decomposition capabilities to integrate in a single volume all the available information contained in the acquired views.
- Evaluate the proposed fusion methodology on multi-view 3D datasets obtained by means of Light-Sheet Fluorescence Microscopy (LSFM). The method should be able to deal with these very big datasets.
- Adapt and apply the proposed fusion method to multi-view non-uniformly sampled 3D Delayed Enhancement Magnetic Resonance Imaging (DE-MRI) datasets

of the heart. Study the clinical applicability of the resulting volumes to the diagnostic and therapeutic process of Ventricular Tachycardia (VT).

1.4 Document Structure

This manuscript is divided in seven chapters, starting with the current one in which the motivation and main objective of the work is stated, together with a description of the underlying problem. In Chapter 2 we give an introduction to the problem and we present the two application fields on which the newly developed methodologies will be applied. We will then go through the current state of the art in Chapter 3, reviewing different image enhancement techniques which are somehow related to the problem we are dealing with, in order to establish the starting point of this work.

The proposed methodology itself will be described in detail in Chapter 4, including all the necessary implementation details and some examples to help understanding each of the steps involved. We will then follow with two chapters devoted to the applications in which the proposed methodology has been tested. Chapter 5 will cover the application of the fusion method in developmental biology, more precisely to a specific fluorescence microscopy modality named Light-Sheet Fluorescence Microscopy (LSFM), and we will show how the fusion helps in the identification of the individual cells of live embryos imaged using this type of microscopy. In Chapter 6 we will explain how the proposed technique can be useful for certain Magnetic Resonance Imaging (MRI) acquisition protocols, with focus in Delayed Enhancement MRI (DE-MRI) using gadolinium-based contrast agents. Moreover, a very specific application for visualizing the substrate of Ventricular Tachycardia (VT) is presented on which the proposed fusion methodology can be of great help.

Then in Chapter 7 we will summarize the main contributions of this work, and finally we will provide a list of publications related to the work we are presenting in this manuscript.

Science... never solves a problem without creating ten more.

George Bernard Shaw

Chapter 2

Applications

In this chapter we describe the two specific fields on which we will apply and test the developed methodologies. Of course these are just two specific application fields, while the methods developed could be useful for any other application that could benefit of a multi-view integration approach of 3D images with complementary information.

2.1 Fluorescence Microscopy

Fluorescence Microscopy [Lichtman 2005] is a branch of optical microscopy in which the objects of interest exhibit fluorescence. Fluorescence is the emission of light by a molecule shortly (within nanoseconds) after absorbing light of a different (shorter) wavelength. It was George G. Stokes who first observed this phenomenon in 1852, noticing that the mineral fluorspar emitted red light when being illuminated with ultraviolet excitation. He also noticed that emitted light had always longer wavelength that the excitation light.

This wavelength difference between absorbed and emitted light, which is due to dissipation of part of the excitation energy during the process, is usually called Stoke's Shift and it is essential for making fluorescence microscopy possible as it allows the use of filters which discriminate illumination light while allowing the transmission of the light emitted by the fluorophores (molecules with fluorescent properties).

The basic scheme of a fluorescence microscope is depicted in Figure 2.1. The specimen under study is illuminated with one wavelength (excitation wavelength), and the return light is filtered to see only the wavelength-shifted (emission wavelength) fluo-

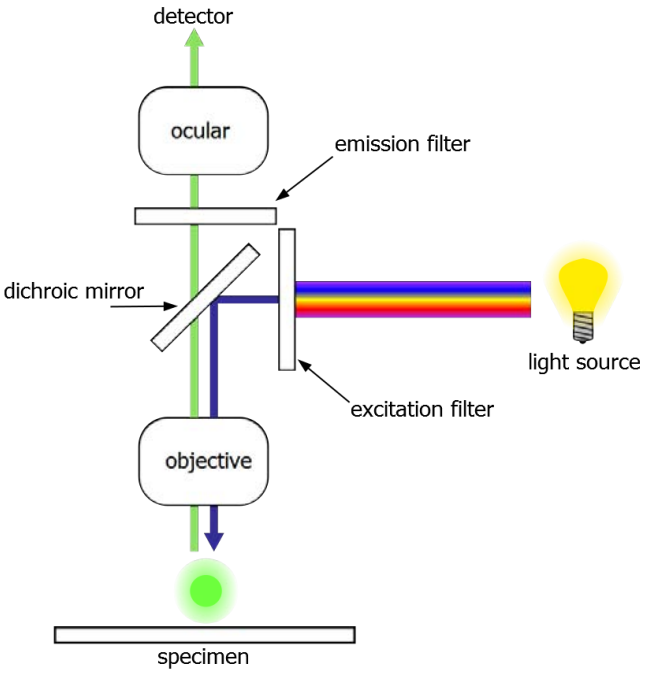

Figure 2.1: Basic scheme of a Fluorescence Microscope

rescence. The dichroic mirror has the role of separating the excitation and emission light-paths, which share the same objective.

The recent breakthrough in fluorescence microscopy applied to biology has followed the development of bio-compatible fluorescent markers. Since 1961 it was known [Shimomura 1972] that the jellyfish *Aequorea victoria* naturally produced a green fluorescent protein (GFP), but it was in 1992 when Praser et al. determined the genetic sequence of the gene responsible of producing GFP [Prasher 1992], opening the possibility of modifying the genome of other organisms to make them also produce GFP [Chalfie 1994] and starting what is now known as The Green Revolution [Stearns 1995]. O. Shimomura, M. Chalfie and R.Y. Tsien were awarded the Nobel prize in Chemistry in 2008 for the discovery and development of GFP.

Since then a number of similar fluorophores have been discovered or developed, broadening the range of colors available and improving some of the features of wild type GFP, like stronger fluorescence and higher resistance. Currently the large spectral range of available fluorophores allows simultaneous imaging of different cellular, subcellular or molecular components.

2.1.1 Fluorescence Microscopy in Developmental Biology

These developments in fluorescent probes technology together with the advances in microscopes have meant a great progress for biological imaging, as they open the possibility of imaging developmental mechanisms that were previously impossible to image [Megason 2003]. One of the main goals of current developmental biology is being able to achieve what is known as "*in toto* imaging" for embryos of animal models such as zebrafish. *In toto* imaging refers to the ability to image every single cell in an embryo and do it not only at a single time point, but over a period of the embryo's development. If that is done with enough temporal and spatial resolution, tracking of every single cell, its mitosis, linage, etc will become possible.

In order to achieve the kind of 4D imaging needed for this purpose, we need imaging techniques which can obtain 3D images of a live embryo, and which are able to do that with enough spatial resolution as to identify individual cells and fast enough to build the full 4D dataset with enough time resolution as to follow those individual cells and their descendants over time.

Optical sectioning, i.e. the ability to reject out-of-focus light, is needed in order to obtain 3D images with fluorescence microscopy, and it can be achieved in several ways. In confocal microscopy [Minsky 1961, Lichtman 2005, Vonesch 2006] a pinhole aperture located in the optical conjugate position to the focal plane blocks out-of-focus light. Two-photon microscopy [Denk 1990, Helmchen 2005] makes use of a non-linear optical property of the fluorophore in which two photons with double the wavelength are absorbed at nearly the same time, as the probability of two-photon excitation drops with the fourth power of the distance to the focal plane, there is very little excitation outside from the focal plane and optical sectioning is achieved. Other techniques which rely on computational post-processing to achieve optical sectioning have also appeared, such as deconvolution microscopy and structured illumination. Recently a new kind of fluorescence microscope called LSFM (Light Sheet Fluorescence Microscopy) has appeared which combines lateral illumination of the sample, achieving optical sectioning, with wide-field collection of the emitted light [Huisken 2004, Engelbrecht 2006].

2.1.2 Light Sheet Fluorescence Microscopy

The fundamental idea behind LSFM is to combine optical sectioning and wide-field fluorescence microscopy by illuminating a thin section of a fluorescent specimen from the side. The object is thus illuminated along a separate optical path orthogonal to the

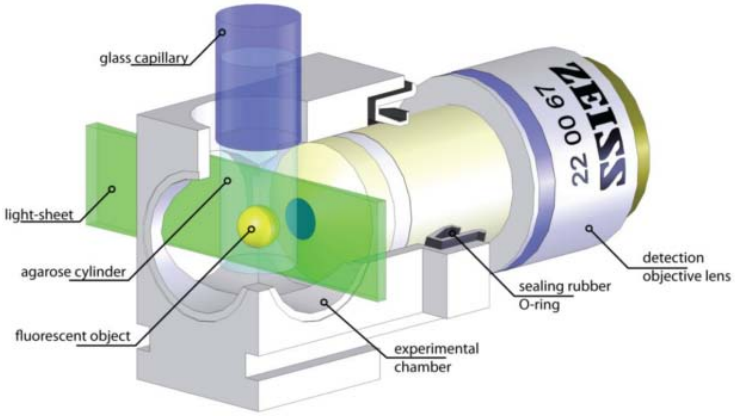

Figure 2.2: Light-Sheet Fluorescence Microscopy

detection axis, while the emitted fluorescence light is detected with a standard wide-field fluorescence microscope (Figure 2.2). In a properly aligned LSFM, the illuminating light-sheet overlaps with the focal plane of the detection objective lens. The fluorescence emission generated in the light-sheet will therefore originate from a volume close to the focal plane of the detection lens and will form a focused, high-contrast image.

In one of LSFM's implementations, known as Selective Plane Illumination Microscopy (SPIM) [Huisken 2004], the light sheet was generated via a set of apertures and by focusing a collimated beam with a cylindrical lens (Figure 2.3a). Later another implementation called Digitally Scanned Light-sheet Microscopy (DSLM) appeared [Keller 2008a, Keller 2008b], in which the light sheet is formed by a rapidly scanned laser beam (Figure 2.3b), and was claimed to present several advantages over SPIM approach:

- Each line in the specimen is illuminated with the same intensity
- DLSM does not rely on apertures to form the laser profile, reducing optical aberrations
- Higher illumination efficiency is achieved by focusing all the illumination power into a single line
- Intensity-modulated illumination patterns can be generated with DSLM, enabling the use of techniques such as structured illumination.

This scheme of microscopy has several advantages over typical Laser Scanning Microscopes (LSM) such as Confocal LSM and Two-Photon LSM. Instead of acquiring

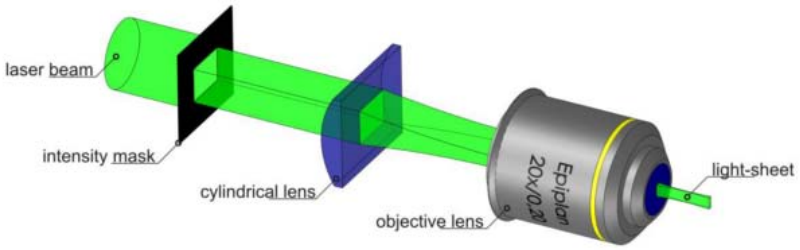

(a) SPIM Illumination Scheme (extractred from [Krzic 2009])

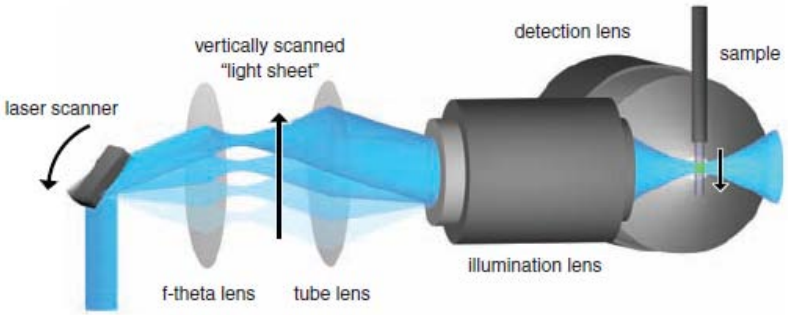

(b) DSLM Illumination Scheme (extracted from [Keller 2008a])

Figure 2.3

one point at a time as in LSM, in LSFM a full slice of the specimen is acquired at once, greatly improving the acquisition speed. Moreover, as the optical sectioning is achieved by illuminating only one slice at a time, the out-of-focus light is reduced and thus contrast is enhanced and noise level is lower. Photo-bleaching is also reduced, as only the slice being imaged at each time gets excited while in LSM the full sample gets illuminated.

But there are also some drawbacks with LSFM. For instance when the illumination beam is scattered or absorbed by features in the sample, shadowing can appear behind (with respect to the illumination direction) those features (Figure 2.4). While it is true that this effect is present in any optical microscope, the collimated nature of the illumination in LSFM makes it more pronounced.

LSFM implementations allow the rotation of the sample in order to perform multi-view acquisitions. By combining several views of the specimen, two problems that typically appear in each of the individual views can be addressed:

- Lateral resolution is typically better than axial resolution. By rotating the sample,

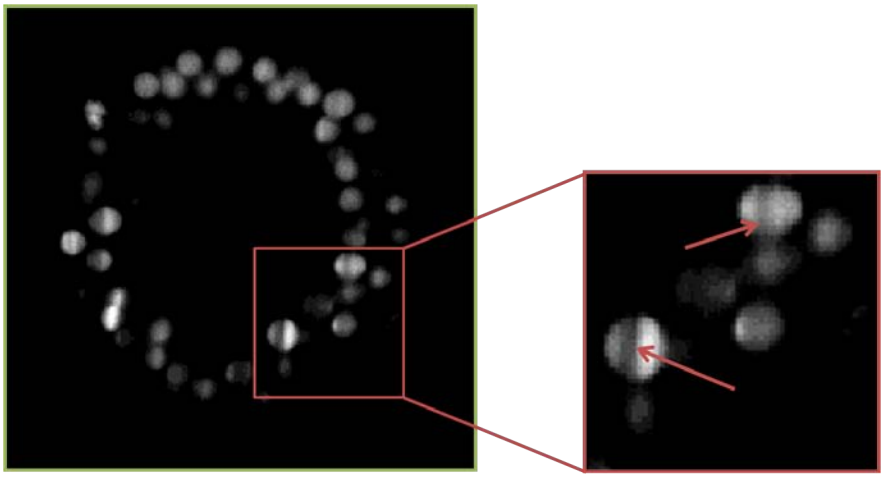

Figure 2.4: Shadowing effect in a sea urchin embryo nuclei image

lateral and axial direction swap roles (Figure 2.5a), and the fusion of the different views can combine the information on all of them.

- As the Light Sheet moves away from the imaging objective, the emitted light has to travel a longer path through the sample, which results in greater scattering and absorbtion leading to increased blur and eventually to lack of information. Rotation and multiple view acquisition provides volumes with complementary blurred&dark and bright&sharp regions (Figure 2.5b) which can be later combined together.

2.1.3 Latest Trends and Open Issues

It is clear that in order to exploit the intrinsic multi-view imaging ability of LSFM and to take full advantage of its benefits for studying the development of live embryos, it is desirable to have all the information fused into a single volume. Apart from making possible to visualize the whole embryo at once, this will also enable, for instance, the use of existing cell segmentation and tracking methods for gaining better knowledge of the embryos' morphogenesis.

For this reason several techniques have appeared during the last years which deal with the problem of combining several LSFM *views* into a single volume. Examples of such methods include multi-view deconvolution, using either invariant [Swoger 2007, Krzic 2009] or spatially variant [Temerinac-Ott 2011] Point Spread Function (PSF), and content-based fusion [Preibisch 2008]. However, as it will be explained in greater

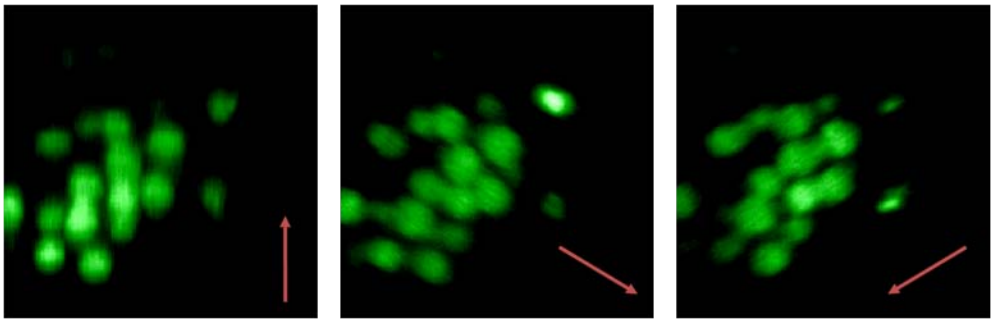

(a) Three views of a sea urchin embryo nuclei acquisition

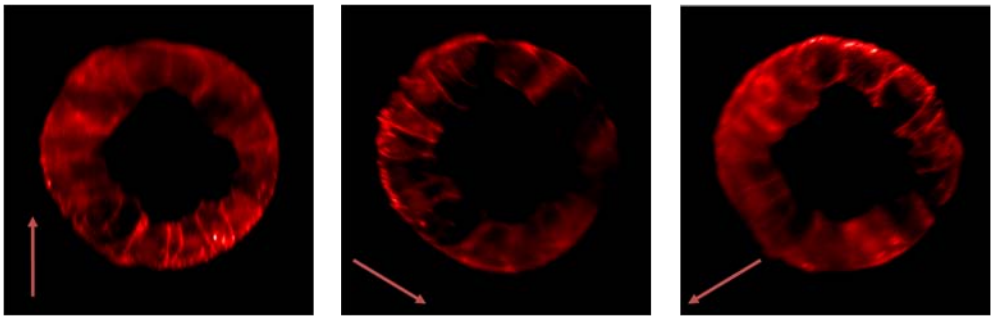

(b) Three views of a sea urchin embryo membranes acquisition

Figure 2.5: Examples of LSFM Multi-View Acquisition. Arrows show illumination direction for each view.

detail in chapter 5, there is still need for methods which not rely on good estimates of the system's PSF and which exploit the directional information provided by different *views*.

For all these reasons, we consider that LSFM imaging could benefit clearly of the methods developed in this work, and we will detail the contribution to this application in chapter 5.

2.2 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is a tomographic imaging technique capable of producing images representing some physical or chemical properties of an object based on externally measured signals.

The main physical phenomenon behind MRI is the Nuclear Magnetic Resonance effect, by which magnetic nuclei placed inside a magnetic field can absorb radio-frequency (RF) energy at a specific frequency (known as the *Larmor* frequency) which depends

on the nuclei and the magnetic field strength. This effect was observed in 1946 independently by F. Bloch and E. Purcell, and both of them were awarded with the Nobel prize in physics in 1952 for their developments in that field.

The Nuclear Magnetic Resonance (NMR) phenomenon allows us to learn about the composition of an object by placing it under a magnetic field, sending RF energy to the object and measuring its reply or *echo*. However, forming an image of the interior of the object needs one more step, the ability to localize the different signals contributing to the *echo*. It took some years to figure out how to do that and take advantage of the NMR phenomenon to produce images of the inside of an object. It was in 1972 when P. Lauterbur developed the spatial encoding principles that make MRI possible, and in 2003, the huge impact that MRI has achieved in medicine was acknowledged by the award of the Nobel prize in medicine to P. Lauterbur and P. Mansfield, who also made a great contribution to the development of MRI by introducing the mathematical formalism and developing techniques for faster imaging.

We will not discuss in detail the image formation process of MRI, there are already some very good books describing it ([Liang 1999] for instance). However, it is worth mentioning that the final image obtained from a given MRI sequence depends on intrinsic parameters of the materials under study (that includes any external contrast agents which might be used in a given application), such as the proton density, the spin-lattice relaxation time T_1 and the spin-spin relaxation time T_2. The effect of these parameters can be suppressed or enhanced depending on a set of imaging parameters such as the timing between the different RF pulses and gradients used in the imaging sequence. This means that two MRI images from the same anatomical site can look completely different depending on the specific acquisition sequence used (Figure 2.6).

2.2.1 Delayed-Enhancement MRI

During the past decades MRI has proven very useful for diagnosis of cardiovascular diseases, giving rise to the subfield of Cardiac MRI [Earls 2002, Finn 2006]. Among the applications of Cardiac MRI we can find coronary arteries morphology evaluation by black-blood imaging [Fayad 2000], cardiac function assessment by dynamic imaging like Cine sequences [Sechtem 1987, Sakuma 1993] or more recently by MR-tagging techniques like HARP [McVeigh 1996, Osman 1999, Ozturk 2003], perfusion imaging [Al-Saadi 2000], coronary angiography [Weber 2003, Deshpande 2001] and myocardial viability assessment by Delayed-Enhancement MRI (DE-MRI) [Kim 1999, Kim 2000, Kim 2003].

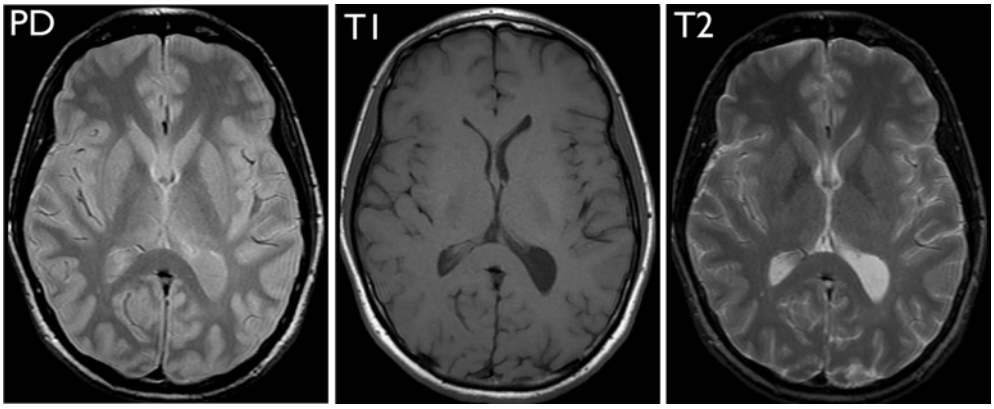

Figure 2.6: Cross-sectional head image obtained by MRI with different contrasts. From left to right: Proton-density weighted contrast, T1-weighted contrast and T2-weighted contrast.

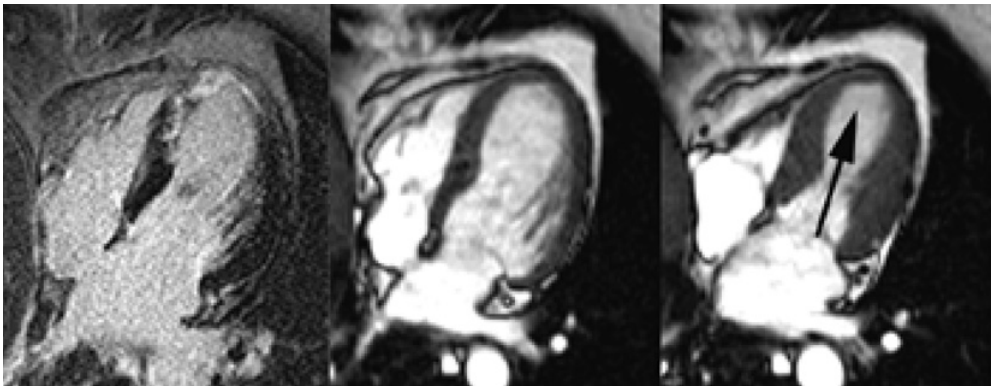

Figure 2.7: (Left) DE-MRI in acute myocardial infarction. (Middle) Cine image at end-diastole. (Right) Cine image at end-systole. DE-MRI shows transmural irreversible damage of distal septum and apex, while cine images reveal corresponding dysfunctional myocardium. (Images extracted from [Earls 2002])

Several studies have proven DE-MRI to be a great tool for differentiating between viable and nonviable myocardium both in patients with ischemic [Kim 1999, Kim 2000, Amado 2004, Kim 2008] and non-ischemic [Tandri 2005, Teraoka 2004] cardiomyopathies. Before the appearance of DE-MRI, the main tools for myocardial viability assessment were Positron Emission Tomography (PET) [Klein 2002] and Single Photon Emission Computed Tomography (SPECT) [Wagner 2003], over which DE-MRI has improved spatial resolution plus the advantage of not involving the use of ionizing radiation.

In order to perform DE-MRI, an intravenous Gadolinium-based contrast (typically Gd-DPTA) is administered around 10-20 min before imaging takes place, and then the use of an inversion-recovery sequence with an appropriate inversion delay

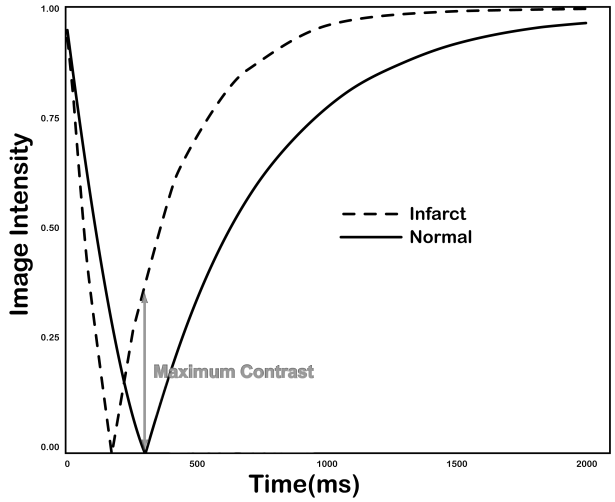

Figure 2.8: Image Intensity for Normal and Infarcted tissue vs. Inversion Time (extracted from [Kim 2003])

time (TI) results in images where viable tissue appears almost black, while nonviable, fibrotic or scarred tissue is greatly enhanced (Figure 2.7). As described in [Simonetti 2001, Kim 2003], selecting the appropriate inversion delay time is critical to obtain the best imaging results. The TI is chosen so that the signal from normal myocardium is *null*, because at that point the contrast between normal and infarcted myocardium is maximized (Figure 2.8).

2.2.2 DE-MRI and Arrhythmias

Heart and circulatory system diseases cause 1.9 million deaths every year in the European Union, representing about half of all the deaths that occur in the European countries [1]. Among this group of diseases, ischemic cardiopathy is a major cause of death.

In patients with some kinds of cardiopathies, both ischemic and non-ischemic, complications can arise in the form of rhythm abnormalities such as ventricular tachycardia (VT), atrial tachycardia (AT) or atrial fibrilation (AF).

Ventricular arrhythmias represent the main cause of sudden deaths. Although the use of implantable cardioverter-defibrillator (ICD) prevents sudden deaths, discharges increase non-arrhythmic mortality, and for this reason ICDs are frequently combined

[1] `http://ec.europa.eu/enlargement/5th_enlargement/europe_at_heart_en.htm`

with other treatments. Arrhythmic substrate ablation (ASA) of the VT significantly reduces the incidence of shock in patients with ICDs.

The annual implant rate in Spain is close to 80 per million inhabitants, thus more than 3000 devices are implanted per year [Peinado Peinado 2008]. If we consider that around 65% of ICDs are related to ischemic cardiopathies and 6% of patients suffer at least one discharge during the first year after the implant [Thijssen 2011], we get an idea of the problem's magnitude. This explains why there is an increasing interest for non-invasive ways of identifying the myocardial substrate of VTs.

Currently, a major diagnostic and treatment technique used in patients with rhythm disorders is Electro-anatomic mapping (EAM), which consists on carrying out through catheterization contact measurements of the electrical activity of the heart at various points and representing these measurements on a three-dimensional map. This technique for electro-anatomical map visualization is very useful, however the construction of detailed maps involves measuring several hundred locations within the patient's heart. This means that the duration of this type of procedure can be up to several hours, during which the patient is exposed to certain doses of ionizing radiation (X rays) due to the fluoroscopy used during the procedure, as well as risks inherent to any catheterization procedure.

On the other hand, it is now well known that DE-MRI images of the patients' heart contain valuable diagnostic information for identifying the substrate of VTs on both ischemic [Bello 2005, Ashikaga 2007] and non-ischemic [Nazarian 2005] cardiomyopathies. Moreover, it has been shown [Reddy 2004, Bogun 2009] how the information obtained in DE-MRI studies can be combined with electro-anatomic maps. In [Reddy 2004] the three-dimensional DE-MRI volumes are registered to EAM maps, and discussion is provided on how this can help the interpretation of those maps by providing additional anatomic information. Additionally both [Reddy 2004] and [Bogun 2009], for ischemic and non-ischemic cardiomyopathy respectively, describe how to manually segment the scar tissue and superimpose the segmentation with EAM information. Given the known correlation between scar tissue identified in DE-MRI images and the VT substrate, such joint visualization is of great value for guiding the arrhythmia substrate ablation procedure.

2.2.3 Latest Trends and Open Issues

In patients with VT, the ability to identify the optimal ablation zone is of key relevance for the success of ablation procedures. The re-entry circuits [de Bakker 1988]

causing VT in patients with chronic ischemic heart disease are known to be related with slow conducting zones [Stevenson 1993] adjacent to the infarct scar tissue. Although recent developments in three-dimensional electro-anatomical mapping systems provide a great help for understanding and identifying those re-entry circuits during ablation procedures, there is currently great interest in extracting similar information from the non-invasive images acquired with DE-MRI.

In this sense, several studies [Yan 2006, Roes 2009, Schmidt 2007] have hypothesized that heterogeneous tissue identified in DE-MRI images is related to slow-conduction zones found in EAM and therefore can be used to predict risk of spontaneous VT. Later in this manuscript (section 6.4) we will discuss a method, developed as part of this work, enabling the visualization of "viability" maps based on the information contained in DE-MRI alone, which has been used to provide additional evidence [Perez-David 2011, Halperin 2011] on the correlation between slow-conduction channels causing SMVT and the heterogeneous tissue corridors that can be identified in such "viability" maps.

It is clear that if DE-MRI images are to be used as an information source for identifying potential VT substrate, it is desirable to get the best DE-MRI images possible. However, while current standard DE-MRI protocols provide a reasonable in-plane resolution of $\sim 1.5mm$, the axial resolution is quite worse, with typical values ranging from $5mm$ to $10mm$ both for slice separation and thickness [Kim 1999, Kim 2003, Simonetti 2001, Reddy 2004, Kim 2008, Bogun 2009, Peters 2009]. While higher resolution DE-MRI images have been demonstrated on ex-vivo swine model [Ashikaga 2007], mechanically ventilated pigs and respiratory navigator-gated MRI sequences [McGann 2008, Kino 2009, Viallon 2011, Moghari 2012], are still not broadly available.

On the other hand it is common that acquisition protocols include several DE-MRI anatomical views of the patients' heart. Short-axis (with slices orthogonal to the left-ventricle axis), 2-chamber long-axis (with slices crossing left ventricle and atrium) and 4-chamber long-axis (with slices crossing the four chambers) are usually acquired. This situation matches the problem statement in section 1.2 and this application constitutes a very appealing target for the methods developed in this work. The application of three-dimensional wavelet fusion to DE-MRI images will be covered in chapter 6.

Those who cannot remember the past are condemned to repeat it.

George Santayana

Chapter 3

State of the Art: Image Enhancement

The objective of this work, as we stated in chapter 1, is to propose, develop and test new image processing methodologies for dealing with multi-view image fusion problems. In this kind of problems the imaging hardware is able to produce several views of a single object, but each of the views alone does not contain enough information as to represent the object being imaged within the quality parameters (sharpness, sampling frequency, ...) required by the final application.

When this happens, the available information about the object under study is spread among the set of acquired views, while it is desirable to have them all merged or *fused* into one single volume with a richer information content. As we explained in section 1.2, a typical multi-view imaging scheme, each individual views suffers from a combination of two different issues:

Spatial Non-Uniformity The image formation process could lead to differences in image sharpness along the whole object, meaning that the object will not be acquired with uniform quality. On the contrary, some regions could appear bright and sharp, while others will be dimmed and blurred. Moreover, in some cases the object will not be completely imaged in any of the individual views alone.

Frequency Non-Uniformity The frequency content of each of the views could be unequal for different orientations. This is frequently the case when the image formation process has an intrinsic sampling anisotropy, where each view has an

"axial direction" with sparser sampling and therefore worse representation of the object details, i.e., less high-frequency content.

Before we explain, in chapter 4, the proposed methodology for dealing with this kind of problems, in this chapter we review the state of the art by going through several families of algorithms which are somehow related to these or other similar problems. For this purpose, we will first set a common notation framework and then we will cover several existing single-view and multi-view algorithms. Finally we will summarize the reasons why we consider that a new multi-view methodology could be of help in situations where none of the existing algorithms provide a fully satisfactory solution.

3.1 Definitions and Notation

In the following sections we will describe several algorithms and methods related to single- or multi-view image restoration problems. As different methods come from different fields of knowledge, sometimes the notation used might vary significatively. Therefore, an effort has been made to use a uniform notation along the chapter.

We will be talking mostly about *images* as most of the methods were initially proposed for two-dimensional data. However, we will be using the same notation both for *images* and *volumes*, as all the operations involved in the methods are similarly defined for two- and three-dimensional data.

Finally, it is worth mentioning that we will be using two different notations to refer to an image (or volume) depending on how the image data is arranged.

- We use lower case letters $(f, g, \ldots)$ to represent images stored in a $M \times N$ matrix, with the $\cdot$ operator meaning element-wise multiplication and the $*$ operator representing the discrete convolution: $(f * h)(i, j) = \sum_{u,v}(f(u, v) \cdot h(i - u, j - v))$

- On the other hand, we will use upper case letters $(F, G, \ldots)$ to represent images stored as a one dimensional $(M \cdot N) \times 1$ array using lexicographical order, with the $\cdot$ operator meaning matrix multiplication and linear transformations represented by two-dimensional matrices: $G = H \cdot F \Longrightarrow G(i) = \sum_j H(i, j) \cdot F(j)$

3.2 Single Image Restoration Methods

3.2.1 Image Deconvolution

In many imaging applications, the obtained image has suffered some kind of degradation or *blur* which can be modeled as a linear process. Usually the most important blur is the one due to the detector's Point Spread Function (PSF), which models the response of the detector to a point source. Deconvolution can be actually considered as a particular case of Super-Resolution 3.3.1, when there is only one observed image, so it can be posed as a linear inverse problem as in eq. 3.15. When the PSF is space-invariant, the image degradation can be also modeled as a convolution with the PSF h (eq. 3.1), which usually acts as a low-pass filter.

$$g = h * f + n \tag{3.1}$$

Deconvolution algorithms (a.k.a. deblurring or image restoration) try to improve image quality by undoing the degradation suffered during the image formation process. A lot of different approaches to deconvolution exist, many of them coming from the field of optical imaging (microscopy, astronomy, etc) but there are applications in other fields like reflection seismology or almost any situation where a measurement process can be modeled like eq. 3.1. We will now describe some well-known algorithms for the deconvolution problem.

