Statistical Methods for Analysis and Processing of Medical Ultrasound

Applications to Segmentation and Restoration

Martino Alessandrini

SUPERVISOR
Prof. Guido Masetti

COORDINATOR
Prof. Claudio Fiegna

Contents

Abstract

MEDICAL ULTRASOUND is nowadays commonly employed in the clinical practice for assessing possible abnormalities in several parts of the human body. Due to the simple image formation process, ultrasound systems are reduced size reasonably priced machines, easily affordable even for small low budget ambulatories widespread distributed on the territory. Besides that ultrasound is a non ionizing radiation and practically harmless to the patient. Despite these desirable features, ultrasound scans have a considerably low image quality as compared to that of other popular techniques as X-ray, Magnetic Resonance Imaging and Computed Tomography, and, as a consequence, their interpretation is often subjective and unclear, and their diagnostic reliability low. In this context, the development of software tools providing the physician with assistance in the image reading and analysis process is a continuous challenge for many researchers working in the field.

Due to the noisy nature of ultrasound frames, successful techniques cannot be directly borrowed from the densely populated literature on the processing of natural images, but rather tailored to the ultrasound very peculiar nature. In particular, applications relying on speckle noise statistics as the driving criterion have led to effective solutions for a wide range of problems.

In this thesis two major topics encountered in medical ultrasound are addressed *i.e.* the problems of image restoration and segmentation. For both, original author contributions will be discussed and their performance compared to the state of the art. In the former case, a novel deconvolution technique will be presented, expressively designed for an improved tissue characterization. Phantom studies will show a relevant increase in classification accuracy, along with the superiority of the proposed algorithm over standard ones. The problem of myocardium detection from 2D echocardiography will be then considered and

an original active contour based solution proposed. The result-
ing algorithm will be validated on a set of clinical ultrasound
sequences. Results will show the proposed method on realistic
clinical data to be feasible and accurate.

Sommario

E SAMI ECOGRAFICI sono comunemente prescritti nella pratica clinica per la diagnosi di possibili patologie di diversi organi del corpo umano. La semplicità del processo di formazione dell'immagine fa sì che le apparecchiature ecografiche siano macchinari di dimensioni relativamente piccole e a basso costo, facilmente annoverabili nel parco macchine di piccoli ambulatori a basso budget capillarmente distribuiti nel territorio. A ciò si aggiunga che l'ecografia è innocua per il paziente, essendo gli ultrasuoni onde non ionizzanti. Malgrado questi vantaggi, le immagini ecografiche hanno una qualità considerevolmente inferiore rispetto a quella ottenuta con altre tecniche di imaging, quali raggi X, risonanza magnetica e tomografia computerizzata, il che ne rende la lettura spesso soggettiva e non chiara, e la validità diagnostica assai limitata. In quest'ottica la comunità scientifica ha proposto numerosi software per supportare il medico nella lettura e l'analisi delle immagini ecografiche.

Sfortunatamente, l'elevata rumorosità delle immagini eco preclude l'adozione di algoritmi proposti per l'elaborazione di immagini naturali, sui quali esiste un'ampia letteratura, e impone invece lo sviluppo di tecniche *ad hoc*. In particolare, le applicazioni più efficaci si basano sull'utilizzo della distribuzione statistica del rumore. In questo contesto si inserisce il presente manoscritto, nel quale vengono affrontano due problematiche relative all'elaborazione di immagini ecografiche, nella fattispecie la loro deconvoluzione e segmentazione. Per entrambe verranno descritte soluzioni originali proposte dall'autore, le prestazioni delle quali saranno confrontate con lo stato dell'arte.

Nel primo caso, verrà descritta una nuova tecnica di deconvoluzione in grado di migliorare la caratterizzazione di tessuti tramite ultrasuoni. Da risultati ottenuti su fantoccio verrà messo in luce come l'algoritmo proposto sia in grado di determinare

una riduzione sostanziale dell'errore di classificazione e insieme superare le performance ottenute tramite tecniche standard.

Verrà quindi presentato un nuovo algoritmo per la segmentazione del miocardio da immagini eco bidimensionali, basato sulla tecnica dei contorni attivi. Una valutazione dell'algoritmo su sequenze ecocardiografiche rivelerà le sue potenzialità come strumento utile di supporto in ambito clinico.

Introduction

S INCE the early 1950s, ultrasound use in medicine has been the basis for several procedures that are widespread in to-day's clinical practice. The principal application is in the field of medical imaging. Medical ultrasound imaging relies on the same principles as sonar or radar units: the ultrasound probe produces a (pulsed) acoustic pressure field; the field propagates through the tissue and is partially reflected and scattered due to the inherent inhomogeneity of most tissues. The backscattered signal is received by the same probe and converted into a gray scale image of the organ.

Medical ultrasound has several advantages over other popu-lar imaging modalities as Magnetic Resonance Imaging (MRI), X-ray and Computed Tomography (CT). At first, unlike X-ray and CT, ultrasound is a non ionizing radiation and hence prac-tically harmless to the human body. Moreover, the simple phe-nomena involved in the signal generation and acquisition pro-cess along with the little computation needed for the image cre-ation (fundamentally an amplitude to gray scale conversion), al-low ultrasound systems to work at extremely high frame rates, easily of the order of 100 frames/sec. This makes ultrasound the standard tool for diagnosis of disease based on organs dynam-ics, as it is in echocardiography. Further advantages connected with ultrasound systems are the their cost effectiveness and re-duced size, making their availability possible even in small local low budget ambulatories. This is instead not the case for X-ray, CT and MRI, whose installation, besides relevant costs, requires extended dedicated areas. In Figure 1 a standard ultrasound ima-ging system is illustrated.

Unfortunately, all these advantages come at a price. *i.e.* the reduced image quality as compared to that of X-ray, CT or MRI. This is principally due to the low spatial resolution, directly con-nected with the finite bandwidth of the transducer and the non-

Figure 1: External parts of an ultrasound imaging system. Image from [1, pp. 298].

negligible width of the acoustic beam, along with the characteristic granular texture referred as speckle noise. All these factors make sometimes the interpretation of ultrasound frames highly ambiguous and subjective, even for expert physicians, thereby limiting their diagnostic reliability. In Figure 2 a short axis slice of the left ventricle is represented from 2D echocariography and cardiac MRI: it is evident how image reading is fairly less straightforward in the ultrasound case.

Improving the quality and diagnostic reliability of medical ultrasound has always represented a continuous challenge for many scientists working in the field. The efforts are mainly spent in two directions: on one side improving the acquisition modality itself, by designing new transducers, more sophisticated beam-forming, compounding or apodization schemes; on the other, acting on the acquired image at a post processing step with suitable signal and image processing tools. In particular, in this thesis the second approach is considered.

Unfortunately, the noisy nature of ultrasound images degrades the performance of standard algorithms developed for natural scenes and implies the design of *ad hoc* techniques. They usually rely on considering the speckle not as a noise term to be eliminated, but rather as a precious source of information. It is indeed well established that speckle distribution is strictly correlated with the micro-structure of the underlying tissue. This

Figure 2: Short axis view of the left ventricle from 2D
echocardiography (a) and cardiac MRI (b).

principle is exploited in a wide range of applications as classification, segmentation, deconvolution and tracking, to make some examples. In this context, this manuscript addresses two major problems encountered in medical ultrasound, both tackled with statistically inspired approaches: the problem of image deconvolution and segmentation.

Deconvolution in medical imaging is commonly employed in the purpose of a visual quality improvement, so to provide the physician with better contrasted and resolved data, suitable for easier interpretation. This is done by removing to the maximum extent possible the blurring effect associated with the non-ideal Point Spread Function of the ultrasound system and restoring an estimate of the tissue response, also called reflectivity. This problem has been diffusively addressed in literature, the most successful solutions being based on Wiener filtering and l^1-norm optimization.

Conversely, this manuscript describes a conceptually novel application of deconvolution: from the observation that the reflectivity of a tissue carries cleaner information on its structure than the raw echo signal does, the possibility of exploiting deconvolution not merely as an image enhancement tool, but rather as a pre-processing step for easing ultrasonic tissue characterization is investigated. In this perspective, standard deconvolution algorithms reveal strong limitations, principally ascribed to the simplified tissue models they make use of. Indeed, though these are sufficient for producing appreciable image quality improve-

ment, they otherwise induce a statistical bias in the statistics of the restored reflectivity, making it unusable for characterization purposes.

Hereto, a novel deconvolution method for ultrasound images has been developed and will be described in this thesis. The algorithm is derived on the base of a non-standard statistical model for the tissue response, defined by the Generalized Gaussian Distribution. By means of two distinct parameters, called scale and shape parameter, this distribution allows sequences of arbitrary energy and sparsity to be generated, and is therefore adequate for providing an accurate description to the most general tissue structures. Deconvolution is then tackled as a *maximum a posteriori* estimate, in which tissue reflectivity is restored along with an estimate of the associated scale and shape parameter. An Expectation Maximization framework is designed to address this task. An evaluation of performance will be presented on experimental data from several tissue-mimicking phantoms having a well-defined particle concentration. These studies will show improvements in classification accuracy of up to the 20% and the superiority of the proposed algorithm over standard ones.

The second problem considered is myocardium segmentation from 2D echocardiography. Echocardiography is one of the leading applications of ultrasound in medicine, indeed this is the standard technique to examine myocardial function in patients with known or suspected heart disease. In clinical practice, the analysis mainly relies on visual inspection and manual segmentation by experienced cardiologists. This approach, besides being tedious and time consuming, suffers from a subjective bias, due to the inherent low signal-to-noise ratio (SNR) of ultrasound scans. An automated procedure is therefore desirable both to reduce intra- and inter-observer variability in border detection and to speed up the segmentation process.

While great attention has been given to the segmentation of the endocardium (the innermost layer of tissue surrounding the ventricular cavity), very limited literature addresses the detection of the epicardium (the outer layer). This is due to the fact that signal dropouts and complex interactions between the ultrasonic pulse and the tissue make the epicardial contour appear highly heterogeneous and discontinuous. Nevertheless, a trustful detection of both structures would have a high clinical relevance, as it would make the computation of fundamental parameters possible.

Hereto, an original segmentation algorithm based on level

sets, specifically designed for the detection of the whole myocardium, has been developed and is described in this thesis. The segmentation flow proceeds by seeking the maximum statistical separation between target, *i.e.* the myocardium, and background. In order to deal with low-contrast or missing boundaries, a localized version of standard region-based methods is adopted. Moreover, shape prior information is efficiently embedded in the evolution equation, forcing the active contour to be approximately annular. This prevents the detection of undesired small structures, like papillary muscles. With the resulting formalism the detection of both endo- and epicardium is addressed efficiently with a single level-set function. A validation will be presented from a set of 59 images acquired from 5 patients. These results will show that the proposed method on realistic clinical data is feasible and accurate.

<div align="center">✧   ✧   ✧</div>

The manuscript is organized as follows. Some background material will be presented in Part I. Ultrasound imaging systems will be briefly described in Chapter 1 and their main features examined. The focus will then move to the description of the ultrasound echo signal. In particular modeling the echo acquisition as a linear time variant system will be discussed in Chapter 2, which will be essential in the formulation of the restoration problem. Statistical inspired algorithms typically rely on modeling the ultrasound echo amplitude by means of parametric probability density functions. Hereto, some popular models will be reviewed in Chapter 3. They will be exploited both in image restoration and segmentation.

The problem of image restoration will be addressed in Part II. The problem will be formulated in Chapter 4 where most common solutions, as Wiener filter and l^1-norm optimization, will be presented as well. The original restoration scheme will be finally derived and evaluated in Chapter 5.

Myocardium segmentation will be discussed in Part III. In Chapter 6 all the theoretical elements needed for a full understanding of the algorithm will be given. Then Chapter 7 will illustrate the main aspects of heart morphology and functioning and address the clinical validity of the segmentation phase. The proposed segmentation framework will be finally presented and evaluated in Chapter 8.

I - Background on Medical Ultrasound

Summary

U LTRASOUND, because of its efficacy, low cost, real-time capability and safety, is often the preferred medical imaging modality and finds a number of applications for different parts of the human body. In obstetrics it is commonly employed for 2D or, more recently, 3D *in vivo* imaging of the fetus. Female (usually) breast ultrasound is primarily used for determining the nature of breast abnormality or as a breast cancer screening test supplemental to mammography. Cardiac ultrasound is employed for early diagnosis and of heart disease and prevention of stroke. Vascular and cardiovascular ultrasound is used for monitoring blood flux through veins and arteries. Besides, the possibility of real time visualization is exploited in ultrasound guided biopsy and catheter placement. As an example of ultrasound diffusion, 5 millions exams were estimated to be given weekly worldwide in 2000 [1]. Figure A illustrates the incidence of ultrasound on the overall medical imaging exams.

When compared to other imaging modalities, ultrasound presents several advantages and shortcomings. These aspects will be examined in Chapter 1, where main features of ultrasound systems will be reviewed and put in relation with those of alternative available techniques. The principal limitation associated to ultrasound imaging is a reduced image quality. In order to compensate this lack of image information many computer tools have been and are currently presented to the scientific community. These algorithms commonly rely on models of the acquisition process as well as of the image content. In this context, in Chapter 2 a time-variant linear model of the image formation process will be derived. It will be obtained from deterministically solving the acoustic waves equation for the case of soft tissues. Conversely, in Chapter 3 several parametric statistical models of the echo signal amplitude distribution will be reviewed.
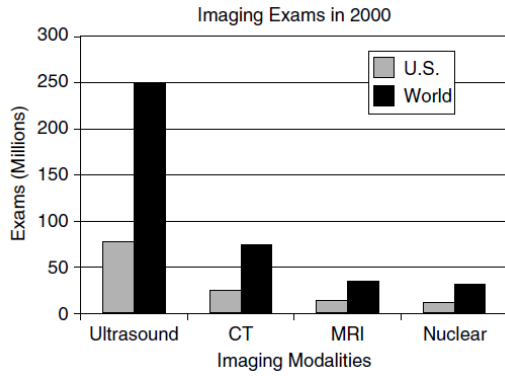
Figure A: Estimated number of imaging exams in 2000. Image from [1, pp. 22].

Ultrasound imaging systems

I N THIS CHAPTER general aspects of medical ultrasound imaging are reviewed. In §1.1 an overview is presented on the main properties of ultrasound imaging systems, with an attention to resolution, penetration and safety. In §1.2 the architecture of an ultrasound imaging system in explained through a simplified block diagram. In §1.3 some concluding remarks are addressed.

1.1 OVERVIEW

The present clinical ultrasound scanners process signals in a manner similar to that of sonar or radar units. To interrogate a tissue, the ultrasound probe produces a (pulsed) acoustic pressure field. The field propagates through the tissue and is partially reflected and scattered due to the inherent inhomogeneity of most tissues. The backscattered signal is received usually by the same probe, supplying useful information about the locations of tissue inhomogeneities and their relative strengths.

Echo amplitude is proportional to the reflection coefficient of those tissue-tissue interfaces encountered along the propagation path. Values of the reflection coefficient for some common interfaces encountered in the medical field are reported in Table 1.1. A high reflection coefficient halts the pulse propagation and makes underneath structures invisible. This fact makes ultrasound unusable in certain situations, e.g. it is not an ideal imaging technique for the bowel or organs obscured by the bowel (CT scanning and MRI are the methods of choice in this setting); for the same reason ultrasound can only see the

INTERFACE	REFLECTION [0,1]
soft tissue - air	0.99
soft tissue - lung	0.52
soft tissue - bone	0.43
vitreous humor - eye lens	0.01
fat - liver	0.79
soft tissue - fat	0.0069
soft tissue - muscle	0.0004
water - lucote	0.13
oil - soft tissue	0.0043

Table 1.1: Reflection coefficient for common interfaces [2].

outer surface of bony structures and not what lies within (MRI is typically preferred here). In cardiac ultrasound, the major application of ultrasound in medicine, the physician is expected obtain clean views of the heart by suitably orientating the transducer through the small *acoustic windows* offered by the rib cage. A big expertise is requested here because of the number of access windows, the differences in anatomy, and the many possible planes of view. Experience is required besides to find relevant planes and targets of diagnostic significance and to optimize instrumentation, as well as to recognize, interpret, and measure images for diagnosis. Such an operation dependency is a first limitation of medical ultrasound.