3.2.1.1 Linear Inverse Filtering

This family of methods solve the deconvolution problem by directly inverting the linear filtering in eq. 3.1. This can be done in a very efficient way by using the Fast Fourier Transform (FFT) as in eq. 3.2 (where $\hat{g} = \mathfrak{F}(g)$,$\hat{h} = \mathfrak{F}(h)$,$\hat{f} = \mathfrak{F}(f)$ and $\hat{n} = \mathfrak{F}(n)$ are the Fourier transforms of g, h, f and n respectively). The main drawback of these methods is that, as the PSF h is a low-pass filter, the high frequencies in the acquired image are dominated by noise, which is greatly amplified if direct inverse filtering is applied.

$$\begin{aligned} g &= h * f + n \\ \hat{g} &= \hat{h} \cdot \hat{f} + \hat{n} \\ f &\approx \mathfrak{F}^{-1}(\hat{h}^{-1} \cdot \hat{g}) \end{aligned} \tag{3.2}$$

Wiener deconvolution tries to minimize the noise amplification problem in frequencies with poor signal-to-noise ratio. This is achieved by implementing a Wiener filter adapted to the problem (eq. 3.3) which takes into account the signal's power spectral density S and the noise mean power spectral density N.

$$\hat{h}_W = \frac{\hat{h}^* \cdot S}{|\hat{h}|^2 \cdot S + N} \tag{3.3}$$
$$f \approx \mathfrak{F}^{-1}(\hat{h}_W \cdot \hat{g})$$

3.2.1.2 Lucy-Richardson Deconvolution

A very well-known and popular method in astronomy and microscopy is the Lucy-Richardson algorithm [Richardson 1972, Lucy 1974]. It is basically a Expectation Maximization (EM) [] solver for the Maximum Likelihood (ML) problem with Poisson noise. In astronomy, microscopy and other optical imaging modalities, the detection process typically suffers Poisson noise, so the deconvolution problem can be posed as the inversion of a linear system with noise like in eq. 3.4, where $N_p(\cdot)$ is a Poisson noise process, F and G are the ideal and measured images, respectively, and are represented as lexicographically ordered one-dimensional vectors, and H is a two-dimensional matrix representing a linear transformation that model the image formation process.

Lucy-Richardson's algorithm tries to find the estimate F_{ML} satisfying eq. 3.5, where $P(G/F)$ is the likelihood of observing G if the real image was F under the image formation model in eq. 3.4. Usually it is more convenient to minimize the negative logarithm of the likelihood $L(G,F) = -ln(P(G/F)$, which for the case of poisson noise is given by eq. 3.6.

$$G = N_p(H \cdot F) \tag{3.4}$$

$$F_{ML} = \underset{F}{\operatorname{argmax}}(P(G/F)) \tag{3.5}$$

$$L(G,F) = -ln(P(G/F)) = \sum_{i=1}^{M}[H \cdot F]_i - G^T \cdot ln(H \cdot F) \tag{3.6}$$

By following the iterative EM equation in eq. 3.7 (where $\cdot$ and $\circ$ represent matrix multiplication and element-wise multiplication, respectively, and H^T represents the transpose of H), the Lucy-Richardon algorithm converges to the ML solution. The

main disadvantages of this algorithm are its slow convergence rate and the lack of use of *a priori* information about the object being observed f.

$$F^{n+1} = F^n \circ \left(H^T \cdot \left(G \circ \frac{1}{H \cdot F^n} \right) \right) \tag{3.7}$$

3.2.1.3 Maximum A Posteriori (MAP)

While ML methods search for the solution which just maximizes the likelihood of observing the acquired image, using only information about the image formation process, there is a related family of methods which incorporate *a priori* information about the object being observed. Such methods are usually known as Maximum A Posteriori (MAP).

$$G = H \cdot f + N$$

$$F_{MAP} = \underset{F}{\operatorname{argmax}}(P(F/G)) = \underset{F}{\operatorname{argmax}} \left(\frac{P(G/F) \cdot P(F)}{P(G)} \right) \tag{3.8}$$

$$F_{MAP} = \underset{F}{\operatorname{argmax}}(P(G/F) \cdot P(G)) \tag{3.9}$$

In eq. 3.9 $P(G/F)$ represents the likelihood of observing G if the actual object was F, while the second term $P(F)$ stands for the *a priori* probability of F, for which a model needs to be assumed regarding the kind of objects we expect to observe. Maximizing the product in eq. 3.9, or minimizing its negative logarithm, is equivalent to maximizing the *a posteriori* probability in eq. 3.8, giving name to this kind of methods.

Equations eq. 3.10 and eq. 3.11 show examples of the resulting functionals for poisson and gaussian noise models respectively, both of them assuming $\|A \cdot F\|^2$ as the smoothness prior (with A being a high-pass linear operator) for the original image. The scalar γ weights relative influence of the likelihood and the *a priori* probability in the final solution. Optimization approaches like gradient descent, conjugate gradients, etc are typically used for minimizing those functionals [].

$$-ln(P(G/F)) - \gamma ln(P(F)) = \sum_{i=1}^{M} [H \cdot F]_i - G^T \cdot ln(H \cdot F) + \gamma \|A \cdot F\|^2 \tag{3.10}$$

$$-ln(P(G/F)) - \gamma ln(P(F)) = \|H \cdot F - G\|^2 + \gamma \|A \cdot F\|^2 \tag{3.11}$$

3.2.2 Wavelet Image Denoising

Another classical group of image restoration methods is image denoising, where the purpose is to reduce the noise level in an image or volume, without losing too much information along the process. There are countless approaches to the image denoising problem (like Median Filtering [Tukey 1976], Anisotropic Diffusion Filtering [Perona 1990, Catte 1992], Total Variation Filtering [Rudin 1992, Rudin 1994], Bilateral Filtering [Paris 2009, Bennett 2007] or Wavelet Denoising [Donoho 1994, Chang 2000]), and it is not our goal to carry out a deep study of their advantages and drawbacks.

However we consider that wavelet denoising (a.k.a. wavelet thresholding or wavelet shrinkage) deserves an special mention in this document. The reason is that it is based on a wavelet decomposition of the image being denoised, and therefore it seems very nicely suited for a potential future integration with the wavelet fusion method we tacke in the present PhD thesis.

The basic workflow of wavelet denoising methods is the following:

1. Apply the N-level wavelet transform of choice to the input data (signal, image or volume) X^0. In this way the original data is decomposed onto coarse approximation coefficients X^N and detail wavelet coefficients $Y^{(n,m)}$ where n is the scale index and m the orientation index.

$$\mathcal{W}(X^0) = \left\{ Y^{(1,1)}, \ldots, Y^{(n,m)}, Y^{(N,7)}, X^N \right\}$$

2. To each individual detail wavelet coefficient, a threshold [Donoho 1994] rule is applied. This can be either a hard threshold (eq. 3.12), or a soft threshold (eq. 3.13).

$$\eta_H(y,t) = \begin{cases} y - t & y \geq t \\ 0 & |y| < t \\ y + t & y \leq -t \end{cases} \tag{3.12}$$

$$\eta_S(y,t) = \begin{cases} y & |y| \geq t \\ 0 & |y| < t \end{cases} \tag{3.13}$$

3. Apply the inverse wavelet transform $\mathcal{W}^{-1}$ to the coarse approximation coefficients and the thresholded wavelet coefficients to obtain the denoised data.

3.3 Multiple Image Restoration Methods

3.3.1 Super-Resolution

Under the name of Super-Resolution we can find several different techniques, all of them with the common objective of enhancing resolution by combining multiple observed low-resolution images to obtain a reconstructed high-resolution image [Park 2003].

Although the Super-Resolution (SR) field was born in the Computer Vision community, many applications adapting SR algorithms for their use in medical imaging applications have appeared in the last years [Greenspan 2008]. One example is the application of SR to MRI stacks [Greenspan 2002] to improve resolution in the slice direction (Figure 3.1). Applications to Positron Emission Tomography (PET) have also been demonstrated [Kennedy 2006, Kennedy 2007] (Figure 3.2).

Super-Resolution belongs to the family of inverse problems, i.e., finding some unobserved variables which are related to the observed data through a known process (known as the direct problem). In Super-Resolution the direct problem is given by the image formation process model, which can usually be expressed as a linear system like eq. 3.14, where F is the unobserved high-resolution image we want to recover, G_i are each of the observed low-resolution images, H_i is the linear transformation matrix containing the observation model which includes the spatial transformation and other effects such as blurring for the corresponding image G_i and N_i is the additive noise of each observation. By using this matrix representation, we can group the N equations and model the whole process for all the observed images simultaneously as in eq. 3.15, where G, H and N are the combination of all the observed images G_i, observation matrices H_i and additive noises N_i, respectively.

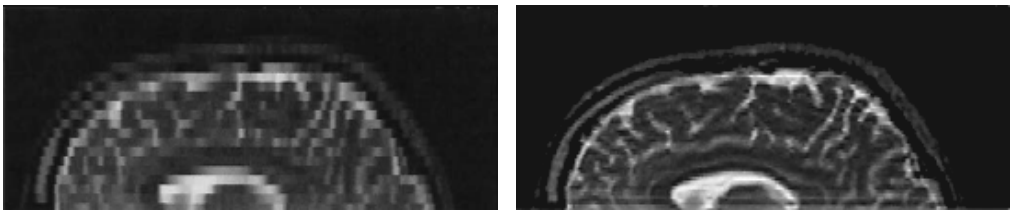

(a) One of the 3 acquired low-resolution stacks

(b) Inter-slice Super-Resolution result

Figure 3.1: Example of MRI Inter-slice SR (extracted from [Greenspan 2002]).

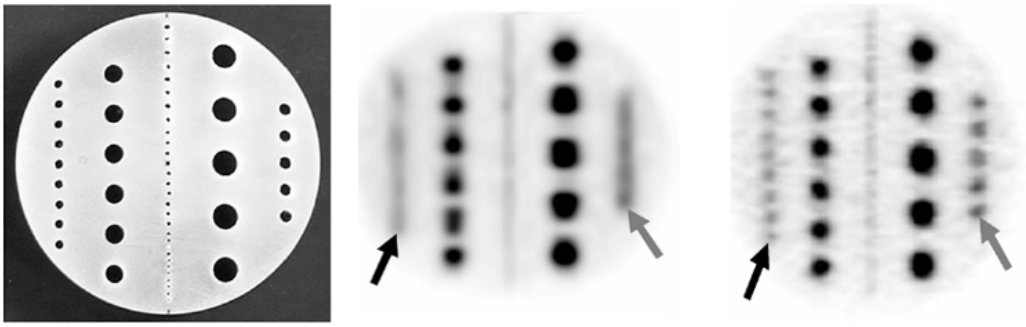

(a) Phantom photograph

(b) One of the 4 acquired images (shifted along vertical direction)

(c) Super-Resolutin result from the 4 acquired images

Figure 3.2: Example of SR in PET (extracted from [Kennedy 2006])

$$G_i = H_i \cdot F + N_i, i = 1...V \tag{3.14}$$

$$G = H \cdot F + N \tag{3.15}$$

The obtained model equation 3.15 is a classic image restoration problem, and can be solved using several approaches like Maximum Likelihood (ML) estimator, Maximum A Posteriori (MAP) or Projection Onto Convex Sets (POCS) [Elad 1997]. Some of this inverse linear problem methods were already discussed in sections 3.2.1.2 and 3.2.1.3.

3.3.2 Multi-View Deconvolution

Recently, extensions to classical deconvolution algorithms have been proposed for dealing with multi-view problems. In [Swoger 2007] and [Krzic 2009] extensions of the Lucy-Richardson and Maximum A Posteriori algorithms, respectively, are described for approaching multiple-view situations in fluorescence microscopy.

The basic idea is the same as in the original Lucy-Richardson and MAP algorithms, the main difference being that on each step the current estimate F^n is compared to each of the available views, and the multiplicative update is thus the mean of the updates suggested by all of the views. For instance, the update equation in [Krzic 2009] is shown in eq. 3.16.

$$F^{n+1} = F^n \circ \frac{1}{V} \sum_{i=1}^{V} \left(PSF_i^T \cdot \left(G_i \circ \frac{1}{F^n \cdot PSF_i} \right) \right) \tag{3.16}$$

Where $G_i, i = 1 \dots V$ are the observed images or views, $PSF_i, i = 1 \dots V$ are the PSFs associated to each of those views and PSF_i^* represents the adjoint of PSF_i.

We mentioned before how Super-Resolution and Deconvolution are two closely related problems, and most of the generic techniques used to solve them are common (ML, MAP, etc). This statement becomes evident when we talk about multi-view deconvolution, which is essentially the same problem as Super-Resolution. In fact, the iterative algorithm in eq. 3.16 is equivalent to the *Iterative Back-Projection* algorithm described in [Irani 1991], just changing the additive update for a multiplicative update, due to the use of a poisson noise model in the former in contrast with the gaussian noise model of the later.

Further improvements to such multi-view deconvolution approaches have been proposed in [Temerinac-Ott 2011], extending the algorithm in order to enable modeling of spatially-variant PSFs, which are commonly obtained in some imaging techniques like Light-Sheet Fluorescence Microscopy (LSFM), as described in section 2.1.2.

3.4 Image Fusion Methods

Image fusion refers to a family of techniques, typically used with 2D images, that can effectively combine information from different sources into a single composite image [Piella 2003]. Some examples of applications are:

Multi-focus Image Fusion Several images acquired with different depth of focus are combined, giving an image where all the scene appears in-focus [Li 1995] (Figure 3.3).

Satellite Image Fusion In current satellite imaging of earth, design constraints usually result in an inverse relationship between spectral and spatial resolution. This means that we can get both panchromatic (no color information) images with high spatial resolution, or multispectral (full color information) images with lower spatial resolution. Image Fusion is widely used in this situation to merge the spatial and color information into a single image [González-Audícana 2005] (Figure 3.4).

Medical Image Fusion Applications have also appeared in the field of medical imaging both for multi-modality fusion, as in [Guihong 2001] where the goal is to create a single image combining different anatomical information from CT and MRI, and for single-modality fusion, as in [Rajpoot 2009] where several images obtained by

Figure 3.3: Fusion of two images with different depth of focus (extracted from [Pajares 2004]).

Figure 3.4: Fusion of panchromatic and multispectral satellite images (extracted from [Pajares 2004]).

3D echocardiography are combined for filling missing anatomical information, extending the field of view and increasing image contrast.

Over the past years several approaches to Image Fusion have been developed, but lately the most popular one is based on multi-resolution transformations, like laplacian pyramids [Burt 1983, Burt 1984], morphological pyramids [Goutsias 2000], steerable pyramids [Liu 2001] and wavelet decompositions [Pajares 2004, Li 1995], which transform the original images to combine the coefficients on the transformed space and then apply the inverse transform to get the final result. Figure 3.5 shows the general scheme of such methods, involving the following steps:

1. Apply the multi-resolution transformation to each of the original images. Depending on the specific transformation, this will give a decomposition of each image into different scales and/or orientations.

2. Compute some kind of *activity* or *saliency* measure for each image's transformation coefficients. This can be simply the absolute value of the coefficients, or it can be more complex by including each coefficient's neighbors in the computation of the activity measure.

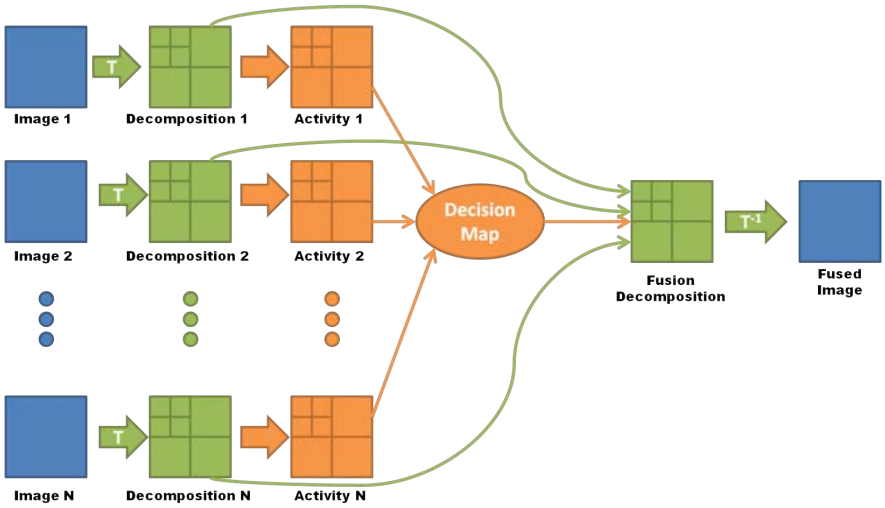

Figure 3.5: General Multi-Resolution Decomposition and Fusion Scheme

3. Based on the activity measures computed in the previous step, a *decision map* is generated which assigns weights which control, for every scale, orientation and location, how the corresponding coefficients from all original images should contribute to the coefficients of the composite image.

4. Using the coefficients from each original image together with the decision map, composite coefficients are generated for every scale, orientation and location in the transformed space.

5. Finally, the inverse transform is applied to the composite coefficients to get the final fused image.

3.4.1 Application to 3D Imaging

To the best of our knowledge, the application of wavelet-based fusion method to 3D multi-view imaging is very limited. We are only aware of the work presented in [Rajpoot 2009, Rajpoot 2011], where such an approach is applied to 3D echocardiography images for combining several anatomical views of the heart, extending the field of view and enhancing the contrast with respecto to individual views. Interestingly, previous work in fusion of multi-view 3D echocardiography was based on multi-view deconvolution techniques [Soler 2005].

In figure 3.6 we reproduce the fusion scheme used in [Rajpoot 2009, Rajpoot 2011], which is a particular case of the general scheme from figure 3.5, where the decision maps

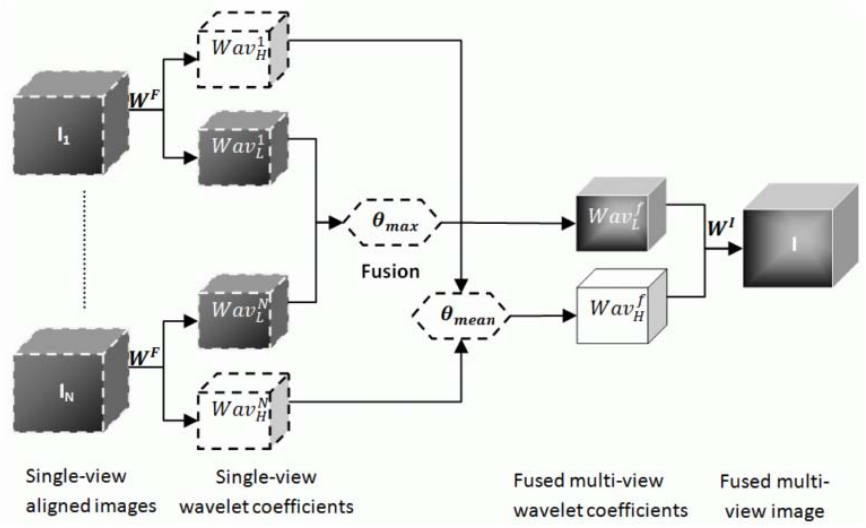

Figure 3.6: 3D Wavelet Fusion Scheme for Echocardiography. (reproduced from [Rajpoot 2009])

are computed so that the high information in the wavelet decomposition is averaged among views, and the maximum from the low frequency information is passed through to the final fusion. This particular choice of fusion rules is designed to reduce speckle noise in the original images and maximize the contrast of the big structures. Figure 3.7 shows an example of the results obtained in [Rajpoot 2009].

3.5 Where are we?

After reviewing all the relevant state of the art we conclude that there is still a need for methods which allow the combination of several three-dimensional views of an object or scene into a single volume which contains all the available and complementary information, instead of having such information spread among the acquired views.

Among existing methods, the ones closest to this type of problems are super-resolution and multi-view deconvolution, which are in fact two families of methods coming from different fields of study but addressing essentially the same problem. The drawback of existing methods in these families is the need for an extremely good knowledge of the image formation process, which is not always feasible to possess.

Therefore, we believe that for a broad range of problems (among which are the two specific applications explored in chapters 5 and 6) a new methodological approach to multi-view fusion which is based on the information available in the acquired data

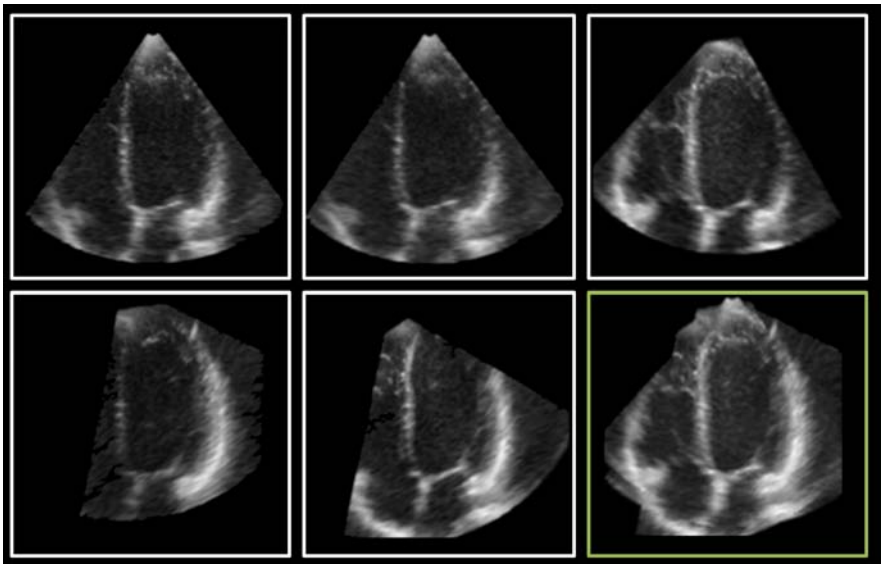

Figure 3.7: Fusion example of five 3D echocardiography images. A representative 2D slice from each volume is shown, followed by the same slice in the fused volume, which is shown in the lower-right corner. (reproduced from [Rajpoot 2009])

alone, without prior knowledge of the imaging hardware, can prove to be an extremely useful tool.

For this goal we have chosen to work in the wavelet transform domain, which seems appropriate for developing such a multi-view fusion method given the well-known capabilities of wavelet transforms for separating a volume's information content onto its different scale and orientation components. This capability will allow us to select the desired features on the transform space and combine them in the fused volume.

In fact, wavelet-based fusion methods have proven their performance on other image processing tasks like multi-focus or multi-modality image fusion, and they have even been applied to three-dimensional medical images for solving related but somehow different multi-view fusion problems [Rajpoot 2009, Rajpoot 2011].

In relation to those works it should be noted that, even though the basic fusion framework used coincides with the one we will be using, the specifics of the problem are quite different. In particular the images used in those works are almost isotropic, while the ones we will be applying our method to present highly anisotropic sampling. Additionally, the main goal in that case was to maximize the contrast for the main structures, while ours will be to preserve all the fine details visible in any of the original views. This mismatch in the fusion goals will translate into a completely different choice of the wavelet coefficient fusion rules, as we will explain later.

In the next chapter we will describe the details of the proposed fusion methodology, and later in chapters 5 and 6 we will cover its application to Fluorescence Microscopy and Magnetic Resonance Imaging, respectively, and results demonstrating the method's performance (including, when applicable, comparison with other existing methods) will be provided.

There is nothing insignificant in the world. It all depends on the point of view.

Johann Wolfgang Von Goethe

Chapter 4

3D Wavelet-Based Multi-View Fusion

4.1 Problem Description

Before describing the proposed method, we will now give a more formal definition to the general problem introduced in section 1.2, settling the notation we will be using along this chapter.

We model the object being studied as a three-dimensional scalar function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$, where $f(x, y, z)$ is the value of some scalar magnitude characterizing the object. The precise nature of this magnitude depends on the imaging modality used, for instance it could be the concentration of some fluorescent molecular probe (Fluorescence Microscopy) or some relaxation time of protons when exited under a strong magnetic field (Magnetic Resonance Imaging).

Next we define a *view* of the object as a discretely sampled version of the scalar function f. Although in principle any imaginable sampling pattern is possible, for this work we will be focusing on orthogonal grid sampling patterns, and so we represent the set of V acquired views from the same object as:

$$X_v : \Omega_v \rightarrow \mathbb{Z}; v = 1 \ldots V$$

where the size of each view is $N_{v_i} \times N_{v_j} \times N_{v_k}$ and its domain is defined as:

$$\Omega_v \equiv \{(i, j, k) | i, j, k \in \mathbb{N}, 0 \leq i < N_{v_i}, 0 \leq j < N_{v_j}, 0 \leq k < N_{v_k}\}$$

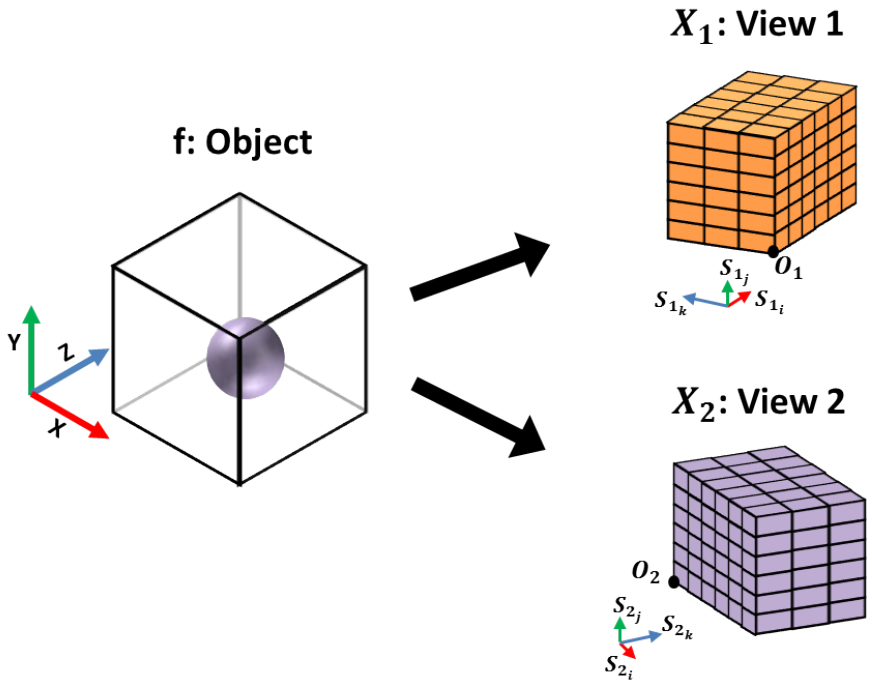

Figure 4.1: Representation of an object by different views with complementary information.

and the relation between the samples in each view X_v and the scalar values f from the imaged object is given by:

$$X_v(i,j,k) = f(O_v + (i,j,k) \cdot S_v) \quad where \quad S_v = \begin{pmatrix} s_{v_i} \\ s_{v_j} \\ s_{v_k} \end{pmatrix} \tag{4.1}$$

with $O_v = (o_{v_i}, o_{v_j}, o_{v_k}) \in \mathbb{R}^3$ representing the origin of the sampling grid and the orientation and dimensions of the sampling grid given by the three vectors $s_{v_i}, s_{v_j}, s_{v_k} \in \mathbb{R}^3$ which are grouped into the transformation matrix S_v. As we said before, we will focus on orthogonal grid, so the sampling vectors $\{s_{v_i}, s_{v_j}, s_{v_k}\}$ must form an orthogonal system, and the sampling distance along each direction is given by $\Delta_{v_i} = \|s_{v_i}\|$, $\Delta_{v_j} = \|s_{v_j}\|$ and $\Delta_{v_k} = \|s_{v_k}\|$.

In figure 4.1 we illustrate the problem description, showing the object being imaged together with the sampling patterns of two views. It can be appreciated how each view has a different sampling grid orientation and, as it will be typically the case in the applications we will focus on, the sampling distance is bigger in the k directions ($\Delta_{1_k} > \Delta_{1_i} = \Delta_{1_j}$ and $\Delta_{2_k} > \Delta_{2_i} = \Delta_{2_j}$).

There is another important aspect of the imaging process that must be taken into account. Up to now we are assuming that the sampling is ideal and there is no distortion due to the imaging process, which is clearly not the case on real imaging techniques. In general, we should assume that each imaging process has an intrinsic Point Spread Function (PSF) which alters the underlying object before the sampling takes place. Of course the PSF will in general be spatial-dependent and also different for each of the acquired *views* ($PSF_v = PSF_v(x, y, z)$). Therefore the imaging equation 4.1 should be modified and becomes:

$$X_v(i, j, k) = \tilde{f}_v(O_v + (i, j, k) \cdot S_v) \tag{4.2}$$

where

$$\tilde{f}_v(x, y, z) = (f * PSF_v)(x, y, z) \tag{4.3}$$

Given the above problem description, the goal of our method is to take a set of input volumes, each one corresponding to a different view of the same object (e.g. a patient's heart or a live embryo), and combine the complementary information available in all of them to create a new fused volume of a better overall quality than any original view alone:

$$\{X_1, X_2, \ldots, X_V\} \xrightarrow{Fusion} X_F$$

4.2 Proposed Methodology

In figure 4.2 we summarize the overall workflow of the proposed method, starting with the set of acquired views $\{X_v\}$, $v = 1 \ldots V$ and all the way through the final fused volume X_F. There are two distinct stages, *Image Registration* and *Fusion*.

4.2.1 Image Registration

In this first stage, all the views are aligned and brought into a common frame of reference. We arbitrarily choose the first view as the reference to which all the remaining views will be aligned. It can be appreciated in the figure, that together with the registration, the aligned volumes X_v^0 might be padded in order to ensure that even those parts of the object which appear in just one of the views are preserved in the final fused volume. Another important detail is the generation of *masks* K_v^0 which keep

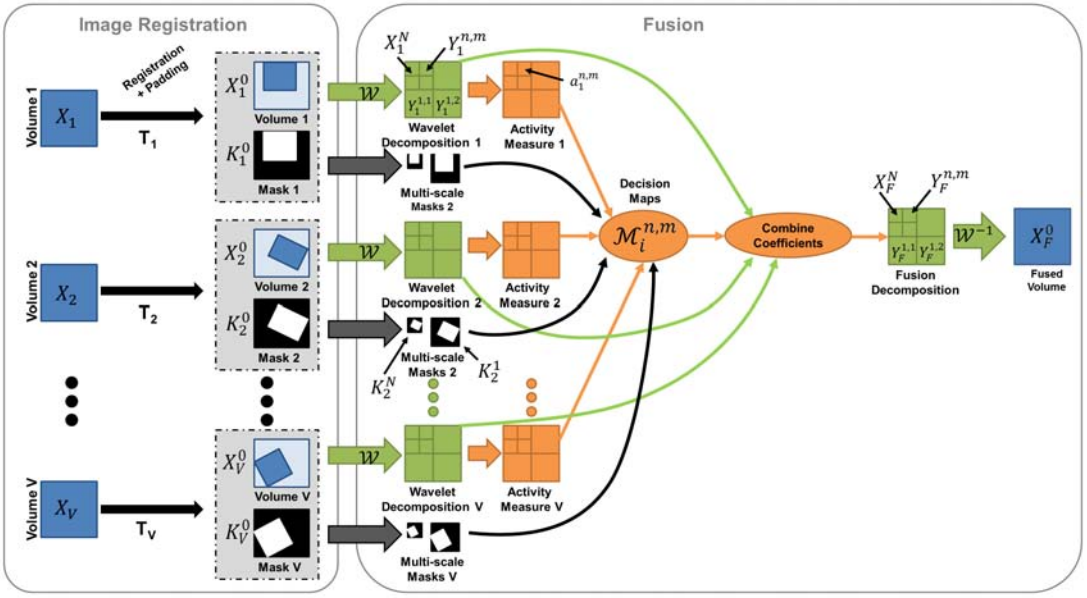

Figure 4.2: Starting with the input volumes $X_v, v = 1 \dots V$, a first registration stage aligns them to a common frame of reference which we arbitrarily choose to be the first view. After the registration we get a set of aligned volumes X_v^0, $v = 1 \dots V$ and a corresponding set of masks K_v^0 which are all passed to the fusion process itself, which we will explain along the following sections.

track of which voxels from X_v^0 contain actual information from the original views X_v. Later in section 4.2.2.2 we will describe how these *masks* are used during the fusion process for preventing the appearance of border artifacts.

The registration method is of key importance for the final output, as any miss-registration might lead to artifacts in the final fused volume. Depending on the source of the multi-view volumes, it might be possible to have some knowledge on the relative position of the views (for instance, the machine coordinates in MRI scanners, or the rotation angle of the micro-positioning stage in some fluorescence microscopes) which allows us to perform a pre-alignment without the need for a registration algorithm. However, due to practical problems like miss-calibration of the imaging hardware or movement of the object during the imaging process, this pre-alignment is usually not good enough and a posterior fine registration is necessary.

In general, and considering that one of the conditions enumerated in section 1.2 was that the object being imaged should not change significantly during the time it takes to acquire the set of views, the registration method we will be using a rigid transformation model. In certain situations an affine transformation model might be useful not for compensating deformations in the object, but for correcting distortions

caused by the imaging hardware.