1.1.1   Resolution

Axial spatial resolution of ultrasound systems is typically chosen equal to 2 wavelength:

$$\lambda_{\text{ax}}(mm) = \frac{2c}{f_c} \tag{1.1}$$

where f_c is the transducer center frequency and c is the sound velocity. Sound speed is highly medium dependent, cf. Table 1.2 and [3]. A mean value commonly agreed for soft tissues is c_{av} = 1540 m/s. For typical frequencies in use ranging from 1 to 15 MHz axial resolution roughly varies from 0.3 to 3 mm. Conversely, lateral spatial resolution is highly depth variant principally due to the beamforming. In particular, this resolution is best at the focal length distance and widens away from this distance in a nonuniform way because of diffraction effects caused

INTERFACE	SOUND VELOCITY (m/s)	DENSITY (Kg/m^3)
Muscles	1580	1070
Liver	1550	1060
Fat	1459	920
Brain	1560	1028
Kidney	1560	1040
Spleen	1570	1059
Blood	1575	1060
Bones	4080	1620
Eye	1670	1135
Lungs	650	430

Table 1.2: Speed of sound and acoustic impedance in some common tissues.
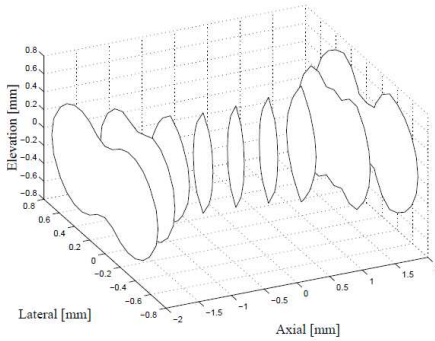
Figure 1.1: Spatial resolution of the acoustic pulse in the lateral and elevation dimensions. Acoustic pressure amplitude contours are -6 dB relative to the peak amplitude within each slice of the point-spread function (PSF) as it propagates. Image from [4, pp. 6]

by apertures on the order of a few to tens of wavelengths. This fact is illustrated in Figure 1.1.

Another factor in determining resolution is attenuation. Attenuation steals energy from the ultrasound field as it propagates and effectively lowers the center frequency of the remaining signals. From (1.1) this implies that axial resolution decreases with the distance from the emitting aperture.

Finally, ultrasound image quality is highly degraded by the presence of granular texture patterns normally referred as speckle noise. Speckle is generated by the constructive and destructive

Modality	Ultrasound	X-ray	CT	MRI
Principle	mechanical properties	mean tissue absorption	tissue absorption	biochemistry
Resolution	frequency and spatially dependent 0.3–3 mm	~1 mm	~1 mm	~1 mm
Penetration	frequency dependent 3–25 cm	Excellent	Excellent	Excellent
Safety	very good	ionizing radiation	ionizing radiation	very good
Speed	100 frames/s	minutes	½ to 1 min	10 frames/s
Cost	$	$	$$$$	$$$$$$$$
Portability	excellent	good	poor	poor

Table 1.3: Comparison of imaging modalities [1, pp. 23].

interference of waves diffused by collections of sub-resolvable scatterers. Where the term diffusive refers to scattering centers of size much smaller than the signal wavelength and sub-resolvable means contained inside one resolution cell of the system. Several speckle models will be presented in Chapter 3.

1.1.2   Penetration

Penetration is defined as the maximum distance reached by the ultrasonic beam inside the tissue. This is fundamentally determined by tissue attenuation, indeed attenuation increases with higher center frequencies; therefore, penetration decreases correspondingly so that fine resolution is difficult to achieve at deeper depths. Depending on the frequency penetration can vary between 3 and 25 cm. Reduced penetration capability makes large patients more difficult to image by ultrasound. In certain cases, this limitation can be offset by specialized probes such as transesophageal (down the throat) and intracardiac (inside the heart) that provide access to regions inside the body.

1.1.3   Frame rate

In a typical brightness mode (B-mode) acquisition modality, the image is build as a collection of adjacent scan lines acquired in
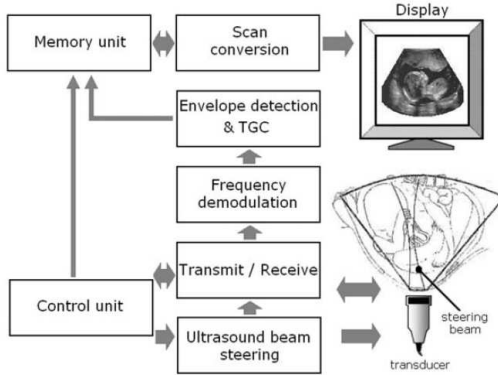
Figure 1.2: Block diagram of an ultrasound imaging system.

sequence. Assuming a number N of scan lines, a maximum depth D and a sound speed c, the frame rate FR is:

$$RF = \frac{c}{2DN} \qquad (1.2)$$

For an image 15 cm deep and 50 scan lines the frame rate is approximately 100 frames/sec, *i.e.* a real time visualization of the organ is allowed.

1.1.4 Safety

As acoustic waves are employed, diagnostic ultrasound does not have any cumulative side effects. For very particular ultrasound applications as hyperthermia, lithotripsy or HIFU, high power regime is involved which is proved to induce bioeffects. Nevertheless these cases will not be discussed here. The interested reader is otherwise addressed to [1, chap. 15]. Besides, diagnostic ultrasound is also generally non invasive, excluding of course "trans" and "intra" families of transducers.

All the features reviewed in this section are compared to the one of other popular medical imaging modalities in Table 1.3.

1.2 Ultrasound system architecture

In this section B-mode imaging will be considered only. Other modes exists, as A-mode and M-mode, sometimes employed in the diagnostic practice, for which the reader is addressed to [1,
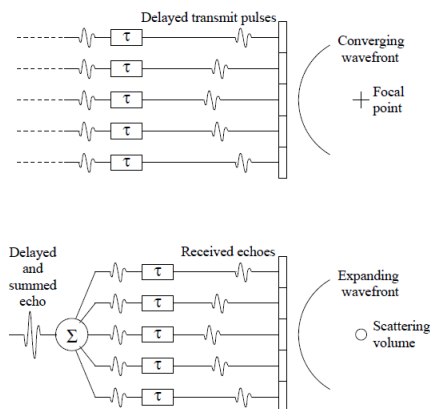
Figure 1.3: Conceptual diagram of phased array beamforming.

chap. 10]. A simplified block diagram of a B-mode ultrasound imaging system is illustrated in Figure 1.2.

The controller sets the excitation sequence of the transducer elements for the beam steering. This process is illustrated in Figure 1.3 where the most common case of phased array beamforming is considered. By planar sweeping the ultrasound beam with a certain increment within a predefined plane, a set of RF-lines are collected and, subsequently, used for visualization of the tissue. The RF signal demodulated is called IQ signal. The envelope is then computed as the absolute value of the IQ signal and logarithmic compression is applied. The compressed signal is often referred as brightness mode (B-mode) signal.

The set of scan lines are normally collected into a matrix. As each scan line is the result of an angular rotation of the ultrasound beam, pixel coordinates in the matrix are related to the true position in the body by a polar to Cartesian scan conversion, allowing the typical B-mode sector image to be displayed on screen. A gray scale color map is commonly adopted: stronger echoes have an higher brightness as weak echoes appear as darker.

As mentioned, tissue absorption induces a depth dependent attenuation, which is compensated at a post processing step through the Time Gain Compensation (TCG) block, to be suitably tuned by the user by moving slides placed on the control panel of the ultrasound system.

1.3 CONCLUSION

Principal features of ultrasound imaging systems have been reviewed and compared to those of other popular imaging modalities. Major advantages of ultrasound are portability, cost effectiveness, safety and real-time capability. Main limitations are low depth-dependent resolution, limited penetration, high user dependency and low image quality due to speckle.

Deterministic description of ultrasound

T HE SOLUTION of the acoustic wave equations for pulse-echo ultrasound leads to complex integral expressions requiring cumbersome computation. Fortunately, as long as weak scattering is concerned, several simplifications can be reasonably made in the calculation of the received pressure field, letting the whole acquisition process to be accounted for by a simple linear shift-variant model. As the linear model is consistent with the majority of diagnostic situations in which soft tissues are imaged, it is otherwise largely violated in some others as in contrast agents imaging or harmonic imaging [1]. Nevertheless these techniques will not be considered in this work.

In this chapter the linear model for the echo image formation will be derived. Such a representation will be largely exploited when image restoration will be discussed. The derivation presented in this chapter closely follows the one proposed in [5], which is at the base of the popular FieldII simulation software [6]. The chapter proceeds as follows. In §2.1 the wave equation for the acoustic pressure field is derived. In §2.2 the expression for the scattered field is presented and the Born expansion is introduced. In §2.3 the incident field is calculated. In §2.4 the linear representation of the received echo signal is presented, which represents the fundamental core of the chapter. The terms reflectivity function and Point Spread Function (PSF) will be there formally defined. In §2.5 the principal properties of the PSF of a medical ultrasound device will be illustrated. In §2.6 the main conclusions will be drawn.

2.1   WAVE EQUATION

Lat's assume a propagating medium state is perturbed by an acoustic pressure field and let's call $P_{ins}(\mathbf{r}, t)$ and $\rho_{ins}(\mathbf{r}, t)$ the instantaneous pressure and density measured in $\mathbf{r}$ at time t. We assume the following to hold:

$$P_{ins}(\mathbf{r}, t) = P + p_1(\mathbf{r}, t)$$
$$\rho(\mathbf{r}, t)_{ins} = \rho(\mathbf{r}) + \rho_1(\mathbf{r}, t) \tag{2.1}$$

where p_1 and ρ_1 are small first order variations due to the propagation of the ultrasound wave. If the transformation is adiabatic (invariant entropy), then the wave equation can be obtained from coupling the dynamic equation with the equation of continuity:

$$\nabla^2 p_1 - \frac{1}{c^2} \frac{\partial^2 p_1}{\partial t^2} = \frac{1}{\rho} \nabla \rho \cdot \nabla p_1 \tag{2.2}$$

where c is the sound speed of the considered medium. If the perturbation is small, then one can expect c and ρ to vary little from their average value c_0 and ρ_0, so that:

$$c(\mathbf{r}) = c_0 + \Delta c(\mathbf{r})$$
$$\rho(\mathbf{r}) = \rho_0 + \Delta\rho(\mathbf{r}). \tag{2.3}$$

By substituting (2.3) into (2.2) and neglecting second order effects, the final equation for the pressure can be derived:

$$\nabla^2 p_1 - \frac{1}{c_0^2} \frac{\partial^2 p_1}{\partial t^2} = -\frac{2\Delta c}{c_0^3} \frac{\partial^2 p_1}{\partial t^2} + \frac{1}{\rho_0} \nabla(\Delta\rho) \cdot \nabla p_1 \tag{2.4}$$

where the right end side represents the source of the scattered field.

2.2   SCATTERED FIELD

The scattered field generated from a distributed region of volume V, measured at $\mathbf{r}_2$, is obtained by integrating all the spherical waves originating from V:

$$p_s(\mathbf{r}_2, t) = \int_V \int_T \left[-\frac{2\Delta c(\mathbf{r}_1)}{c_0^3} \frac{\partial^2 p_1(\mathbf{r}_1, t_1)}{\partial t_1^2} + \right.$$
$$\left. + \frac{1}{\rho_0} \nabla(\Delta\rho(\mathbf{r}_1)) \cdot \nabla p_1(\mathbf{r}_1, t_1) \right] G(\mathbf{r}_1, t_1 | \mathbf{r}_2, t_2) dt_1 d\mathbf{r}_1^3 \tag{2.5}$$

where

$$G(\mathbf{r}_1, t_1 | \mathbf{r}_2, t_2) = \frac{\delta(t - t_1 - |\mathbf{r}_2 - \mathbf{r}_1|/c_0)}{4\pi|\mathbf{r}_2 - \mathbf{r}_1|} \tag{2.6}$$

is the free space Green function.

In general, the overall acoustic field p_1 in the scattering region is given by the sum of the incident and the scattered pressure fields, denoted by p_i and p_s respectively:

$$p_1(\mathbf{r}, t) = p_i(\mathbf{r}, t) + p_s(\mathbf{r}, t) \tag{2.7}$$

It is evident that with this general expression (2.5) cannot be solved for p_s. Hereto it is useful here to apply the Born-Neumann expansion:

$$\begin{aligned}
p_s(\mathbf{r}_2, t) = & \left[G_i F_{op}\right] p_i(\mathbf{r}_1, t_1) + \\
& \left[G_i F_{op}\right]^2 p_i(\mathbf{r}_1, t_1) + \\
& \left[G_i F_{op}\right]^3 p_i(\mathbf{r}_1, t_1) + \cdots \\
& \left[G_i F_{op}\right]^N p_i(\mathbf{r}_1, t_1)
\end{aligned} \tag{2.8}$$

where G_i is the Green operator over $\mathbf{r}_1$ and t_1 while:

$$F_{op} = -\frac{2\Delta c(\mathbf{r}_1)}{c_0^3} \frac{\partial^2 p_1(\mathbf{r}_1, t_1)}{\partial t_1^2} + \frac{1}{\rho_0} \nabla(\Delta\rho(\mathbf{r}_1)) \cdot \nabla p_1(\mathbf{r}_1, t_1) \tag{2.9}$$

in such a way (2.5) can be rewritten as $p_s = G_i F_{op} p_1$. Now, the terms of (2.8) involving $[\bullet]^N$, where $N > 1$, refer to multiple scattering of order N. In the weak scattering approximation, multiple scattering events can be neglected and the first order Born approximation holds, so that p_1 in (2.5) can be in practice replaced by p_i:

$$\begin{aligned}
p_s(\mathbf{r}_2, t) = & \int_V \int_T \left[-\frac{2\Delta c(\mathbf{r}_1)}{c_0^3} \frac{\partial^2 p_i(\mathbf{r}_1, t_1)}{\partial t_1^2} + \right. \\
& \left. + \frac{1}{\rho_0} \nabla(\Delta\rho(\mathbf{r}_1)) \cdot \nabla p_i(\mathbf{r}_1, t_1) \right] G(\mathbf{r}_1, t_1 | \mathbf{r}_2, t_2) dt_1 d\mathbf{r}_1^3
\end{aligned} \tag{2.10}$$

2.3   INCIDENT FIELD

The incident field is generated from the ultrasonic transducer. The coordinate system to be used for this calculation is represented in Figure 2.1, where $\mathbf{r}_3$ is the transducer position, $\mathbf{r}_4$ denotes a point on the transducer surface in the local coordinate
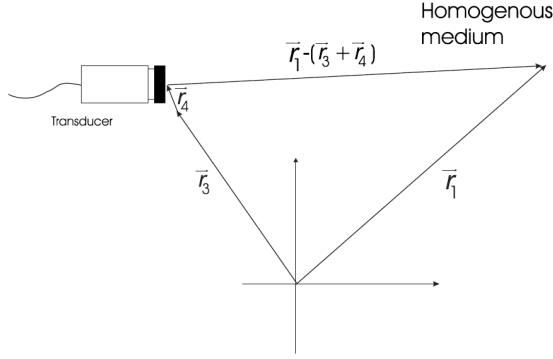
Figure 2.1: Coordinate system for calculation of the incident field.
Image taken from [5, pp. 23].

system centered in $\mathbf{r}_3$ and $\mathbf{r}_1$ denotes a point in the scattering volume.