The specific registration method selected is highly application-dependant, and a broad discussion of the available registration methods is beyond the scope of this work. We will instead describe the particular registration methods used for each of the studied applications in their respective chapters.

4.2.2 Wavelet-Based Fusion

As we mentioned earlier in section 3.4, the use of multi-resolution transformations has become one of the most popular approaches to image fusion [Burt 1983, Burt 1984, Piella 2003]. The ability of such transformations to decompose the information contained in an image onto several scales and/or orientations has proven quite useful for a wide range of image processing tasks, and in particular for image fusion.

In particular, the method we are about to describe falls in the general framework of wavelet-based image fusion [Li 1995, Pajares 2004, Rajpoot 2009, Rajpoot 2011]. The main idea behind such approach is to take advantage of the properties of multi-band image decomposition schemes like the discrete wavelet transform.

In this kind of decompositions, an input image (or volume) X^0 is decomposed onto several bands (each one corresponding to a specific scale and orientation, for instance) $\mathcal{W}(X^0) = \{Y^1, \ldots, Y^k, \ldots, Y^K, X^K\}$, where $Y^k, k = 1 \ldots K$ are the different bands in which the image is decomposed and X^K is the residual low-pass approximation of X^0.

Salient features in the input image X^0 become high-energy coefficients in at least one of the decomposition bands Y^k, while smooth regions become low- or zero-energy coefficients. So, roughly speaking, by applying a common multi-band decomposition to all the input images available, and then choosing the highest energy coefficients, a fused image combining salient features from all input images can be obtained.

The detailed process is outlined in the right side of figure 4.2 and, before explaining the different steps involved, we recall the inputs of the fusion method:

- A set of I volumes $X_v^0 (v = 1 \ldots V)$, each of them containing a different three-dimensional view of the object, which have already been registered, using the first volume X_1^0 as reference.
- The corresponding set of binary masks K_v^0, each of them indicating which voxels on each of the volumes X_v^0 contain relevant information. These masks are used, for instance, to tell which voxels in a registered volume come from the actual

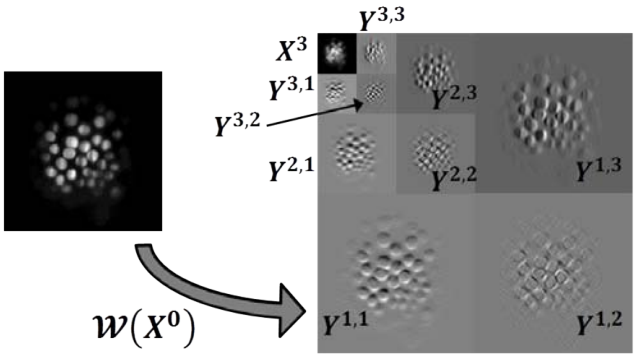

Figure 4.3: Wavelet decomposition example of a two-dimensional slice from a fluorescence microscopy image of a sea urchin embryo. The decomposition consists of three scales and three orientations for each scale, resulting in nine bands (N=3 and K=9), plus a remaining coarse aproximation. Although in this work we use undecimated wavelet transform, for illustration purposes this figure represents the classical decimated wavelet transform.

transformed volume, and which ones are just padding, allowing us to avoid border artifacts due to incomplete object coverage in some of the views.

We will now cover in detail each of the steps involved in the fusion process.

4.2.2.1 Transform input volumes

The first step consists on applying the wavelet transform of choice to each of the input volumes: $\mathcal{W}(X_v^0) = \{Y_v^1, \ldots, Y_v^K, X_v^K\}$, where $v = 1 \ldots V$ is the view index. Depending on the specific wavelet transform chosen, there will be a different number and arrangement of the bands. If three-dimensional Discrete Wavelet Transform (DWT) [Mallat 2008] is selected, we get 7 different orientations on each scale, so when N scales are used during the decomposition, we end up with $K = 7 \cdot N$ bands. In this case it is more practical to index the bands by scale ($n = 1 \ldots N$) and orientation ($m = 1 \ldots 7$) as in eq. 4.4

$$\mathcal{W}(X_v^0) = \left\{ Y_v^{(1,1)}, \ldots, Y_v^{n,m}, Y_v^{N,7}, X_v^N \right\} \tag{4.4}$$

Figure 4.3 shows an example of the decomposition of a single slice of a Light-Sheet Fluorescence Microscopy volume, using a two-dimensional DWT with three scales (N=3) and three orientations on each scale, resulting in $K = 3 \cdot N = 9$ bands.

The size of each of the bands will depend on the precise wavelet transform used. When using decimated DWT, for instance, the bands' size is halved on each level, so

that the total number of elements in the whole decomposition $\mathcal{W}(X_v^0)$ is always equal to that of X_v^0. Undecimated implementations of DWT, on the other hand, keep the size of each $Y_v^{n,m}$ and X_v^N equal to X_v^0s size, so each additional level increases the total number of elements in the decomposition.

For this work we have chosen to work with an undecimated implementation of the DWT, more precisely a custom three dimensional extension of Matlab's implementation (`swt2` and `iswt2` functions) of the undecimated DWT which is based on the *à trous* algorithm. The undecimated *à trous* implementation of DWT is a redundant transformation and, although more consuming in terms of computation time and computer memory, shows the shift-invariant property (which decimated DWT lacks) that makes it more robust when it comes to image processing tasks like de-noising [Coifman 1995] or image fusion [Amolins 2007, Gonzalez-Audicana 2004, Redondo-Tejedor 2007].

4.2.2.2 Generating multi-scale masks

In parallel to the transformation of the input volumes, the corresponding masks are processed in order to prepare adequate multi-scale masks: $K_v^n, (v = 1 \dots V)$. These multi-scale masks are designed to ensure that only valid voxels from the input volumes are taken into account when generating the fused volume.

For this purpose we define the support of any given wavelet coefficient ($Y_v^{n,m}(\cdot)$ or $X_v^n(\cdot)$), as the set of voxels (i, j, k) in the corresponding input volume which get involved in the computation of said coefficient. The size of the support will depend on the scale n and on the wavelet family used. If we use `haar` wavelets, for instance, the support will be a $2 \times 2 \times 2$ sub-volume for coefficients in scale 1, a $4 \times 4 \times 4$ sub-volume for scale 2, and so on.

If we denote by $\mathcal{S}\{X_v^n(\cdot)\}$ the support of the coefficient $X_v^n(\cdot)$, the multi-scale masks are built so that $K_v^n(\cdot)$ takes the value 1 if and only if the original mask K_v^0 is equal to 1 at all the voxels (i, j, k) belonging to $\mathcal{S}\left(X_v^n(\cdot)\right)$, taking the value 0 otherwise, as expressed by eq. 4.5.

$$K_v^n(\cdot) = \begin{cases} 1 & if \quad K_v^0(i,j,k) = 1, \forall (i,j,k) \in \mathcal{S}\{X_v^n(\cdot)\} \\ 0 & otherwise \end{cases} \tag{4.5}$$

The multilevel masks K_v^n ($v = 1 \dots V$ and $n = 1 \dots N$) depend only on the original masks K_v^0, the wavelet family used and the decomposition level n, but not on the band

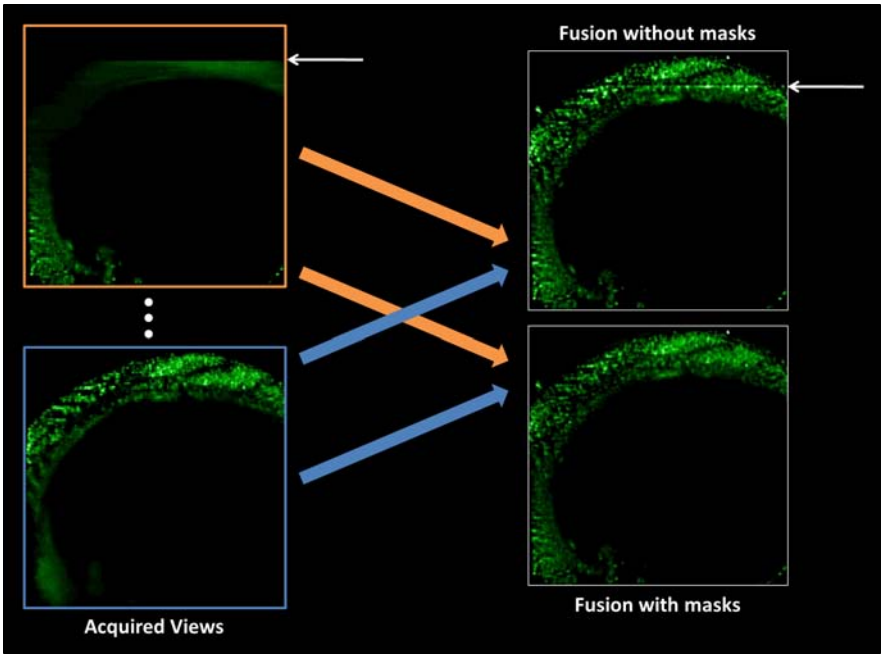

Figure 4.4: To the left we show a blow-up of 2D slices of two different LSFM views of a live zebrafish embryo, where it can be appreciated how one of the views does not cover the full embryo and presents a sharp border where the acquired volume ends (marked with a thin white arrow). To the right we show the result of applying the fusion method without (top-right) and with (bottom-right) the use of masks to indicate the acquired data limits. The fusion without masks shows a clear artifact (marked with a thin white arrow) while the fusion with masks is free of such border artifacts.

m, in other words K_v^n will be applied both to the low-pass approximation X_v^n and the decomposition bands $Y_v^{n,m}$.

In order to illustrate the motivation behind the use of this multilevel masks, figure 4.4 shows an example of how their use prevents the appearance of artifacts when some of the acquired views do not fully cover the object.

4.2.2.3 Computing Saliency Measure

The next step is to compute an activity or saliency measure at every location of each band from all the wavelet decompositions. This will later be used to decide which of the wavelet coefficients are related to salient features in the input volumes, so that they should be included in the fused volume. A common and simple choice, providing very good results, is to use the energy or the module of wavelet coefficients. In this context, the saliency measure for any given wavelet coefficient $Y_v^{n,m}(\cdot)$ is defined as:

$$a_v^{n,m}(\cdot) = \|Y_v^{n,m}(\cdot)\| \tag{4.6}$$

4.2.2.4 Computation of Decision Maps

This is the most important step in the process depicted in figure 4.2, and it might be considered the core of the fusion method. We mean by decision maps the weights that define how, for each location in every band of the wavelet decompositions (including the residual low-pass approximation), the coefficients from all the input volumes are combined to create the fused volume.

As already mentioned, this kind of methods aim at keeping the most salient features, at different scales and orientations, among the input volumes. Having the saliency measure from the previous step, a straightforward way to implement such behavior is to create decision maps $M_i^{n,m}$ that, for each location, scale and orientation, select the maximum saliency measure from the different input volumes. The multi-level masks are taken into account in order to avoid selecting coefficients whose support includes non-valid voxels from the input volumes. This can be achieved with the following decision map equation:

$$\mathcal{M}_v^{n,m}(\cdot) = \begin{cases} 1 & \text{if} \quad a_v^{n,m}(\cdot) \cdot K_v^n(\cdot) = \max_{j=1\ldots V} \left[a_j^{n,m}(\cdot) \cdot K_v^n(\cdot) \right] \\ 0 & \text{otherwise} \end{cases} \tag{4.7}$$

However, this definition is not adequate to be used on every decomposition band. For instance, depending on the size of the objects present in the input volumes, some scales of their wavelet decomposition might be dominated by noise, so selecting the maximum coefficients on those bands would lead to increased noise in the fused volume. As an alternative for bands featuring noisy components, selecting the minimum saliency coefficients instead of the maximum leads to better results. For instance, figure 5.14 in chapter 5 shows a comparison demonstrating how using minimum saliency selection on some bands can improve the fusion results.

With respect to the low-pass approximations corresponding to each input volume, they should contain similar information after extracting the relevant features into the different bands. For this reason it is usually preferred to average all the low-pass approximations, taking into account only those coefficients whose associate multi-level

mask equals 1, which can be implemented by defining a different decision map equation, which is simply:

$$\mathcal{M}_v^{n,m}(\cdot) = \frac{K_v^n(\cdot)}{\sum_{j=1}^{V} K_v^n(\cdot)} \tag{4.8}$$

Other possible strategies are easy to incorporate into the framework by just creating alternative definitions of the decision maps.

4.2.2.5 Combination of Wavelet Coefficients

Once the decision maps have been computed, the next step consists on using them for combining the coefficients from all the input volume decompositions, generating the wavelet decomposition of the fused volume:

$$Y_F^{n,m}(\cdot) = \frac{\sum_{j=1}^{V} \left(\mathcal{M}_j^{n,m}(\cdot) \cdot Y_j^{n,m}(\cdot) \right)}{\sum_{j=1}^{V} \mathcal{M}_j^{n,m}(\cdot)} \tag{4.9}$$

As for the residual low-pass approximation of the fused volume X_F^N, it is computed in the same way by averaging the low-pass approximations of all the available volumes $X_v^N, i = 1 \dots V$.

4.2.2.6 Inverse Transform

Finally, the wavelet decomposition generated in the previous step is inverted to get the final fused volume.

$$\mathcal{W}^{-1}\left(\left\{ Y_F^{1,1}, \dots, Y_F^{n,m}, Y_F^{N,7}, X_F^N \right\} \right) = X_F^0 \tag{4.10}$$

4.2.3 Implementation Details

The whole fusion framework has been implemented in Matlab. Until version `R2010a` there was no 3D implementation of the Discrete Wavelet Transform in Matlab, and starting with version `R2010a` the decimated version of the DWT was included (functions `dwt3` and `idwt3`), but not the un-decimated DWT. Therefore, a custom implementation of the un-decimated DWT (with two new functions, `swt3` and `iswt3` was developed and used for the fusion framework.

However, as we mentioned before, one drawback of the un-decimated DWT is the fact that each band has the same size as the original volumes. This means that when we compute a full 3-level decomposition of a given volume we end up with a wavelet decomposition which takes 22 times (3×7 bands $+1$ low-pass approximation) more memory than the original volume (in fact it is even worse, as input data voxels are usually stored as one or two byte unsigned integer, while wavelet coefficients are stored as 8-byte double precision float numbers). A quick calculation tells us that for a $256 \times 256 \times 100$ volume (usual in DE-MRI images, for instance) we need 1.1GB of memory for the full wavelet decomposition.

This is too big when we consider that in order to perform the fusion process we need to decompose all the available views and also store the fusion decomposition while we compute it. In order to deal with this problem, an *off-line* implementation of the fusion process was developed, in which each of the partial steps we described in the previous sections is saved to the filesystem and erased from the RAM, and it is later loaded back from the filesystem when needed. While this has an impact on the method's performance and on its storage usage, it allows us to perform the fusion on bigger datasets than a regular implementation that would store everything in the RAM.

While this approach was enough for allowing us to work with DE-MRI images, datasets from LSFM are even bigger, with $600 \times 600 \times 600$ volumes being quite common and with some datasets reaching $1000 \times 1000 \times 1000$ volumes. If we repeat the previous computation we see that a single 3-level wavelet decomposition for one of such volumes would need $\sim$ 35 and $\sim$ 160 GB of RAM, respectively, with every individual band taking ~ 1.5 and ~ 7.5 GB.

Clearly for this kind of datasets even the *off-line* implementation is not enough, as it might be impossible to keep even a single band of the wavelet decomposition in memory. In order to solve the problem and enable the use of the fusion method for arbitrarily large volumes, we decided to take advantage of the fact that all the steps taking place along the process are spatially-local.

For this purpose we developed yet another implementation, in which the input volumes are broken onto small pieces (with a size that can be adapted to the available RAM) and all the fusion process steps are performed for each piece separately. We can then stitch the pieces back together and get the same result as if we had processed the full volumes at once. An additional benefit of such approach is that it is very well suited for parallelization, which would allow us to take advantage of the multi-core capabilities of current CPUs.

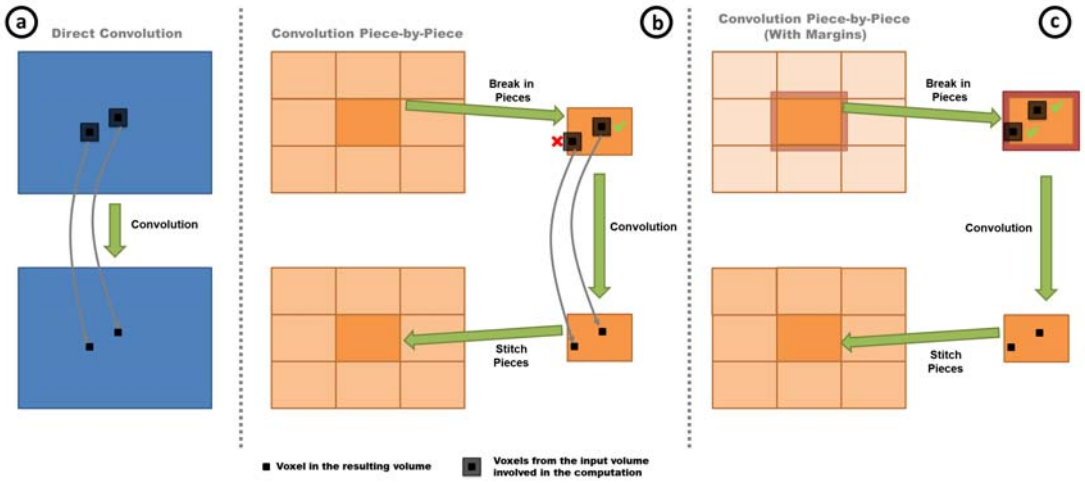

Figure 4.5: Explanation of the need to use additional margins when performing a convolution in a piece-by-piece basis. (a) Regular convolution performed directly on the full volume. (b) Piece-by-piece convolution. The volume is broken onto several pieces, the convolution is performed on each piece separately and the result is stitched back together to get the final result. It can be appreciated how the convolution result on voxels close to the piece border will be different from the result of direct convolution for the same voxel. (c) Piece-by-piece convolution, with an overlap among different pieces, so that the voxels close to the border of each piece are *wasted* and do not get stitched in the final result. When the amount of overlap is chosen properly, this approach yields results identical to those of the direct convolution.

There is however one pitfall when such an approach is used: special care must be taken in order to avoid problems with the borders of the pieces, which come from the use of convolutions in the implementation of the wavelet decomposition operations (`swt3` and `iswt3`). As explained in figure 4.5, when we perform a convolution operation on a piece of a bigger volume, not all the resulting voxels will be equivalent to the result we would obtain if the convolution took place on the original volume itself.

This means that each time we carry out a convolution we *waste* the border voxels of the small pieces (with that border having a size related to the support of the filter we are convolving with). Along the fusion process, each piece gets convolved $2 \times N$ times (N being the number of levels), N times for the decomposition and N times more for the inverse wavelet transform. This means that once we know the number of levels and the wavelet family (which defines the filters for the convolutions) we can compute the cumulative border that will be wasted in each piece, and overlap the pieces by that amount. In that way we can later discard the border of each piece before carrying out the stitching and get the correct result, as shown in figure 4.5. Clearly this has an impact on performance, as some extra computations are performed that will be discarded in the end, but it is a small price to pay when compared to the benefits of

this approach.

So quick bright things come to confusion.

William Shakespeare

Chapter 5

Multi-View 3D Microscopy on Live Embryos

5.1 Introduction

3D+time images of fluorescently labeled cells in live model organisms are essential to developmental biology [Muzzey 2009, Truong 2011, Dzyubachyk 2010]. When this kind of images is acquired with enough temporal and spatial resolution, tracking of every single cell and reconstructing the cell lineage tree becomes possible [Olivier 2010, Swoger 2010]. Typical optical microscopy techniques (like confocal or two-photon laser scanning microscopy) provide images with resolution along the optical axis considerably worse than lateral resolution. Moreover, image quality gets progressively worse as light travels deeper inside the specimen, meaning that for relatively large specimens (such as zebrafish embryos or larvae) it is impossible to get good images of the whole embryo.

To overcome this kind of limitations other microscopy techniques have been developed. Spinning-disk microscopy [Graf 2005], for instance, provides much higher imaging speed compared to scanned techniques. In theta-microscopy [Stelzer 1994], separate illumination and detection optical paths along orthogonal directions are used in order to improve the axial resolution. In a more recent approach called light-sheet based fluorescence microscopy (LSFM) [Huisken 2004, Huisken 2009] wide-field detection is combined with a lateral light-sheet illumination along the focal plane of the detection objective. LSFM is a powerful tool for 3D+t studies of live embryos in developmental biology, and is claimed to offer several advantages over conventional confocal laser scanning microscopy (CLSM) technique:

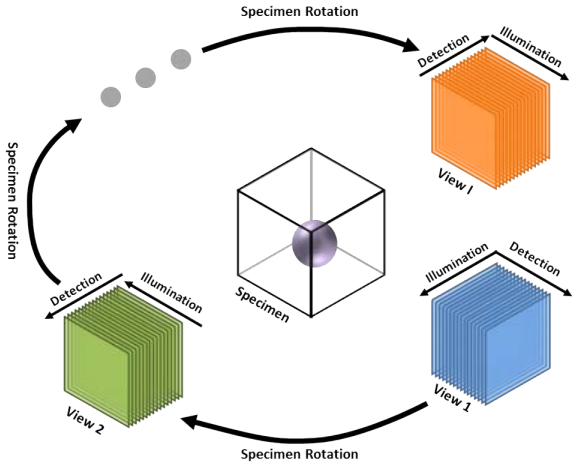

Figure 5.1: General scheme of the acquisition of several volumes or views of the specimen, each of them consisting on a stack of slices or planes orthogonal to the optical detection axis

- Unlike CLSM where a large part of the specimen is illuminated when recording a single point or plane, LSFM illuminates only the imaged plane, greatly reducing the overall photo-damage of the specimen.

- The system's point spread function (PSF) is the combination of the light-sheet shape and the detection objective's PSF. For this reason, when low numerical aperture detection lenses are used, LSFM achieves significantly better axial resolution than other techniques such as confocal or two-photon fluorescence microscopy.

- In LSFM wide-field detection is used, meaning that a full plane of the volume is acquired at once by an array of detectors (typically a CCD camera). In contrast, laser scanning microscopy (either confocal or two-photon) acquires one pixel at a time. This means that, for identical frame rates, LSFM has much longer per-pixel measurement time, which translates in a higher number of collected photons per-pixel, higher signal-to-noise ratio and dynamic range.

Another important feature of typical LSFM implementations is their ability to perform multi-view imaging, i.e. to obtain several volumes of the sample with different orientations. This is usually achieved by mounting the specimen under study on a rotating stage, and it takes advantage of the large working distance of water dipping lens objectives or of relatively low NA objectives, which makes the implementation of the rotation stage easier. Each view is a three-dimensional volume formed as a stack of

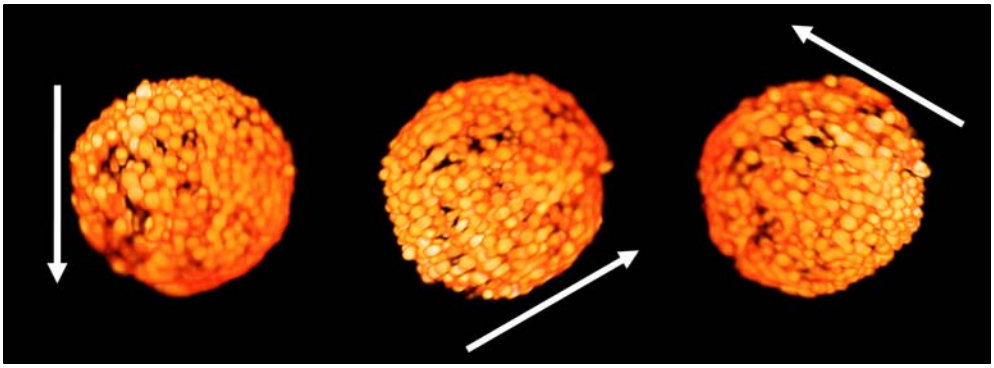

Figure 5.2: Example of multi-view acquisition from a sea urchin embryo, at approximately 20 hours post fertilization (hpf).

planes acquired sequentially. General schemes of LSFM can be found in the literature [Huisken 2009] and in figure 5.1.

Such multi-view acquisition capability is desirable because, despite the previously mentioned advantages of LSFM over other fluorescence microscopy techniques, single views still show a series of problems (see figure 5.2), like an axial resolution still worse than in plane resolution, slices increasingly dim and blurred with depth, etc, which will be further described below. An example is given in figure 5.2, which shows 3D renders of the three different views acquired from a sea urchin embryo at 20 hours post fertilization (the embryo has $\sim$ 800 cells at that time). White arrows indicate the detectors optical axis, with arrow heads pointing away from detection objective. Note that the regions close to the arrows heads are quite blurred. For this dataset three views were acquired with a 120 rotation between consecutive views.

5.1.1 Problem description

Combining the different views is still an open problem in LSFM, despite recent contributions in the field. In this context, this chapter deals with the proposal and validation of a method designed to deal with the multi-view volume fusion problem in LSFM. There are mainly two ways in which different three dimensional fluorescence microscopy acquisitions can be complementary in a multi-view setting [Swoger 2007]:

- Even though LSFM offers a relatively small axial PSF, due to the use of orthogonal illumination and detection axes, axial resolution remains usually worse than in-plane resolution. The in-plane and axial resolution throughout the specimen's volume change roles while acquiring the different views, as the microscope rotates

around an axis orthogonal to the detection axis. Two orthogonal views offer the most complementary information in terms of PSF

- For highly diffusive samples (which applies to living tissues even when relatively transparent) excitation and emission of light far from the excitation and detection objectives respectively gets severely attenuated/absorbed as it travels through the sample, which results in increasingly degraded image quality deeper into the sample. For large specimens such as the zebrafish embryo or larva, each view will contain useful information for less than half of the sample. The complementarity of views taken from different angles to compensate for this effect is quite clear.

The goal of our method is to take a set of volumes, each one corresponding to a different view of the same sample (e.g. live embryo or larva of chosen animal models), and combine the complementary information available in all of them to create a new fused volume of a better overall quality than any original view alone.

5.1.2 Multi-view fusion in LSFM

The multi-view imaging capability of LSFM provides extra information with respect to a single-view approach, as those regions of the specimen which are acquired with lower quality in one view will appear sharper and brighter in a different view.

To be fully useful, the information of the sample distributed among several volumes by multi-view imaging should be combined into a single volume (see figure 5.3). Several techniques have been recently proposed for this task [Krzic 2009, Swoger 2007, Preibisch 2008, Temerinac-Ott 2011].

One approach is to pose the problem as a multi-view deconvolution, which can be solved using extensions of classic iterative deconvolution algorithms like Richardson-Lucy [Krzic 2009] or Maximum A Posteriori [Swoger 2007] to the multi-view situation, by updating the estimate using one view at a time. This kind of algorithms performs fine when the actual PSF is well approximated and used in the deconvolution process. Its main drawback is precisely the need for a good estimate of the PSF, which is particularly difficult to obtain when the PSF is spatially variant as observed in the LSFM imaged volume due to the shape of light sheet and thickness of specimens. In [Temerinac-Ott 2011] an spatially-variant Richardson-Lucy is proposed, where the PSF is dynamically estimated based on fluorescent beads embedded in the sample. Finally, it is worth mentioning that these PSF based algorithms require an extremely precise registration of the views.

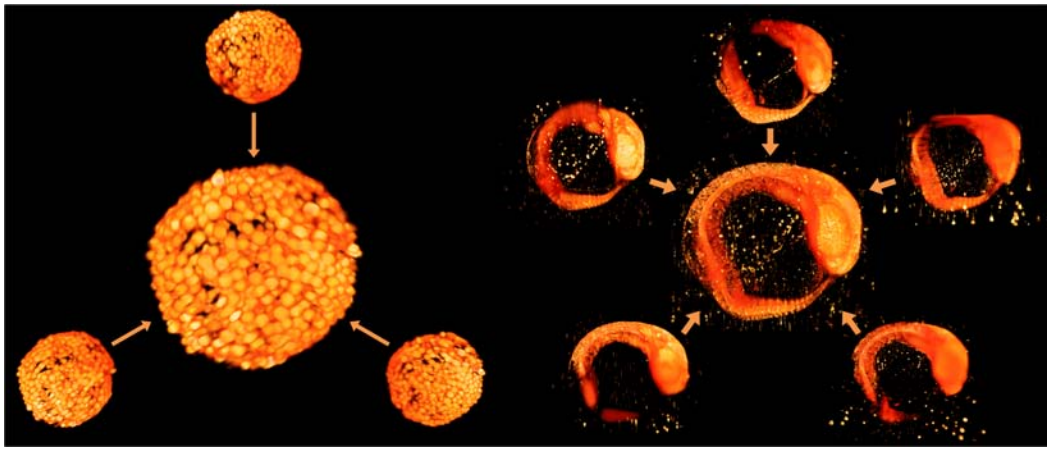

Figure 5.3: Fusion renderings for embryos from two different species. Each one is surrounded by smaller scale renderings of the original views.

A different approach called content-based fusion [Preibisch 2008], which has no need for information about the system's PSF, was recently proposed in which the fused volume is the weighted average of all the available views, using the local entropy of each view as weights. This method is very fast and the results provided show clear improvement with respect to each of the acquired views.

Here we propose an alternative method for multi-view fusion on the Discrete Wavelet Transform (DWT) [Mallat 2008] space, as described in chapter 4, which is also independent of the system's PSF. This represents an advantage when compared to multi-view deconvolution methods, which do require a PSF estimate, in situations where good knowledge of the system's PSF is often difficult to obtain, and in particular for imaging systems like LSFM microscopes which are characterized by a spatially-variant PSF [Temerinac-Ott 2011].

With respect to the existing PSF-free method, the so called content-based fusion [Preibisch 2008], the proposed method has the advantage of working in the DWT space instead of in the image intensity space. The use of such multi-band transforms, which decompose each acquired view onto several scales and orientations, means that our method is able to select the best view independently along each band, and can thus be considered as a multi-scale and multi-orientation content-based fusion. This enhanced orientation discrimination feature is very valuable for LSFM data, as in a given region of the sample two or more views often provide useful information, each one along a different orientation. In such situations, the image-space weighted average performed by the content-based fusion method is unable to independently discriminate directional

information components on each view, while the proposed method can do so by working in the wavelet-space.

5.2 Methods

5.2.1 Pre-processing

Several pre-processing steps might be performed before applying the fusion method to the data, e.g. cropping the volumes and, for very large datasets, down-sampling the in-plane resolution to lower the computation time. However the only necessary step that must be performed before the fusion is the alignment or registration of all the acquired views, because any miss-registration will lead to artifacts in the final fused volume.

The rotation angle of the sample mount is known from the acquisition setup. But it only provides a coarse estimation of the actual transformation between views. Several artifacts including the lack of precision and accuracy in the rotation calibration and operation [Krzic 2009] make a fine registration necessary.

Fine registration of LSFM images is quite challenging, as the depth dependent blurring means that the registration procedure is complicated by information differences between the views. Furthermore, for large specimens, some regions can be totally obscured because of light scattering as described above. To face these problems, more views might be required to ensure sufficient overlap of usable information and external cues such as fluorescent beads in the mounting medium might be required [Preibisch 2010]. Two different approaches have been used during this work, depending on the characteristics of the dataset:

- A bead-based affine registration algorithm [Preibisch 2010] was used for datasets with fluorescent beads embedded in the mounting medium around the specimen.
- A gray-level affine registration algorithm [Thevenaz 1998] for datasets with no fluorescent beads, which uses just the information from the image intensity values.

For small embryos such as sea urchin, bead-less sample preparation and thus gray-level registration is preferred. However, 3D volumes acquired from large embryos like zebrafish show extensive blurring and eventually little overlap between views, and we

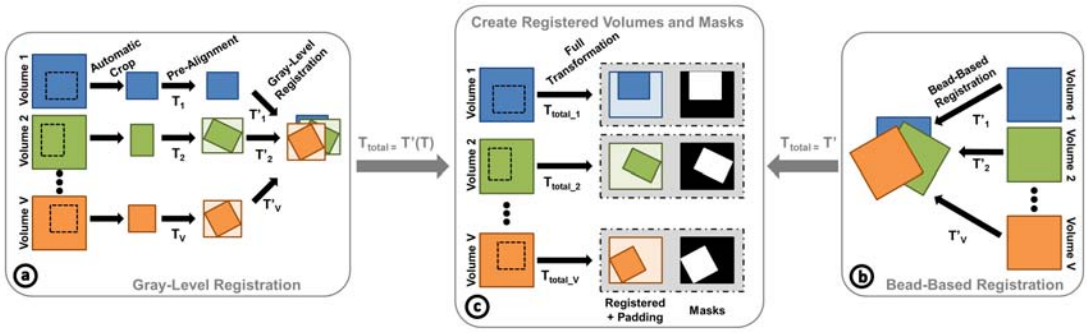

Figure 5.4: Overview of the pre-processing workflow. On the left (a) and on the right (b) we find the two possible registration methods, while the final step is depicted in the middle (c).

confirm that the bead-based registration proposed previously [Preibisch 2010] is a useful strategy. The detailed pre-processing scheme used for each of the registration approaches is depicted in figure 5.4.

For the gray-level registration, each of the acquired volumes is first automatically cropped in order to reduce as much as possible the data size and thus the time and memory consumption of the registration. Due to the low noise level of LSFM images, their back-ground is quite uniform and such automatic cropping works well. Each cropped volume is then pre-aligned according to the rotation step configured during the acquisition, and the pre-aligned volumes are then passed to the gray-level affine registration algorithm. The transformation matrices corresponding to pre-alignment and fine registration are then composed to produce the total transformation matrix of each input volume.

On the other hand, when bead-based registration is selected no initial cropping is used, mainly because it would not make sense to remove the beads needed for the registration algorithm to perform its task. No pre-alignment step is needed either for this kind of registration, as it performs well directly on the acquired volumes.