The incident field can be written as:

$$p_i(\mathbf{r}, t) = \rho_0 \frac{\partial \Psi(\mathbf{r}, t)}{\partial t} \qquad (2.11)$$

where Ψ is the velocity potential, satisfying the homogeneous wave equation:

$$\nabla^2 \Psi - \frac{1}{c_0^2}\frac{\partial^2 \Psi}{\partial t^2} = 0. \qquad (2.12)$$

By calling v the particle velocity normal to the transducer surface, and assuming an uniform velocity distribution on the surface itself, then it is:

$$\Psi(\mathbf{r}_1, \mathbf{r}_3, t) = \int_T v(t_3) \int_S g(\mathbf{r}_1, t | \mathbf{r}_1 + \mathbf{r}_4, t_3) d\mathbf{r}_4^2 dt_3 \qquad (2.13)$$

where g is the bounded Green function

$$g(\mathbf{r}_1, t | \mathbf{r}_1 + \mathbf{r}_4, t_3) = \frac{\delta(t - t_3 - |\mathbf{r}_1 - \mathbf{r}_3 - \mathbf{r}_4|/c_0)}{2\pi |\mathbf{r}_1 - \mathbf{r}_3 - \mathbf{r}_4|} \qquad (2.14)$$

By substituting (2.13) into (2.11), the following expression for the incident pressure field can be derived:

$$p_i(\mathbf{r}_1, \mathbf{r}_3, t) = \rho_0 \frac{dv}{dt} \underset{t}{*} h(\mathbf{r}_1, \mathbf{r}_3, t) \qquad (2.15)$$

Figure 2.2: Coordinate system for calculation of the received field. Image taken from [5, pp. 24].

where h is the spatial impulse response, depending exclusively on the geometry of the radiating aperture:

$$h(\mathbf{r}_1, \mathbf{r}_3, t) = \int_S \frac{\delta(t - |\mathbf{r}_1 - \mathbf{r}_3 - \mathbf{r}_4|/c_0)}{2\pi|\mathbf{r}_1 - \mathbf{r}_3 - \mathbf{r}_4|} d\mathbf{r}_4^2 \qquad (2.16)$$

2.4 RF ECHO SIGNAL

The received echo signal is the scattered pressure field integrated over the transducer surface, convolved with the transducer electromechanical impulse response:

$$p_r(\mathbf{r}_5, t) = E_m(t) \underset{t}{*} \int_S p_s(\mathbf{r}_6 + \mathbf{r}_5, t) d\mathbf{r}_6^2 \qquad (2.17)$$

where the coordinate system used for this calculation is illustrated in Figure 2.2. The final expression for p_r is then obtained by substituting p_s in (2.10) and p_i expressed in (2.15). As the interest is here only to present the equations for the linear model of the echo generation process, the final expression is directly presented, addressing the interested reader to [5, pp. 24–26] for a complete description of all steps.

The final expression for the received pressure field is:

$$p_r(\mathbf{r}_5, t) = v_{pe}(t) \underset{t}{*} f_m(\mathbf{r}_1) \underset{\mathbf{r}}{*} h_{pe}(\mathbf{r}_1, \mathbf{r}_5, t) \qquad (2.18)$$

Note that the output voltage trace produced by the transducer is exactly equal (2.18), a part of a multiplication by a constant factor determined by the piezoelectric crystals gain. The term f_m

is commonly referred as *reflectivity function* and represents the tissue signature contained in the echo signal. In particular it accounts for the inhomogeneities in the tissue due to density and propagation velocity perturbations which give rise to the scattered field. The term h_{pe} is the modified pulse-echo spatial impulse response and is exclusively dependent on the geometry of the problem: it relates the transducer geometry to the spatial extent of the scattered field. Finally, the term v_{pe} represents electro mechanical stimulation modality of the transducer elements.

Explicitly written out these terms are:

$$v_{pe}(t) = \frac{\rho_0}{2c_0^2} E_m(t) \underset{t}{*} \frac{d^3 v(t)}{dt^3} \tag{2.19}$$

$$f_m(\mathbf{r}_1) = \frac{\Delta\rho(\mathbf{r}_1)}{\rho_0} - \frac{2\Delta c(\mathbf{r}_1)}{c_0} \tag{2.20}$$

$$h_{pe}(\mathbf{r}_1, \mathbf{r}_5, t) = h(\mathbf{r}_1, \mathbf{r}_5, t) \underset{t}{*} h(\mathbf{r}_5, \mathbf{r}_1, t) \tag{2.21}$$

It is common to rewrite (2.18) by separating tissue and transducer dependent effects in two separate terms as:

$$p_r(\mathbf{r}_5, t) = H(\mathbf{r}_1, \mathbf{r}_5, t) \underset{\mathbf{r}}{*} f_m(\mathbf{r}_1) \tag{2.22}$$

$$H = v_{pe}(t) \underset{t}{*} h_{pe}(\mathbf{r}_1, \mathbf{r}_5, t) \tag{2.23}$$

where H is the transducer Point Spread Function Equation (PSF) and represents the echo signal acquired from a single ideal point scatterer. Equation (2.22) shows that, in the weak scattering assumption, which is typically satisfied in the case of soft tissues, the echo image formation process can be approximated as a linear filtering operation. This approximation will be largely exploited in the rest of the manuscript, in particular when the deconvolution problem will be discussed.

2.4.1   IQ echo signal

Equation (2.22) represents the linear model for the RF echo signal. For computational reasons it is sometimes useful to work on the baseband equivalent rather then on the RF directly. Fortunately, it can be shown that, due to the linearity of the demodulation process, an expression completely analogous to (2.22) exists for the complex envelope as well [7].

Given a real valued signal x with center frequency $2\pi\omega_0$, the *in phase/quadrature* signal x_{IQ} is computed as:

$$x_{IQ} = \left[x - j\mathcal{H}(x) \right] e^{-j\omega_0 t} \tag{2.24}$$

Figure 2.3: PSF shape at different depths corresponding do different focusing and apodization schemes. Images obtained with the Field II software.

where $\mathcal{H}(\cdot)$ is the Hilbert transform and $x_{AN} = [x - j\mathcal{H}(x)]$ is the analytic signal [8]. x_{IQ} is sometimes referred as complex envelope and its magnitude is the envelope of x. With this notation it can be shown that [7]:

$$p_{r,IQ} = H_{IQ} \underset{\mathbf{r}}{*} \tilde{f}_m \tag{2.25}$$

where $\tilde{f}_m = f_m e^{-2jk_0 z}$, z is the axis of axial propagation and $k_0 = \omega_0/c_0$ is the wavenumber.

2.5   POINT SPREAD FUNCTION

It is important to examine here the main factors which determine the shape of the ultrasonic PSF (H), since they must be taken into account in the derivation of an accurate deconvolution framework. In this sense the peculiar property of an ultrasound PSF is its spatial variability. We will see later in the manuscript how dealing with spatially variant blurring kernels involves relevant computational issues in the implementation of a restoration algorithm, while in this section our interest is to briefly examine the phenomena at the origin of this variability. These factors are both system and medium dependent.

Figure 2.4: Changes in pressure-pulse shape of an initially Gaussian pulse propagating in a medium with a 1dB/MHz1.5-cm absorption coefficient for three different increasing propagation distances (z). Image taken from [1, pp. 76]

The main system dependent factor is beamforming. As mentioned, beamforming consists in suitably delaying the piezoelectric crystals excitation in order to have an acoustic field of minimum width in correspondence with the region of interest. The most common beamforming technique is represented by static single focus beamforming, in which the acoustic beam is the narrowest at the focus and widens progressively before and after that spot, cf. Figure 2.3(a). In this context, more complicated dynamic focusing schemes or apodization techniques have be designed in order to have the most uniform PSFs on the imaged plane, but they are not the norm on commercial scanners, and neither completely prevent spatial variations.

Tissue dependent effects are instead connected with tissue attenuation. Indeed real data indicate that absorption and dispersion of the ultrasonic beam due to the presence of the propagating medium have a power law dependency. As a result, acoustic pulses not only become smaller in amplitude as they propagate, but they also change shape, cf. Figure 2.4. Tissue dependent effects can be embedded in the PSF expression by introducing a modified spatial impulse response of the kind:

$$h_{att}(t, \mathbf{r}) = \int_T \int_S a(t - \tau, |\mathbf{r} + \mathbf{r}_1|) \frac{\delta(t - |\mathbf{r} + \mathbf{r}_1|/c)}{|\mathbf{r} + \mathbf{r}_1|} dS d\tau \quad (2.26)$$

where a denotes the attenuation term, and is often defined through its Fourier Transform $A(f, |\mathbf{r}|)$. A simplified model for A is given by:

$$A(f, |\mathbf{r}|) = \exp(-\alpha|\mathbf{r}|) \exp(-\beta(f - f_0)|\mathbf{r}|) \qquad (2.27)$$

where f_0 is the transducer center frequency. For more complicated models see [5, pp. 26–29] and references therein.

One have to note that these tissue dependent effects cannot be a priori known. As a consequence, every estimate of the PSF deriving from modeling or measurements will be necessarily inaccurate and, hence, the PSF is an unknown term in the acquisition equation. We will explore in the next part of this manuscript how this issue further complicates the problem of restoration of medical ultrasound images.

2.6 Conclusion

The acoustic wave equations have been solved for the calculation of the echo signal backscattered to the transducer from an inhomogeneous region. Under weak scattering assumption the acquisition process is modeled as a linear shift-variant system. Reflectivity function and Point Spread Function have been here introduced and formally defined. These terms will be extensively used throughout the rest of the manuscript. The PSF spatial variability has been discussed and the mechanism at its origin illustrated.

While the echo signal has been here derived by formally solving the acoustic wave equations, the next chapter is instead dedicated to the description of the echo generation as a statistical process.

Statistical description of ultrasound

ULTRASOUND images exhibit characteristic speckle patterns. Speckle is a granular texture generated by the constructive and destructive interference of waves diffused by collections of sub-resolvable scatterers, where the term diffusive refers to scattering centers much smaller than the signal wavelength and sub-resolvable means contained inside one resolution cell of the ultrasound equipment.

Despite its noise like resemblance, speckle patterns are not random but deterministic, as they can be exactly reproduced if the transducer is returned in the same position. This feature of speckle is used to track movements and displacements of imaged organs, which could have a fundamental clinical relevance in certain circumstances, like in the study of heart dynamics from echocardiography. Besides, speckle has been shown to be strictly correlated with tissue structure. This is exploited for discriminating among different tissues on US images (segmentation) and is otherwise at the base of Ultrasonic Tissue Characterization (UTC).

As seen, if speckle may be considered a noise source in some particular cases, for which *ad hoc* despeckling techniques have been designed, it is otherwise common to consider the speckle itself a fundamental source of information, to be not only preserved, but exploited for driving certain applications. This is indeed the case for image segmentation, deconvolution and classification, to be discussed next. In this context, the common paradigm is to describe speckle statistical distribution through suitable parametric probability density functions.

A number of different distributions have been proposed such
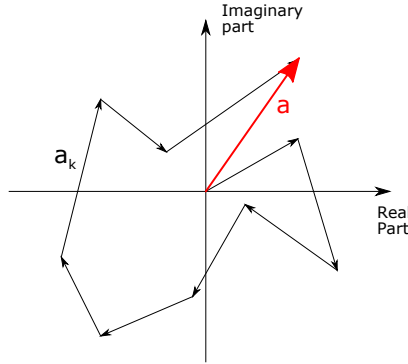
Figure 3.1: Random walk. In red the final phasors sum.

as Rician [9, pp.50–52], K [10], Homodyne-K [11], and Nakagami [12]. When the RF signal is considered instead, Krf distribution [13] and Generalized Gaussian distribution (GGD) [14] are the most appropriate. An exhaustive description of all these distributions is beyond the scope of this work, rather, in this chapter the emphasis will be put on the only statistical models exploited in the applications to be discussed next: these are Rayleigh and Nakagami distribution for the envelope signal (§3.1 and §3.2) and Generalized Gaussian for the RF (§3.3). For all these distributions a parameters estimate technique will be presented. Insights on the goodness of fit of these parametric distributions will be given in §3.4.

3.1   RAYLEIGH DISTRIBUTION

The Rayleigh distribution is derived here for the first order statistic of the ultrasound envelope signal. The envelope is computed as the magnitude of the analytic signal of the RF echo, as discussed in §2.4.1. Here we refer the RF echo as $\mathbf{y}$, the analytic signal as $\mathbf{a}$ and the real envelope of $\mathbf{y}$ as $\rho = |\mathbf{a}|$.

In this section we will closely follow the derivation of the Rayleigh distribution presented in [9, chap. 2]. The analytic echo signal $\mathbf{a}$ returned from a collection of N unresolvable diffusive scatterers, can be expressed by the random phasors sum:

$$\mathbf{a} = \sum_{k=1}^{N} \mathbf{a}_k; \quad \text{with } \mathbf{a} = \rho \cdot e^{j\theta}, \quad \mathbf{a}_k = \frac{1}{\sqrt{N}} \cdot \rho_k \cdot e^{j\theta_k} \qquad (3.1)$$

where the effect of each scatter is taken into account as a deviation of the incident wave along a random direction. This process is represented in Figure 3.1. The following assumptions are made:

- amplitude ρ_k and phase θ_k of each elementary phasor $\mathbf{a}_k$ are statistically independent;

- each phase θ_k is uniformly distributed in the interval $[-\pi, \pi]$;

- all amplitudes ρ_k are identically distributed, with mean value $< \rho >$ and second moment $< \rho^2 >$.

If the number of scatterers N per resolution cell is high, then the Central Limit Theorem holds and the joint probability density function (pdf) of $r \triangleq \Re\{\mathbf{a}\}$ and $i \triangleq \Im\{\mathbf{a}\}$ becomes:

$$p_{RI}(r, i) = \frac{1}{2\pi\sigma^2} \exp\left\{-\frac{r^2 + i^2}{2\sigma^2}\right\}, \quad \sigma^2 = \frac{< \rho^2 >}{2} \qquad (3.2)$$

The distribution of ρ and θ in (3.1) can be derived from (3.2) after the variable substitution:

$$r = \rho \cos\theta$$
$$i = \rho \sin\theta \qquad\qquad (3.3)$$

so that

$$p_\rho(\rho) = \frac{\rho}{\sigma^2} \exp\left\{-\frac{\rho^2}{2\sigma^2}\right\}, \qquad (3.4)$$

while θ is uniformly distributed on $[-\pi, \pi]$. The modulus $\rho = \sqrt{r^2 + i^2}$ represents the envelope signal and (3.4) defines the Rayleigh distribution.

The principal assumptions employed in the derivation of the Rayleigh distribution are that scatterers within the resolution cell must have random location and must be in a big number. The first hypothesis precludes the Rayleigh pdf from modeling tissues exhibiting regularity patterns, as for instance skin, while the second makes the modeling of low scatterers density impossible. More specifically, the condition of high N values is referred as *fully developed speckle* condition. Although fully developed speckle well represents the response of certain tissues as blood [15], in the most general case scatter distribution may deviate substantially from the fully developed model making the

Rayleigh distribution highly inaccurate [10]. This can be intuitively justified from the fact that the Rayleigh pdf posses a single parameters, directly related to the signal energy. As a consequence, if scatterers strength can be suitably modeled, that is not the case for other features which may characterize the scatterer distribution, as their concentrations, or the pattern of their location. For these situations more complicated multi-parametric distributions have been proposed, as we will discuss later in this chapter.

3.1.1 Rayleigh parameter estimate

As these parametric distributions are used for fitting the echo signal histogram, their parameters have to be set so to obtain the best fit. In this case it is common to adopt Maximum Likelihood (ML) parameters estimators which can be shown to be asymptotically unbiased and minimum variance [16].

Given $\{x_i\}_{i=1}^{N}$ a set independent identically distributed (i.i.d.) samples obeying a Rayleigh distribution, the ML estimate $\hat{\sigma^2}$ of the Rayleigh parameter σ^2 is:

$$\hat{\sigma} = \frac{1}{2N} \sum_{i=1}^{N} x_i^2. \tag{3.5}$$

A explicit derivation can be found in [17].

3.2 NAKAGAMI DISTRIBUTION

As mentioned before and largely documented in literature [18, 9, 19], Rayleigh model is not adequate for modeling the most general tissue scattering conditions. In this context a wide range of distributions have been proposed, among which Nakagami is the most commonly adopted. Indeed, as stated in [12], this model can describe the statistics of the envelope of the backscattered echo from an ensemble of scatterers with varying number densities, varying cross sections, and the presence or absence of regularly spaced scatterers.

Nakagami pdf writes as:

$$p(\rho) = \frac{2m^m \rho^{2m-1}}{\Gamma(m)\Omega^m} \exp\left(-\frac{m}{\Omega}\rho^2\right) \tag{3.6}$$

where m is the Nakagami parameter and Ω is the scale parameter. See Figure 3.2 for an illustration of the Nakagami pdf relative to different values of m. Note that for $m = 1$ one obtains
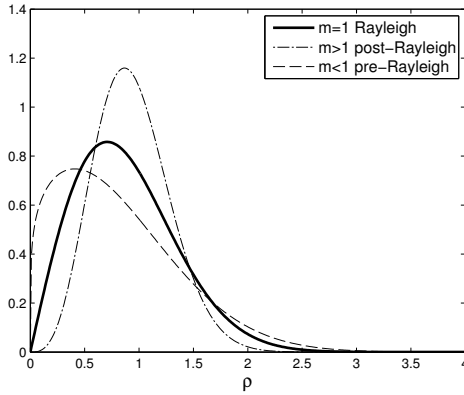
Figure 3.2: Nalagami pdf for different values of m

the Rayleigh pdf. Values of m smaller or bigger than one correspond to scattering conditions referred as pre-Rayleigh and post-Rayleigh [12]. While Ω is connected with signal energy, m accounts for other tissue properties. Due to this flexibility, Nakagami distribution has been widely exploited for ultrasonic tissue characterization [18].