Once the total transformation matrix has been computed for each of the views (by composing pre-alignment and fine registration matrices in the gray-level case, or directly obtained from the bead-based registration otherwise) they are applied to the original views in order to obtain the registered volumes, which are padded as needed in order to make all of them have the same dimensions. In parallel to this transformation, binary masks are also computed which indicate which voxels of each registered volume contain actual acquired data. These masks will be later used to prevent border artifacts, as it will be further explained later.

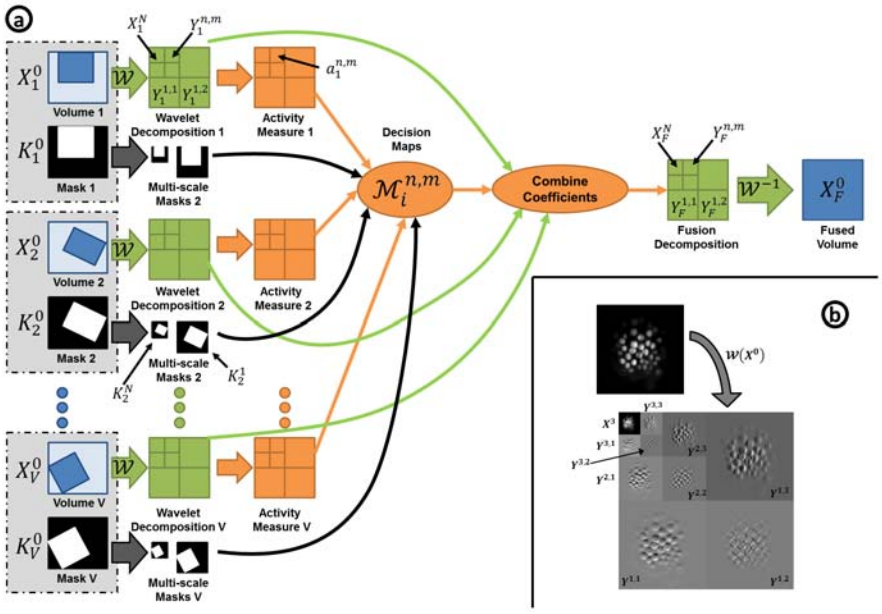

Figure 5.5: (a) Overview of the wavelet-based fusion method. (b) Wavelet decomposition example of a two-dimensional slice, using three scales and three orientations, resulting in nine bands ($N = 3$ and $K = 9$). Although in this work we use undecimated wavelet transform, for illustration purposes this figure represents the classical decimated wavelet transform.

In both approaches, the first view is arbitrarily chosen as a reference, and either registration method is used to compute the affine transform matching it with the other views. When all the views have been brought to the reference frame coordinate system, the fusion process can start.

5.2.2 Wavelet-Based Multi-View Fusion

Once all the available views have been registered and the masks have been generated, the wavelet-based fusion process can start. The details of the fusion method were explained in the previous chapter (section 4.2.2), but for convenience we reproduce in figure 5.5 the fusion workflow and we recall its major steps:

- The inputs are a set of volumes containing the registered views X_v^0 $(v = 1 \dots V)$ and a corresponding set of masks K_v^0 $(v = 1 \dots V)$.
- Each of the registered views is then passed to the wavelet transform of choice, generating the wavelet decomposition bands $Y_v^{n,m}$ (where $n = 1 \dots N$ is the scale

index and $m = 1 \dots 7$ the orientation index) and the low-pass approximation residuals X_v^N.

- In parallel the masks are also processed to get the multi-level masks K_v^n ($n = 1 \dots N$) needed later.
- A saliency measure $a_v^{n,m}$ is computed for each of the wavelet coefficients.
- By looking at those saliency measures together with the multi-level masks, the decision maps $M_v^{n,m}$ are computed.
- The wavelet coefficients are combined according to the decision maps, giving place to the fusion decomposition $\{Y_F^{1,1}, \dots, Y_F^{N,7}, X_F^N\}$.
- Finally we apply the inverse transform to get the fused volume X_F^0 as output.

5.3 Experiments and Results

Datasets used in this work were acquired using a Digital Scanned Lightsheet (DSLM) version of the LSFM microscope [Keller 2008a] where the light sheet is generated by scanning a light beam instead of using a cylindrical lens as in the Selective Plane Illumi-nation Microscope (SPIM) [Huisken 2004]. Our method was applied to various kinds of datasets. The results presented here correspond to two datasets obtained from developing sea urchin (Paracentrotus lividus) and zebrafish (Danio rerio) embryos. Acquisition parameters are summarized in Table 5.1.

Animal [Dataset ID]	Objective	Volume Size (voxels)	Voxel Size (μm)	#Views (angle)	#Time Steps
Sea Urchin [090916eS]	40X/1.0NA W	972x972x82	0.185x0.185x1.85	3 (120deg)	406
Zebrafish [100928aS]	10X/0.3NA W	600x600x111	1.48x1.48x7.4	5 (72deg)	281

Table 5.1: Acquisition parameters for the two datasets used for this work

5.3.1 Live Sea Urchin embryo

A live sea urchin embryo was imaged for 20 hours, from 5 hours post fertilization (hpf) until 25hpf. The whole acquisition time was divided in 180s intervals, leading to over 400 time steps. At each time step 3 volumes (i.e. views) were acquired, using a 120°

rotation of the sample between consecutive views. On all the acquired volumes, signal intensity relates to the local concentration of fluorescent protein accumulated in the cell nucleus.

The sample did not contain fluorescent beads in this study, so the image registration method used before the fusion is the aforementioned gray-level affine registration [Thevenaz 1998].

After views registration, we applied the wavelet fusion method as described in section 4.2.2. A three-dimensional undecimated DWT was used as $\mathcal{W}(X_V^0)$, with Haar wavelet filters and three scales (N=3). Regarding the different possibilities for computation of decision maps (as discussed in section 4.2.2.4), the following configuration parameters were used for this dataset:

- Minimum saliency selection for bands in scale 1 (the finest) which is dominated by noise
- Maximum saliency selection for bands in scales 2 and 3, which contain the most relevant detail information
- Averaging of the low-pass approximations

Figure 5.6 shows an example of the fusion performance, by comparing slices and profiles of the fused volume and the acquired views corresponding to a time step around 20hpf. Each column shows three orthogonal slices from a single volume. Columns 1-3 contain slices from each of the three acquired views, while column 4 contains slices from the volume obtained by our fusion method. Rows 1-3 show slices along XY, XZ and YZ planes, respectively.

Several things are apparent on figure 5.6. First of all each of the three views provides an incomplete picture of the embryo, with some regions showing sharp and bright nuclei, while others exhibit heavy blurring and even some missing nucleus. Nevertheless, slices from the fused volume present overall sharp nuclei with good contrast.

At the bottom of the figure, three line profiles compare each of the individual views with the fusion. On each profile, the blue plot represents intensity values on the corresponding view, while the yellow line follows intensity on the fused volume. Line segments have been overlaid to the XZ slices (middle row), in parallel to the group of five cells represented by the profiles. These profiles show that the fusion provides the best representation of cell populations, because even though each view might show

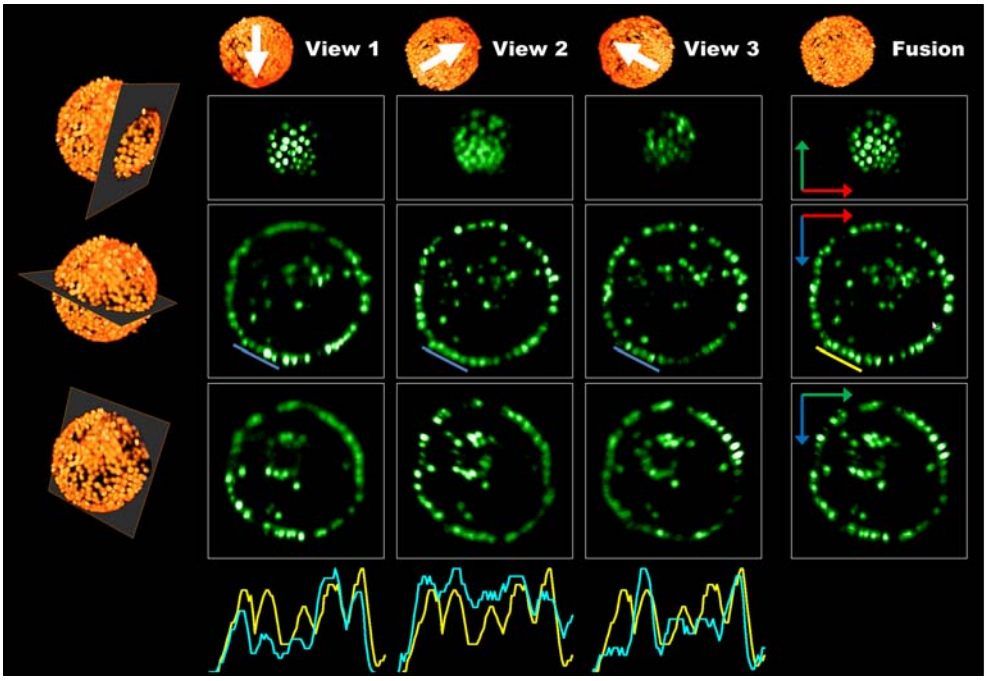

Figure 5.6: Results of the proposed method on a sea urchin embryo, imaged at 20 hours-post-fertilization (hpf).

increased brightness and contrast for some cells, they might completely miss other ones that are indeed recovered in the fused volume.

The whole dataset (406 time steps) was processed similarly. A complete view of the results for this dataset is showed in a video displaying the evolution of the embryo during the 20 hours of the imaging procedure [1]. The movie shows the improvement of individual cell identification in the fused volume that becomes even more obvious as the embryo grows and the individual views are unable to properly capture the details of the whole embryo.

5.3.2 Live Zebrafish embryo

The second example is a live zebrafish embryo imaged during 12 hours, starting at 5hpf. The whole acquisition time was divided in 155s intervals, leading to over 280 time steps. At each time step 5 volumes (i.e. views) were acquired, using a 72° rotation of the sample between consecutive views. In all the acquired volumes, signal intensity correlates with the fluorescent protein concentration in the cell nucleus.

[1] `http://www.die.upm.es/im/videos/SPIM/Movie_S11.avi`

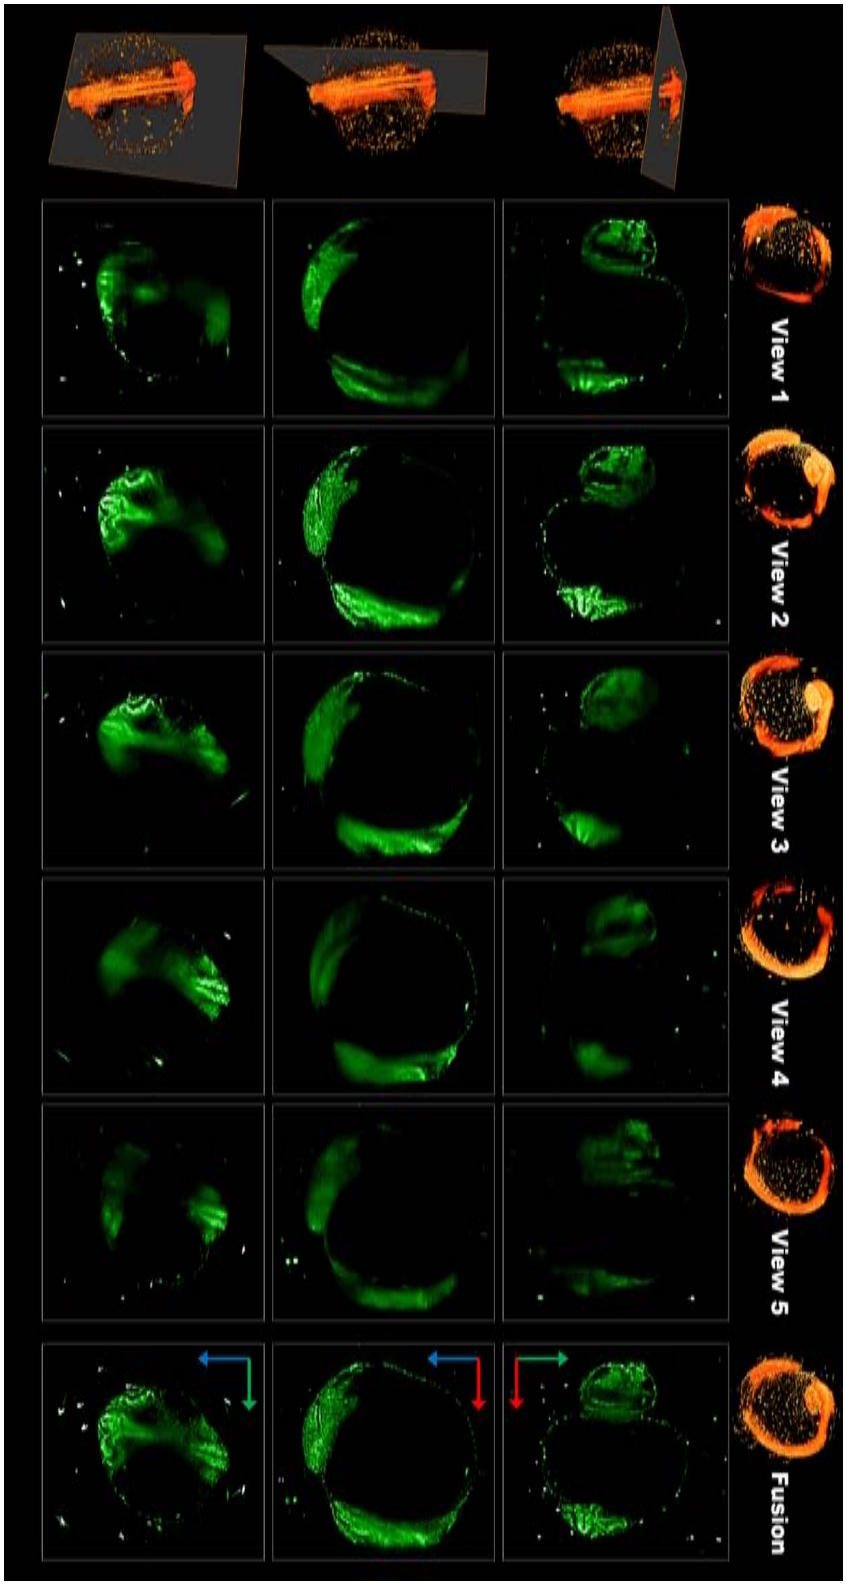

Figure 5.7: Results of the proposed fusion method on a zebrafish embryo, imaged at 17 hpf.

In this study, fluorescent beads were mixed in the sample preparation, so that the bead-based registration method could be used. Due to the relatively big size of the embryo, only a small region of it is imaged with reasonable quality in each view. This means that the five acquired views show little overlapping information among them, as it can be appreciated in figure 5.3. For this reason, standard gray-based registration would be very challenging for this dataset, and bead-based registration is preferred.

In figure 5.7 each column shows three orthogonal slices from a single volume. Columns 1-5 show slices from each of the five acquired views, and column 6 contains slices from the volume obtained by our fusion method. Rows 1-3 show slices along XY, XZ and YZ planes, respectively.

The configuration parameters used for this dataset, including wavelet family, number of scales and decision maps computation scheme, were identical to the ones described in section 5.3.1 for the sea urchin dataset.

The fusion represents the full embryo better than any of the individual volumes, and, while some regions of the embryo are brighter in one specific view, the fusion captures all the relevant information, keeping good contrast whenever such contrast exists in any of the views. A video showing the embryo 360deg around and comparing all the acquired views and the fusion simultaneously is also available for download from [1].

It is worth mentioning that, due to the size of this embryo, each view appears highly blurred in some regions, but the fusion method is able to deal with this situation and gives preference to non-blurred information from other views. For visualization purposes a hard threshold was applied to the original and fused volumes in order to reduce the number of visible beads in the video.

5.3.3 Comparison with existing methods

We have compared our results with those of the previously mentioned content-based weighted averaging [Preibisch 2008], which is also PSF independent and freely available as part of Fiji [2] open source image processing package.

We mentioned in section 5.1.2 that the use of wavelet decomposition provides multi-orientation discrimination capabilities on the information content of each view. This feature is very valuable for LSFM data, as in a given region of the sample two or

[1] `http://www.die.upm.es/im/videos/SPIM/Movie_S12.avi`
[2] `http://fiji.sc/wiki/index.php/Fiji`

more views often provide useful information, each one along a different orientation. In this kind of situation, the image-space weighted average performed by the content-based fusion method is unable to independently discriminate directional information components on each view, while the proposed method can do so by working in the wavelet-space.

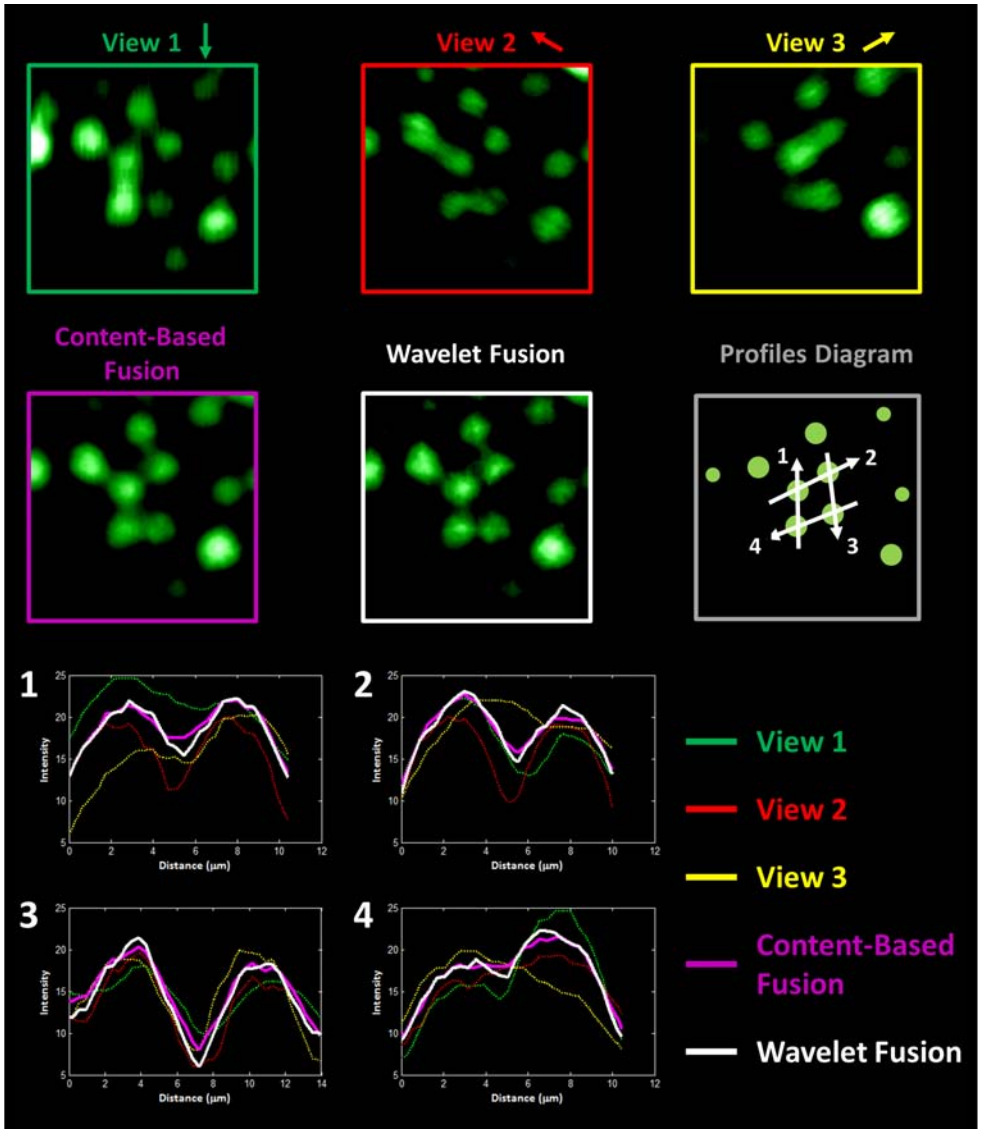

Figure 5.8: Comparison of enhanced orientation discrimination of the proposed Wavelet Fusion, compared to Content-Based Fusion.

An example of such situation is provided in figure 5.8, in which we show slices and line profiles extracted from the dataset described in section 5.3.1 (more precisely, from

the same timestep shown in figure 5.6).

The upper half of figure 5.8 shows 2D slices obtained orthogonally to the samples rotation axis. In the first row, slices from each of the three acquired views are shown, with the arrow beside each views title indicating the direction of the detection axis for that view. On the second row, the same slice from the content-based fusion and the wavelet fusion volumes are shown, together with a diagram depicting where each of the line profiles has been extracted. The line profiles themselves are shown on the bottom half of the figure. Each profile passes through a pair of adjacent cells and compares both fusions methods and the three original views.

We also provide a qualitative comparison, based on 2D slices of the full embryos, between our method and Preibisch et al.'s method. In figure 5.9a we show the comparison for the sea urchin embryo, and in figure 5.9b for the zebrafish. In the latter, the results of the wavelet-based fusion method look less blurred, and more individual nuclei can be distinguished, as it can be appreciated at the bottom of the figure thanks to a line profile comparison through a group of three nuclei located inside a blurred region (figure 5.9b). For the sea urchin, the difference is less striking, but the line profiles in figure 5.9a reveal slightly better peak contrast with our wavelet fusion method (in blue) compared to the content-based fusion (in yellow).

Additionally, a quantitative comparison of the proposed method and content-based fusion was performed for both datasets, by measuring the contrast of several line profiles, each one passing through two adjacent nuclei. Two timesteps, separated several hours, were selected for each dataset (sea urchin and zebrafish). The complete comparison process is depicted in figure 5.10, and consists on the following steps:

(a) A set of cell nuclei (between 12 and 16) was manually selected for both fused volumes.

(b) 12 line profiles were automatically traced, each one passing through two adjacent nuclei.

(c) A mixture of two gaussians was then fitted to each line profile.

(d) A normalized contrast measure was obtained from each fitting, ranging from 0 (no contrast) to 1 (maximum contrast).

Table 5.2 summarizes the contrast measures for each case. The contrast is higher in the volumes fused using the proposed wavelet method in 47 out of 48 (98%) of

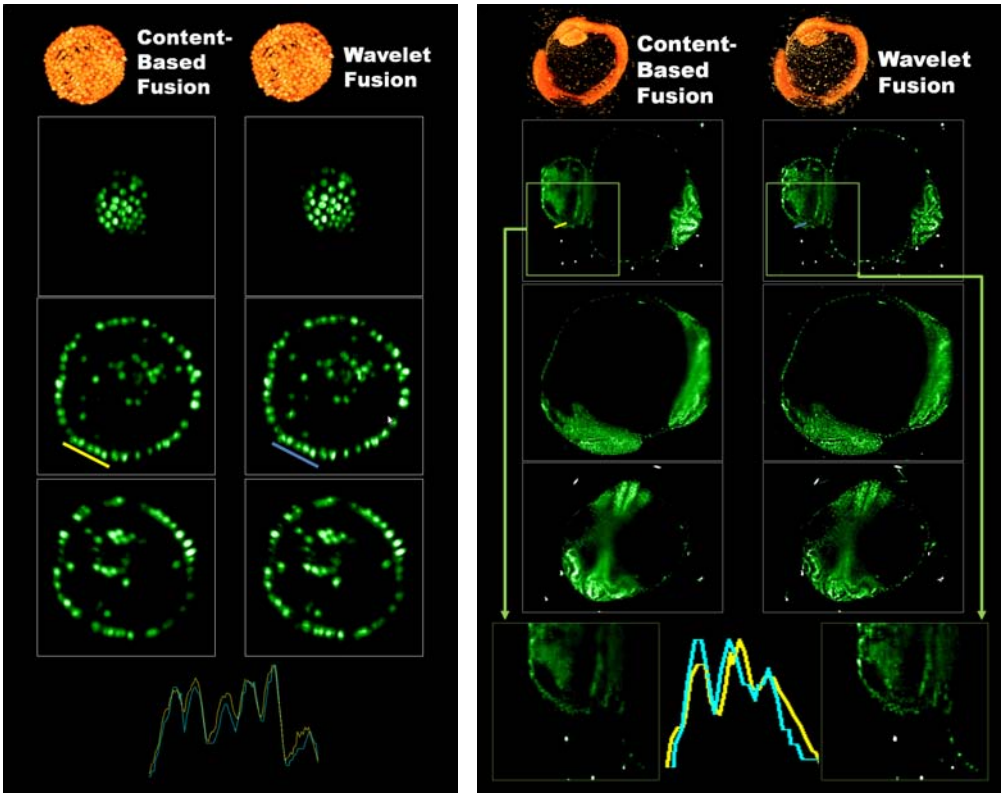

(a) Sea Urchin **(b)** Zebrafish

Figure 5.9: Qualitative comparison of fusion methods.

the line profiles. The overall average improvement is 0.08, and it is higher in the zebrafish volumes (0.11 average improvement) than in the sea urchin volumes (0.05 average improvement). In percentage, this means an average relative improvement of 31% ($\sim$ 20% for sea urchin volumes and $\sim$ 42% for zebrafish volumes). The statistical significance of the contrast difference between both fusion methods was checked using a paired sample two-tailed t-test, obtaining p-values well below 0.01 for all the datasets and timesteps.

Finally, the full set of profiles used for the comparison are shown in figures 5.11, and 5.12.

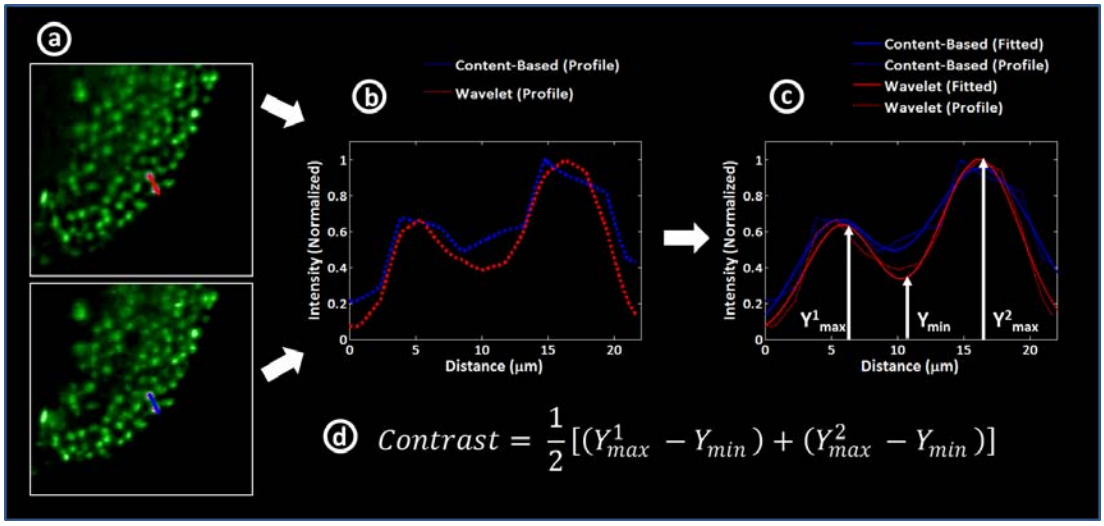

Figure 5.10: Quantitative comparison process for the two datasets described in section 5.3

	Timestep	Fusion Method	Profile #											
			1	2	3	4	5	6	7	8	9	10	11	12
Sea Urchin	150 [12hpf]	Content-Based	0.19	0.25	0.18	**0.20**	0.21	0.36	0.32	0.50	0.36	0.16	0.21	0.44
		Wavelet	**0.22**	**0.30**	**0.21**	0.16	**0.29**	**0.41**	**0.48**	**0.54**	**0.42**	**0.19**	**0.24**	**0.51**
	300 [20hpf]	Content-Based	0.15	0.13	0.06	0.25	0.24	0.18	0.22	0.21	0.12	0.39	0.44	0.29
		Wavelet	**0.20**	**0.18**	**0.07**	**0.27**	**0.27**	**0.18**	**0.27**	**0.24**	**0.15**	**0.47**	**0.55**	**0.36**
Zebrafish	000 [5hpf]	Content-Based	0.29	0.36	0.30	0.34	0.22	0.16	0.53	0.44	0.39	0.44	0.35	0.50
		Wavelet	**0.33**	**0.52**	**0.44**	**0.45**	**0.30**	**0.24**	**0.66**	**0.53**	**0.45**	**0.59**	**0.44**	**0.59**
	280 [17hpf]	Content-Based	0.08	0.34	0.32	0.32	0.38	0.09	0.31	0.19	0.17	0.23	0.36	0.31
		Wavelet	**0.23**	**0.37**	**0.46**	**0.47**	**0.53**	**0.09**	**0.47**	**0.35**	**0.25**	**0.31**	**0.58**	**0.48**

Table 5.2: Comparison of contrast measures between the proposed method and content-based fusion

5.4 Practical Issues

5.4.1 Incomplete Views

When multi-view datasets are acquired with LSFM, it is not unusual to get some incomplete views. This can be due to non-perfect rotation calibration or even due to the embryo's movement during the acquisition of the time-lapse.

As we mentioned in section 4.2.2.2 in order to avoid such incomplete views causing artifacts in the fused volume, a set of masks is generated when the views are registered. These masks help keeping track of which voxels contain actual information from the original view and which ones come from the necessary padding that we introduce during the registration. The masks are taken into account during the fusion, in the way

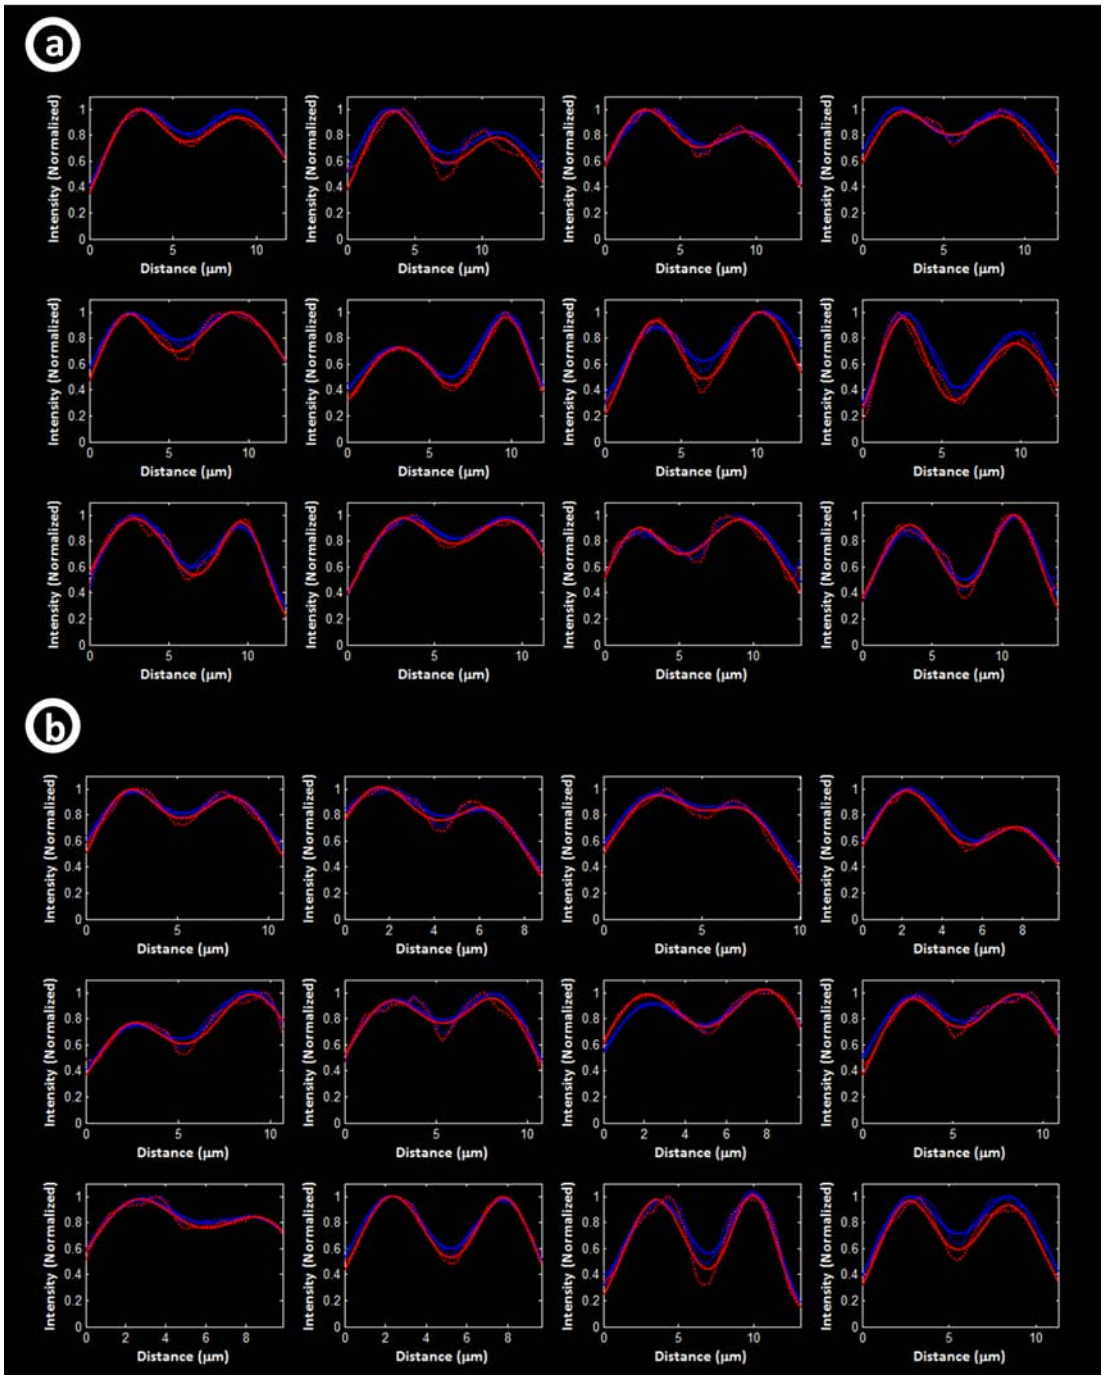

Figure 5.11: Line profiles and gaussian fittings for timesteps 150 [12hpf] (a) and 300 [20hpf] (b) of the sea urchin dataset. The line profiles are plotted with dotted lines and the gaussians in solid lines. The profiles and fittings corresponding to content-based fusion appear in blue, while those corresponding to wavelet fusion appear in red.

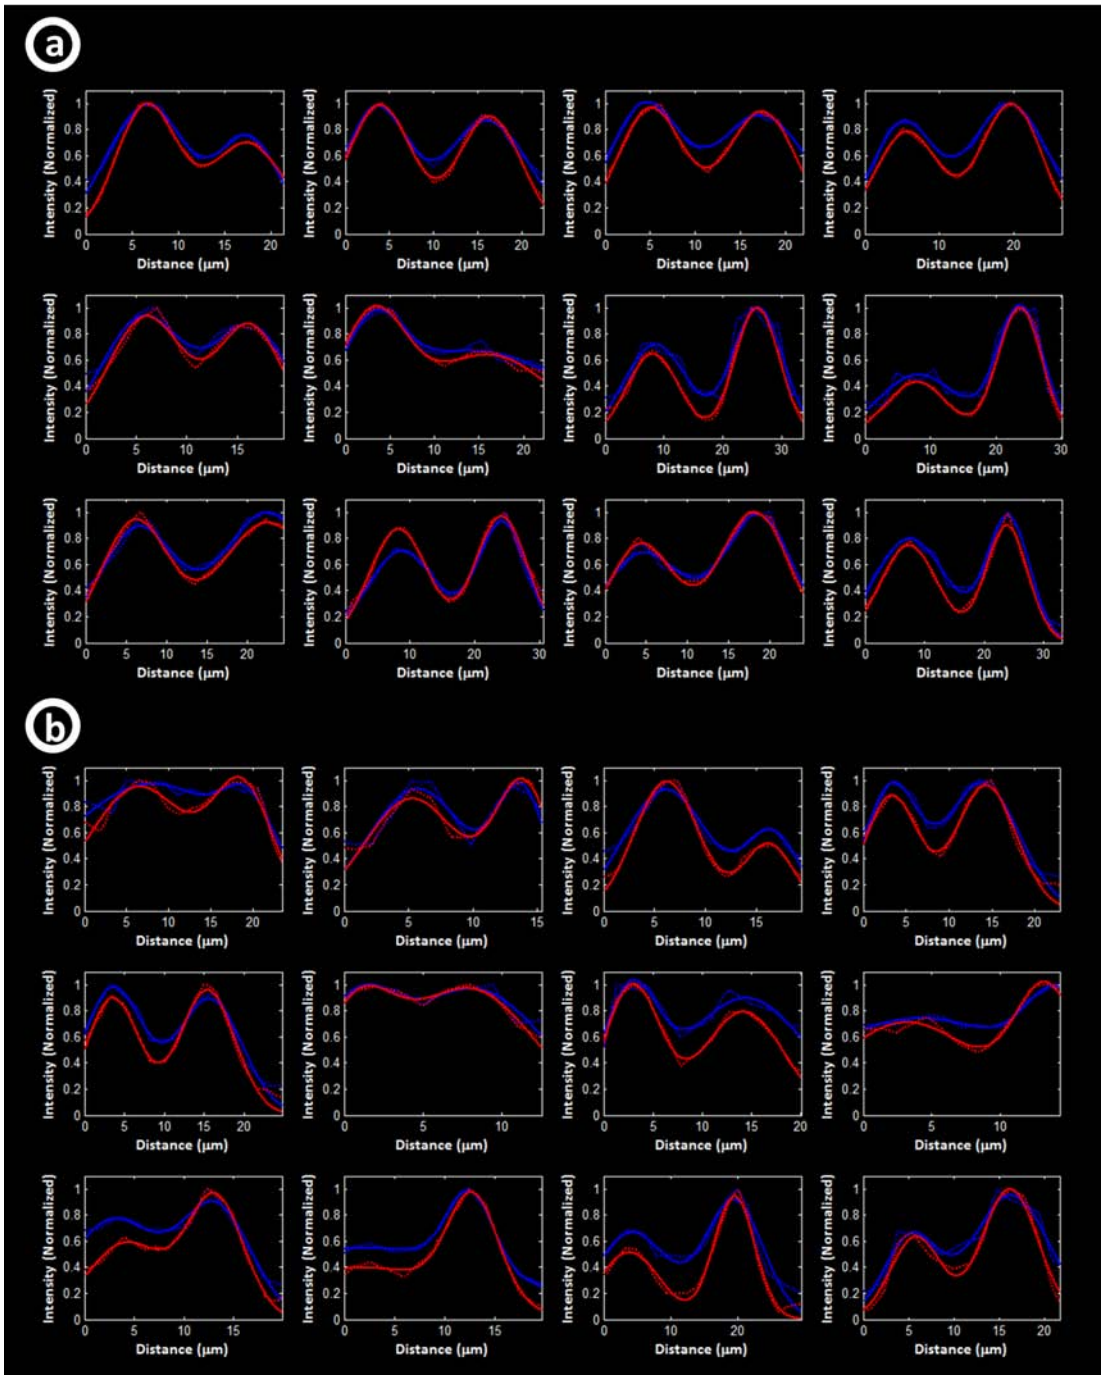

Figure 5.12: Line profiles and gaussian fittings for timesteps 000 [5hpf] (a) and 280 [17hpf] (b) of the zebrafish dataset. The line profiles are plotted with dotted lines and the gaussians in solid lines. The profiles and fittings corresponding to content-based fusion appear in blue, while those corresponding to wavelet fusion appear in red.

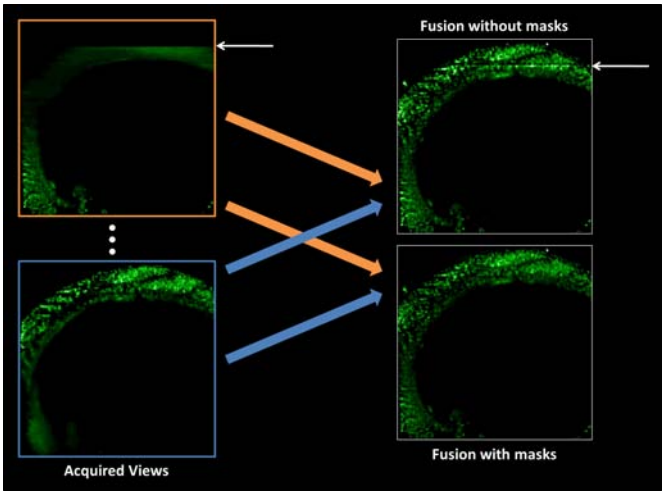

Figure 5.13: Example of how the use of masks can help us avoid fusion artifacts when some of the views are incomplete

explained in section 4.2.2.2, so that the final fusion is free of artifacts.

In figure 5.13 we reproduce a practical example of this kind of situation extracted from the zebrafish dataset (in which one of the five views was incomplete). To the left we show a blow-up of 2D slices of two different LSFM views of a live zebrafish embryo, where it can be appreciated how one of the views does not cover the full embryo and presents a sharp border where the acquired volume ends (marked with a thin white arrow). To the right we show the result of applying the fusion method without (top-right) and with (bottom-right) the use of masks to indicate the acquired data limits. The fusion without masks shows a clear artifact (marked with a thin white arrow) while the fusion with masks is free of such border artifacts.

5.4.2 Fusion decision rules

We mentioned in section 4.2.2.4 the fact that several decision rules can be used when computing the decision maps that decide how wavelet coefficients from each input view are combined into the fused volume. In particular we described how it might be useful to use *minimum* saliency instead of *maximum* saliency selection for some bands because depending on the size of the objects present in the input volumes, some scales of their wavelet decomposition might be dominated by noise.

That is indeed the situation for both datasets described in sections 5.3.1 and 5.3.2, and in figure 5.14 we support the choice of decision rules used for their fusions. The

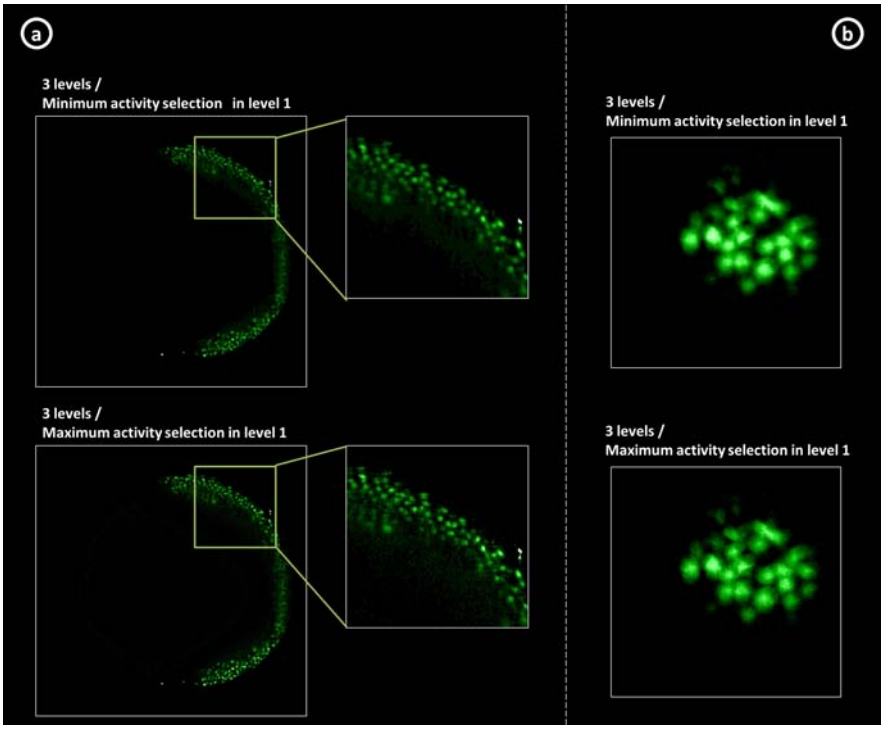

Figure 5.14: Demonstration of the effect of using the "minimum saliency coefficients" selection for the finest scale of the wavelet decomposition.

figure compares slices from fused volumes corresponding to both embryos: (a) The zebrafish described on the left and (b) the sea urchin on the right. The slices shown in the top half of the figures correspond to fused volumes obtained using the configuration mentioned in sections 5.3.1 and 5.3.2, that is:

- Three-dimensional undecimated DWT as $\mathcal{W}(X_v^0)$,
- `haar` family wavelet filters,
- three scales (N=3)
- minimum saliency selection for bands in scale 1 (the finest) which is dominated by noise,
- maximum saliency selection for bands in scales 2 and 3, which contain the most relevant detail information,
- Averaging of the low-pass approximations.

The slices shown in the bottom half of the figures were extracted from fused volumes obtained using mostly the same parameters, except the fact that "maximum saliency" selection was used for bands in scales 1, 2 and 3. It can be appreciated that the slices on the top half have a lower noise level than those on the bottom half, which as mentioned in section 5.3.1 can be explained by the fact that scale 1 (the finest) is dominated by noise

5.5 Discussion

We proposed a novel methodology to fuse multi-view images, such as those obtained from light-sheet based fluorescence microscopy, but not limited to that particular technique. Almost any imaging technique in which several volumes of the same object are acquired, each of them providing incomplete and complementary information, could benefit from the proposed fusion scheme, including recent light-sheet techniques like DSLM using structured illumination [Keller 2010].

The proposed method is based on wavelet multi-scale and multi-orientation decomposition of the input volumes, and does not rely on a priori knowledge of the system PSF. While PSF-aware methods have more potential for reducing PSF related artifacts, the proposed method can be used on those situations in which such knowledge is not available as is the case with LSFM acquisition where the PSF varies over the volume due to light sheet shape and to distortion caused by the specimen itself.

We demonstrated the method on LSFM acquisitions from live sea urchin and zebrafish embryos. The fusion results show improved overall contrast and details when compared to any of the acquired volumes alone. With respect to the other PSF independent method mentioned in this work, content-based weighted averaging, the new method has improved orientation and scale discrimination as a direct consequence of performing the fusion in the wavelet domain instead of the intensity space. This improved discrimination is useful for instance in regions of the sample where two views contribute sharp details in different orientations.

This work focuses on the fusion method itself, and does not consider the effect of other complementary methods like denoising techniques which could be used in a preprocessing step. For instance wavelet denoising [Donoho 1994, Chang 2000] and bilateral filtering [Paris 2009, Yang 2009] seem good choices for future tests. Some variations (cross/joint and dual bilateral filtering) have been used for image fusion in certain situations [Bennett 2007], but they are not adequate in situations where

different views do not cover the full specimen, as uniform regions in the image being filtered will stay uniform regardless of the information provided by the other views.

A possible way to further improve the usefulness of the proposed method would be to automatically adapt the wavelet fusion parameters, e.g. number of scales and fusion scheme, to the content of the input volumes. Moreover, the general scheme can be accommodated to any kind of wavelet transform (different base wavelet basis, non-separable decompositions, etc), making the proposed method highly flexible and thus adaptable to any kind of input volumes.

I have not yet lost a feeling of wonder, and of delight, that this delicate motion should reside in all the things around us, revealing itself only to him who looks for it. I remember, in the winter of our first experiments, just seven years ago, looking on snow with new eyes. There the snow lay around my doorstep - great heaps of protons quietly precessing in the earth's magnetic field. To see the world for a moment as something rich and strange is the private reward of many a discovery.

Nobel Lecture, 11th December 1952
Edward Mills Purcell

Chapter 6

Cardiac DE-MRI Multi-View Imaging

6.1 Introduction

Delayed-Enhancement Magnetic Resonance Imaging (DE-MRI) is a Contrast-Enhanced MRI (ceMRI) modality using gadolinium based contrast agents, and has proven to be a major tool for the assessment of myocardial tissue viability. Delayed enhancement MRI images are acquired 10-15 minutes after the injection of the contrast agent, and several studies prove that it is a precise tool for the identification of injured myocardium tissue after myocardial infarction [Kim 1999, Fieno 2000, Ricciardi 2001, Simonetti 2001, Wu 2001, Amado 2004, Bello 2005, Kim 2008].

By classifying myocardial tissue depending on its signal intensity in gadolinium-based DE-MRI images, three categories can be defined: healthy, heterogeneous and scar tissue. The extent and morphology of the heterogeneous tissue (areas of necrosis interspersed with bundles of viable myocytes) is receiving increasing interest due to its correlation with inducible ventricular tachycardia (VT) [Bello 2005, Schmidt 2007]. In this line of work several studies, some of the most relevant ones related to this PhD Thesis [Perez-David 2011], have tracked down such correlation to the correspondence of heterogeneous tissue in DE-MRI data with slow conducting zones detected during electrophysiology studies [Yan 2006, Roes 2009].

Given such correlation between the information contained in DE-MRI images and the voltage maps obtained through electro-anatomic mapping (EAM), there is great potential in using the information acquired in a non-invasive way through DE-MRI for

pre-planning ablation procedures, and even for evaluating or predicting the patient's susceptibility of suffering VT and therefore deciding the most adequate treatment.

However, the interpretation of the information contained in DE-MRI images in a slice-by-slice basis is non-trivial, as the topology of the scar and the heterogeneous tissue is three-dimensional and it can be quite complex to analyze it by means of two-dimensional slices. For this reason, it seems highly useful to obtain an alternative representation of the information contained in the DE-MRI images which allows an easier and better understanding of the 3D structure in the underlying tissue.

But there is a limitation in typical DE-MRI images, which usually provide good in-plane resolution ($\sim 1.5mm$) but poor axial resolution ($5mm$ to $10mm$ slice thickness) [Kim 1999, Kim 2003, Simonetti 2001, Reddy 2004, Kim 2008, Bogun 2009, Peters 2009]. While higher resolution DE-MRI images has been demonstrated on ex-vivo swine models [Ashikaga 2007], mechanically ventilated pigs and respiratory navigator-gated MRI sequences [McGann 2008, Kino 2009, Viallon 2011, Moghari 2012], they are still not broadly available.

To overcome this limitation, DE-MRI studies typically include at least three volumes acquired on three different anatomically-related views:

Short Axis Image stack where the slices are orthogonal to left ventricle's axis

Four Chambers Image stack where slices contain the left ventricle's axis, and pass through the heart's 4 chambers. It is also known as Horizontal Long Axis view.

Two Chambers Image stack where slices contain the left ventricle's axis, and passes only through the left ventricle and atrium. It is also known as Vertical Long Axis view.

Figure 6.1 shows a diagram of the three different views, together with examples of slices acquired for each of them.

Therefore the problem we approach in this chapter is to adapt and apply the proposed fusion method (chapter 4) to multi-view non-uniformly sampled 3D Delayed Enhancement Magnetic Resonance Imaging (DE-MRI) datasets of the heart. The fused volumes will then be used to generate a 3D visualization of myocardial tissue viability for aiding the diagnostic and therapeutic process of Ventricular Tachycardia.

Interestingly a wavelet-based approach has been recently proposed [Aganj 2012] for the fusion of orthogonal views on T2-weighted MRI brain studies, confirming the interest and suitability of this kind of approach for combining multi-view MRI datasets.

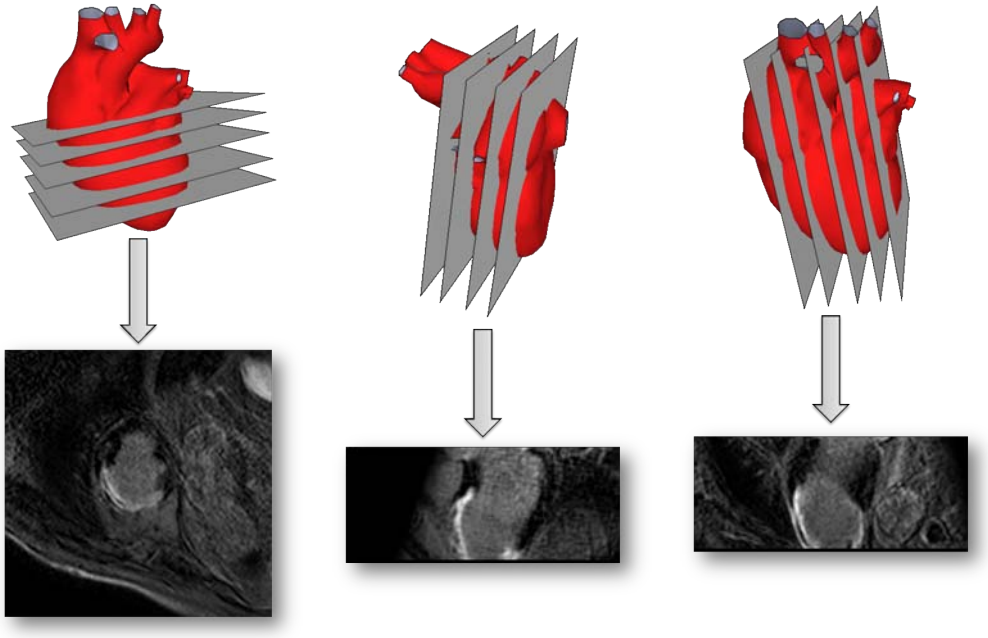

(a) Short Axis Slice Stack

(b) Four Chambers Slice Stack

(c) Two Chambers Slice Stack

Figure 6.1: Three anatomical views of the heart

In section 6.2 we explain how the fusion method was adapted to this particular problem and later in section 6.3 we describe the datasets on which the method was applied and we show some of the results obtained. In section 6.4 we explain the clinical application of the resulting volumes for the diagnostic and therapeutic process of Ventricular Tachycardia (VT) and finally we include a brief discussion in section 6.5.

6.2 Methods

6.2.1 Pre-Processing

There are some pre-processing steps which need to be followed before the fusion method might be applied to the DE-MRI datasets. Their aim is to compensate for some issues present in the original acquired datasets, which generally speaking fall into one of the two following categories:

Registration It was already mentioned in section 4.2.1 how important a correct alignment of the input volumes is for the fusion process. In this application the mis-

alignment of the volumes comes from the way in which they are acquired. In the first place, the planes in each of the three volumes are acquired along different orientations, as shown in figure 6.1, so in order to align all the volumes to a common frame of reference we need to use the information contained in the DICOM headers of the DE-MRI volumes, which specify the position of each volume's voxels with respect to the MRI scanner's *machine coordinates*. However there is an additional source of misalignment, coming from the fact that the volumes are acquired during breath-holds, with three or four breath-holds normally used for a full dataset, which means that there will be some variation in the patient's position between volumes.

Normalization As it was explained in chapter 4 the wavelet-based fusion works by selecting the most salient features among all input volumes. This features are in fact coefficients from different scales and orientations of a discrete wavelet transform, and therefore are just the result of applying some linear filter to the volumes.

This means that, in order to compare coefficients or "features" from different input volumes, they must be normalized, i.e., if any relevant feature in the object under study appears in more than one volume, it should be represented with the same magnitude in all of them.

Unlike other medical imaging modalities, like CT which is calibrated so that the pixel values they provide have a physical meaning (Hounsfield Units), in MRI the pixel values have arbitrary units and can actually vary with many of the configurable sequence parameters. Therefore, it is to be expected that pixel values in the three anatomical views of the heart described in section 6.1 are not directly comparable.

In figure 6.2 we show the whole pre-processing process, with the three volumes and the manual segmentation as inputs, and the normalized and registered volumes as outputs. As it can be appreciated in the figure, the inputs for the pre-processing are the set of DE-MRI views and segmentation contours of the left-ventricle's endocardium and epicardium borders. For this work, the segmentations were performed manually using QMass MR [1] and were defined on the Short-Axis view. However, the workflow is equally valid for any alternative segmentation method (manual, semi-automatic or fully-automatic).

[1] `http://www.medis.nl/products/qmass-mr.html`

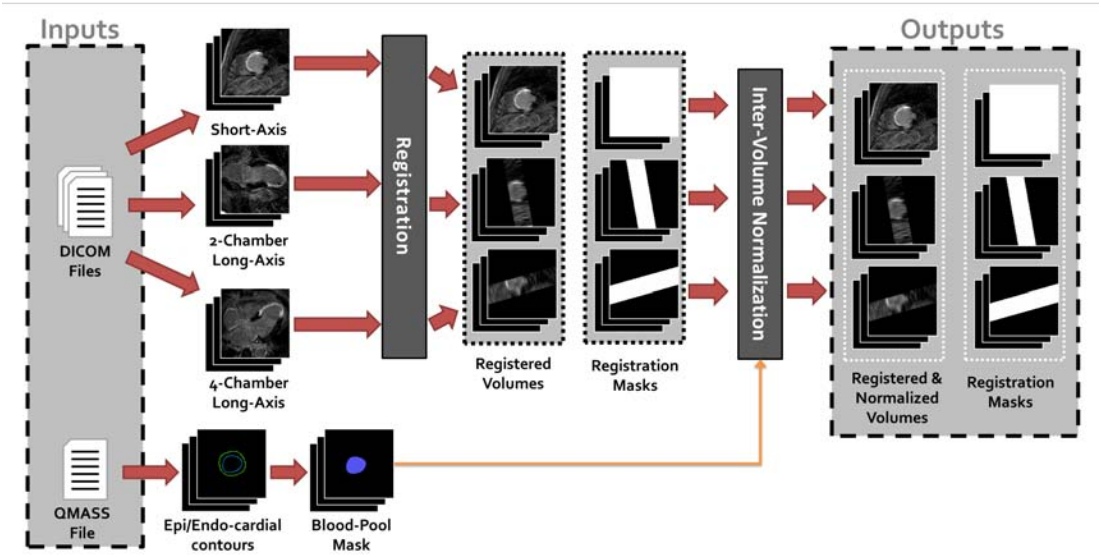

Figure 6.2: Detailed workflow for pre-processing the DE-MRI datasets before they different views can be fused together.

6.2.1.1 Registration

The MR scanner stores the volumes in DICOM files, usually one file for each acquired slice. If we just read the slices from each file and stack then we get the full acquired volume in its own intrinsic xyz coordinate system, as shown in figure 6.3.

Among the DICOM file headers, however, we can find two tags ("ImagePositionPatient" and "ImageOrientationPatient") which can be used to recover the transformation relating this intrinsic coordinate system with the fixed coordinate system associated to

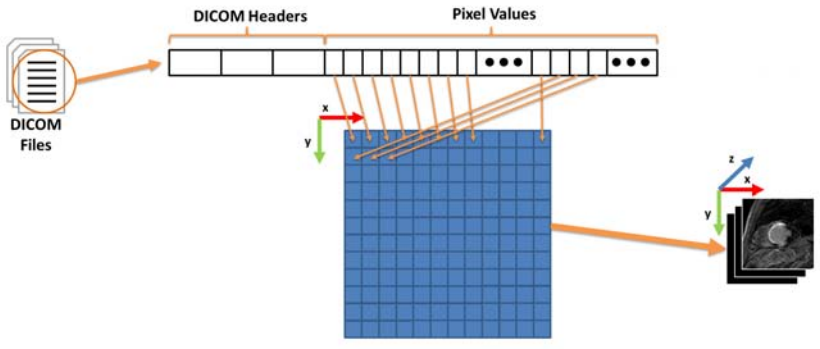

Figure 6.3: Reading DE-MRI volumes one slice at a time.

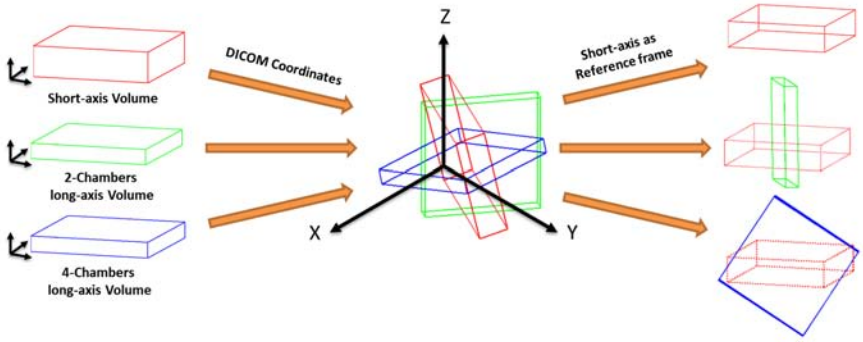

Figure 6.4: Pre-alignment of the views using DICOM headers. (left) Each view is read from the DICOM files into its own "intrinsic coordinates". (center) Using DICOM's *ImagePositionPatient* and *ImageOrientationPatient* headers we can find the location and orientation of each view in the *Machine Coordinates* reference frame. (right) We arbitrarily choose the Short-Axis volume as reference for aligning the remaining views.

the MR scanner (*machine coordinates*). By computing this transformation for each of the available volumes (short-axis, 2-chamber and 4-chamber views) we can find the relative positions and orientations of each of them, as shown in figure 6.4.

In this way, we have all the volumes aligned in the *machine coordinates* reference frame (center part of figure 6.4). However, in order to avoid the interpolation step for at least one of the volumes, we arbitrarily choose the short-axis volume as reference and then, using combinations of those transformations, we transform the 2-chamber and 4-chamber volumes to the short-axis' intrinsic reference frame, as depicted in the rightmost part of figure 6.4.

However, as each of the DE-MRI volumes is acquired during a different patient breath-hold, the patient's heart position, with respect to the scanner's *machine coordinates*, will vary slightly on each volume. This means that, even after using the information available in the DICOM headers from each volume in order to align all the acquired volumes in the *machine coordinates* reference frame, there will still be a small (typically $<< 1cm$) misregistration among them.

For all these reasons, we use a fine gray-level registration step after the DICOM-based alignment, by which we compensate the patient's heart misalignment among different views. A rigid transformation model is used for this purpose, and cross-correlation is used as metric for the registration [Zitová 2003, Gonzalez 2007], given its quite robust behavior when faced with intensity differences among the volumes under registration.

6.2.1.2 Normalization

A convenient way of normalizing the three different views available for each patient is to use a common structure present in all of them. The biggest and most uniform structure available in this case is the left ventricle's blood-pool.

So in order to correct intensity level mismatches among views, we compute the mean intensity level of the blood-pool on each of them and then we arbitrarily choose the short-axis view as a reference and scale the two remaining views by the necessary scale factor, so that after the scaling the mean intensity level of the blood-pool is the same on all the volumes.

In order to select the voxels from the blood-pool on each volume and compute the average intensity for normalization purposes, we use the manual segmentations mentioned earlier. By filling the endocardial contour we get a mask of the interior of the left ventricle, which contains the blood-pool voxels but also some undesired structures like the papillary muscles. For this reason, an erosion is applied to the mask so that only the central part of the mask, which contains blood-pool voxels only, remains.

As we only have the endocardial segmentation for the short-axis volume, for the normalization of the remaining volumes (the 2-chambers and 4-chambers long-axis) we need to apply the necessary transformations to them, including the transformation obtained from the DICOM headers and the one obtained from the fine registration described above, before we can use the blood-pool masks on those volumes and get the normalization factors, as shown in figure 6.2.

6.2.2 Wavelet-Based Multi-View Fusion

Once all the available views have been registered and the masks have been generated, the wavelet-based fusion process can start. The details of the fusion method were explained in the previous chapter (section 4.2.2), but for convenience we reproduce in figure 6.5 the fusion workflow and we recall its major steps:

- The inputs are a set of volumes containing the registered views X_v^0 $(v = 1 \dots V)$ and a corresponding set of masks K_v^0 $(v = 1 \dots V)$.
- The wavelet transform of choice is then applied to each of the registered views, generating the wavelet decomposition bands $Y_v^{n,m}$ (where $n = 1 \dots N$ is the scale

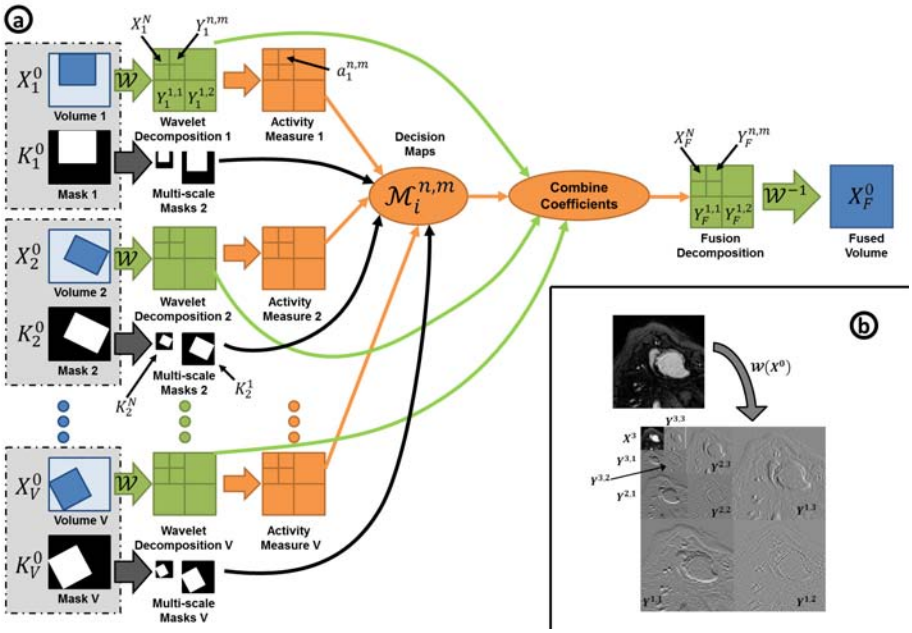

Figure 6.5: (a) Overview of the wavelet-based fusion method. (b) Wavelet decomposition example of a two-dimensional slice, using three scales and three orientations, resulting in nine bands ($N = 3$ and $K = 9$). Although in this work we use undecimated wavelet transform, for illustration purposes this figure represents the classical decimated wavelet transform.

index and $m = 1 \dots 7$ the orientation index) and the low-pass approximation residuals X_v^N.

- In parallel the masks are also processed to get the multi-level masks K_v^n ($n = 1 \dots N$) needed later.
- An saliency measure $a_v^{n,m}$ is computed for each of the wavelet coefficients.
- By looking at those saliency measures together with the multi-level masks, the decision maps $M_v^{n,m}$ are computed.
- The wavelet coefficients are combined according to the decision maps, giving place to the fusion decomposition $\{Y_F^{1,1}, \dots, Y_F^{N,7}, X_F^N\}$.
- Finally we apply the inverse transform to get the fused volume X_F^0 as output.

6.3 Experiments and Results

In this section we present the datasets used for testing and validating the proposed method, which include DE-MRI acquisitions from both swine and human studies.

On the human studies the fused volumes were compared directly to their respective set of low axial resolution views in order to asses the improvement on the fused volume. However, no further validation was possible for these datasets as DE-MRI volumes with higher axial resolution were not available due to acquisition timing constraints, as we mentioned earlier in this chapter (which is actually one of the reasons behind the choice of this particular application as a target for our fusion method).

Therefore, the method validation has been carried out on a swine model of myocardial infarction, on which thanks to the use of mechanical ventilation we were able to obtain datasets with pixel size $1.29x1.29x5.0mm^3$, similar to what is achievable in human studies, but also with higher axial resolution (pixel size $1.29x1.29x2.5mm^3$). In this way we can apply the fusion method to the lower axial resolution data and compare the results against the higher axial resolution images.