3.2.1 Parameter estimation

Maximum likelihood parameter estimation does not have a closed for solution. As a consequence simpler moment matching based estimates are often preferred:

$$\hat{m} = \frac{[\mathrm{E}(\rho^2)]^2}{\mathrm{E}[\rho^2 - \mathrm{E}(\rho^2)]^2}$$

$$\hat{\Omega} = \mathrm{E}(\rho^2)$$

(3.7)

where the statistical expectation $\mathrm{E}\{\bullet\}$ is clearly implemented as samples average when a realization of (3.6) is available.

3.3 GENERALIZED GAUSSIAN DISTRIBUTION

As the RF is not an easily accessed output on commercial ultrasound equipments, limited attention has been paid on modeling its distribution. In this sense the simplest model is given by the Gaussian distribution. In particular, Rayleigh distribution for the envelope directly derives from the Gaussian model for the

Figure 3.3: GGD pdf for different values of ξ

RF, as expressed in (3.2). As a consequence Gaussian shares the same limitations associated with the Rayleigh model.

Most sophisticated models in this sense are given by the Krf [13] and Generalized Gaussian Distribution (GGD) [14]. In particular the latter has been shown to be superior in terms of accuracy and feasibility [14]. The GGD pdf writes as:

$$p_Y(y) = a \exp\left(- \left| \frac{y}{b} \right|^{\xi} \right) \tag{3.8}$$

where $a = \xi/(2b\Gamma(1/\xi))$, $b = \sigma\sqrt{\Gamma(1/\xi)/\Gamma(3/\xi)}$ is the scale parameter, σ is the standard deviation, ξ is the shape parameter and $\Gamma(\cdot)$ is the Gamma function [20]. As σ explains signal energy, ξ is directly connected to signal sparsity and scatterer concentration. In this sense it plays the same role as m for the Nakagami pdf. Examples of a GGD distribution for different values of ξ are given in Figure 3.3. Note that Gaussian and Laplacian distribution are both special cases of the GGD.

In the following of this manuscript we will show how GGD, besides being a valuable model for the RF signal, it can realistically describe tissue reflectivity as well.

3.3.1 GGD parameter estimate

No closed form solution exists for GGD shape parameter estimate, neither based on ML nor moments matching. The ML estimator is presented here since, as will see later, it will allow

Blood Region   Muscle Region

(a)   (b)

Figure 3.4: Selected regions of interest: (a) blood pool internal to the ventricle, (b) portion of the cardiac muscle.

us to design a deconvolution framework consistent with the Expectation Maximization framework.

The ML $\hat{\xi}$ estimate of ξ is obtained from solving:

$$\frac{N}{\xi} + \frac{N\Psi(1/\xi)}{\xi^2} - \sum_{k=1}^{N}\left(\frac{|x_k|}{b}\right)^{\xi}\log\left(\frac{|x_k|}{b}\right) = 0 \qquad (3.9)$$

and then

$$\hat{b} = \left(\frac{\xi}{N}\sum_{k=1}^{N}|x_k|^{\xi}\right)^{1/\xi}. \qquad (3.10)$$

The Newton method in [21] can be used to solve (3.9).

3.4   COMPARISON

We present here an example of the fitting capability of the distributions described in this chapter for the case of cardiac ultrasound data. In particular, data from two regions are considered, corresponding to the blood pool inside the left ventricle and the cardiac muscle, see Figure 3.4. These two kind of tissues exhibit different scattering properties, in particular blood regions are close to the fully developed model, while muscle areas are pre-Rayleigh. From this information one can anticipate that Gaussian and Rayleigh distributions will provide accurate fits for the
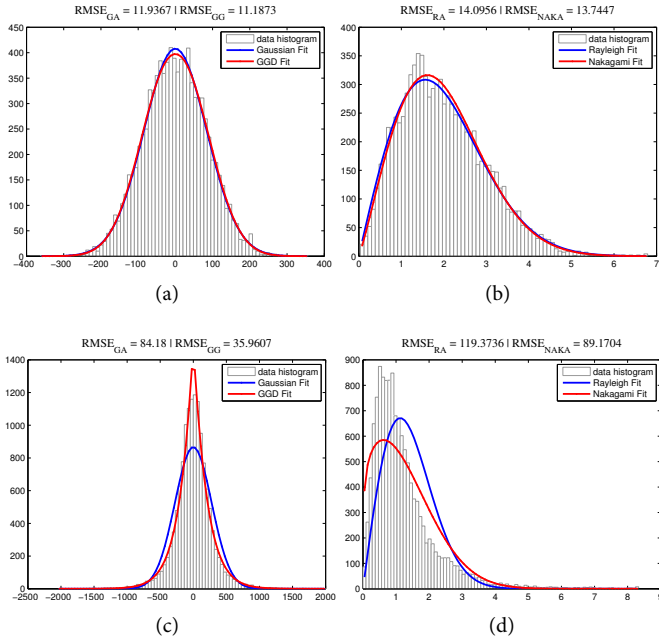
Figure 3.5: Goodness of fit of different parametric models when applied to a blood region ((a)–(b)) or a muscle region ((c)–(d)). Images (a) and (c) are relative to the RF signal, (b) and (d) to the envelope.

former region, while GGD and Nakagami will be more precise for the latter. This guess is indeed confirmed experimentally, cf. Figure 3.5.

In Figure 3.5 the goodness of fit has been measured by the root mean square error (RMSE) between the data histogram and the parametric model. Parameters have been estimated as described in this chapter.

3.5 CONCLUSION

In this chapter some popular parametric distribution for modeling the statistics of ultrasound signals have been reviewed. Techniques for parameters estimation have been presented for each distribution. We will show how these models can be exploited in the design of ultrasound specific image processing applications. In particular Rayleigh speckle model will be used for driving the

evolution of an active contour addressing the recognition of my-ocardial boundaries from echocardiography; Generalized Gaussian distribution will be used as prior distribution for modeling the tissue reflectivity in the derivation of a Maximum a Posteriori deconvolution framework; Nakagami parameters will be adopted, along with GGD ones, for characterizing different scatterer concentrations from tissue mimicking phantoms.

II - Restoration of Ultrasound Images

Summary

I T HAS BEEN shown that, at least in the weak scattering approximation, the ultrasound image can be expressed as the linear convolution of the tissue response with the system PSF. The non-negligible spatially-variant extent of the PSF is one of the main causes to the limited quality of ultrasound frames, which makes sometimes opt for more onerous potentially harmful but more reliable imaging modalities. In this context, deconvolution is commonly proposed in literature as a post processing tool for improving resolution and contrast of ultrasound images by removing to the maximum extent possible the blurring effect associated with the PSF. The problem is often tackled as an l^2-norm or l^1-norm constrained optimization task.

In this part of the thesis, an alternative application of deconvolution is considered instead, i.e. its use as a pre-processing step for potentiating ultrasonic tissue characterization. This procedure relies on the quantitative analysis of the echo signal to infer information about the tissue structure. The feasibility of these approaches for discriminating healthy tissues against cancerous ones from breast and prostate ultrasound has been largely documented. From the observation that the reflectivity of a tissue carries cleaner information on its structure than the raw echo signal does, the possibility of exploiting the deconvolved image rather then the unprocessed one for tissue characterization is investigated. From this perspective, standard deconvolution algorithms reveal strong limitations, principally ascribed to the simplified tissue models they make use of. Indeed, though these are sufficient for producing appreciable image quality improvement, they otherwise induce a statistical bias in the statistics of the restored reflectivity, making it unusable for characterization purposes.

Hereto, a novel deconvolution method for ultrasound images has been developed and is described. The algorithm is de-

rived on the base of a non-standard statistical model for the tissue response, defined by the Generalized Gaussian Distribution. By means of two distinct parameters, called scale and shape parameter, this distribution allows sequences of arbitrary energy and sparsity to be generated, and is therefore adequate for providing an accurate description to the most general tissue structures. Deconvolution is then tackled as a Maximum a Posteriori estimate, in which tissue reflectivity is restored along with an estimate of the associated scale and shape parameter. An Expectation Maximization framework is designed to address this task. An evaluation of performance will be presented on experimental data from several tissue-mimicking phantoms having a well-defined particle concentration. These studies will show improvements in classification accuracy of up to the 20% and the superiority of the proposed algorithm over standard ones.

This thesis part proceeds as follows. The deconvolution problem is formulated in Chapter 4 where most common solutions, as Wiener filter and l^1-norm optimization, are presented. The original restoration scheme will be finally derived and evaluated in Chapter 5.

Deconvolution problem

CONNATURATED with every imaging technique is a discrepancy between the true scene and the imaged one. The entity of these errors is dependent on the features of the sensing devices as well as on the physical phenomena underlying the image formation process. Many of these elements can be reasonably accounted for by linear models. Although not rigorous, they represent a commonly accepted approximation for the largest majority of applications. Under this assumption of linearity, the *observed image* can be expressed as the linear convolution of the *true image* with a linear *blurring kernel*. The problem of image deconvolution, or, equivalently, restoration or deblurring, naturally arises from this scenario, and has its goal in the enhancement of image resolution and contrast by the restoration of an estimate of the true image. Image restoration is a very common problem in image processing, encountered in a wide variety of technical areas as astronomy, seismology, microscopy and medical imaging. See [22, chap. 1] for a review.

In this chapter the problem of medical ultrasound restoration is addressed. In §4.1 the problem is defined and some major issues implied by the peculiar nature of ultrasound are discussed. In §4.2 the simple predictive deconvolution scheme is presented along with its main limitations. In §4.3 it is shown how deconvolution can be formalized as a Bayesian inference problem and the most popular related solutions (Wiener filtering and l^1-norm deconvolution) are presented. As predictive deconvolution is a completely blind technique, Bayesian techniques require the availability of an estimated PSF instead. Hereto in §4.4 several techniques for estimating ultrasound PSFs from the

backscattered echo are discussed. The chapter concludes in §4.5.

4.1   Problem statement

As shown in Chapter 2, as long as imaging of soft tissues is concerned, the first order Born approximation can be applied and the radio-frequency (RF) image y can be expressed as:

$$y[n.m] = \sum_{k,l} h[n, k; m, l]x[k, l] + v[n, m] \qquad (4.1)$$

where x represents the tissue *reflectivity function*, h is the transducer's Point Spread Function (PSF) and v is Gaussian measurement noise. The problem of image restoration is then defined as producing an estimate of x from the observation y. In the sequel we will address some very important issues arising in the case of ultrasound images, making the restoration problem for this class of images particularly challenging.

At first we note that the model (4.1) is intrinsically non stationary, indeed h is allowed to vary its shape when measured at different locations. This is indeed the case of medical ultrasound, where the system PSF may change considerably on the imaged plane. This spatial variability is principally connected with the beamforming performed by the array in order to have better resolved images, but is also due to the presence of the tissue itself, whose attenuating and dispersive action on the acoustic field propagation is known to be highly frequency dependent. These issues have been more exhaustively explained earlier in §2.5 of this manuscript.

These considerations not only confirm that a spatially variant PSF must be considered in a realistic signal model, but introduce a further issue. This consists in the fact that, although spatial variations associated with system dependent effects could theoretically be *a priori* modeled given a perfect knowledge of the acquisition modality, instead this is not the case for tissue dependent ones. Indeed they are intrinsically depend on the structural properties of the insonified medium, which, excluding very controlled experimental setups, are completely unpredictable. As a result, the blurring term h in (4.1) is also an unknown of the problem. The class of deconvolution problems in which none or little information is available on the blurring kernel is referred in literature as *blind deconvolution* [22, 23].

The blind deconvolution problem can be tackled in two ways. The first consists in estimating the tissue reflectivity and the trans-

ducer PSF simultaneously. The second approach is instead to estimate the PSF first, followed by using such an estimate to solve the deconvolution problem in a non-blind manner. We anticipate that this second solution is fairly the most common. Indeed, as we will see, these methods are free from strong hypothesis on the signal statistics implied by the former ones, and allow for more realistic tissue models to be taken into account.

Among the completely blind deconvolution strategies *predictive deconvolution* represents a very popular alternative. This technique is described in the next section.

4.2   Predictive deconvolution

Predictive deconvolution was first introduced in [24] for deconvolving seismograms obtained in reflection seismology. Seismology and medical ultrasound represent two techniques inspired by the same intent, that is to say to characterize a propagating medium through the echo returned after an acoustic solicitation of the same. In the former case air guns are employed as sources and the medium is the earth crust, in the latter piezo-electric transducers are used to image biological tissues. These similarities motivate the tentative to apply tools developed for the processing of seismic signals to the treatment of medical ones. In particular, predictive deconvolution for medical ultrasound has been investigated in few works, as [25, 26]. The reader is addressed to [27] for an extensive dissertation on the employment of predictive techniques, and adaptive filters in general, for medical ultrasound restoration.

Predictive deconvolution assumes an autoregressive (AR) model for the echo signal y, i.e.

$$y[n] = \sum_{k=1}^{P} a[k] y[n-k] + x[n] \qquad (4.2)$$

where y and x have the same meaning as in (4.1). The coefficients $a[k]$ are called auto-regressive coefficients of y, and P is the order of the model. In this context x is also called excitation of the autoregressive model, and must be white noise. Note that a one dimensional formulation has been used. Indeed, in medical ultrasound, the common practice is to apply blind deconvolution techniques of this kind only along the axial direction of propagation [22, 25].

Figure 4.1: Predictive deconvolution.

Note that (4.2) is equivalent to $y[n] = (h * x)[n]$ if:

$$H(z) = \frac{Y(z)}{X(z)} = \frac{1}{1 - \sum_{k=1}^{p} a[k]z^{-k}} \qquad (4.3)$$

where $H(z) = \sum_n h[n]z^{-n}$ is the Z-transform [8, chap. 3] of the PSF h. That is to say when h is an all-pole, minimum-phase filter [8, pp. 280–290].

The scheme for predictive deconvolution is represented in Figure 4.1. The function of the predictive block is to produce a guess $\hat{y}[n]$ of $y[n]$ given a set of its delayed samples. In practice it is implemented as a linear finite impulse response (FIR) filter with taps $\mathbf{w} = [w_1, \cdots, w_M]$, so that:

$$\hat{y}[n] = \sum_{k=1}^{M} w[k]y[n-k] \qquad (4.4)$$

The values of the coefficients is computed so to satisfy the optimal prediction condition:

$$\mathbf{w}^* = \underset{\mathbf{w}}{\mathrm{argmin}}\{E(e^2[n])\} \qquad (4.5)$$

where $e[n] = y[n] - \hat{y}[n]$ is the prediction error and $E\{\cdot\}$ represents the statistical expectation.

It is immediate to show that, when y is an autoregressive process, the so obtained linear prediction coefficients provide in effect an estimate of the autoregressive parameters a. As a consequence, the prediction error $e[n]$ is an estimation of the excitation process x, that is to say the deconvolved trace.

A common way to compute the prediction coefficients in (4.5) is by solving the associated Wiener-Hopf equations [28, chap. 6]. This solution is optimal provided that a consistent estimate of the autocorrelation function of the input trace $y[n]$ can be computed. This implies that a stationary time realization of sufficient duration is available. Unfortunately, medical ultrasound signals cannot be considered stationary or, at least,

stationarity hypothesis can be made only on small segments. Indeed, though the time invariant autoregressive model has been introduced in (4.2) for conformity with the literature on AR processes, a much more realistic model would instead take time-variant coefficients into account, i.e.

$$y[n] = \sum_{k=1}^{P} a[n;k]y[n-k] + x[n] \qquad (4.6)$$

As a consequence not a fixed set of AR coefficients must be estimated from an entire ultrasonic scan line, otherwise they should be adaptively updated so to second to the time-variant properties of the signal. A consolidated alternative for dealing with this class of signals is represented by adaptive filtering techniques.

4.2.1   RLS algorithm for linear prediction

Adaptive filters are filters in which taps are recursively updated throughout all the acquisition process in order to compensate for time varying changes of acquisition system or propagating channel. The literature on adaptive filters is extremely wide, see [28] for an exhaustive dissertation. The most classical solutions for the updating rule are the Least Mean Square (LMS) algorithm [28, chap. 9] and the Recursive Least Squares (RLS) algorithm [28, chap. 13]. In particular, the second is known to ensure better performance both in terms of convergence and tracking capabilities.

The RLS algorithm relies on the following approximation of the cost function $E\{e^2[n]\}$:

$$\Gamma\{e^2[n]\} = \sum_{k=1}^{n} \lambda^{n-k} e^2[n] \qquad (4.7)$$

where $\lambda \in [0,1]$ is called the *forgetting factor* and represents the memory of the filter. In particular, values of λ close to 1 correspond to considering segments of higher duration: this improves the coefficients estimate when stationarity assumptions can be made but limits the tracking capabilities for highly non stationarity signals. Symmetrically, small values of λ improve the dynamic behavior on one side, but make the approximation given by Γ more dramatic on the other, so degrading the consistency of the estimated coefficients. As a consequence, λ has to be carefully tuned in relation with the application at hand.

Table 4.1: Summary of the RLS algorithm

Initialization: $\mathbf{P}[0] = \delta^{-1} \cdot \mathbf{I}$, $\hat{\mathbf{w}}[0] = \mathbf{0}$,

---

For n= 1,2,..., compute:

$$\mathbf{k}[n] = \frac{\lambda^{-1}\mathbf{P}[n-1]\mathbf{u}[n]}{1 + \lambda^{-1}\mathbf{u}^T[n]\mathbf{P}[n-1]\mathbf{u}[n]},$$

$$\xi[n] = y[n] - \hat{\mathbf{w}}^T[n-1]\mathbf{u}[n],$$

$$\hat{\mathbf{w}}[n] = \hat{\mathbf{w}}[n-1] + \mathbf{k}[n]\xi[n],$$

$$\mathbf{P}[n] = \lambda^{-1}\mathbf{P}[n-1] - \lambda^{-1}\mathbf{k}[n]\mathbf{u}^T[n]\mathbf{P}[n-1],$$

Before going through the description of the algorithm it is useful to introduce here the some variables. We define

$$\mathbf{u}[n] = [y[n-1], \cdots, y[n-M]]^T \tag{4.8}$$

the input to the predictor, so that $\hat{y}[n] = \mathbf{w}^T\mathbf{u}[n]$;

$$\Phi[n] = \sum_{i=1}^{n} \lambda^{n-i}\mathbf{u}[i]\mathbf{u}^T[i] \tag{4.9}$$

the estimate of the autocorrelation matrix and

$$\mathbf{z}[n] = \sum_{i=1}^{n} \lambda^{n-i}\mathbf{u}[i]y[i] \tag{4.10}$$

is the estimate of the crosscorrelation vector between the input sequence $\mathbf{u}$ and y.

The solution $\mathbf{w}^*$, that minimizes (4.7), is then obtained from the relation

$$\mathbf{w}^*(n) = \Phi^{-1}(n)\mathbf{z}(n) \tag{4.11}$$

In practice no matrix inversion is needed and the matrix inversion lemma is used instead [28, pp. 565]. A summary of the RLS algorithm is reported in Table 4.1. The RLS algorithm is a very consolidated framework in adaptive filter theory, the interested reader may find an exhaustive derivation e.g. in [28].

4.2.2 *Limitations of the approach*

Several pros and cons go along with the predictive deconvolution technique when employed in the context of medical ultra-

sound. They have been mentioned throughout the text but it's useful to summarize them here. The main advantages are:

- The small computational complexity. In fact, they relies on simple filtering techniques for which several simple hardware implementations could be thought or borrowed from the telecommunications literature. Moreover each scan line is typically processed independently from the adjacent one, as a consequence this intrinsic parallelism could be exploited for considerably reducing the total cost associated to the deconvolution of the entire image. For these reasons, adaptive filtering deconvolution techniques are with no doubts the most promising candidates as long as real-time processing is pursued.

- They are completely blind. As a consequence no PSF estimation strategy is needed. We will see in the following that several complications are associated with these estimation techniques.

- The intrinsic non-stationarity of the system is directly taken into account as the model coefficients are continuously updated.

Limitations are mainly due to the assumptions at the base of the predictive deconvolution scheme. Indeed, if they can be adequate in reflection seismology or in some telecommunication application, they are instead too severe for medical ultrasound. These are:

- The PSF cannot be well represented by an all-pole system as the one in (4.3). Indeed the AR model is useful for modeling random signals that possess peaky power spectral densities. When more complex situations with broad peaks or sharp nulls exist, which is the case in medical ultrasound, more general models should be alternatively used, as autoregressive moving average (ARMA). These concept is illustrated in Figure 4.2, where the inadequacy of the AR model for representing the actual shape of the spectrum of an experimental PSF is evident, along with the better approximation offered by an ARMA model. Unfortunately, several shortcomings accompany the employment of ARMA schemes, which limit their use as well. At first much onerous model estimation procedures are

Figure 4.2: (a) Experimental PSF obtained as the response of a metal wire immerse in a water tank. (b) Power spectrum of the experimental pulse. (c) Power spectrum of the AR estimate. (d) Power spectrum of the ARMA estimate. For the AR the best fit has been obtained for an filter length of $M=4$. For the ARMA both number of zeros and poles was equal to 5.

needed. Moreover, adaptive strategies cannot be formalized [22, pp.185–190], which preclude the time variance of the system from being correctly taken into account.

- Experimental PSFs are known to have a non-negligible lateral extent, which makes the 1-D model highly inadequate. For this reason, modern deconvolution techniques all exploit more realistic 2-D [29, 30, 31], when not 3-D [32, 33], PSF models.

- Assuming the tissue reflectivity to be a white Gaussian process is an unacceptable simplification for biological tissues. At first, variations in the echogenicity profiles are the immediate consequence of the simultaneous presence of different biological structures encountered inside realistic images, which make a white process a highly inadequate

model. Moreover more realistic scattering conditions can
be accounted for by more flexible distributions than the
Gaussian, as detailed in Chapter 3. These concepts will be
further addresses in the next sections and more diffusely
in the next Chapter.

The aforementioned limitations justify the fact that very few
authors spent big efforts in the investigation of predictive de-
convolution schemes for medical ultrasound. The most com-
mon strategy is instead to exploit different approaches, relying
upon statistical estimation and convex-optimization theory. At
the same time we further stress that techniques based on adap-
tive filters are instead the most promising alternative when on-
line processing is needed.

4.3   Maximum a posteriori deconvolution

Unlike predictive deconvolution, this kind of techniques address
the blind deconvolution problem in a two step scheme: the PSF
h is estimated first and subsequently the task is tackled in a non-
blind manner. In this section we describe the most common
techniques for the latter step. Possible ways for estimating the
PSF from the echo image are described later in this chapter.

4.3.1   Problem formulation

For compactness of notation, let's rewrite the signal model in
(4.1) with the common matrix-vector notation [34]:

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{v} \tag{4.12}$$

where $\mathbf{y}$, $\mathbf{x}$ and $\mathbf{v}$ are vectors obtained after lexicographical or-
dering from the corresponding images and $\mathbf{H}$ is the matrix as-
sociated to the blur [34, Appendix A]. Specifically, assuming a
generic $N \times M$ RF-image, then $\mathbf{y}$ is the $NM \times 1$ vector obtained
by stacking each column under the previous one and $\mathbf{H}$ is an
$NM \times NM$ matrix. With this formalism, deconvolution trans-
lates into providing a reasonable solution $\hat{\mathbf{x}}$ to (4.12).

We remark that PSF shape variation requires a *spatially vari-
ant* blur to be taken into account within an accurate deconvolu-
tion framework. Though a pointwise variant kernel could theo-
retically be embedded in (4.12), this would require products $\mathbf{H}\mathbf{x}$
and $\mathbf{H}^T\mathbf{x}$ to be computed explicitly, thereby making deconvolu-
tion computationally unfeasible. Hereto, the simplest approach

is to assume a slowly varying PSF [31, 29, 35], and hence well approximated by a *piecewise constant* function. In this case, the data image can be divided into a number of segments, whose size is chosen to be small enough to guarantee that each of the segments is formed by a stationary convolution with a different PSF.

With this model, the following expressions holds [36]:

$$\mathbf{H} = \sum_{k=1}^{K} D_k \mathbf{H}_k, \qquad \mathbf{H}^T = \sum_{k=1}^{K} \mathbf{H}_k^T D_k \qquad (4.13)$$

where $\mathbf{H}_k$ denotes the *spatially invariant* PSF associated to the kth segment and D_k is the diagonal matrix determining the interpolation between the kth region and the neighboring ones (cf. Appendix B). Since each of the $\mathbf{H}_k$ is Block Toeplitz with Toeplitz Blocks, then products of the type $\mathbf{Hx}$ and $\mathbf{H}^T\mathbf{x}$ can be efficiently implemented via Fast Fourier Transform [36]. Details about structure of $\mathbf{H}$ are given in Appendix B of this manuscript.

4.3.2 Bayesian framework

Image deconvolution is often tackled as a statistical inference problem [37, 38, 35]. In this case, vectors $\mathbf{y}$, $\mathbf{x}$ and $\mathbf{v}$ of the model in (4.12) are treated as single realizations of random processes, to which a suitable probability density function (pdf) $p(\cdot)$ is assigned. In this context, the two most common estimation paradigms are Maximum Likelihood (ML) and Maximum a Posteriori (MAP), respectively seeking the realization of $\mathbf{x}$ maximizing the *log-likelihood* $\ln[p(\mathbf{y}|\mathbf{x})]$ and the *log-posterior* $\ln[p(\mathbf{x}|\mathbf{y})]$, see [16].

If the noise is modeled as zero mean white Gaussian with variance σ_n^2:

$$p_v(\mathbf{v}) = \frac{1}{(2\pi\sigma_n^2)^N} \exp\left(-\frac{1}{2\sigma_n^2}||\mathbf{v}||_2^2\right) \qquad (4.14)$$

then $p(\mathbf{y}|\mathbf{x}) = p_v(\mathbf{y} - \mathbf{Hx})$:

$$p(\mathbf{y}|\mathbf{x}) = \frac{1}{(2\pi\sigma_n^2)^N} \exp\left(-\frac{1}{2\sigma_n^2}||\mathbf{y} - \mathbf{Hx}||_2^2\right) \qquad (4.15)$$

where N is now the total number of samples in the image. From (4.15) one can derive that ML estimate corresponds to the least squares solution to (4.12), minimizing the error norm $||\mathbf{y} - \mathbf{Hx}||_2^2$.

As known, this problem is directly solved by pre-multiplying $\mathbf{y}$ with the pseudo-inverse of $\mathbf{H}$, i.e. $(\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T$. Unfortunately, in medical ultrasounds, like in any other measurement technique, the system matrix $\mathbf{H}$ is necessarily ill-conditioned. This fact is directly connected to the band-limitedness of acquisition system impulse response. Hence, inverse filtering results in an unacceptable amplification of the out-of-band noise components. In this context, MAP technique provides a way to cope with the ill-posedness of the restoration problem by including further information in the estimation task. In particular, it requires a *prior* model $p(\mathbf{x})$ to be assigned to the desired solution. Using the Bayes rule $p(\mathbf{x}|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{x})p(\mathbf{x})$, the MAP estimate of $\mathbf{x}$ then writes as:

$$\hat{\mathbf{x}}_{\text{MAP}} = \underset{\mathbf{x}}{\text{argmin}}\left\{\frac{1}{2\sigma_n^2}\|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 - \log p(\mathbf{x})\right\} \qquad (4.16)$$

where $p(\mathbf{x})$ reflects the prior belief about $\mathbf{x}$, acting as a regularizing constraint in the inversion of (4.12). In practice, reflectivity samples x_i are modeled as independent identically distributed random variables, so that $p(\mathbf{x}) = \prod p(x_i)$. The two reflectivity models by far the most common in deconvolution literature, i.e. Gaussian and Laplacian, are discussed in the next Section.

4.3.3   Gaussian model

If $\mathbf{x}$ obeys a multivariate Gaussian distribution with zero mean and co-variance matrix $\text{E}\{\mathbf{x}^T\mathbf{x}\} = \Sigma_{\mathbf{x}}$:

$$p_{\mathbf{x}}(\mathbf{x}) = \frac{1}{(2\pi)^{1/2}|\Sigma_{\mathbf{x}}|^{1/2}} \exp\{\frac{1}{2}\mathbf{x}^T\Sigma_{\mathbf{x}}^{-1}\mathbf{x}\} \qquad (4.17)$$

then the restoration problem maps into a (weighted) l^2-norm optimization task:

$$\log p(\mathbf{x}|\mathbf{y}) \propto \left\{\frac{1}{2\sigma_n^2}\|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \frac{1}{2}\mathbf{x}^T\Sigma_x^{-1}\mathbf{x}\right\} \qquad (4.18)$$

The solution $\nabla_{\mathbf{x}}\log p(\mathbf{x}|\mathbf{y}) = 0$ is given by:

$$\hat{\mathbf{x}} = (\mathbf{H}^T\mathbf{H} + \sigma_n^2\Sigma_{\mathbf{x}}^{-1})^{-1}\mathbf{H}^T\mathbf{y} = \mathbf{W}\mathbf{y} \qquad (4.19)$$

where $\mathbf{W}$ is the well known Wiener filter [34, Ch.8]. Due to the size of the problem the matrix inversion involved by (4.19) is not feasible in a direct way. It is hereto common to adopt iterative

Table 4.2: Conjugate Gradients Algorithm

Initialization: $\mathbf{r}_0 = \mathbf{b} - \mathbf{Ax}$, $\mathbf{p}_0 = \mathbf{r}_0$, $k = 0$

repeat

$\qquad \alpha_k = \mathbf{r}_k^\dagger \mathbf{r}_k / (\mathbf{r}_k^\dagger \mathbf{Ar}_k)$,

$\qquad \mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{x}_k$,

$\qquad \mathbf{r}_{k+1} = \mathbf{r}_k - \alpha_k \mathbf{Ar}_k$,

$\qquad \beta_k = \mathbf{r}_{k+1}^\dagger \mathbf{r}_{k+1} / (\mathbf{r}_k^\dagger \mathbf{r}_k)$,

$\qquad \mathbf{p}_{k+1} = \mathbf{r}_{k+1} + \beta_k \mathbf{p}_k$,

$\qquad k = k + 1$

until: $\|\mathbf{r}_k\|_2 < \tau_{CG}$

---

linear system solvers as Conjugate Gradients [39]. A schematic of this technique is reported in Table 4.2, where $\mathbf{A} = (\sigma_n^2 \Sigma_\mathbf{x}^{-1} + \mathbf{H}^T\mathbf{H})$ and $\mathbf{b} = \mathbf{H}^T\mathbf{y}$.

Despite the multivariate pdf in (4.17) represents an accurate model, allowing variations in echogenicity profiles to be taken into account, and has been explored in few works [35], it otherwise involves some tedious issues as the estimate of $\Sigma_\mathbf{x}$ and its inversion. In this context it is fairly more common to assume white models for $\mathbf{x}$ so that $\Sigma_\mathbf{x} = \sigma_x^2\mathbf{I}$, where $\mathbf{I}$ is the identity matrix. As a consequence (4.19) becomes:

$$\hat{\mathbf{x}} = \left(\mathbf{H}^T\mathbf{H} + \frac{\sigma_n^2}{\sigma_x^2}\right)^{-1} \mathbf{H}^T\mathbf{y} = \mathbf{Wy} \qquad (4.20)$$

This simplified model was adopted in [31].

Moreover, as far as spatially-invariant PSFs are considered, then $\mathbf{H}$ can be immediately diagonalized by means of its Fourier coefficients and the restoration problem can efficiently be solved in the Fourier domain as:

$$\hat{X} = \frac{H^*}{|H|^2 + \sigma_n^2/\sigma_x^2} \qquad (4.21)$$

where the transformed signals have been denoted by capital letters. Equation (4.21) represents the most popular expression for Wiener filtering. This DFT domain solution has been widely exploited in literature for its computational feasibility, mainly by Taxt *et al.*[29, 30, 40, 32]. When (4.21) is adopted, the common

way do deal with non stationarity of the PSF is to deconvolve single sub-images and then to merge the results. Taking care of avoiding block artifacts by suitably overlapping the image segments.