6.3.1 Dataset Description

For the swine model of myocardial infarction, four pigs underwent DE-MRI with a 1.5-T scanner five weeks after the creation of the myocardial infarction (Intera, Philips, Best, the Netherlands). A 5-element dedicated cardiac coil was used. All images were obtained with electrocardiogram gating and breath-holding and the pigs were maintained under general anesthesia and mechanical ventilation while performing the study.

The MRI study consisted of delayed enhancement imaging of myocardial scar tissue (3D inversion-recovery turbo gradient echo sequence, pre-pulse delay optimized for maximal myocardial signal suppression; SENSE $\times$ 2, flip angle 15°, repetition time: 4.7 ms, echo time: 1.6 ms, actual pixel size: $1.55 \times 1.74mm^2$, interpolated pixel size $1.29 \times 1.29mm^2$, inversion time: 200 to 300 ms; acquisition window was set to 150 to 170 ms and breath-holding length was 50 to 60s depending on heart rate). The images were obtained in short-axis, 4- and 2- chamber views. Two different sets of views were acquired; one with 5.0 mm slice thickness (corresponding to approximately 20 contiguous slices for each view) in order to have the same acquisition as in the human studies, and another set with 2.5 mm slice thickness (corresponding to approximately 40 contiguous slices for each view) to be used as a reference for comparison. Delayed enhancement imaging was performed 10 to 15 min after a total injection of 0.2 mmol/kg gadodiamide (Omniscan, GE Healthcare, Oslo, Norway).

As for the human DE-MRI studies, we processed datasets of 16 patients from the SMVT (Sustained Monomorphic Ventricular Tachycardia) group in the study presented

in [Perez-David 2011], all of them suffering from chronic myocardial infarction, with clinically documented SMVT and no contraindication for MRI.

All the patients underwent DE-MRI in a 1.5-T scanner (Intera, Philips, Best, The Netherlands). A 5-element dedicated cardiac coil was used. All images were obtained with electrocardiogram gating and breath-holding. The DE-MRI imaging of myocardial scar tissue was performed with the following parameters: 3D inversion-recovery turbo gradient echo sequence, pre-pulse delay optimized for maximal myocardial signal suppression; SENSE × 2, flip angle 15°, repetition time 3.4 ms, echo time 1.3 ms, actual pixel size $1.48 \times 1.66 mm^2$, interpolated pixel size $1.29 \times 1.29 mm^2$, 5 mm actual slice thickness, inversion time 200 to 300 ms, acquisition window set to 150 to 170 ms and breath-holding length of 13 to 14 s depending on heart rate. The images were obtained in the short-axis, 4- and 2-chamber views. Delayed enhancement was performed 10 to 15 min after a total injection of 0.2 mmol/kg gadodiamide (Omniscan, GE Healthcare, Oslo, Norway).

6.3.2 Validation Methodology

For each dataset, the three available views (short-axis, 2- and 4- chamber long-axis) were registered and normalized as described in section 6.2.1 and figure 6.2. Bi-cubic interpolation was used while transforming the views in order to obtain volumes with isotropic voxels ($1.29 \times 1.29 \times 1.29 mm^3$).

After views registration, we applied the wavelet fusion method as described in section 6.2.2. A three-dimensional undecimated DWT was used as $\mathcal{W}(X_V^0)$, with Haar wavelet filters and three scales (N=3). Regarding the different possibilities for computation of decision maps (as discussed in section 4.2.2.4), the following configuration parameters were used for this dataset:

- Maximum saliency selection for decomposition bands in the three scales
- Averaging of the low-pass approximations

In order to asses the improvement in the resulting volumes, the fusion obtained from the $1.29 \times 1.29 \times 5.00 mm^3$ views is then compared to the three original views. This comparison is carried out by means of representative two-dimensional slices and line profiles passing through relevant features.

Apart from the visual comparison, it is also interesting to compute some quantitative estimation of the improvement achieved by the fusion method. Estimating the quality

and information content of an image or volume is an open problem for which several solutions have been proposed in the literature. Given that what we claim to achieve is an increase in the information on relevant or salient features in the fused volumes, we have selected the following measure which combines Michelson visibility definition and second-derivative estimation [DelMarco 2009, Strong 1959]:

$$SDVO\{V\} = \frac{1}{N}\sum_{i,j,k}\left|\frac{V_{MAX} - 2\cdot V(i,j,k) + V_{MIN}}{V_{MAX} + 2\cdot V(i,j,k) + V_{MIN}}\right|$$

where V is the volume for which the Second Derivative Visibility Operator (SDMO) is being computed, N is the number of pixels in the volume, $V(i,j,k)$ is an individual pixel and V_{MAX} and V_{MIN} are the maximum and minimum, respectively, values over a window around the central pixel $V(i,j,k)$. A $3 \times 3 \times 3$ window was used for results presented in the following sections.

Additionally, for the swine infarction model datasets the fusion result is also compared to the short-axis view from the $1.29 \times 1.29 \times 2.50mm^3$ study, which is used as *gold standard*. The goal is therefore to check how the fusion result is closer to such *gold standard* than any of the original ($5.00mm$ thickness) views alone.

6.3.3 Results

6.3.3.1 Swine infarction model datasets

In figure 6.6 we show, for one of the pig datasets, a qualitative comparison of the three original $1.29 \times 1.29 \times 5.00mm^3$ views, the fused volume and the $2.50mm$ short-axis view. The comparison is done by selecting representative two-dimensional slices, and it is apparent how the fused volume captures the most salient information from each of the input volumes and is very close to the *gold standard* (the $2.50mm$ short-axis). It is also noteworthy how the sharp border edges in the registered input volumes, which arise from the incomplete overlap among them, do not translate in artifacts in the fused volume, thanks to the use of masks (see section 4.2.2.2).

In all cases, the short-axis volume was selected as a reference for the views alignment, but the fused volume was defined to be slightly larger than the short-axis along its axial direction. The reason behind this decision is the fact that short-axis often show limited coverage of the apical region of the heart, while long-axis do provide useful information in that same region. By defining a target fusion region larger than the reference short-axis, we can include this additional information in the final fused volume. This should

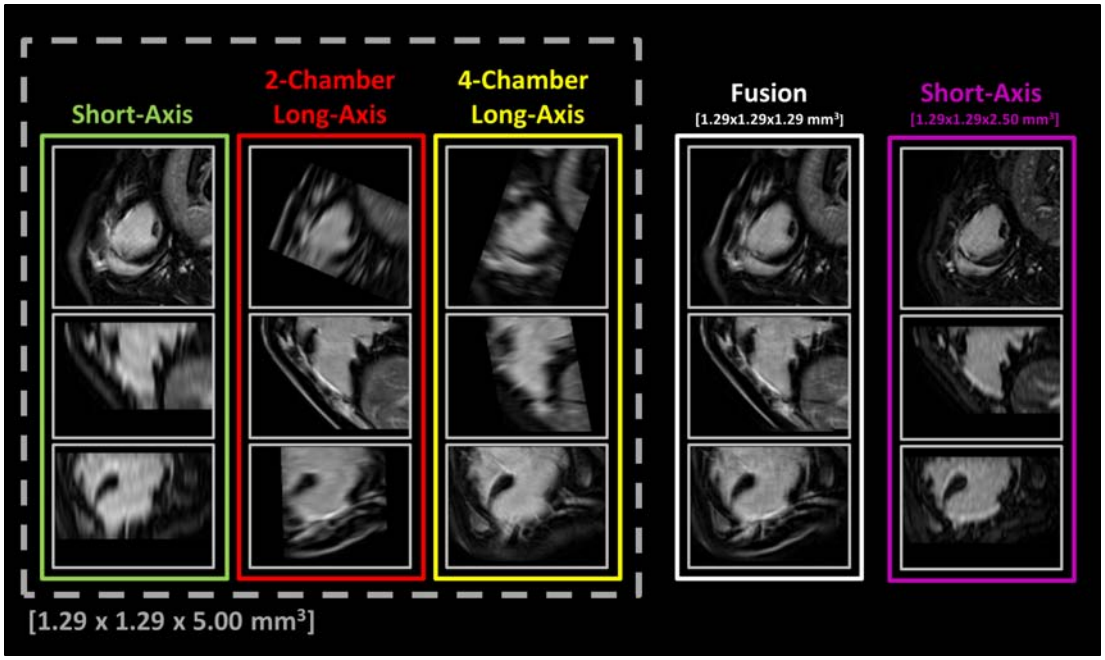

Figure 6.6: Qualitative comparison of the results for one of the datasets from the swine model of myocardial infarction. The three $1.29 \times 1.29 \times 5.00mm^3$ acquired DE-MRI volumes (short-axis, 2-chamber long-axis and 4-chamber long-axis) are compared to the volume resulting of their fusion. A $1.29 \times 1.29 \times 2.50mm^3$ short-axis volume acquired for the same subject is also shown in order to help with the evaluation of the fusion performance.

be taken into consideration when looking at the results in figure 6.6, as the fusion result has a different field of view than the *gold standard* (the fused field of view is larger in the apical and basal regions).

Figure 6.7 shows several line profiles extracted from the same slices and passing through the infarct scar tissue, which allows us to analyze in more detail the behavior of the fusion method in some relevant areas of the myocardial wall. We can appreciate how the profiles extracted from the fused volume are most similar to that view with the highest saliency (Short Axis in profile 1 and 2-Chamber Long Axis in profiles 2 and 3), while taking little influence from those other views with low-frequency information. Apart from the fusion and the three original $5.00mm$ views, the figure also includes the corresponding profiles from the $2.50mm$ short-axis view used as *gold standard*. The comparison of the fused volume with the $2.50mm$ shows that the former is always quite close to the latter.

In figure 6.8 we show the three high-resolution ($1.29 \times 1.29 \times 2.50mm^3$) views from the same dataset together with the corresponding fused volume. The same two-

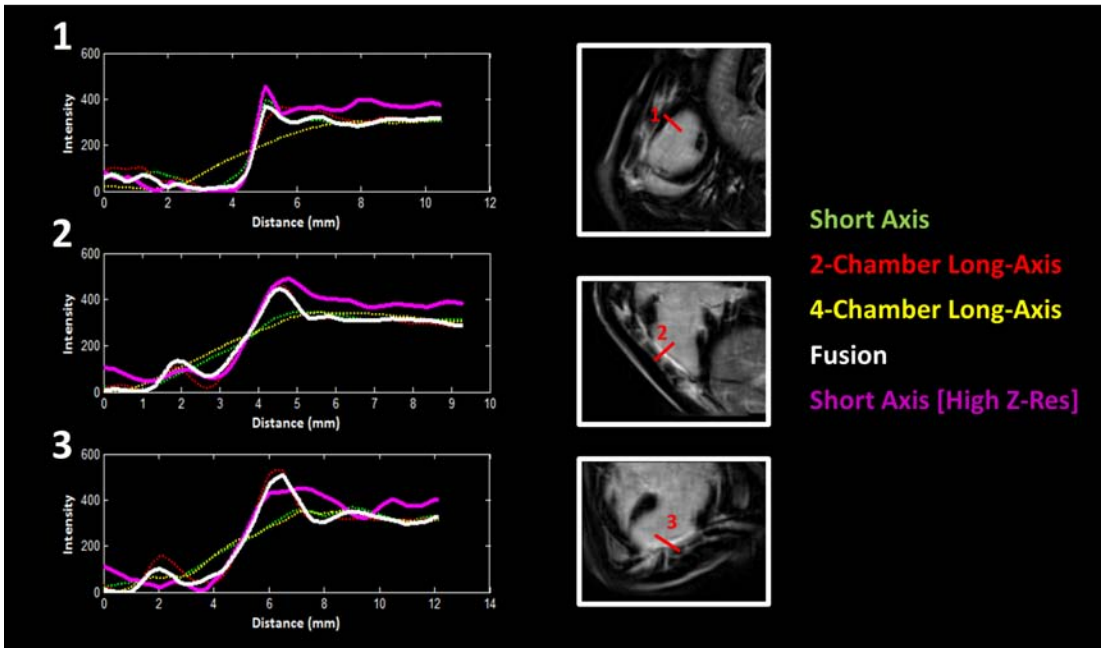

Figure 6.7: Line profile comparison of the results for one of the datasets from the swine model of myocardial infarction. Relevant line profiles were extracted from the original volumes and the fusion.

dimensional slices as in figure 6.6 are shown and again it can be appreciated how the fused volume captures the most relevant features from all the three views, although the difference between the fusion and the original views is less significant than in the previous figure given the thinner slices acquired in this case.

As for the quantitative comparison, in table 6.1 we summarize the results obtained by applying the aforementioned second derivative visibility operator to the four datasets from the swine infarction model. The first thing we notice is that in every case, the high-resolution (2.50mm slices) views obtain higher scores than the corresponding low-resolution (5.00mm slices) ones, as one would expect. More interestingly, we can also appreciate that the fused volumes always get higher scores than all the corresponding views from which they are computed. It is also noteworthy the fact that the score obtained by the 5.0mm fusion is comparable to the scores of individual 2.5mm views.

The table also includes the SDVO relative improvement of the fused volumes with respect to each of their input views ($(SDVO_{Fusion} - SDVO_{View})/SDVO_{View}$). The average relative improvement was 41.2% for the fusions of 2.5mm-slice views and 51.6% for the 5.0mm-slice ones.

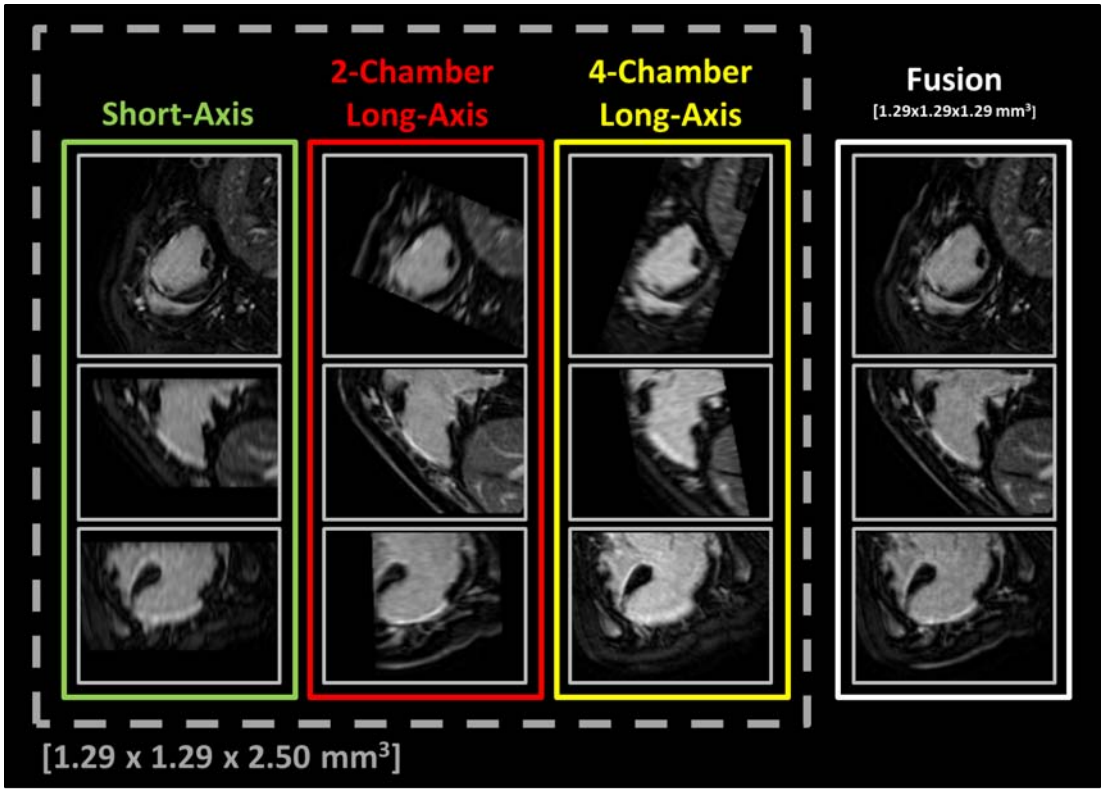

Figure 6.8: Qualitative comparison of the results for one of the datasets from the swine model of myocardial infarction. The three $1.29 \times 1.29 \times 2.50mm^3$ acquired DE-MRI volumes (short-axis, 2-chamber long-axis and 4-chamber long-axis) are compared to the volume resulting of their fusion.

6.3.3.2 Human datasets

The fusion method was applied to the three $1.29 \times 1.29 \times 5.0mm^3$ volumes acquired for each patient. In figure 6.9 we show a qualitative comparison of the original views and the fused volume for one of the patients in the study. The comparison is done by selecting representative two-dimensional slices, and it is apparent how the fused volume captures the most salient information from each of the input volumes. The short-axis volume was selected as a reference for the views alignment, but the fused volume was defined to be slightly larger than the short-axis along its axial direction. The reason behind this decision is the fact that short-axis often show limited coverage of the apical region of the heart, while long-axis do provide useful information in that same region. By defining a target fusion region larger than the reference short-axis, we can include this additional information in the final fused volume. This should be taken into consideration when looking at the results in the following figures.

Patient ID	Measure	Short-Axis	2-Chamber Long-Axis	4-Chamber Long-Axis	Fusion
P12-389	SDVO	0.18	0.19	0.15	0.23
$[1.29 \times 1.29 \times 2.50mm^3]$	Relative fusion improvement	30%	25%	53%	
P12-389	SDVO	0.14	0.16	0.10	0.20
$[1.29 \times 1.29 \times 5.00mm^3]$	Relative fusion improvement	45%	20%	91%	
P12-392	SDVO	0.19	0.18	0.15	0.24
$[1.29 \times 1.29 \times 2.50mm^3]$	Relative fusion improvement	23%	33%	56%	
P12-392	SDVO	0.13	0.15	0.10	0.19
$[1.29 \times 1.29 \times 5.00mm^3]$	Relative fusion improvement	44%	28%	98%	
P12-430	SDVO	0.14	0.14	0.13	0.19
$[1.29 \times 1.29 \times 2.50mm^3]$	Relative fusion improvement	35%	39%	49%	
P12-430	SDVO	0.08	0.07	0.07	0.11
$[1.29 \times 1.29 \times 5.00mm^3]$	Relative fusion improvement	38%	55%	56%	
P12-433	SDVO	0.13	0.10	0.15	0.17
$[1.29 \times 1.29 \times 2.50mm^3]$	Relative fusion improvement	36%	76%	15%	
P12-433	SDVO	0.08	0.11	0.06	0.12
$[1.29 \times 1.29 \times 5.00mm^3]$	Relative fusion improvement	43%	3%	97%	

Table 6.1: Second Derivative based visibility operator results for all the swine datasets and their corresponding fused volumes.

It is also noteworthy how the sharp edges in the registered input volumes, which arise from the incomplete overlap among them due to the fact that individual views do not fully cover the left ventricle, do not translate in artifacts in the fused volume, thanks to the use of masks (see section 4.2.2.2).

Figure 6.10 shows several line profiles extracted from the same slices, which allows us to analyze in more detail the behavior of the fusion method in some relevant areas of the myocardial wall and the infarcted tissue. We can appreciate how the profiles extracted from the fused volume are always very similar to the best of the original views, and it is also apparent how the $5.0mm$ views do not cover all the left ventricle and therefore their corresponding profiles appear incomplete in some of the profiles

As for the quantitative comparison, in table 6.2 we summarize the results obtained by applying the aforementioned second derivative visibility operator to the sixteen human datasets. We can also appreciate that the fused volumes always get higher scores than all the corresponding views from which they are computed. The table also includes the SDVO relative improvement of the fused volumes with respect to each of their input views ($(SDVO_{Fusion} - SDVO_{View})/SDVO_{View}$). The average relative improvement was 52.1%, which is consistent with the 51.6% improvement presented

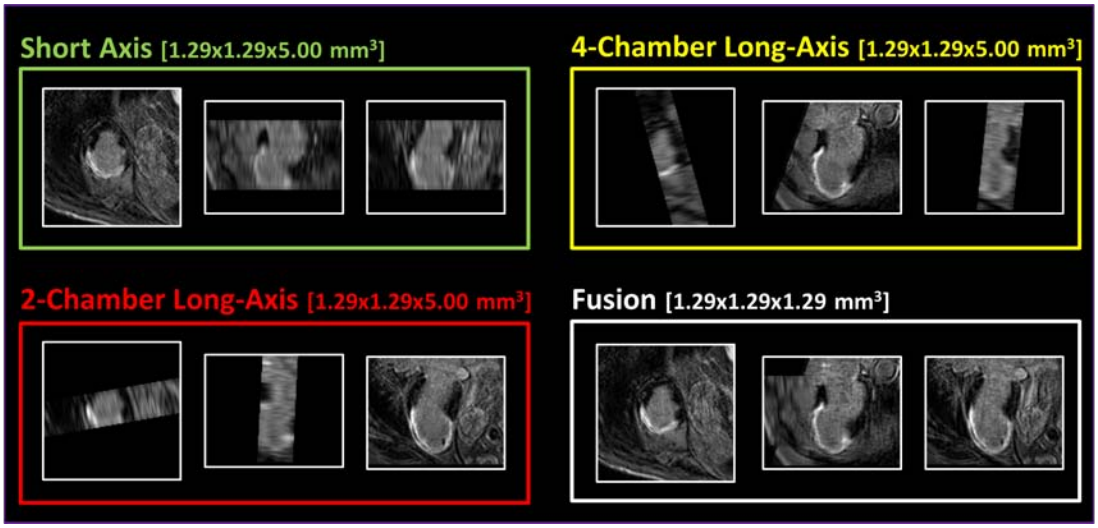

Figure 6.9: Qualitative comparison of the results for one of the datasets from the human study. The three acquired DE-MRI volumes (short-axis, 2-chamber long-axis and 4-chamber long-axis) are compared to the volume resulting of their fusion.

earlier for the 5.0mm-slice swine datasets.

6.4 Application: Non-Invasive Identification of Ventricular Tachycardia Substrate

6.4.1 Introduction

As we mentioned earlier in this chapter (section 6.1) several recent studies have reported a strong relationship between the extent of contrast-enhanced myocardial regions in DE-MRI images and inducibility of Ventricular Tachycardia (VT) [Bello 2005, Schmidt 2007], and it has been hypothesized that heterogeneous tissue in the surroundings of deep scar identified by DE-MRI is related with slow conduction zones found during electro-anatomic mapping procedures in patients with ischemic cardiomyopathy [Yan 2006, Roes 2009].

A case-control study was recently carried out [Perez-David 2011] with the purpose of demonstrating the connection between slow conduction channels [Arenal 2004, Hsia 2006] measured during electrophysiology studies, which are known to be related to re-entry circuits causing Sustained Monomorphic Ventricular Tachycardia (SMVT), and heterogeneous tissue corridors found in DE-MRI.

Patient ID	Measure	Short-Axis	2-Chamber Long-Axis	4-Chamber Long-Axis	Fusion
P1	SDVO	0.053	0.049	0.079	0.085
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	59.3%	73.6%	8.2%	
P2	SDVO	0.096	0.040	0.072	0.103
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	6.7%	157.8%	42.5%	
P3	SDVO	0.037	0.037	0.036	0.052
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	40.8%	38.5%	45.1%	
P4	SDVO	0.079	0.071	0.073	0.114
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	44.4%	62.3%	56.1%	
P5	SDVO	0.066	0.038	0.074	0.088
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	34.0%	130.1%	19.1%	
P6	SDVO	0.060	0.050	0.100	0.102
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	69.2%	103.8%	2.0%	
P7	SDVO	0.091	0.055	0.079	0.112
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	22.7%	104.0%	42.2%	
P8	SDVO	0.062	0.058	0.064	0.087
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	40.9%	49.9%	36.3%	
P9	SDVO	0.051	0.061	0.042	0.072
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	40.7%	18.1%	70.5%	
P10	SDVO	0.096	0.083	0.080	0.130
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	34.9%	55.8%	62.2%	
P11	SDVO	0.056	0.030	0.041	0.063
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	12.5%	113.2%	53.0%	
P12	SDVO	0.149	0.105	0.124	0.188
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	26.4%	78.8%	52.2%	
P13	SDVO	0.036	0.029	0.036	0.046
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	28.1%	59.4%	26.0%	
P14	SDVO	0.076	0.072	0.082	0.113
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	48.4%	55.6%	36.7%	
P15	SDVO	0.065	0.041	0.039	0.076
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	16.8%	84.4%	96.4%	
P16	SDVO	0.060	0.044	0.055	0.077
$[1.29 \times 1.29 \times 5.00 mm^3]$	Relative fusion improvement	28.1%	74.5%	39.8%	

Table 6.2: Second Derivative based visibility operator results for all the human datasets and their corresponding fused volumes.

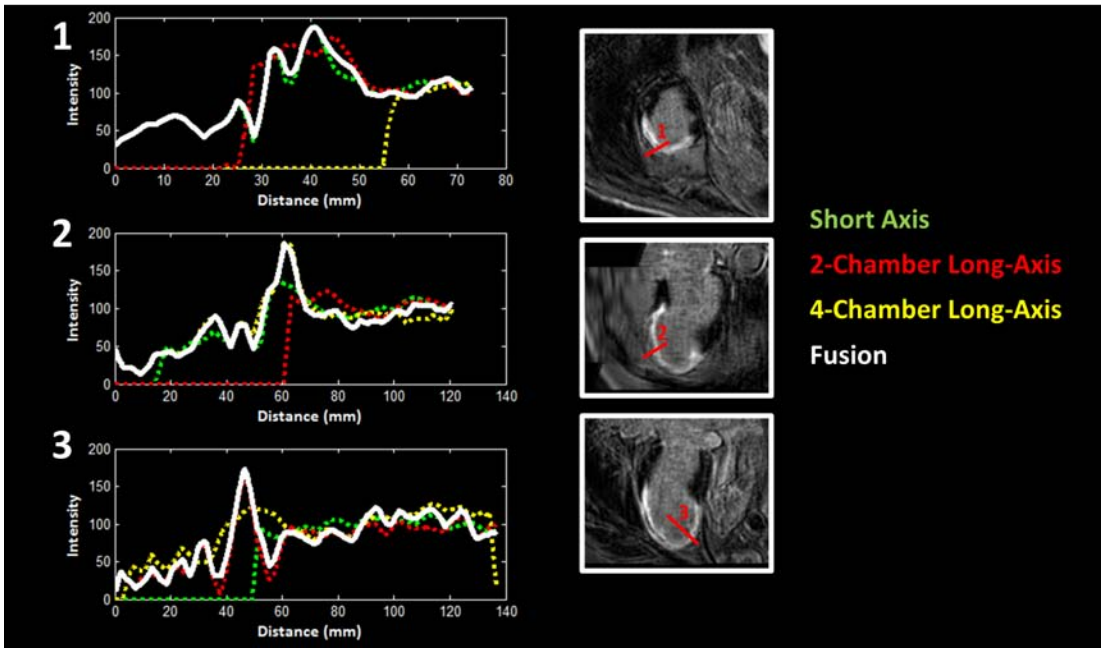

Figure 6.10: Line profile comparison of the results for one of the datasets from the human study. Relevant line profiles were extracted from the original volumes and the fusion.

As part of this PhD Thesis we developed image processing and visualization means for their use in the aforementioned study [Perez-David 2011]. The goal was to project the information from DE-MRI datasets into a three-dimensional "viability map" in which healthy, heterogeneous and scar tissue, as well as their shapes and relative position, could be observed at once. This kind of "viability maps" will facilitate the identification of heterogeneous tissue corridors in the patients' myocardium, in the same way as existing three-dimensional electro-anatomic mapping systems (like CARTO, Biosense-Webster, Diamond Bar, CA, USA and EnSite NavX, St. Jude Medical, St. Paul, MN, USA) help identifying the slow conduction channels.

There are some clear advantages of a MRI-based "viability map" over existing EAM methods, arising from the non-invasive and pre-interventional nature of the former. On the one hand, by having the viability maps *a priori* they can be used for pre-planing and guidance of ablation procedures, therefore shortening the total intervention time. On the other hand, their non-invasive nature enables the possibility to obtain such MRI-based maps for any patient suffering from chronic myocardial infarction, paving the way for further studies aimed to find out whether the information in the viability maps can be used for making decisions about the best therapy for each patient.

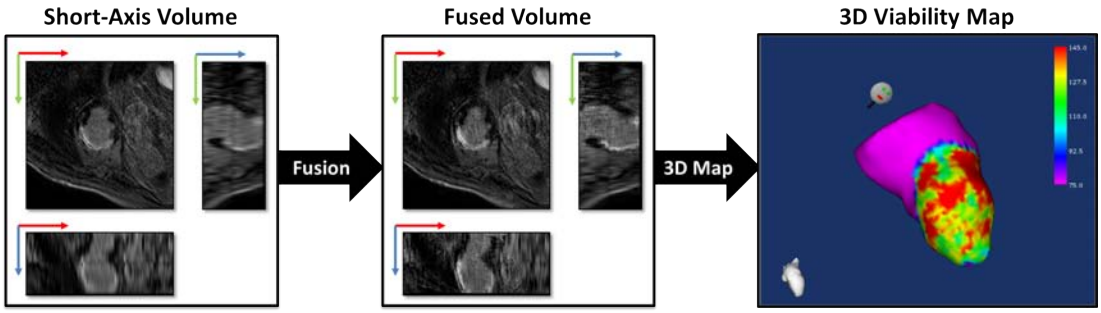

Figure 6.11: DE-MRI Fusion and Viability Map Example. (left) Orthogonal slices extracted from a patient's Short-Axis DE-MRI volume are shown. (middle) The same three orthogonal slices in the fused volume computed from the Short-Axis, 2-Chamber Long-Axis and 4-Chamber Long-Axis DE-MRI volumes of the same patient. (right) 3D Viability Map extracted from the information contained in the fused volume.

It is clear that if DE-MRI images are to be used as an information source for identifying the VT substrate it is desirable to get the best DE-MRI images possible, therefore the fusion method was applied to the DE-MRI volumes of each patient (as described in section 6.2, except for the modifications described in appendix A) in order to get a single fused volume, which was then used to generate the desired 3D viability map as described in section 6.4.2.1. Figure 6.11 shows the overall workflow and an example of the resulting fused volume and viability map.

In order to demonstrate the benefit of using the fused volume when generating the *viability map* (as described in the following sections), in figure 6.12 we show a comparison of the *viability map* generated from the fused DE-MRI volume and the same map obtained from the acquired short-axis volume alone.

Apart from the increased detail in the map when using the fused volume, we see that the short-axis based map does not cover the apical region of the heart. The reason is that in short-axis acquisitions, specially when thick slices are used, the apex is heavily blurred because of partial volume effect, so it is usually not even acquired. In the fused volume, on the other hand, we can see the apex because we get that information from the 2- and 4-chamber long axis.

6.4.2 DE-MRI Based Myocardium Viability Map

In this section we describe the methodology used for creating the 3D triangle mesh representing the patients' myocardium, as well as the way in which the signal intensity information from a DE-MRI volume is then projected into that 3D mesh, the method

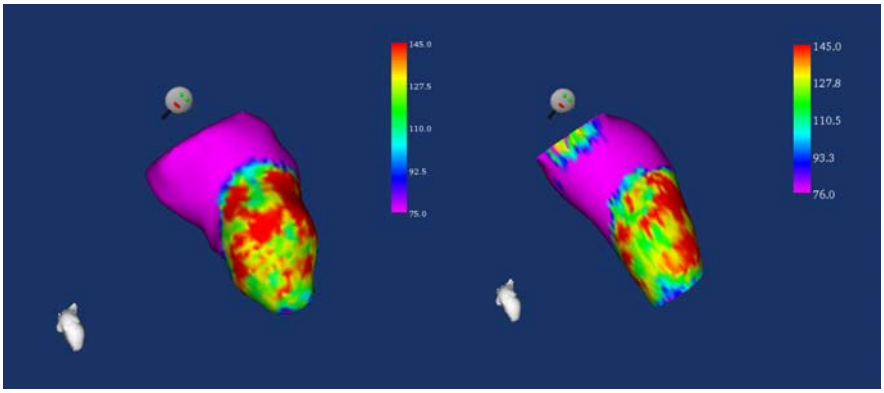

Figure 6.12: Comparison of *viability maps* obtained from the fused DE-MRI volume (left) and from the short-axis volume alone (right).

has been filed as a patent application [Ledesma-Carbayo 2010] which is currently under examination. While describing the general method, we will also mention the specific choices used when generating the maps for the study presented in [Perez-David 2011].