Despite the computational efficiency, DFT-domain Wiener filter yields sub-optimal results because the complex exponentials that comprise the basis of the DFT have poor spatial localization, so that modifications to individual DFT coefficients of $\mathbf{x}$ result in global changes to the entire image that are often manifested as Gibbs ringing artifacts [41]. Moreover, as a linear filter, Wiener solution is incapable of interpolating the out of band information that was lost during the image formation process, so producing oversmoothed solutions [31]. The reader may note that these shortcomings are shared with the predictive deconvolution scheme described at the beginning of this Chapter. Indeed, although derived in different ways, these two techniques have much in common: in both cases a Gaussian model for the reflectivity is assumed, which, as we will discuss later, in fairly too simplified for a real tissue, moreover, in both cases a linear regularized approximation of inverse filter is produced, which goes along with the oversmoothing and the ringing phenomena.

These considerations motivate the need for alternative reflectivity models. In this context Laplacian distribution has been the most popular.

4.3.4   Laplacian model

If a Laplacian model is assumed for the reflectivity:

$$p(x_i) = \frac{1}{\sqrt{2\sigma_x^2}} \exp\left(-\sqrt{\frac{2}{\sigma_x^2}} |x_i| \right) \tag{4.22}$$

then the restoration problem in (4.16) translates into an l^1-norm optimization task [31]:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \left\{ \frac{1}{2\sigma_n^2} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1 \right\} \tag{4.23}$$

where $\lambda = \sqrt{2}/\sigma_x$ and $\|\mathbf{x}\|_1 = \sum |x_i|$ denotes the l^1-norm. The optimization problem defined by (4.23) does not have a closed form solution, and must be handled with iterative convex optimization schemes [31, 42]. A possible solution to (4.23) will be addressed in the next Chapter, where a general framework for the solution of l^p-norm problems will be described.

Figure 4.3: Comparison between Gaussian and Laplacian pdfs.

The literature on l^1-norm optimizations is wide, indeed it is the standard procedure solving linear systems when the solution is known to be sparse. In the context of medical ultrasound deconvolution, Laplacian prior was first introduced by Michailovich *et al.*[31, 43]. It has to be noted that Laplacian pdf possesses heavier tails than the Gaussian one and, hence, there is a greater allowance for a few occasional large-amplitude samples to occur, cf. Figure 4.3. For this reason, such a model is better suited for representing more realistic tissues with diffusive scatterers superimposed with sparse specular reflectors. A comparison between Wiener and l^1-norm deconvolution is given in Figure 4.4, where the increased sparsity associated to the Laplacian prior is evident, along with the oversmooting effect caused by Wiener filtering.

4.3.5   *Advantages and shortcomings*

These maximum a posteriori techniques represent the most common restoration strategies for medical ultrasound. An overview of the literature concerning ultrasound images restoration produced in the last 15 years is summarized in Table 4.3. Those methods are directly derived on the base of the optimization tasks derived in the previous section, or on slight modifications. From the second column of the Table is evident how in some case working on the *in-phase/quadrature* (IQ) signal is preferred than on the raw RF. This is due to the fact that, after demodulation, the IQ signal can be downsampled without introducing aliasing effects. The downasmpling factor is dependent on the ratio between the transducer's bandwidth and the sampling frequency (factors of 4-5 are often feasible with standard acquisi-

Figure 4.4: Comparison between wiener and laplacian deconvolution relative to a prostate scan acquired with Trans Rectal Ultrasound (TRUS).

tion setups). This is highly desirable as standard deconvolution algorithm as the one discussed above mostly have a computational complexity which raises with the square of the number of samples.

Main advantages of MAP techniques include:

- Two dimensional [30] or even three dimensional [32] PSFs can be accounted for.

- No assumption is made on the PSF, neither on the number of zeros or poles, neither on the position of the zeros on the complex plane. In particular both minimum phase (all zeros and poles inside the unit circle) and mixed phase [8] PSF can be considered. In this context we will see that this second hypothesis is the more realistic for ultrasound pulses, which involve several complications in the PSF estimation procedure.

- More advanced models than white Gaussian can be assumed for the tissue reflectivity. In particular a degree of correlation among samples can be introduced, as in [35],

Table 4.3: Literature Overview on Medical Ultrasound
Deconvolution

WORK	METHOD		
	Signal	PSF [1]	Optimization
Abeyratne [44] (1995)	1D-RF	Estimated (HOS)	Wiener DFT (4.21)
Taxt [40] (1999)	1D-RF	Estimated (H)	Wiener Block (4.21)
Taxt [30] (2001)	2D-IQ	Estimated (H)	Wiener Block (4.21)
Taxt [32] (2001)	3D-RF	Estimated (H)	Wiener Block (4.21)
Michailovich [45] (2003)	1D-RF	Estimated (GH)	Wiener Block (4.21)
Taxt [29] (2004)	2D-IQ	Estimated (H)	Wiener Block (4.21)
Michailovich [31] (2005)	2D-IQ	Estimated (GH)	l^1-norm (4.23), Wiener (4.20)
Jiřík [46] (2006)	2D-RF	Estimated (H)	Wiener DFT (4.21)
Ng [35] (2007)	2D-IQ	Simulated: FieldII [6]	Iterative Wiener (4.19)
Michailovich [43] (2007)	2D-IQ	Iterative refinement of GH estimate	Hybrid. l^1-norm.

[1] When an estimated PSF is used, the following notation is adopted: HOS
for the cumulants-based estimate in §4.4.1, H for the cepstrum estimate
in §4.4.2, GH for the generalized homomorphic estimate in §4.4.3.

where piecewise-smooth echogenicity variations due to
the presence of multiple tissues in the image are taken into
account by considering a non-white multivariate distribu-
tion for the reflectivity. Moreover, the hypothesis of Gaus-
sianity can be disregarded as well. This is done for instance
with the Laplacian model, but this concept will be further
expanded in the next Chapter.

Main limitations include:

- An increased computational cost, in particular when more
complex scheme than Wiener are employed. As a conse-
quence, though considerable speed ups could be obtained
thanks to efficient implementations on dedicated hard-
ware platforms or on GPUs, these tools are not promising

candidate as long as real-time processing is the principal goal.

- Although PSF estimation strategies have received a great attention among several authors [29, 31, 45] this step still remain tedious. In particular this is motivated by the need of phase unwrapping procedures when non-minimum phase PSFs are considered. As long as ultrasounds are considered this operation is particularly challenging due to the high and rapid variations of the signal phase, which make standard phase unwrapping algorithms developed for other image modalities as MRI [47] of little use.

4.4 PSF ESTIMATION

As mentioned, deconvolution of ultrasound images is a blind problem and MAP techniques cope with this fact by estimating the PSF from the RF image, and then use the obtained estimate within a non-blind setting. As the image formation is known to be non stationary, the common solution is to partition the image axially into segments small enough to consider the process locally stationary and then to estimate a different PSF on each segment. The global blurring operator is then to be built as in (4.13).

The literature on PSF estimation techniques is wide. These methods can be grouped into three main classes, which we briefly resume in this section. The reader is addressed to [22, chap. 5] for a more complete review on the subject. In order to keep the following discussion fluent let's introduce here the variables to be used throughout the rest of the section, see Table 4.4.

4.4.1 Higher-Order Statistic technique

This method was first proposed in [44] for 1D signals. It relies on the following hypotheses: the psf h is deterministic and mixed-phase, the reflectivity x is stationary zero-mean and non-Gaussian, the noise v is stationary white Gaussian. Under these assumptions the *bispectrum* of y can be written as [48]:

$$C_y(\omega_1, \omega_2) = C_x(\omega_1, \omega_2)H(\omega_1)H(\omega_2)H^*(\omega_1+\omega_2)+C_v(\omega_1, \omega_2)$$

It can be shown that, if v is white Gaussian then $C_v = 0$. Moreover, if x is white non-Gaussian then its bispectrum is simply a

Table 4.4: List of Symbols for PSF Estimate[1]

SYMBOL	DEFINITION	DESCRIPTION
$G[\omega]$	$\mathfrak{F}\left(g[n]\right)$	spectrum
$G[\omega_1, \omega_2]$	$\mathfrak{F}_{2D}\left(g[n, m]\right)$	2D spectrum
$\hat{G}[\omega_1, \omega_2]$	$\log G[\omega_1, \omega_2]$	log-spectrum
$c_g[\tau_1, \tau_2]$	$E\{g[n]g[n+\tau_1]g[n+\tau_2]\}$	3-rd order cumulant
$G_g[\omega_1, \omega_2]$	$\mathfrak{F}_{2D}\left(c_g[\tau_1, \tau_2]\right)$	bispectrum
$\hat{g}[n, m]$	$\mathfrak{F}_{2D}^{-1}\left(\hat{G}[\omega_1, \omega_2]\right)$	complex cepstrum

[1] In this table $\mathfrak{F}$, $\mathfrak{F}_{2D}$ and $\mathfrak{F}_{2D}^{-1}$ represent Fourier transform, two-dimensional Fourier transform and inverse two-dimensional Fourier transform respectively. In practice, as long as discrete time series are considered, they are implemented by the Fast Fourier Transform (FFT) algorithm.

constant. Therefore is evident how C_g is, a part of the multiplication by a constant, exclusively determined by the spectrum of h, which suggests how the blur can be estimated from the bispectrum of the rf signal.

A strong limitation of this approach is that consistent estimates of higher order statistics require long stationary time realizations to be available. Otherwise all the desirable properties described above may not hold and the resulting PSF estimate be inaccurate. As mentioned, this condition is never satisfied in a real world setting, where the echo rapidly varies its properties. Moreover 2D extension to these techniques is not straightforward and may become fairly too computation demanding. A much more feasible and popular alternative is represented by homomorphic techniques.

4.4.2   Homomorphic technique

Homomorphic deconvolution has been widely explored with a certain success by Taxt *et al.*, see Table 4.3. The log-spectrum $\hat{Y}$

of y is:

$$\hat{Y}(\omega_1, \omega_2) = \log|X(\omega_1, \omega_2)| + \log|H(\omega_1, \omega_2)| + \\ +j\left\{\angle H(\omega_1, \omega_2) + \angle X(\omega_1, \omega_2)\right\} \tag{4.24}$$

It can be shown that, at least withing the transducer bandwidth, both $\log|H|$ and $\angle H$ are reasonably smooth signals, while both the power spectrum and the phase of x are spiky noise-like sequences [49]. As a consequence, the spectrum of H can be retrieved from $\hat{Y}$ by filtering both its real and imaginary parts.

Homomorphic techniques accomplish this separation in the cepstrum domain [8]. In particular, from (4.24) it is straightforward to derive $c_y[n, m] = c_x[n, m] + c_h[n, m]$, where, for the considerations above, c_h is a fast converging sequences, with its energy confined at the origin of the cepstrum domain, while the energy of y is spread over the entire cepstrum domain due to the non regularity of $\hat{Y}$. Although the two sequences c_h and c_y may have a certain overlapping, this is usually neglected and the cepstrum of the PSF c_h is simply obtained by truncating the cepstrum c_y of the echo signal [50, 29, 51].

4.4.3   Generalized homomorphic technique

Generalized homomorphic technique [31] also exploits different smoothing properties of H and X in order to retrieve a PSF estimate from (4.24). The main difference is that smoothing is not performed by truncation in the cepstrum domain, as in the traditional homomorphic scheme, but by projection of both the real and imaginary parts of (4.24) on a reduced resolution space. This is done by means of wavelet filtering, where the signal statistics are taken into consideration in order to define the optimal basis and threshold. This technique is documented to outperform the traditional one in [50] in terms of mean square reconstruction error and will be adopted in the following Chapter.

We conclude this section by noting that a major complication goes along with homomorphic or generalized homomorphic techniques, represented by the need of unwrapping the phase of Y before the PSF estimation procedure. As the phase of ultrasound signals present very peculiar properties, making standard unwrapping tools developed for alternative imaging modalities like MRI of little or no use at all [47], several *ad hoc* phase unwrapping algorithms have been proposed in literature. A discussion on phase unwrapping is beyond the scope

of this manuscript. Major contributors in this context are Strand and Taxt [52, 53] and Michailovich [49, 54]. Although good results are documented, details about their implementation are often left quite obscure. In the following chapter, as a compromise between feasibility of implementation and accuracy of results, we adopt the phase unwrapping technique described in [55].

4.5   CONCLUSION

In this chapter the problem of deconvolution for medical ultrasound has been introduced. The principal techniques for tackling this task have been reviewed, with an emphasis on predictive deconvolution and Maximum a Posteriori approaches. We mentioned that predictive deconvolution, despite its computational efficiency, introduces several simplifications which are too severe for medical ultrasound and, for this reason, it has received only a limited attention. In particular, we have obtained some interesting results relative to the employment of predictive deconvolution within a Trans Rectal Ultrasound (TRUS) based Computer Aided Detection (CAD) tool for classification of prostate cancer [26, 56]. Otherwise, that project was marginal to the primary doctoral research activity and wont be further detailed in this manuscript.

The focus then drifted to Maximum a Posteriori restoration techniques, which represent the state of the art in medical ultrasound and in image processing in general. Indeed, due to their flexibility, they allow suitably modeling the majority of tissue and system dependent effects involved in a real acquisition. In particular most common Wiener filter and l^1-norm optimization have been defined. Those MAP techniques assume a priori knowledge of the blurring kernel, i.e. the system PSF. As the problem of deconvolution of medical ultrasound is blind, the common solution is to use estimates of the PSF obtained from the echo signal. Hereto the most classical estimation strategies have been reviewed at the end of the chapter.

All the concepts here introduced will be exploited in the next chapter, where a novel algorithm for deconvolution of medical ultrasound will be derived.

Deconvolution and tissue characterization

MEDICAL ULTRASOUND represents a non-invasive and cost-efficient imaging modality. Beyond that, ultrasound provides a precious tool for tissue characterization since, from the analysis of the backscattered echo signal, relevant information on structural properties of the insonified tissue can be inferred, otherwise unaccessible from simple visual inspection.

This principle is at the base of a consolidated research branch referred as Ultrasonic Tissue Characterization (UTC). In particular, it has been effectively exploited in a clinical context in the development of ultrasound-based Computer Aided Detection (CAD) schemes. The purpose of CAD is to support the physician in the interpretation of ultrasound images and to assist him in the decision-making process when suspicious situations are encountered. Significant results have been documented on a range of clinical applications, like prostate cancer recognition on TRUS [57], detection of suspicious masses in breast ultrasounds [58] and diagnosis of hepatic steatosis [59]. The computer output is derived from quantitative analysis of the echo signal. In this context, a large amount of features of different nature have been proposed in literature, which can be subdivided according to their contribution in highlighting specific tissue properties.

Tissue characterization based on the acoustic parameters such as attenuation and backscattering coefficients extracted from radio-frequency (RF) echo signals has been widely studied. These quantities are commonly estimated by using one-dimensional [60] or two-dimensional [61] spectrum analysis of the RF signal. Spectral features have proven to provide useful outputs for

diagnosis of diseases in various organs, like eye, prostate, breast and liver; see [62] for a comprehensive review. Besides the RF-spectrum analysis, many researchers have used texture features extracted from ultrasound B-scan images for characterization purposes. The initiative is that the speckle pattern in the ultrasonic image can reveal structural information about the tissue. Usefulness of textural features within clinical settings has been largely documented, e.g. for prostate carcinoma diagnosis [63, 26], evaluation of liver diseases like hepatoma and cirrhosis [64] and detection of atherosclerotic plaques in the carotid artery [65].

A third class of features are derived from statistical modeling of the ultrasound echo distribution, allowing to bring information on scatterer concentrations and distribution pattern. As said in Chapter 3, many statistical models have been developed in literature associated with different scattering conditions. The most popular is represented by the Rayleigh distribution for the envelope signal, which can be analytically derived when diffuse scattering occurs [9, pp.48–50],[15]. Nevertheless, diffuse scattering conditions are often violated in biological tissues, either because the number of scatterers per resolution cell may not be large enough, or because of the presence of regular patterns in the scatterers location. In these cases Rayleigh model is inadequate and more complex models have to be adopted. In this context, a variety of distributions have been used, such as Rician [9, pp.50–52], K [10], Homodyne-K [11], and Nakagami [12]. In particular, Nakagami represents the most popular model for fitting envelope amplitude histograms [12]: its two parameters, called scale and shape parameter, are strictly correlated to scatterer strength and concentration, and are widely used for tissue characterization purposes [18]. Besides, Nakagami parameters have been shown to assess good discriminating performances on B-mode, logarithmically compressed, data as well [66]. When the RF signal is considered instead, Krf distribution [13] and Generalized Gaussian distribution (GGD) [14] have been shown to provide a precise fit under the most general scattering conditions.

Due to their computational feasibility and attested diagnostic relevance [13, 18], statistical features are particularly well suited for those applications in which online assistance has to be provided, like for biopsy guidance. Besides, their ability of enhancing different scatterer concentrations makes them a precious support tool for those clinical situations in which healthy and dis-

eased tissues are discriminated according to density, but exhibit similar echogenicity, and result therefore invisible from visual inspection.

These features are normally extracted on the echo signal. Nevertheless, due to the blurring associated to the system Point Spread Function (PSF), the backscattered echo is known to carry degraded information on the tissue. This fact suggests that a deconvolution step could tentatively be employed in order to mitigate this effect and restore an estimate of the true tissue response, and, correspondingly, features computed on deconvolved data be exploited for an improved characterization. In this context, we already obtained some interesting result in [26], where a simple predictive deconvolution scheme was employed and slight improvement in classification accuracy is documented for prostate cancer detection. Due to the strong limitations related to the predictive deconvolution scheme delineated in the previous chapter, we were motivated to further improve those results by designing a more sophisticated deconvolution framework, which will be the subject of this chapter.

As mentioned, deconvolution in medical imaging is commonly employed in the only purpose of a visual quality improvement, so to provide the physician with better contrasted and resolved data, suitable for easier interpretation. As long as an image quality improvement is only concerned, standard deconvolution schemes exploit simplified models for the tissue reflectivity, prevalently Gaussian or Laplacian, which are typically chosen so to produce the most appealing images. The statistics of the restored solutions are hence badly biased by these prior assumptions, and are not well suited for characterization.

In this context, a novel deconvolution method for ultrasound images is presented in this chapter, expressively designed for an improved tissue characterization. The algorithm is derived on the base of a non-standard, more general statistical model for the tissue response, defined by a GGD [45]. By means of two distinct parameters, called scale and shape parameter, this distribution allows sequences of arbitrary energy and sparsity to be generated, and is therefore adequate for providing an accurate description to the most general tissue structures. Deconvolution is then tackled as a Maximum a Posteriori (MAP) estimate, in which tissue reflectivity is restored along with an estimate of the associated scale and shape parameters. An Expectation Maximization (EM) framework is designed to tackle the estimation problem.

The remaining of the chapter is structured as follows. In §5.1 the proposed Generalized Gaussian model for the tissue reflectivity as well as the solution of the related optimization problem are introduced, and the substantial elements of difference with respect to traditional models are explained. The algorithm was validated on phantoms containing identical scatterers at different concentration. Hereto §5.2 is dedicated to the description of the experimental setup and of the processing flow: feature set, feature-extraction and classification paradigms are discussed. The obtained results are presented §5.3 and the superiority of the proposed deconvolution scheme over traditional ones is motivated. Conclusive considerations and hints on future research directions are provided in §5.4.

5.1   REFLECTIVITY MODEL AND OPTIMIZATION SCHEME

As stated, deconvolution is not employed here for a mere visual enhancement, but rather to access higher quality information on the tissue structure for characterization. In this sense, any prior assumptions on the reflectivity distribution could badly bias the statistics of the restored signal, thereby degrading its information content, and should therefore be avoided. To this purpose, a more flexible model for the tissue reflectivity is needed, allowing to accurately describe the most general scattering conditions. In this sense, the GGD [45] is particularly well suited. By means of two distinct parameters, called *scale* and *shape* parameter, this distribution allows sequences of arbitrary energy and sparsity to be generated. In particular, low values of shape parameter correspond to sparser sequences while high values correspond to dense ones (see Figure 5.1).

Since energy is directly connected to the scatterers average strength and sparsity is intuitively related to their concentration, any kind of tissues could theoretically be well modeled as a GGD process. Moreover, the suitability of the GGD for providing accurate fitting to the RF signal histograms, documented in [14], further proves the adequacy of this model for describing tissue reflectivity as well. Indeed, the linear filtering operations associated to the acquisition process does not modify the qualitative data distribution, a part of decreasing its sparsity [67]. This fact is qualitatively illustrated in Figure 5.2

Figure 5.1: Examples of random sequences drawn from zero-mean unit variance GGD, having unit variance and shape parameter ξ equal to $\xi = .5$ (a), $\xi = 1$ (b) and $\xi = 2$ (c). Let's note how for low values of ξ tissues with isolated reflectors are well modeled. For values close to 2, instead, the model is well suited for describing diffusive regions.

The GGD pdf writes as:

$$p(x_i) = a \exp\left(-\left|\frac{x_i}{b}\right|^{\xi} \right) \tag{5.1}$$

where $a = \xi/(2b\Gamma(1/\xi))$, $b = \sigma_x \sqrt{\Gamma(1/\xi)/\Gamma(3/\xi)}$ is the scale parameter, σ_x is the standard deviation, ξ is the shape parameter and $\Gamma(\cdot)$ is the Gamma function [20]. Note that Gaussian and Laplacian distributions belong to the GGD family: in particular, they are obtained from (5.1) by substituting $\xi = 2$ and $\xi = 1$ respectively.

Under the GGD assumption, the log-posterior becomes:

$$L = -\frac{1}{2\sigma_n^2}\|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 - \lambda\|\mathbf{x}\|_{\xi}^{\xi} + N \cdot a \tag{5.2}$$

Figure 5.2: Distribution of a synthetic reflectivity generated as a Generalized Gaussian process in (a), distribution of the correspondent RF signal generated by filtering the reflectivity with an experimental PSF in (b). In blue the histograms and in red the GGD obtained from their fitting. Note that for both signals fitting is accurate. The shape parameter before and after filtering is 0.3 and 1.6 respectively.

where $\lambda = b^{-\xi}$, $\|\mathbf{x}\|_p = (\sum |x_i|^p)^{1/p}$ denotes the l^p-norm and N is the total number of samples in the image.

It must be noted here that classical restoration techniques, (4.18) and (4.23), are special cases of the maximization of (5.2), in which the sparsity parameter ξ is *a priori* fixed to 2 and 1 respectively. Otherwise, in the present work, any prior knowledge on tissue sparsity is avoided, and correspondingly an EM scheme is designed in which this information is jointly restored along with the tissue response.

The EM algorithm is an iterative procedure for the solution of ML or MAP statistical estimation problems whose solution is not analytically tractable [68]. Each iterations of the EM algorithm consists in an Expectation step (E-step), in which the expectation function Q is defined, and a Maximization step (M-step), in which the new estimate of the interesting quantity is produced. It can be proved that EM algorithm monotonically converges toward a local maximum of the log-likelihood or log-posterior function.

In [38] an EM framework is derived for the solution of MAP image restoration problems in which the prior distribution $p(\mathbf{x})$ can be expressed as a Gaussian Scale Mixture (GSM). Since the GGD belongs to this family of distributions, that framework can be adopted here to tackle the deconvolution task at hand. As-

suming a shape parameter value ξ, the E-step and M-step write respectively as:

$$\text{E-STEP} \quad Q(\mathbf{x}, \hat{\mathbf{x}}_k) = -\frac{\|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2}{2\sigma_n^2} - \frac{1}{2}\mathbf{x}^T\mathbf{D}_k\mathbf{x}, \tag{5.3}$$

$$\text{M-STEP} \quad \hat{\mathbf{x}}_{k+1} = \underset{\mathbf{x}}{\text{argmax}}\{Q(\mathbf{x}, \hat{\mathbf{x}}_k)\} = (\sigma_n^2\mathbf{D}_k + \mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{y} \tag{5.4}$$

where $\mathbf{D}_k$ is a diagonal matrix of size $N \times N$ with entries $[\mathbf{D}_k]_{ii} = \lambda\xi/|\hat{x}_i^k|^{2-\xi}$, and $\hat{x}_i^k$ denotes the ith sample of the vector $\hat{\mathbf{x}}_k$.

The iterative scheme defined by (5.4) guarantees the maximization of the log-posterior $L(\mathbf{x}, \xi)$ defined in (5.2) with respect to $\mathbf{x}$, in such a way that $L(\hat{\mathbf{x}}_{k+1}, \xi) > L(\hat{\mathbf{x}}_k, \xi)$. As mentioned, in the present work, the shape parameter ξ is also a quantity to be estimated, and therefore (5.2) must be optimized w.r.t. $\mathbf{x}$ and ξ jointly. The simplest solution is to introduce a parameter update step after (5.4):

$$\hat{\xi}_{k+1} = \underset{\xi}{\text{argmax}}\{L(\hat{\mathbf{x}}_{k+1}, \xi)\} \tag{5.5}$$

As known, (5.5) is solved by the ML estimate of the shape parameter ξ from the sequence $\hat{\mathbf{x}}_{k+1}$. For a possible solution see §3.3.1.

The complete algorithm is reported in Table 5.1. It should be noted here that, because of the ML parameter estimation step in (5.5), the global framework is no longer rigorously an EM and therefore convergence is not theoretically guaranteed. Nevertheless, (5.5) ensures that $L(\mathbf{x}, \hat{\xi}_{k+1}) > L(\mathbf{x}, \hat{\xi}_k)$ and monotonic growth of the log-posterior is preserved. Besides, convergence is observed in all practical experimentations.

The most time consuming step in the proposed framework is the signal update step in (5.4), which is equivalent to solving the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$, defined by $\mathbf{A} = (\sigma_n^2\mathbf{D}_k + \mathbf{H}^T\mathbf{H})$ and $\mathbf{b} = \mathbf{H}\mathbf{y}$. Since the direct solution is computationally unfeasible, an iterative scheme must be adopted; specifically the Conjugate Gradients (CG) algorithm in Table 4.2 is here adopted. In this context a very common strategy for improving the convergence rate of the CG algorithm is preconditioning [39], which consists in pre-multiplying the computed residuals for a suitable approximation of the inverse $\mathbf{P}$ of the system matrix $\mathbf{A}$. In particular, in this work we adopt Jacobi pre-conditioning, cf. Appendix B.

Table 5.1: MAP restoration under GGD prior

Initialization: $\hat{\mathbf{x}}_0 = \mathbf{y}$, $k = 0$

---

repeat

$\qquad \hat{\xi}_k = \arg\max_{\xi} L(\hat{\mathbf{x}}_k, \xi)$,

$\qquad d_i^k = \lambda \hat{\xi}_k / |\hat{x}_i^k|^{2 - \hat{\xi}_k}$,

$\qquad \mathbf{D}_k = \text{diag}(\mathbf{d}_k)^1$; $\quad \mathbf{d}_k = [d_1^k, \cdots, d_N^k]$

$\qquad \hat{\mathbf{x}}_{k+1} = (\sigma_n^2 \mathbf{D}_k + \mathbf{H}^T \mathbf{H})^{-1} \mathbf{H} \mathbf{y}$

$\qquad \boxtimes = |L(\hat{\mathbf{x}}_{k+1}, \hat{\xi}_k) - L(\hat{\mathbf{x}}_k, \hat{\xi}_k)| / |L(\hat{\mathbf{x}}_k, \hat{\xi}_k)|$

$\qquad k = k + 1$

until: $\boxtimes < \tau_{\text{EM}}$

---

¹ $diag(\mathbf{a})$ denotes the diagonal matrix with vector $\mathbf{a}$ as entries.

In order to make (5.2) differentiable in the origin, a smooth-convex approximation of $|x|$ is often adopted [38, 31]. Possible choices are e.g. $\eta_1(x) = \{|x/\alpha| - \log(1 + |x/\alpha|)\}$ and $\eta_2(x) = \sqrt{x^2 + \alpha}$, with $\alpha \ll 1$ [43]. In this work, we adopt the approximation defined by η_2. Since the weight of the approximation is clearly dependent on the variance of the signal, then we propose here to take $\alpha = \gamma \sigma_x^2$, with γ a small positive proportionality constant.

In the remaining of this section other important issues are addressed. In particular the slight modifications implied by considering the complex envelope signal are described, together with the adopted techniques for the estimation of the system PSF along with signal and noise variances σ_x^2 and σ_n^2.

5.1.1 IQ Signal

In this work we do not process directly the RF signal, but its baseband equivalent, i.e. the IQ signal, cf. §2.4.1. As seen, because of the linearity of the demodulation process, the linear model (4.12) is still satisfied in the IQ domain [7], where $\mathbf{y}$, $\mathbf{x}$ and $\mathbf{v}$ represent now the IQ image, the complex reflectivity and the complex noise respectively, while $\mathbf{H}$ is the complex blurring matrix. Let's remark here that Hermitian transpose $\mathbf{H}^\dagger$ must replace $\mathbf{H}^T$ as long as complex valued signals are concerned.

The usefulness of working with the IQ signal is twofold: at

first, since the RF signal is narrowband, then the IQ signal can be subsampled without aliasing problems, which considerably reduces the computational effort; moreover, the downsampled signal has a wide relative bandwidth, which highly improves the conditioning of **H**.

In order to formalize the deconvolution framework for the IQ image, the tissue prior $p(\mathbf{x})$ has to be defined for complex variables, hereto we introduce an extension to the GGD which writes:

$$p(x_i) = \alpha \exp\left(-\left|\frac{x_i}{\beta}\right|^{\xi} \right) \tag{5.6}$$

where x_i is the i-th sample of the complex reflectivity and parameters α and β are defined in Appendix C. Thanks to the new model in (5.6), the deconvolution scheme in Algorithm 5.1 is preserved, with the only exception that parameters α and β must now be used instead of a and b in the definition of L in (5.2).

5.1.2 PSF estimation

As mentioned above and discussed in §4.4, deconvolution of ultrasound images is a blind problem. The common way to deal with it is to estimate a suitable PSF first, followed by using the resultant estimate to solve the deconvolution problem in a nonblind manner. Here we adopt the generalized homomorphic technique presented in §4.4. Figure 5.3 provides an example of PSF estimated from the phantom data used in the present study, which will be described in §5.2; signals were acquired with a 30 MHz transducer and sampled at 500 MHz.

5.1.3 Estimation of σ_n and σ_x

In this work we adopted the wavelet based noise variance estimator presented in [35]. It is based on the reasonable assumption that the finest-scale wavelet coefficients of **y** are mainly associated to noise. By calling **w** these coefficients the noise standard deviation estimate is then:

$$\sigma_n = \frac{median(|\mathbf{w}|)}{\sqrt{\ln 4}}. \tag{5.7}$$

Concerning the estimate of the signal variance σ_x^2, this is straightforward if we assume that all the PSFs have unit energy, so that the variance of **Hx** is approximately equal to the variance
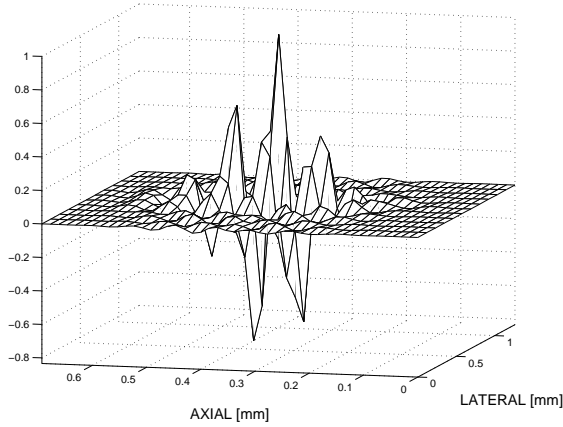
Figure 5.3: Example of estimated PSF.

of $\mathbf{x}$ and then $\sigma_x^2 \approx \sigma_y^2 - \sigma_n^2$, where $\sigma_y^2 = \|\mathbf{y}\|_2^2 / (2N)$. The normalization for $2N$ in the definition of σ_y^2 is due to the fact that complex variables are considered, in the real case a normalization for N should be used.

5.2   Materials and methods

We tested the performance of the proposed deconvolution algorithm on experimental data from several tissue-mimicking phantoms having a particular scatterer concentration. The aim of the study was to evaluate the improvement of medium characterization after deconvolution.

5.2.1   *Experimental Setup*

Ultrafine polyamide particles (Orgasol, Arkema, France) of diameter 10 ± 2 μm (2001 EXD NAT 1) were used as scatterers. The tissue-mimicking phantoms were prepared by mixing a specific concentration of Orgasol particles with distilled water and 1% w/w ICI synperonic N surfactant in order to improve the particle wetting. The Orgasol particles had a density of 1030 kg/m^3 so that the first order Born approximation is acceptable and the linear model is valid. Seven experiments were performed with seven mixtures having Orgasol concentrations 0.25%, 0.5%, 0.75%, 1%, 6%, 12% and 24%. Low concentration ranging from 0.25% to

Table 5.2: RMV-707B probe (Visualsonics Inc.)