6.4.2.1 Map Construction

The map construction process starts with two inputs: a DE-MRI volume (which could be any of the acquired anatomic views, but will preferably be the fused DE-MRI volume) and a segmentation of the endo- and epi-cardial surfaces on that volume. The main steps, summarized in figure 6.13, are the following:

- Using the endocardial segmented surface, a three-dimensional triangle mesh is constructed using any existing triangulation algorithm (for instance, *marching cubes* or *marching tetrahedrons*).
- Both endo- and epi-cardial segmented contours were used to generate a mask of the myocardium which was then applied to the input DE-MRI volume in order to get a new volume in which only the myocardium intensity values from DE-MRI are present. For this work, the contours were manually generated using QMass® MR 7.0 (MEDIS, The Netherlands)
- We now select a region of interest within the myocardium, in this work the inner 50% of its width was used, and we *project* the selected myocardium intensity values into the vertex of the 3D mesh. While there are several ways in which we can choose to *project* the intensity values on the surface (see [Ledesma-Carbayo 2010])

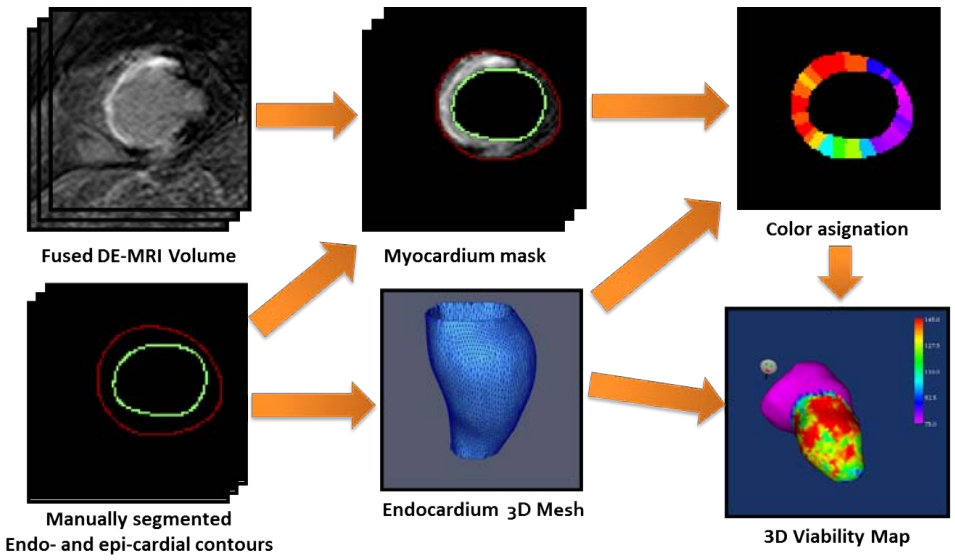

Figure 6.13: 3D Viability Map creation process.

for this work it was done using a minimum-distance criterium, i.e., each voxel in the myocardium was assigned to the closest vertex, and then the average intensity among all the voxels assigned to a given vertex was used as the vertex's *scalar value*.

- Next we need to define a colormap which specifies which color will be used when visualizing each vertex in the 3D mesh, based on the *scalar value* assigned to it in the previous step. For this work we used a colormap similar to the one used in EAM studies (i.e. a colormap ranging from magenta to red with linear behavior in the *hue* value of the HSV space). The colormap range was decided based on intensity values of small healthy and deep scar reference regions: any vertex with a value below the maximum intensity present in the healthy reference region was represented in magenta, any vertex with a value greater than the average intensity in the deep scar healthy reference region was shown in red, and in between the full colormap spectrum was used (figure 6.14).

- Once the 3D mesh, the *scalar values* for each of its vertex and the visualization colormap are available, we just need to visualize it using some adequate library, which in our case was VTK (Visualization ToolKit).

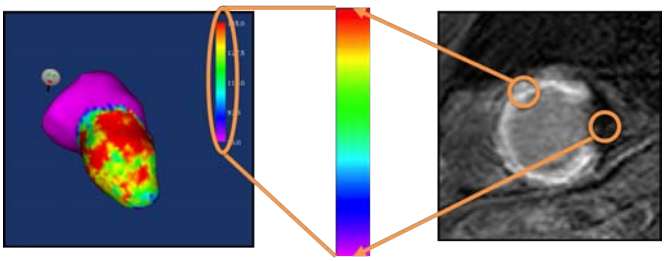

Figure 6.14: Colormap range selection based on healthy and deep-scar regions from the DE-MRI data.

6.4.2.2 Experiments and Results

The *viability maps* were generated from the fused DE-MRI volume as described in the previous sections, and a blind comparison was carried out on the study group between heterogeneous tissue corridors identified in the *viability maps* and slow conduction channels visible in endocardial voltage maps obtained during the electrophysiological study. For the comparison, the left ventricle was divided into 12 segments [Josephson 2002] and two orientations were defined: parallel or perpendicular to the mitral annulus. Both heterogeneous tissue corridors and conduction channels were classified according to those segments and segmentations, and a perfect match was found.

In figure 6.15 we show comparisons between viability maps based on DE-MRI signal intensity and voltage maps obtained during the electrophysiological study for two different patients. It is important to note that this side-by-side comparison was carried out after the channel identification had been performed in order to maintain blindness.

At the time of this study, there was no feasible way of importing the full viability maps into the CARTO® console used for the EAM procedures, so in order to perform a side-by-side comparison in the console itself we had to develop a custom export/import process in which the full colormap range of the viability map was divided into just three regions (healthy, heterogeneous and scar tissue) and those regions were exported as *fake* DICOM study with just four different pixel values (one for each region plus background). In this way we could trick the CARTO MERGE® tool, which is designed to import anatomical data from CT/MRI studies, and import the surfaces of the three regions. In figure 6.16 we show the resulting surfaces, corresponding to the two viability maps examples we presented before.

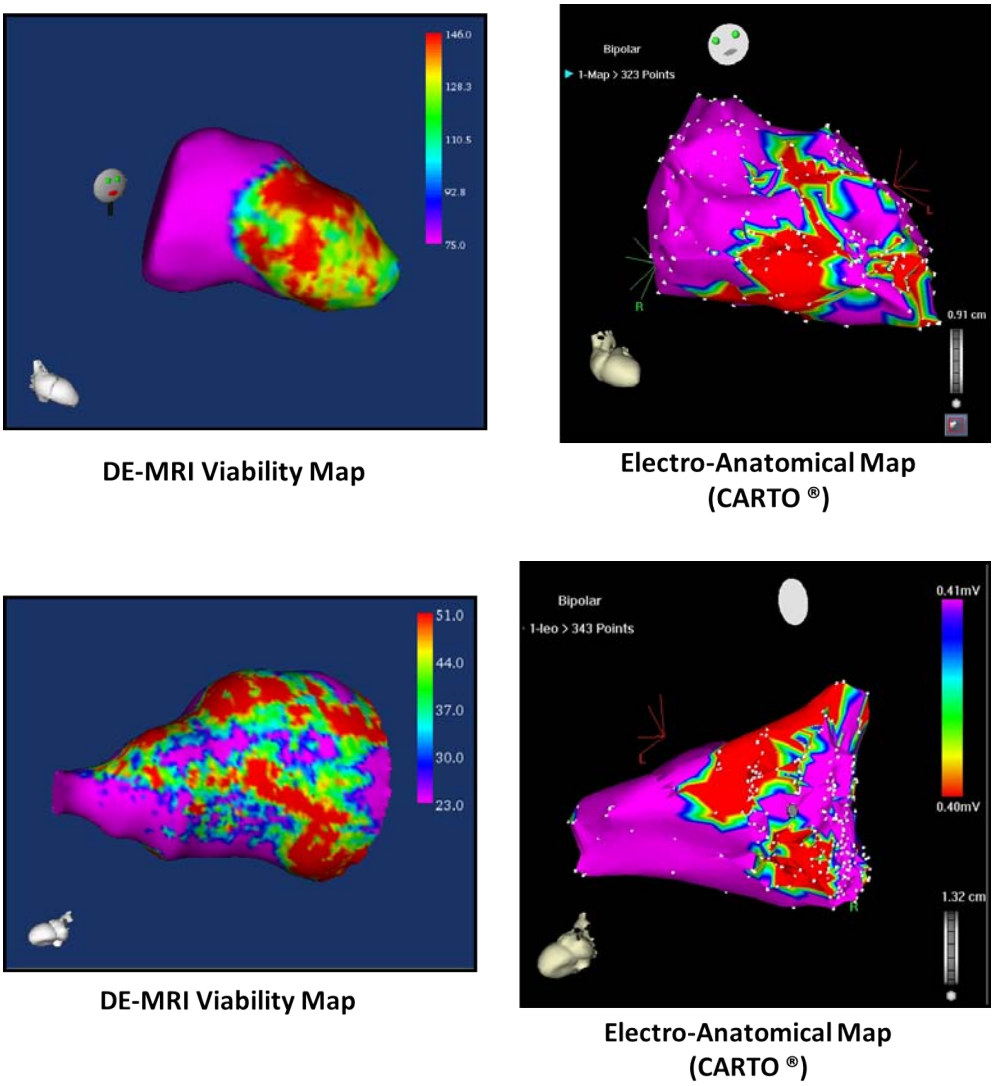

Figure 6.15: Two examples comparing *Viability Maps* (left) and the corresponding endocardial Voltage Maps (right).

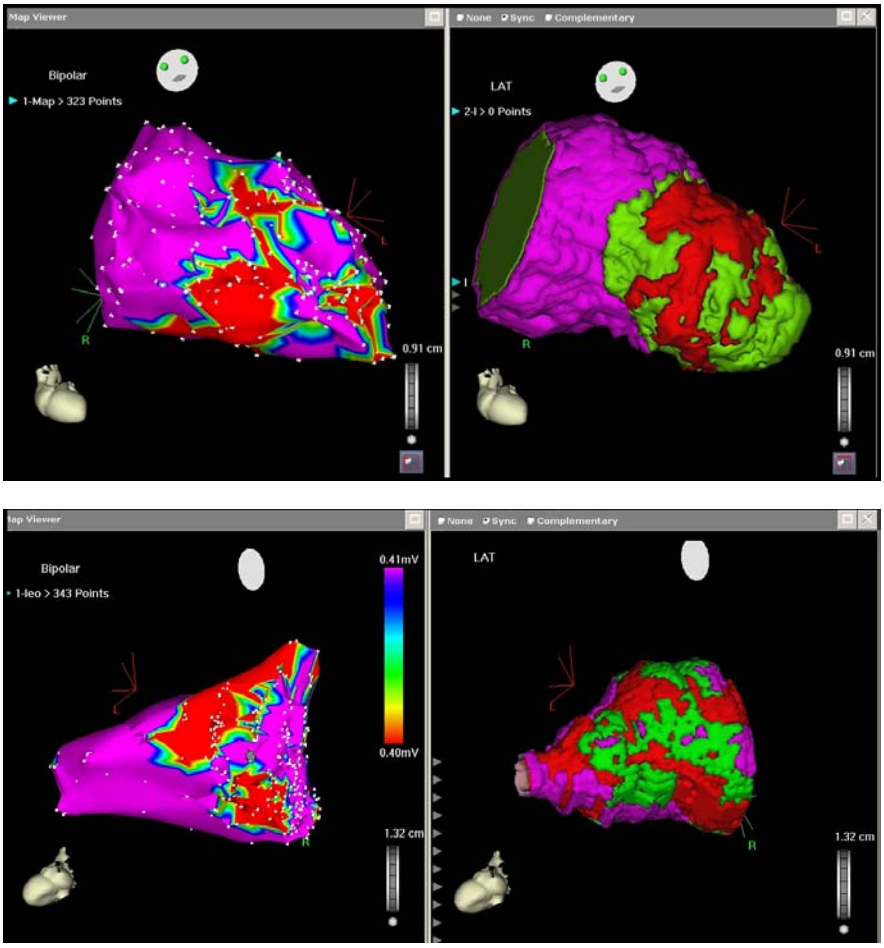

Figure 6.16: Two examples of endocardial voltage maps (left) and the corresponding *viability maps* imported into CARTO® (right).

6.4.3 Graphical User Interface

As part of this PhD Thesis, a Graphical User Interface (GUI) was developed in Java for the generation, visualization and manipulation of the viability maps described in the previous sections. The GUI asks for the DE-MRI study and the endo- and epicardial segmentations and performs the full map generation process. Figure 6.17 shows a screenshot of the user interface.

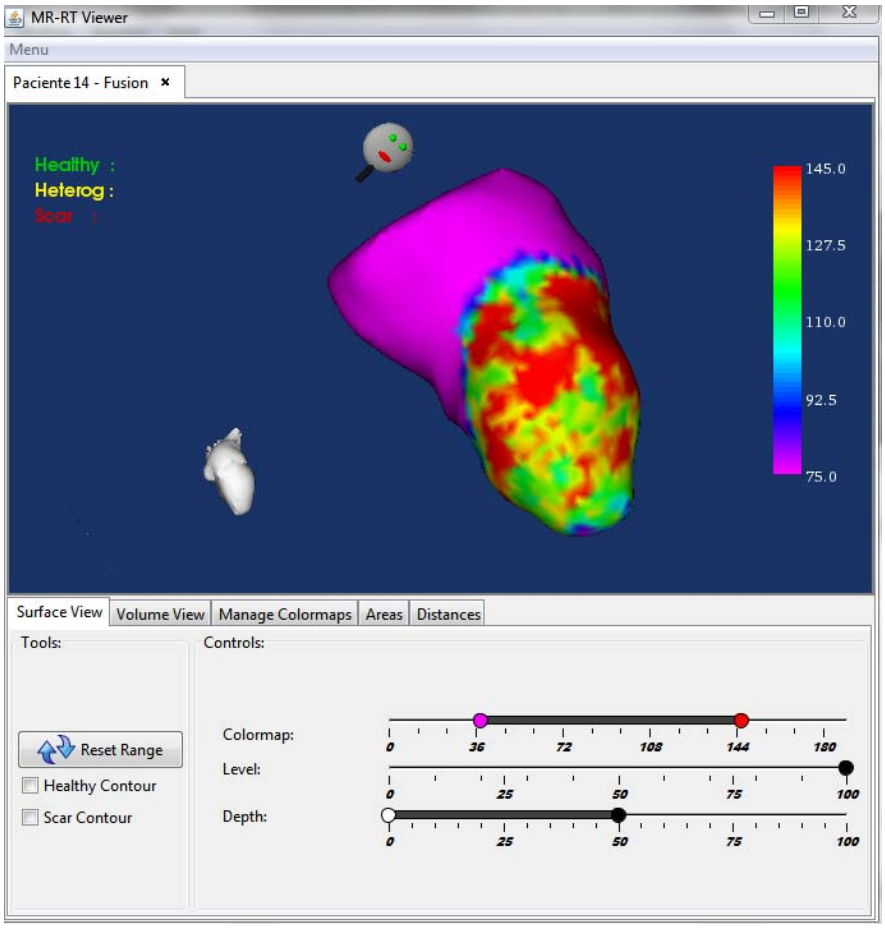

Figure 6.17: Screenshot of the GUI developed for the generation, visualization and manipulation of the viability maps.

The program has been further developed after the version which was used for the aforementioned study ([Perez-David 2011]) and now contains the following set of features:

- Viability Maps can be displayed jointly with the underlying DE-MRI data, allow-

ing a better understanding of the heterogeneous and scar regions when needed.

- It is possible to define the percentage of myocardium used for the computation of the viability map. This allows to display both endocardial (0% to 50% myocardium depth) and epicardium (50% to 100% myodardium depth) maps, or any other custom selection of depth-range.
- Total healthy, heterogeneous and scar areas are automatically computed and displayed.
- Manual measurements of distance and areas over the map surface are possible.
- Direct export to both CARTO and Ensite NavX is available from the GUI.
- Export to CARTO has been enhanced to take advantage of the direct map import capabilities of the latest CARTO release. This means that the full colormap can now be exported from the GUI and imported into CARTO, instead of the reduced 3-color version shown in figure 6.16.
- A given viability map study can be saved and later restored, including all the manually defined measures (distances and areas).

Finally, it is worth mentioning that the software has been transferred to the Non-Invasive Cardiology Department of Hospital General Universitario Gregorio Marañon (Madrid, Spain), where is currently been used for several on-going studies.

6.5 Discussion

In this chapter we have described how the proposed fusion method has been adapted to multi-view anisotropic Delayed Enhancement Magnetic Resonance Imaging (DE-MRI) datasets. The resulting fused volumes show a quantitative improvement in image sharpness and therefore enable an easier and more complete interpretation of complex three-dimensional scar and heterogeneous tissue information in ischemic cardiomyopathy patients.

Additionally, in order to take advantage of the improved resolution in the fused volumes, we have proposed and implemented a novel visualization method which simplifies the interpretation of MRI information for the non-invasive evaluation of myocardial substrate of rhythm abnormalities like Sustained Monomorphic Ventricular Tachycardia

(SMVT) for diagnostic and therapeutic purposes. This visualization method was filed as an Spanish patent in 2010 and the extended as a PCT (Patent Cooperation Treaty) application in 2011 [Ledesma-Carbayo 2010].

The clinical applicability of the novel visualization method has been tested on a set of 18 patients suffering form chronic myocardial infarction and SMVT. The results show that the non-invasive information contained in the DE-MRI data is highly correlated with electrical activation maps of the myocardium acquired through Electro-Anatomical Mapping (EAM), with a perfect match in location and orientation of the slow conduction channels identified by both methods.

Finally, it is worth mentioning that during the development of this thesis, another work on wavelet-based fusion of multi-view MRI datasets was reported in [Aganj 2012], in which the fusion was applied to T2-weighted MRI brain volumes and also to simulated phantoms. The method proposed in that work is similar to the one described in this chapter and is actually a particular case of the general method described in chapter 4, with a particular choice of decision rules which exploit the orthogonality of the brain MRI datasets they work with. The work in [Aganj 2012] confirms the interest and suitability of wavelet space fusion methods as an approach for combining multi-view MRI datasets.

When you cease to make a contribution, you begin to die.

Eleanor Roosevelt

Chapter 7

Contributions

The main contributions of this work are listed below:

- A new methodology has been proposed and developed which enables the combination of the complementary information spread among several three-dimensional *views* on multi-view imaging schemes. Such method makes possible the creation of a fused volume which condenses all the available information, surpassing the sampling and image quality limitations of individual *views*. The method operates in the wavelet domain and therefore exploits the ability of the discrete wavelet transform to distribute the information contained in a given volume into several scales and orientations.

- Efficient implementation has been achieved taking into account practical aspects: the method is able to work on really large datasets by minimizing the amount of information which is necessary to handle at any given point of the process pipeline; a novel solution has been proposed to deal with the problem of uncomplete views, which may introduce artifacts in the fused volume unless specific measures are adopted.

- The proposed method has been applied and validated on fluorescence microscopy multi-view datasets, obtained by means of Light-Sheet Fluorescence Microscopy from live embryos of animal models (sea urchin and zebrafish). The fusion results show improved contrast and detail discrimination when compared to any of the individual views and the method does not rely on prior knowledge of the system's point spread function (PSF). The results have shown improved performance with respect to another PSF independent method (content-based fusion),

with a 31% average relative improvement in the contrast of line profiles passing through adjacent nuclei. The method is currently in use at the Institut de Neurobiologie Alfred Fessard (CNRS, France) and has been made freely available to the scientific community.

- We have adapted the fusion method to multi-view anisotropic Delayed Enhancement Magnetic Resonance Imaging (DE-MRI) datasets. The resulting fused volumes enable an easier and more complete interpretation of complex three-dimensional scar and heterogeneous tissue information in ischemic cardiomyopathy patients. The results show that the fused volumes provide an average relative improvement close to 52% in their Second Derivative Visibility Operator value with respect to the original views when $5.0mm$-slice views are fused, and an average relative improvement close to 41% when $2.5mm$-slice views are fused.

- In order to take advantage of the improved resolution in the fused volumes, we have proposed and implemented a novel visualization method which simplifies the interpretation of MRI information for the non-invasive evaluation of myocardial substrate of rhythm abnormalities like Sustained Monomorphic Ventricular Tachycardia (SMVT) for diagnostic and therapeutic purposes. This visualization method was filed as a Spanish patent in 2010 and extended as a PCT (Patent Cooperation Treaty) application in 2011 [Ledesma-Carbayo 2010].

- The clinical applicability of the novel visualization method has been tested on a set of 18 patients suffering form chronic myocardial infarction and SMVT. The results show that the non-invasive information contained in the DE-MRI data is highly correlated with electrical activation maps of the myocardium acquired through Electro-Anatomical Mapping (EAM), with a good agreement in location and orientation of the slow conduction channels identified by both methods.

Appendix A

Additional Preprocessing Steps Applied to the Human DE-MRI Datasets

The human datasets used for the study reported in chapter 6 have some specific features that require additional care apart from the general pre-processing considerations discussed in section 6.2.1. The reason is that these datasets were acquired during breath-holds, and given the limited time available for each breath-hold ($\sim 14s$) the maximum number of slices that can be acquired during this time is $\sim 10 - 14$. With a slice separation of $5mm$, that number of slices is usually enough to cover the left ventricle when acquiring 2-chamber and 4-chamber long-axis view, but not when acquiring short-axis views. For this reason the short-axis volumes need to be divided in two halves, and acquired during two consecutive breath-holds.

Unfortunately, this division introduces an extra source of misalignment which needs to be corrected before the fusion takes place, as each of the short-axis halves is acquired during a different breath-hold and therefore the heart will not be in the same exact position on both halves. Moreover, due to the use of separate breath-hold and to some peculiarity of the precise MRI sequence used, in all the datasets used for this study we found that the two central slices in the resulting short-axis appear slightly under-exposed, presenting lower intensity values than the slices around them.

Figure A.1 gives an example of the two issues we just mentioned. In the figure, a section of the short-axis along a 4-chamber oriented plane is shown, where the two central slices appear clearly under-exposed and a discontinuity can be appreciated between

Figure A.1: Practical issues with the Short-Axis DE-MRI volumes when they are acquired during two separate breath-holds.

slices 10 and 11 (out of 20) due to the aforementioned misalignment of the sub-volumes acquired during each breath-hold.

In order to correct the "under-exposure" problem on the two central slices, an additional step was introduced in the pre-processing chain. Similarly to the way in which inter-volume normalization was performed (section 6.2.1.2), we use the manually segmented endocardium contours to compute the average blood-pool intensity level for each slice in the short-axis volume. Then by comparing the mean intensity of the under-exposed slices with that of the surrounding slices, a correction factor is computed and applied to the two central slices, as shown in figure A.2. This intra-volume normalization of the short-axis volume should be performed before the inter-volume normalization, as the short-axis will be the used as reference to normalized the other volumes.

Regarding the misalignment between the two halves which compose the short-axis, we solve it by using the same kind of cross-correlation registration mentioned in section 6.2.1.1. However in this case the problem is a bit trickier, as there is no common information among both halves of the short-axis. So in order to align them, we split the short-axis onto its two halves, we register each of the halves with one of the long-axis volumes available, and we use the difference among the transformations obtained from both registrations to align the two halves of the short-axis. Again, this intra-volume registration is performed before the inter-volume registration, and we also use the result of the intra-volume registration to modify the segmentations, which were obtained on the original (and misaligned) short-axis.

In figure A.3 we present an updated version of the pre-processing workflow, taking into account the two additional steps we just described.

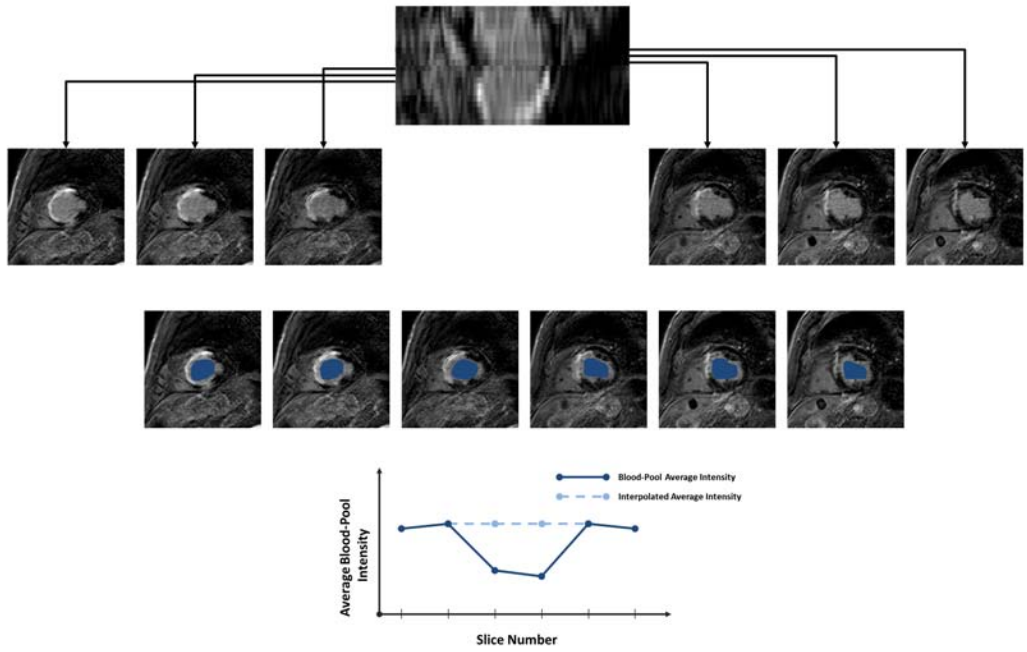

Figure A.2: Normalization of the two central slices in a DE-MRI Short-Axis volume acquired during two different breath-holds.

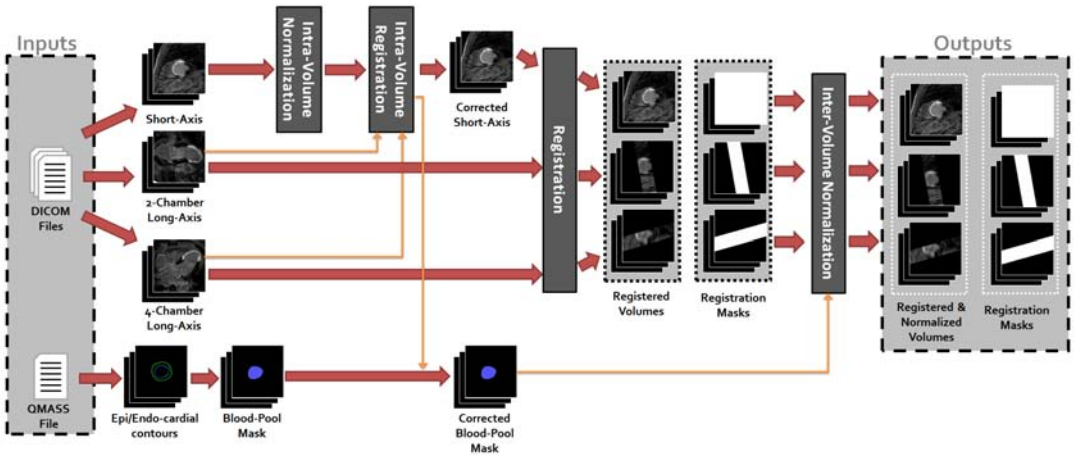

Figure A.3: Detailed workflow for pre-processing the DE-MRI datasets before they different views can be fused together. With respect to figure 6.2, the additional issues related to split-acquisition of the Short-Axis volume are considered here.

References

[Aganj 2012] Iman Aganj, Christophe Lenglet, Essa Yacoub, Guillermo Sapiro and Noam Harel. *A 3D wavelet fusion approach for the reconstruction of isotropic-resolution MR images from orthogonal anisotropic-resolution scans.* Magnetic resonance in medicine, vol. 67, no. 4, pages 1167–72, April 2012.

[Al-Saadi 2000] N Al-Saadi, E Nagel, M Gross, A Bornstedt, B Schnackenburg, C Klein, W Klimek, H Oswald and E Fleck. *Noninvasive detection of myocardial ischemia from perfusion reserve based on cardiovascular magnetic resonance.* Circulation, vol. 101, no. 12, pages 1379–83, March 2000.

[Amado 2004] Luciano C Amado, Bernhard L Gerber, Sandeep N Gupta, Dan W Rettmann, Gilberto Szarf, Robert Schock, Khurram Nasir, Dara L Kraitchman and João a C Lima. *Accurate and objective infarct sizing by contrast-enhanced magnetic resonance imaging in a canine myocardial infarction model.* Journal of the American College of Cardiology, vol. 44, no. 12, pages 2383–9, December 2004.

[Amolins 2007] K Amolins, Y Zhang and P Dare. *Wavelet based image fusion techniques.* ISPRS Journal of Photogrammetry and Remote Sensing, vol. 62, no. 4, pages 249–263, September 2007.

[Arenal 2004] Angel Arenal, Silvia del Castillo, Esteban Gonzalez-Torrecilla, Felipe Atienza, Mercedes Ortiz, Javier Jimenez, Alberto Puchol, Javier García and Jesús Almendral. *Tachycardia-related channel in the scar tissue in patients with sustained monomorphic ventricular tachycardias: influence of the voltage scar definition.* Circulation, vol. 110, no. 17, pages 2568–74, October 2004.

[Ashikaga 2007] Hiroshi Ashikaga, Tetsuo Sasano, Jun Dong, M Muz Zviman, Robert Evers, Bruce Hopenfeld, Valeria Castro, Robert H Helm, Timm Dickfeld, Saman

Nazarian, J Kevin Donahue, Ronald D Berger, Hugh Calkins, M Roselle Abraham, Eduardo Marbán, Albert C Lardo, Elliot R McVeigh and Henry R Halperin. *Magnetic resonance-based anatomical analysis of scar-related ventricular tachycardia: implications for catheter ablation.* Circulation research, vol. 101, no. 9, pages 939–47, October 2007.

[Bello 2005] David Bello, F. Scott Pereles, Rod Passman, Gina Song, Alan H. Kadish, Jeffrey J. Goldberger, Raymond J. Kim and David S. Fieno. *Infarct morphology identifies patients with substrate for sustained ventricular tachycardia.* Journal of the American College of Cardiology, vol. 45, no. 7, pages 1104–1108, April 2005.

[Bennett 2007] Eric P Bennett, John L Mason and Leonard McMillan. *Multispectral bilateral video fusion.* IEEE Transactions on Image Processing, vol. 16, no. 5, pages 1185–94, May 2007.

[Bogun 2009] Frank M. Bogun, Benoit Desjardins, Eric Good, Sanjaya Gupta, Thomas Crawford, Hakan Oral, Matthew Ebinger, Frank Pelosi, Aman Chugh, Krit Jongnarangsin and Fred Morady. *Delayed-enhanced magnetic resonance imaging in nonischemic cardiomyopathy: utility for identifying the ventricular arrhythmia substrate.* Journal of the American College of Cardiology, vol. 53, no. 13, pages 1138–45, March 2009.

[Burt 1983] P. Burt and E. Adelson. *The Laplacian Pyramid as a Compact Image Code.* IEEE Transactions on Communications, vol. 31, no. 4, pages 532–540, 1983.

[Burt 1984] Peter Jeffrey Burt. *The pyramid as a structure for efficient computation.* In Multiresolution Image Processing and Analysis, pages 6–35. Springer-Verlag, 1984.

[Catte 1992] Francine Catte, Pierre-Louis Lions, Jean-Michel Morel and Tomeu Coll. *Image Selective Smoothing and Edge Detection by Nonlinear Diffusion.* SIAM Journal on Numerical Analysis, vol. 29, no. 1, page 182, 1992.

[Chalfie 1994] M Chalfie, Y Tu, G Euskirchen, W W Ward and D C Prasher. *Green fluorescent protein as a marker for gene expression.* Science, vol. 263, no. 5148, pages 802–5, February 1994.

[Chang 2000] S G Chang, B Yu and M Vetterli. *Adaptive wavelet thresholding for image denoising and compression.* IEEE Transactions on Image Processing, vol. 9, no. 9, pages 1532–46, January 2000.

[Coifman 1995] R.R. Coifman and D.L. Donoho. *Translation-Invariant De-Noising.* In A. Antoniadis and G. Oppenheim, editeurs, Wavelets and Statistics, pages 125–150. Springer-Verlag, 1995.

[de Bakker 1988] J M de Bakker, F J van Capelle, M J Janse, A A Wilde, R Coronel, A E Becker, K P Dingemans, N M van Hemel and R N Hauer. *Reentry as a cause of ventricular tachycardia in patients with chronic ischemic heart disease: electrophysiologic and anatomic correlation.* Circulation, vol. 77, no. 3, pages 589–606, March 1988.

[DelMarco 2009] Stephen DelMarco and Sos Agaian. *The design of wavelets for image enhancement and target detection.* In Proceedings of SPIE, pages 735103–735103–12. SPIE, 2009.

[Denk 1990] W Denk, J H Strickler and W W Webb. *Two-photon laser scanning fluorescence microscopy.* Science, vol. 248, no. 4951, pages 73–6, April 1990.

[Deshpande 2001] V S Deshpande, S M Shea, G Laub, O P Simonetti, J P Finn and D Li. *3D magnetization-prepared true-FISP: a new technique for imaging coronary arteries.* Magnetic Resonance in Medicine, vol. 46, no. 3, pages 494–502, September 2001.

[Donoho 1994] D.L. Donoho and I.M. Johnstone. *Threshold selection for wavelet shrinkage of noisy data.* In Proceedings of 16th Annual International Conference of the IEEE EMBS, pages A24–A25. IEEE, 1994.

[Dzyubachyk 2010] Oleh Dzyubachyk, Jeroen Essers, Wiggert A van Cappellen, Céline Baldeyron, Akiko Inagaki, Wiro J Niessen and Erik Meijering. *Automated analysis of time-lapse fluorescence microscopy images: from live cell images to intracellular foci.* Bioinformatics, vol. 26, no. 19, pages 2424–2430, October 2010.

[Earls 2002] James P Earls, Vincent B Ho, Thomas K Foo, Ernesto Castillo and Scott D Flamm. *Cardiac MRI: recent progress and continued challenges.* Journal of magnetic resonance imaging, vol. 16, no. 2, pages 111–27, August 2002.

[Elad 1997] M. Elad and A. Feuer. *Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images.* Image Processing, IEEE Transactions on, vol. 6, no. 12, pages 1646–1658, December 1997.

[Engelbrecht 2006] Christoph J Engelbrecht and Ernst H Stelzer. *Resolution enhancement in a light-sheet-based microscope (SPIM).* Optics letters, vol. 31, no. 10, pages 1477–9, May 2006.

[Fayad 2000] Z A Fayad, V Fuster, J T Fallon, T Jayasundera, S G Worthley, G Helft, J G Aguinaldo, J J Badimon and S K Sharma. *Noninvasive in vivo human coronary artery lumen and wall imaging using black-blood magnetic resonance imaging.* Circulation, vol. 102, no. 5, pages 506–10, August 2000.

[Fieno 2000] David S Fieno, Raymond J Kim, E L Chen, Jon W Lomasney, Francis J Klocke and Robert M Judd. *Contrast-enhanced magnetic resonance imaging of myocardium at risk: distinction between reversible and irreversible injury throughout infarct healing.* Journal of the American College of Cardiology, vol. 36, no. 6, pages 1985–91, November 2000.

[Finn 2006] J Paul Finn, Kambiz Nael, Vibhas Deshpande, Osman Ratib and Gerhard Laub. *Cardiac MR imaging: state of the technology.* Radiology, vol. 241, no. 2, pages 338–54, November 2006.

[Gonzalez-Audicana 2004] M. Gonzalez-Audicana, J.L. Saleta, R.G. Catalan and R. Garcia. *Fusion of multispectral and panchromatic images using improved IHS and PCA mergers based on wavelet decomposition.* IEEE Transactions on Geoscience and Remote Sensing, vol. 42, no. 6, pages 1291–1299, June 2004.

[González-Audícana 2005] M. González-Audícana, X. Otazu, O. Fors and A. Seco. *Comparison between Mallat's and the 'à trous' discrete wavelet transform based algorithms for the fusion of multispectral and panchromatic images.* International Journal of Remote Sensing, vol. 26, no. 3, pages 595–614, 2005.

[Gonzalez 2007] Rafael C. Gonzalez and Richard E. Woods. Digital Image Procesing. Prentice Hall, Reading, MA, 3rd édition, 2007.

[Goutsias 2000] J Goutsias and H M Heijmans. *Nonlinear multiresolution signal decomposition schemes–part I: morphological pyramids.* IEEE Transactions on Image Processing, vol. 9, no. 11, pages 1862–76, January 2000.

[Graf 2005] Ralph Graf, Jens Rietdorf and Timo Zimmermann. *Live cell spinning disk microscopy.* Advances in biochemical engineering/biotechnology, vol. 95, pages 57–75, January 2005.

[Greenspan 2002] H. Greenspan, G. Oz, N. Kiryati and S. Peled. *MRI inter-slice reconstruction using super-resolution.* Magnetic resonance imaging, vol. 20, no. 5, pages 437–46, June 2002.

[Greenspan 2008] H. Greenspan. *Super-Resolution in Medical Imaging.* The Computer Journal, vol. 52, no. 1, pages 43–63, 2008.

[Guihong 2001] Qu Guihong, Zhang Dali and Yan Pingfan. *Medical image fusion by wavelet transform modulus maxima.* Optics Express, vol. 9, no. 4, page 184, August 2001.

[Halperin 2011] Henry R Halperin and Saman Nazarian. *Magnetic resonance identification of the ventricular tachycardia critical isthmus finding the needle in the haystack.* Journal of the American College of Cardiology, vol. 57, no. 2, pages 195–7, January 2011.

[Helmchen 2005] Fritjof Helmchen and Winfried Denk. *Deep tissue two-photon microscopy.* Nature methods, vol. 2, no. 12, pages 932–40, December 2005.

[Hsia 2006] Henry H Hsia, David Lin, William H Sauer, David J Callans and Francis E Marchlinski. *Anatomic characterization of endocardial substrate for hemodynamically stable reentrant ventricular tachycardia: identification of endocardial conducting channels.* Heart rhythm, vol. 3, no. 5, pages 503–12, May 2006.

[Huisken 2004] Jan Huisken, Jim Swoger, Filippo Del Bene, Joachim Wittbrodt and Ernst H. K. Stelzer. *Optical sectioning deep inside live embryos by selective plane illumination microscopy.* Science, vol. 305, no. 5686, pages 1007–9, August 2004.

[Huisken 2009] Jan Huisken and Didier Y. R. Stainier. *Selective plane illumination microscopy techniques in developmental biology.* Development, vol. 136, no. 12, pages 1963–1975, June 2009.

[Irani 1991] M. Irani and S. Peleg. *Improving resolution by image registration.* CVGIP: Graphical Models and Image Processing, vol. 53, no. 3, pages 231–239, 1991.

[Josephson 2002] Mark E. Josephson. Clinical Cardiac Electrophysiology. Lippincott Williams & Wilkins, Philadelphia, PA, 3rd édition, 2002.

[Keller 2008a] Philipp J Keller, Annette D Schmidt, Joachim Wittbrodt and Ernst H K Stelzer. *Reconstruction of zebrafish early embryonic development by scanned light sheet microscopy.* Science, vol. 322, no. 5904, pages 1065–9, November 2008.

[Keller 2008b] Philipp J Keller and Ernst H K Stelzer. *Quantitative in vivo imaging of entire embryos with Digital Scanned Laser Light Sheet Fluorescence Microscopy.* Current opinion in neurobiology, vol. 18, no. 6, pages 624–32, December 2008.

[Keller 2010] Philipp J Keller, Annette D Schmidt, Anthony Santella, Khaled Khairy, Zhirong Bao, Joachim Wittbrodt and Ernst H K Stelzer. *Fast, high-contrast imaging of animal development with scanned light sheet-based structured-illumination microscopy.* Nature Methods, vol. 7, no. 8, pages 637–42, August 2010.

[Kennedy 2006] John A. Kennedy, Ora Israel, Alex Frenkel, Rachel Bar-Shalom and Haim Azhari. *Super-resolution in PET imaging.* IEEE transactions on medical imaging, vol. 25, no. 2, pages 137–47, February 2006.

[Kennedy 2007] John A. Kennedy, Ora Israel, Alex Frenkel, Rachel Bar-Shalom and Haim Azhari. *Improved Image Fusion in PET/CT Using Hybrid Image Reconstruction and Super-Resolution.* International journal of biomedical imaging, vol. 2007, page 46846, January 2007.

[Kim 1999] Raymond J. Kim, David S. Fieno, Todd B. Parrish, Kathleen Harris, E L Chen, Orlando Simonetti, Jeffrey Bundy, J. Paul Finn, Francis J. Klocke and Robert M. Judd. *Relationship of MRI delayed contrast enhancement to irreversible injury, infarct age, and contractile function.* Circulation, vol. 100, no. 19, pages 1992–2002, November 1999.

[Kim 2000] R J Kim, E Wu, A Rafael, E L Chen, M A Parker, O Simonetti, F J Klocke, R O Bonow and R M Judd. *The use of contrast-enhanced magnetic resonance imaging to identify reversible myocardial dysfunction.* The New England journal of medicine, vol. 343, no. 20, pages 1445–53, November 2000.

[Kim 2003] Raymond J. Kim, Dipan J. Shah and Robert M. Judd. *How We Perform Delayed Enhancement Imaging.* Journal of Cardiovascular Magnetic Resonance, vol. 5, no. 3, pages 505–514, 2003.

[Kim 2008] Raymond J Kim, Timothy S E Albert, James H Wible, Michael D Elliott, John C Allen, Jennifer C Lee, Michele Parker, Alicia Napoli and Robert M

Judd. *Performance of delayed-enhancement magnetic resonance imaging with gadoversetamide contrast for the detection and assessment of myocardial infarction: an international, multicenter, double-blinded, randomized trial.* Circulation, vol. 117, no. 5, pages 629–37, February 2008.

[Kino 2009] Aya Kino, Sven Zuehlsdorff, John J Sheehan, Peter J Weale, Timothy J Carroll, Renate Jerecic and James C Carr. *Three-dimensional phase-sensitive inversion-recovery turbo FLASH sequence for the evaluation of left ventricular myocardial scar.* American Journal of Roentgenology, vol. 193, no. 5, pages W381–8, November 2009.

[Klein 2002] C. Klein. *Assessment of Myocardial Viability With Contrast-Enhanced Magnetic Resonance Imaging: Comparison With Positron Emission Tomography.* Circulation, vol. 105, no. 2, pages 162–167, January 2002.

[Krzic 2009] Uros Krzic. *Multiple-View Microscopy with Light-Sheet based Fluorescence Microscope.* PhD thesis, University of Heidelberg, 2009.

[Ledesma-Carbayo 2010] Maria J. Ledesma-Carbayo, Andres Santos, Jose L. Rubio-Guivernau, Angel Arenal, Esther Perez-David, Javier Bermejo, Manuel Desco and Francisco Fernandez-Aviles. *Method for Displaying the Information Contained in Three-Dimensional Images of the Heart*, 2010.

[Li 1995] H Li, B.S. Manjunath and S.K. Mitra. *Multisensor Image Fusion Using the Wavelet Transform.* Graphical Models and Image Processing, vol. 57, no. 3, pages 235–245, May 1995.

[Liang 1999] Zhi-Pei Liang and Paul C. Lauterbur. Principles of Magnetic Resonance Imaging: A Signal Processing Perspective. Wiley-IEEE Press, 1 édition, 1999.

[Lichtman 2005] Jeff W Lichtman and José-Angel Conchello. *Fluorescence microscopy.* Nature methods, vol. 2, no. 12, pages 910–9, December 2005.

[Liu 2001] Z Liu. *Image fusion by using steerable pyramid.* Pattern Recognition Letters, vol. 22, no. 9, pages 929–939, July 2001.

[Lucy 1974] L. B. Lucy. *An iterative technique for the rectification of observed distributions.* The Astronomical Journal, vol. 79, page 745, June 1974.

[Mallat 2008] Stephane Mallat. A Wavelet Tour of Signal Processing. Academic Press, 3rd édition, 2008.

[McGann 2008] Christopher J McGann, Eugene G Kholmovski, Robert S Oakes, Joshua J E Blauer, Marcos Daccarett, Nathan Segerson, Kelly J Airey, Nazem Akoum, Eric Fish, Troy J Badger, Edward V R DiBella, Dennis Parker, Rob S MacLeod and Nassir F Marrouche. *New magnetic resonance imaging-based method for defining the extent of left atrial wall injury after the ablation of atrial fibrillation.* Journal of the American College of Cardiology, vol. 52, no. 15, pages 1263–71, October 2008.

[McVeigh 1996] Elliot R McVeigh. *MRI of myocardial function: motion tracking techniques.* Magnetic Resonance Imaging, vol. 14, no. 2, pages 137–150, January 1996.

[Megason 2003] Sean G. Megason and Scott E. Fraser. *Digitizing life at the level of the cell: high-performance laser-scanning microscopy and image analysis for in toto imaging of development.* Mechanisms of Development, vol. 120, no. 11, pages 1407–1420, November 2003.

[Minsky 1961] Marvin Minsky. *US Patent 3013467: Microscopy Apparatus*, 1961.

[Moghari 2012] Mehdi H. Moghari, Raymond H. Chan, Susie N. Hong, Jaime L. Shaw, Lois A. Goepfert, Kraig V. Kissinger, Beth Goddu, Mark E. Josephson, Warren J. Manning and Reza Nezafat. *Free-breathing cardiac MR with a fixed navigator efficiency using adaptive gating window size.* Magnetic Resonance in Medicine, February 2012.

[Muzzey 2009] Dale Muzzey and Alexander van Oudenaarden. *Quantitative time-lapse fluorescence microscopy in single cells.* Annu Rev Cell Dev Biol, vol. 25, pages 301–27, January 2009.

[Nazarian 2005] Saman Nazarian, David A. Bluemke, Albert C. Lardo, Menekhem M. Zviman, Stanley P. Watkins, Timm L. Dickfeld, Glenn R. Meininger, Ariel Roguin, Hugh Calkins, Gordon F. Tomaselli, Robert G. Weiss, Ronald D. Berger, Joao A.C. Lima and Henry R. Halperin. *Magnetic Resonance Assessment of the Substrate for Inducible Ventricular Tachycardia in Nonischemic Cardiomyopathy.* Circulation, vol. 112, no. 18, pages 2821–2825, November 2005.

[Olivier 2010] N. Olivier, M. A. Luengo-Oroz, L. Duloquin, E. Faure, T. Savy, I. Veilleux, X. Solinas, D. Debarre, P. Bourgine, A. Santos, N. Peyrieras and E. Beaurepaire. *Cell Lineage Reconstruction of Early Zebrafish Embryos Using*

Label-Free Nonlinear Microscopy. Science, vol. 329, no. 5994, pages 967–971, August 2010.

[Osman 1999] N F Osman, W S Kerwin, E R McVeigh and J L Prince. *Cardiac motion tracking using CINE harmonic phase (HARP) magnetic resonance imaging.* Magnetic Resonance in Medicine, vol. 42, no. 6, pages 1048–60, December 1999.

[Ozturk 2003] Cengizhan Ozturk, J Andrew Derbyshire and Elliot R McVeigh. *Estimating Motion From MRI Data.* Proceedings of the IEEE, vol. 9, no. 10, pages 1627–1648, October 2003.

[Pajares 2004] G Pajares and J Manuel de la Cruz. *A wavelet-based image fusion tutorial.* Pattern Recognition, vol. 37, no. 9, pages 1855–1872, September 2004.

[Paris 2009] Sylvain Paris, Pierre Kornprobst, JackTumblin Tumblin and Frédo Durand. Bilateral Filtering: Theory and Applications, volume 4. 2009.

[Park 2003] S. C. Park, M. K. Park and M. G. Kang. *Super-resolution image reconstruction: a technical overview.* IEEE Signal Processing Magazine, vol. 20, no. 3, pages 21–36, 2003.

[Peinado Peinado 2008] Rafael Peinado Peinado, Esteban G Torrecilla, José Ormaetxe and Miguel Alvarez. *[Spanish Implantable Cardioverter-Defibrillator Registry. Fourth Official Report of the Spanish Society of Cardiology Working Group on Implantable Cardioverter-Defibrillators (2007)].* Revista española de cardiología, vol. 61, no. 11, pages 1191–203, November 2008.

[Perez-David 2011] Esther Perez-David, Angel Arenal, José L Rubio-Guivernau, Roberto Del Castillo, Leonardo Atea, Elena Arbelo, Eduardo Caballero, Verónica Celorrio, Tomas Datino, Esteban Gonzalez-Torrecilla, Felipe Atienza, Maria J Ledesma-Carbayo, Javier Bermejo, Alfonso Medina and Francisco Fernández-Avilés. *Noninvasive identification of ventricular tachycardia-related conducting channels using contrast-enhanced magnetic resonance imaging in patients with chronic myocardial infarction comparison of signal intensity scar mapping and endocardial voltage mapping.* Journal of the American College of Cardiology, vol. 57, no. 2, pages 184–94, January 2011.

[Perona 1990] P. Perona and J. Malik. *Scale-space and edge detection using anisotropic diffusion.* IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 12, no. 7, pages 629–639, July 1990.

[Peters 2009] Dana C Peters, Evan A Appelbaum, Reza Nezafat, Basem Dokhan, Yuchi Han, Kraig V Kissinger, Beth Goddu and Warren J Manning. *Left ventricular infarct size, peri-infarct zone, and papillary scar measurements: A comparison of high-resolution 3D and conventional 2D late gadolinium enhancement cardiac MR.* Journal of magnetic resonance imaging, vol. 30, no. 4, pages 794–800, October 2009.

[Piella 2003] Gemma Piella. *A general framework for multiresolution image fusion: from pixels to regions.* Information Fusion, vol. 4, no. 4, pages 259–280, 2003.

[Prasher 1992] D C Prasher, V K Eckenrode, W W Ward, F G Prendergast and M J Cormier. *Primary structure of the Aequorea victoria green-fluorescent protein.* Gene, vol. 111, no. 2, pages 229–33, February 1992.

[Preibisch 2008] Stephan Preibisch, Torsten Rohlfing, Michael P. Hasak and Pavel Tomancak. *Mosaicing of single plane illumination microscopy images using groupwise registration and fast content-based image fusion.* Proc. of SPIE, vol. 6194, pages 69140E 1–8, 2008.

[Preibisch 2010] Stephan Preibisch, Stephan Saalfeld, Johannes Schindelin and Pavel Tomancak. *Software for bead-based registration of selective plane illumination microscopy data.* Nature Methods, vol. 7, no. 6, pages 418–419, June 2010.

[Rajpoot 2009] Kashif Rajpoot, J. Noble, Vicente Grau, Cezary Szmigielski and Harald Becher. *Multiview RT3D Echocardiography Image Fusion.* In Lecture Notes in Computer Science, pages 134–143. Springer, 2009.

[Rajpoot 2011] Kashif Rajpoot, Vicente Grau, J Alison Noble, Cezary Szmigielski and Harald Becher. *Multiview fusion 3-D echocardiography: improving the information and quality of real-time 3-D echocardiography.* Ultrasound in medicine & biology, vol. 37, no. 7, pages 1056–72, July 2011.

[Reddy 2004] Vivek Y. Reddy, Zachary J. Malchano, Godtfred Holmvang, Ehud J. Schmidt, Andre D'Avila, Christopher Houghtaling, Raymond C. Chan and Jeremy N. Ruskin. *Integration of cardiac magnetic resonance imaging with three-dimensional electroanatomic mapping to guide left ventricular catheter manipulation: Feasibility in a porcine model of healed myocardial infarction.* Journal of the American College of Cardiology, vol. 44, no. 11, pages 2202–2213, December 2004.

[Redondo-Tejedor 2007] Rafael Redondo-Tejedor. *New Contributions on Image Fusion and Compression Based on Space-Frequency Representations.* PhD thesis, Universidad Politecnica de Madrid, 2007.

[Ricciardi 2001] M. J. Ricciardi, E. Wu, C. J. Davidson, K. M. Choi, F. J. Klocke, R. O. Bonow, R. M. Judd and R. J. Kim. *Visualization of Discrete Microinfarction After Percutaneous Coronary Intervention Associated With Mild Creatine Kinase-MB Elevation.* Circulation, vol. 103, no. 23, pages 2780–2783, June 2001.

[Richardson 1972] William Hadley Richardson. *Bayesian-Based Iterative Method of Image Restoration.* Journal of the Optical Society of America, vol. 62, no. 1, page 55, January 1972.

[Roes 2009] Stijntje D Roes, C Jan Willem Borleffs, Rob J van der Geest, Jos J M Westenberg, Nina Ajmone Marsan, Theodorus a M Kaandorp, Johan H C Reiber, Katja Zeppenfeld, Hildo J Lamb, Albert de Roos, Martin J Schalij and Jeroen J Bax. *Infarct tissue heterogeneity assessed with contrast-enhanced MRI predicts spontaneous ventricular arrhythmia in patients with ischemic cardiomyopathy and implantable cardioverter-defibrillator.* Circulation. Cardiovascular imaging, vol. 2, no. 3, pages 183–90, May 2009.

[Rudin 1992] L Rudin, S Osher and E Fatemi. *Nonlinear total variation based noise removal algorithms.* Physica D: Nonlinear Phenomena, vol. 60, no. 1-4, pages 259–268, November 1992.

[Rudin 1994] L.I. Rudin and S. Osher. *Total variation based image restoration with free local constraints.* In Proceedings of 1st International Conference on Image Processing, volume 1, pages 31–35. IEEE Comput. Soc. Press, 1994.

[Sakuma 1993] H Sakuma, N Fujita, T K Foo, G R Caputo, S J Nelson, J Hartiala, A Shimakawa and C B Higgins. *Evaluation of left ventricular volume and mass with breath-hold cine MR imaging.* Radiology, vol. 188, no. 2, pages 377–80, August 1993.

[Schmidt 2007] André Schmidt, Clerio F Azevedo, Alan Cheng, Sandeep N Gupta, David a Bluemke, Thomas K Foo, Gary Gerstenblith, Robert G Weiss, Eduardo Marbán, Gordon F Tomaselli, João a C Lima and Katherine C Wu. *Infarct tissue heterogeneity by magnetic resonance imaging identifies enhanced*

cardiac arrhythmia susceptibility in patients with left ventricular dysfunction. Circulation, vol. 115, no. 15, pages 2006–14, April 2007.

[Sechtem 1987] U Sechtem, P W Pflugfelder, R D White, R G Gould, W Holt, M J Lipton and C B Higgins. *Cine MR imaging: potential for the evaluation of cardiovascular function.* American Journal of Roentgenology, vol. 148, no. 2, pages 239–46, February 1987.

[Shimomura 1972] Osamu Shimomura and Frank H. Johnson. *Structure of the light-emitting moiety of aequorin.* Biochemistry, vol. 11, no. 9, pages 1602–1608, April 1972.

[Simonetti 2001] O P Simonetti, R J Kim, D S Fieno, H B Hillenbrand, E Wu, J M Bundy, J P Finn and R M Judd. *An improved MR imaging technique for the visualization of myocardial infarction.* Radiology, vol. 218, no. 1, pages 215–23, January 2001.

[Soler 2005] P. Soler, O. Gerard, P. Allain, E. Saloux, E. Angelini and I. Bloch. *Comparison of fusion techniques for 3D+T echocardiography acquisitions from different acoustic windows.* In Computers in Cardiology, 2005, pages 141–144. IEEE, 2005.

[Stearns 1995] T Stearns. *Green fluorescent protein. The green revolution.* Current Biology, vol. 5, no. 3, pages 262–4, March 1995.

[Stelzer 1994] E Stelzer. *Fundamental reduction of the observation volume in far-field light microscopy by detection orthogonal to the illumination axis: confocal theta microscopy.* Optics Communications, vol. 111, no. 5-6, pages 536–547, October 1994.

[Stevenson 1993] W G Stevenson, H Khan, P Sager, L A Saxon, H R Middlekauff, P D Natterson and I Wiener. *Identification of reentry circuit sites during catheter mapping and radiofrequency ablation of ventricular tachycardia late after myocardial infarction.* Circulation, vol. 88, no. 4 Pt 1, pages 1647–70, October 1993.

[Strong 1959] John Strong and G.A. Vanasse. *Interferometric Spectroscopy in the Far Infrared.* Journal of the Optical Society of America, vol. 49, no. 9, page 844, September 1959.

[Swoger 2007] Jim Swoger, Peter Verveer, Klaus Greger, Jan Huisken and Ernst H. K. Stelzer. *Multi-view image fusion improves resolution in three-dimensional microscopy.* Optics Express, vol. 15, no. 13, page 8029, June 2007.

[Swoger 2010] Jim Swoger, Mariana Muzzopappa, Hernán López-Schier and James Sharpe. *4D retrospective lineage tracing using SPIM for zebrafish organogenesis studies.* Journal of Biophotonics, vol. 4, no. 1-2, pages 122–134, October 2010.

[Tandri 2005] Harikrishna Tandri, Manoj Saranathan, E Rene Rodriguez, Claudia Martinez, Chandra Bomma, Khurram Nasir, Boas Rosen, João A C Lima, Hugh Calkins and David A Bluemke. *Noninvasive detection of myocardial fibrosis in arrhythmogenic right ventricular cardiomyopathy using delayed-enhancement magnetic resonance imaging.* Journal of the American College of Cardiology, vol. 45, no. 1, pages 98–103, January 2005.

[Temerinac-Ott 2011] Maja Temerinac-Ott, Olaf Ronneberger, Roland Nitschke, Wolfgang Driever and Hans Burkhardt. *Spatially-variant Lucy-Richardson deconvolution for multiview fusion of microscopical 3D images.* In IEEE ISBI 2011, pages 899–904, March 2011.

[Teraoka 2004] Kunihiko Teraoka, Masaharu Hirano, Hiroyuki Ookubo, Kazuyoshi Sasaki, Hiroaki Katsuyama, Masayuki Amino, Yimihiko Abe and Akira Yamashina. *Delayed contrast enhancement of MRI in hypertrophic cardiomyopathy.* Magnetic resonance imaging, vol. 22, no. 2, pages 155–61, February 2004.

[Thevenaz 1998] P Thevenaz, U E Ruttimann and M Unser. *A pyramid approach to subpixel registration based on intensity.* IEEE Transactions on Image Processing, vol. 7, no. 1, pages 27–41, January 1998.

[Thijssen 2011] Joep Thijssen, C Jan Willem Borleffs, Johannes B van Rees, Mihály K de Bie, Enno T van der Velde, Lieselot van Erven, Jeroen J Bax, Suzanne C Cannegieter and Martin J Schalij. *Driving restrictions after implantable cardioverter defibrillator implantation: an evidence-based approach.* European heart journal, vol. 32, no. 21, pages 2678–87, November 2011.

[Truong 2011] Thai V Truong and Willy Supatto. *Toward high-content/high-throughput imaging and analysis of embryonic morphogenesis.* Genesis, vol. 49, no. 7, pages 555–569, April 2011.

[Tukey 1976] J.W. Tukey. Exploratory data analysis. Addison-Wesley, Reading, MA, 1976.

[Viallon 2011] Magalie Viallon, Alexis Jacquier, Carmen Rotaru, Bénédicte M A Delattre, Nathan Mewton, Fabrice Vincent and Pierre Croisille. *Head-to-head comparison of eight late gadolinium-enhanced cardiac MR (LGE CMR) sequences at 1.5 tesla: from bench to bedside.* Journal of magnetic resonance imaging, vol. 34, no. 6, pages 1374–87, December 2011.

[Vonesch 2006] C. Vonesch, F. Aguet, J.-L. Vonesch and M. Unser. *The colored revolution of bioimaging.* IEEE Signal Processing Magazine, vol. 23, no. 3, pages 20–31, May 2006.

[Wagner 2003] A Wagner, H Mahrholdt, T Holly, M Elliott, M Regenfus, M Parker, F Klocke, R Bonow, R Kim and R Judd. *Contrast-enhanced MRI and routine single photon emission computed tomography (SPECT) perfusion imaging for detection of subendocardial myocardial infarcts: an imaging study.* The Lancet, vol. 361, no. 9355, pages 374–379, February 2003.

[Weber 2003] Oliver M Weber, Alastair J Martin and Charles B Higgins. *Whole-heart steady-state free precession coronary artery magnetic resonance angiography.* Magnetic Resonance in Medicine, vol. 50, no. 6, pages 1223–8, December 2003.

[Wu 2001] Edwin Wu, Robert M Judd, John D Vargas, Francis J Klocke, Robert O Bonow and Raymond J Kim. *Visualisation of presence, location, and transmural extent of healed Q-wave and non-Q-wave myocardial infarction.* The Lancet, vol. 357, no. 9249, pages 21–8, January 2001.

[Yan 2006] Andrew T Yan, Adolphe J Shayne, Kenneth a Brown, Sandeep N Gupta, Carmen W Chan, Tuan M Luu, Marcelo F Di Carli, H Glenn Reynolds, William G Stevenson and Raymond Y Kwong. *Characterization of the peri-infarct zone by contrast-enhanced cardiac magnetic resonance imaging is a powerful predictor of post-myocardial infarction mortality.* Circulation, vol. 114, no. 1, pages 32–9, July 2006.

[Yang 2009] Q. Yang, K.H. Tan and N. Ahuja. *Real-time O(1) bilateral filtering.* IEEE CVPR 2009, no. 1, pages 557–564, June 2009.

[Zitová 2003] Barbara Zitová and Jan Flusser. *Image registration methods: a survey.* Image and Vision Computing, vol. 21, no. 11, pages 977–1000, October 2003.

Publications

Journal Articles

- J. L. Rubio-Guivernau, V. Gurchenkov, M. A. Luengo-Oroz, L. Duloquin, P. Bourgine, A. Santos, N. Peyrieras, and M. J. Ledesma-Carbayo, "Wavelet-based image fusion in multi-view three-dimensional microscopy.," *Bioinformatics*, vol. 28, pp. 238–45, Jan. 2012.

- A. Arenal, J. Hernández, E. Pérez-David, J. L. Rubio-Guivernau, M. J. Ledesma-Carbayo, and F. Fernández-Avilés, "Do the spatial characteristics of myocardial scar tissue determine the risk of ventricular arrhythmias?," *Cardiovascular research*, Apr. 2012.

- M. A. Luengo-Oroz, J. L. Rubio-Guivernau, E. Faure, T. Savy, L. Duloquin, N. Olivier, D. Pastor, M. Ledesma-Carbayo, D. Debarre, P. Bourgine, E. Beaurepaire, N. Peyrieras, and A. Santos, "Methodology for reconstructing early zebrafish development from in vivo multiphoton microscopy.," *IEEE Transactions on Image Processing*, vol. 21, pp. 2335–40, Apr. 2012.

- M. A. Luengo-Oroz, D. Pastor, C. Castro, E. Faure, T. Savy, B. Lombardot, J. L. Rubio-Guivernau, L. Duloquin, M. Ledesma-Carbayo, P. Bourgine, N. Peyrieras, and A. Santos, "3D+t Morphological Processing: Applications to Embryogenesis Image Analysis.," *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, vol. 21, pp. 3518–3530, May 2012.

- E. Perez-David, A. Arenal, J. L. Rubio-Guivernau, R. Del Castillo, L. Atea, E. Arbelo, E. Caballero, V. Celorrio, T. Datino, E. Gonzalez-Torrecilla, F. Atienza, M. J. Ledesma-Carbayo, J. Bermejo, A. Medina, and F. Fernández-Avilés, "Noninvasive identification of ventricular tachycardia-related conducting channels using contrast-enhanced magnetic resonance imaging in patients with chronic myocardial infarction

comparison of signal intensity scar mapping and endocardial voltage mapping," *Journal of the American College of Cardiology*, vol. 57, pp. 184–94, Jan. 2011.

- J. E. Ortuño, G. Kontaxakis, J. L. Rubio, P. Guerra, and a. Santos, "Efficient methodologies for system matrix modelling in iterative image reconstruction for rotating high-resolution PET.," *Physics in medicine and biology*, vol. 55, no. 7, pp. 1833–1861, 2010.

- A. Iriarte, C. O. S. Sorzano, J. M. Carazo, J. L. Rubio, and R. Marabini, "A theoretical model for EM-ML reconstruction algorithms applied to rotating PET scanners.," *Physics in medicine and biology*, vol. 54, pp. 1909–34, Apr. 2009.

- P. Guerra, J. L. Rubio, J. E. Ortuño, G. Kontaxakis, M. J. Ledesma, and A. Santos, "Performance analysis of a low-cost small animal PET/SPECT scanner," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 571, pp. 98–101, Feb. 2007.

- J. E. Ortuño, P. Guerra-Gutierrez, J. L. Rubio, G. Kontaxakis, and A. Santos, "3D-OSEM iterative image reconstruction for high-resolution PET using precalculated system matrix," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 569, pp. 440–444, Dec. 2006.

- T. Thireou, J. L. Rubio-Guivernau, V. Atlamazoglou, M. J. Ledesma, S. Pavlopoulos, A. Santos, and G. Kontaxakis, "Evaluation of data reduction methods for dynamic PET series based on Monte Carlo techniques and the NCAT phantom," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 569, pp. 389–393, Dec. 2006.

- P. Guerra, J. L. Rubio, G. Kontaxakis, J. E. Ortuño, M. J. Ledesma, and A. Santos, "Programmable electronics for low-cost small animal PET/SPECT imaging," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 569, pp. 162–166, Dec. 2006.

Conference Proceedings

- J. L. Rubio-Guivernau, E. Perez-David, A. Arenal, J. Bermejo, A. Santos, and M. Ledesma-Carbayo, "3D visualization of myocardial substrate using Delayed Enhancement MRI for pre-planning and guidance of ablation procedures of ventricular

tachycardia," *Journal of Cardiovascular Magnetic Resonance*, vol. 13, no. Suppl 1, p. O56, 2011.

- M. A. Luengo-Oroz, T. Savy, J. L. Rubio, L. Duloquin, E. Faure, N. Olivier, M. Ledesma-Carbayo, D. Debarre, P. Bourgine, E. Beaurepaire, N. Peyrieras, and A. Santos, "Processing pipeline for digitalizing the lineage tree of early zebrafish embryogenesis from multiharmonic imaging," in *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pp. 1561–1564, IEEE, Mar. 2011.

- J. L. Rubio-Guivernau, M. A. Luengo-Oroz, L. Duloquin, T. Savy, N. Peyrieras, P. Bourgine, and A. Santos, "Combining sea urchin embryo cell lineages by error-tolerant graph matching.," in *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, vol. 2009, pp. 5918–21, Jan. 2009.

- J. E. Ortuno, J. L. Rubio, P. Guerra, G. Kontaxakis, M. Desco, J. J. Vaquero, and A. Santos, "Efficient methodology for 3D statistical reconstruction of high resolution coplanar PET/CT scanner," in *2008 IEEE Nuclear Science Symposium Conference Record*, pp. 5096–5100, IEEE, Oct. 2008.

- D. Wallach, F. Lamare, J. Rubio, M. J. Ledesma-Carbayo, G. Kontaxakis, A. Santos, P. Marechal, C. Roux, and D. Visvikis, "Super-resolution in 4D positron emission tomography," in *2008 IEEE Nuclear Science Symposium Conference Record*, pp. 4285–4287, IEEE, Oct. 2008.

- J. L. Rubio-Guivernau, M. J. Ledesma-Carbayo, F. Lamare, J. E. Ortuno, P. Guerra, D. Visvikis, A. Santos, and G. Kontaxakis, "Respiratory motion correction in PET with super-resolution techniques and non-rigid registration," in *2007 IEEE Nuclear Science Symposium Conference Record*, pp. 3560–3563, IEEE, 2007.

- J. E. Ortuño, G. Kontaxakis, J. L. Rubio, P. Guerra, and A. Santos, "3D Iterative Reconstruction of High Resolution PET/CT Images using Anatomical Priors and Attenuation Correction," in *9th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, pp. 394–397, 2007.

- N. Sakellios, J. L. Rubio, N. Karakatsanis, G. Kontaxakis, G. Loudos, A. Santos, K. Nikita, and S. Majewski, "GATE simulations for small animal SPECT/PET using voxelized phantoms and rotating-head detectors," in *2006 IEEE Nuclear Science Symposium Conference Record*, pp. 2000–2003, IEEE, 2006.

- J. E. Ortuno, J. L. Rubio, P. Guerra, G. Kontaxakis, and A. Santos, "Multi-grid 3D-OSEM reconstruction technique for high resolution rotating-head PET scanners," in *2006 IEEE Nuclear Science Symposium Conference Record*, pp. 2215–2218, IEEE, 2006.

Patents

- M. J. Ledesma-Carbayo, A. Santos, J. L. Rubio-Guivernau, A. Arenal, E. Perez-David, J. Bermejo, M. Desco, and F. Fernandez-Aviles, "Method for Displaying the Information Contained in Three-Dimensional Images of the Heart," 2010.