PARAMETER		VALUE
center frequency	f_0	30 MHz
transducer radius	R	7 mm
focal distance	D_f	12.7 mm
axial resolution	ry	55 μm
lateral resolution	rx	115 μm
azimuthal resolution	ry	115 μm

6% constituted random medium and the two other concentrations (12% and 24%) mimicked dense medium in order to study more packed distributions of particles, as performed in [69]. A magnetic agitator was employed in order to preserve the solution homogeneous throughout the acquisition.

The RF signal was acquired with a high resolution ultrasound scanner (Vevo 770, Visualsonics, Toronto, Canada), equipped with the single element transducer RMV-707B, whose parameters are reported in Table 5.2. Acquisitions were performed using the US equipment available at the Animage platform (Cermep-Lyon-France). Since the transducer had a focal depth of 12.7 mm, a 6 mm thick layer of agar gel (Prolabo) at 3% was interposed between the probe and the phantom in order to place the probe focus in correspondence with the region of interest. The RF signal was suitably amplified (Pulse/Receiver Sofranel Panametrics 5052pr) and acquired at a sampling frequency of 500 MHz with 8 bit resolution (Gagescope, model CS11G8-1, Acquisys, France). For each phantom 3 acquisitions were made. Data was then processed using MATLAB (R2008b, The MathWorks). A schematic of the experimental setup is reported in Figure 5.4.

Each image was demodulated, downsampled by a factor 8, and deconvolved using Wiener filtering, Laplacian prior, and Generalized Gaussian prior. After deconvolution, data was then modulated back to the transducer center frequency and upsampled so to have an image size identical to the one of unprocessed frames. Each image was then subdivided into 48 non-overlapping Regions of Interest (ROIs) of 405 pixels in the axial direction and 29 in the lateral one, corresponding roughly to 1×3 mm². This size was equivalent to three times the extent of an estimated PSF
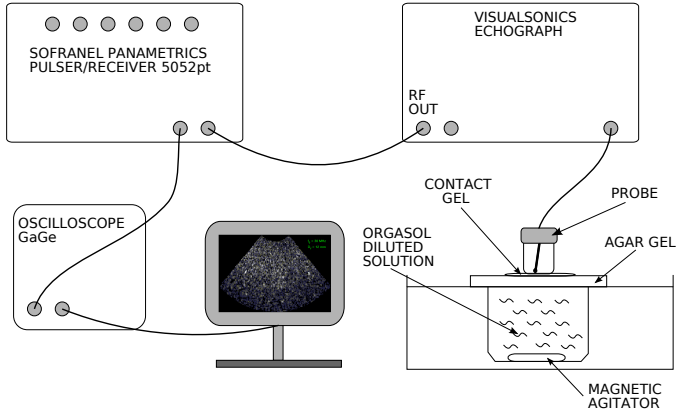
Figure 5.4: Experimental setup used in the experiments.

in each direction, and the number of included samples was verified to be large enough to make biasing effects in the features computation largely negligible.

Considering the 3 images per concentration, the total number of ROIs related to a single concentration was then equal to 144. Several features were extracted on each ROI of the deconvolved and unprocessed data set and classification was performed in the feature space.

5.2.2 Deconvolution

Each deconvolution algorithm was implemented with the Expectation Maximization framework of Table 5.1. In the Wiener and Laplacian cases, a fixed value of shape parameter was adopted and correspondingly the ML update step in (5.5) was skipped. The algorithm was then stopped when a relative decrease in the log-likelihood $\tau_{\text{EM}} = 10^{-4}$ was reached. Such a threshold was normally reached in about 20 iterations.

On sub-sampled data, whose size is 210×382 pixels, each iteration was accomplished in roughly 6 seconds on a laptop equipped with an Intel Core i5 2.27 GHz processor and 4 GB of RAM.

5.2.3 Feature Set

A six-features approach was adopted, defined by the following parameters: the two GGD parameters computed from the RF

Figure 5.5: Processing flowchart

signal, the two Nakagami parameters computed from the envelope and the two Nakagami parameters computed from the B-mode image.

5.2.4 Feature Extraction

The six adopted features are basically associated with two physical properties, echogenicity and density, of the insonified medium. Although considering different data representations (RF, uncompressed envelope and B-mode) ensures a degree of complementarity in the information, such a disparity may otherwise lead to some correlation between features of the same kind.

In this context, feature extraction is a common technique in Pattern Recognition, which combines observed data into a reduced dimensionality space, in which redundancies are minimized and the informative content is preserved [70]. Besides, dimensionality reduction is desirable for computational reasons as well.

Standard solutions to the feature extraction problem are given by Principal Component Analysis (PCA) [70, Ch.9, pp.319–329] and Linear Discriminant Analysis (LDA) [70, Ch.4], which represent an un-supervised and a supervised approach respectively. When a trustful ground-truth is given, like in the present case, where the actual concentrations are known, the LDA algorithm effectively embeds this knowledge and its performances are superior to the ones of the PCA. Therefore, in this study, Linear Discriminant Analysis (LDA) was adopted for feature extraction and the number of extracted features was chosen equal to 2.

5.2.5 Classification

Both linear and non-linear classification schemes were evaluated. Non linear methods were found to be largely preferable in

terms of classification errors and computation time, especially for a number of classes higher than 2. In particular, a Support Vector Machine (SVM) with Radial Basis Functions as kernels was used [70, Ch.5].

Feature extraction and classification part was performed with the Statistical Pattern Recognition Toolbox developed for Matlab [71]. A flow chart of the whole processing scheme is reported in Figure 5.5.

It is worth noting that the main scope of this study was expressively to evaluate the deconvolution contribution to a classification task and not to provide a standard flow for medium characterization through ultrasound. For this reason, the classification scheme proposed in this section might be sub-optimal, but it is nevertheless sufficient to give important and trustful insights on the effect of a deconvolution step.

5.3   Results

5.3.1   Data classification

Two kind of studies were made. In the first set of trials, four datasets were separately considered, consisting respectively in the acquired images before and after deconvolution with the three techniques, i.e. Wiener filtering, Laplacian prior and the proposed scheme. Each data set was then classified with the processing flow previously described and the misclassification error measured.

For each classification task, LDA and SVM were trained with 50% of the ROIs, randomly picked from the entire data set. The remaining 50% was adopted as testing set. For any data type, cross-validation was used to find the optimal kernel parameters, which were then used in the SVM. In order to avoid biasing in the results due to the specific choice of training and testing sets, 20 independent trials were made for each experiment.

In the ideal case, all the concentrations should be simultaneously discriminated, leading to a seven classes classification task. Since classification becomes more challenging as the number of classes increases, several tasks were considered, starting for the simplest binary case, to the most complicated and complete case with 7 classes. The misclassification error was then computed: the complete set of results is reported in Table 5.3.

From Table 5.3 several conclusions can be drawn:

Table 5.3: Classification error on 20 independent trials (mean value ± standard deviation)

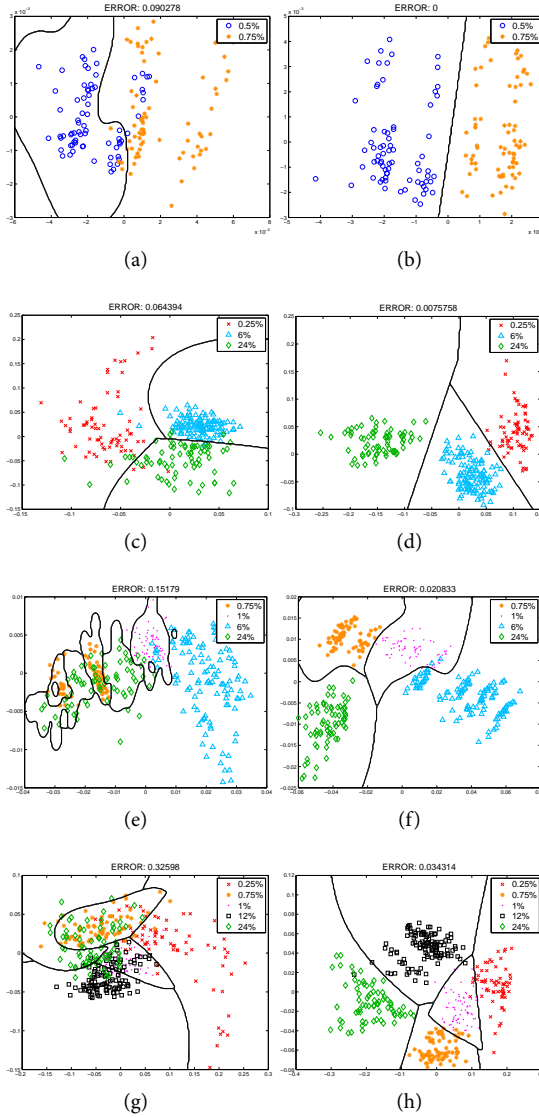|	CLASSES	CLASSIFICATION ERROR			
#	CONCENTRATIONS	ORIGINAL	WIENER	LAPLACE	GGD
2	$[0.5\%, 0.75\%]$	0.1 ± 0.04	0.17 ± 0.04	0.07 ± 0.02	$\mathbf{0.001 \pm 0.004}$
3	$[0.25\%, 6\%, 24\%]$	0.07 ± 0.02	0.03 ± 0.01	0.02 ± 0.01	$\mathbf{0.008 \pm 0.01}$
4	$[0.75\%, 1\%, 6\%, 24\%]$	0.20 ± 0.04	0.22 ± 0.02	0.09 ± 0.02	$\mathbf{0.02 \pm 0.01}$
5	$[0.25\%, 0.75\%, 1\%, 12\%, 24\%]$	0.32 ± 0.02	0.25 ± 0.03	0.08 ± 0.02	$\mathbf{0.03 \pm 0.01}$
6	$[0.25\%, 0.5\%, 1\%, 6\%, 12\%, 24\%]$	0.35 ± 0.02	0.44 ± 0.01	0.11 ± 0.02	$\mathbf{0.08 \pm 0.01}$
7	$[0.25\%, 0.5\%, 0.75\%, 1\%, 6\%, 12\%, 24\%]$	0.42 ± 0.02	0.38 ± 0.02	0.19 ± 0.01	$\mathbf{0.16 \pm 0.01}$

Figure 5.6: Scatter plot relative to the testing set before ((a),(c),(e),(g)) and after ((b),(d),(f),(h)) deconvolution with a GGD prior. The axes correspond to the two features extracted via LDA. Each row corresponds to a particular number of classes, specifically 2, 3, 4 and 5. The black lines correspond to the boundaries of the classes identified by the SVM.

- Wiener filtering did not bring any contribution to the classification task. Sometimes, even worst results were observed with respect to the unprocessed case. This reveals that Wiener filtering, which is by far the most common restoration technique, is not a relevant tool for improving characterization performances. This fact is probably connected to the widely documented over-smoothing associated to Wiener filtering [45, 31], in particular when low concentrations are considered, where filtering of the peaky signal components, due to the presence of isolated scatterers, prevents an effective estimate of their concentration.

- Differently, deconvolution with Laplacian prior was found to positively affect the classification task. This is particularly evident for a high number of classes, where a classification error reduction of up to 20% is associated to this deconvolution procedure. Such a behavior is explained by the fact that a Laplacian model is more effective than Gaussian to provide realistic representations of general reflectivity structures, in particular when isolated reflectors are superimposed to distributed diffusive scattering centers.

- Finally, the proposed deconvolution scheme was found to ensure by far the best classification performances. As in the Laplacian case, the contribution became more incisive as the number of classes increased. In particular, in the 4 classes case, the classification error related to the proposed framework was still at 2%, while it was four times bigger for the Laplacian case and ten times for rough data. Moreover, in the most challenging case, in which all concentrations were simultaneously considered, the classification error was at the satisfactory rate of 16%, while it reached an unacceptable 42% rate on unprocessed data.

In Figure 5.6 several scatter plots are displayed. Each point represents the position of a single ROI of the testing set in the extracted features space; a different marker is associated to each concentration. Black contours represent the inter-classes boundaries computed by the SVM classifier. From the figure it is evident that classes are better clustered and consequently more easily separable after deconvolution; a much faster and more accurate classification is therefore allowed.

Figure 5.7: Generalized Gaussian shape parameter histograms associated to four different concentrations. In 5.7(a) the unprocessed case is illustrated. Images 5.7(b), 5.7(c) and 5.7(d) report the situation after deconvolution is applied, by using Wiener filtering, laplacian prior and the proposed method respectively.

As the proposed deconvolution algorithm was derived in order to adaptively restore an estimate of the true reflectivity shape parameter, it was interesting to examine how concentrations were separated accordingly to this parameter only. This is illustrated in Figure 5.7, where a four classes case is observed. When the unprocessed signal is considered, the system PSF effect implies a mixing of single scatterers contributions, and correspondingly the density information, explained by the shape parameter, results highly deteriorated. On the contrary, the proposed framework allows to retrieve an improved estimate of this important signature of the insonified medium. Such a behavior is not encountered when traditional schemes are applied, which is a direct consequence of the biasing effect associated to these methods.

Figure 5.75.7(d) shows clearly how the direct correlation ex-

Figure 5.8: Boxplot representation of data in Figure 5.75.7(d). On each box the central mark is the median and the edges are the 25th and 75th percentiles, the whiskers are the 5th and 95th percentiles.

isting between scatterers concentration and shape parameter is restored after deconvolution (see also Figure 5.8).

5.3.2   Feature ranking

Data deconvolved with the proposed algorithm were shown to allow for a more precise characterization. A further study was then made to evaluate the effective contribution to the classification task of features computed before and after deconvolution.

To this purpose, a binary classification task was considered, defined by the two classes corresponding to 0.5% and 0.75% concentrations. The 6 considered features were then computed on unprocessed and deconvolved data separately, and collected into an 12 elements feature set. Features were then ranked according to a distance-based criterion. The total number of classified samples in this study was equal to 576; this number derives from counting 48 ROIs per image, 3 images per concentration, 2 concentrations and the fact that unprocessed and deconvolved data are jointly considered.

In particular, Sequential Forward Selection (SFS) criterion was used for the ranking [70, Ch.9, pp.315]. SFS is a bottom-up search procedure that adds one feature at a time until the complete set is reached. Specifically, the feature for which the inter-classes distance J is maximized is added to the feature set

Table 5.4: Features Ranking

Type	Feature Name	Signal	J_{1D}
D	GGD shape	RF	▇▇▇▇▇▇▇▇ 3.50
D	Nakagami scale	Bmode	▇▇▇▇ 1.62
D	Nakagami shape	Bmode	▇▇▇ 1.26
U	Nakagami scale	Bmode	▇ 0.27
D	Nakagami scale	envelope	▇▇ 0.95
U	Nakagami scale	envelope	▇ 0.42
U	Nakagami shape	envelope	▇ 0.18
D	GGD scale	RF	▇▇ 0.73
U	GGD scale	RF	▇ 0.41
U	Nakagami shape	Bmode	▇ 0.44
U	GGD shape	RF	▇ 0.28
D	Nakagami shape	envelope	▇▇ 0.70

obtained at the previous step, where J was defined in this study as the Mahalanobis Distance:

$$J = \left(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\right)^T \left(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2\right)^{-1} \left(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\right) \qquad (5.8)$$

where $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$ ($i \in \{1, 2\}$) denote the intra class mean value and covariance matrix respectively.

The final ranking is reported in Table 5.4, where features are sorted from the first classified to the last. Features computed after deconvolution are tagged with a "D" while a tag "U" is associated to unprocessed data. The fourth column visualizes the Mahalanobis Distance J_{1D}, corresponding to the selection of a single-attribute feature set.

From the ranking results provided in Table 5.4 several observations can be made:

- The first positions in the ranking were occupied by features computed on deconvolved data. This proves that a more characterizing information on the insonified medium can be restored thanks to the proposed deconvolution framework. In particular, GGD shape parameter placed first in the ranking, which confirms the high significance of this feature after deconvolution, cf. Figure 5.7 and Figure 5.8;

- If sorting by the value of J_{1D} is made, the first half of the ranking happens to be filled by "D"-tagged features exclu-

sively. That's to say all features improve their individual discrimination capability after deconvolution.

In order to have more exhaustive evaluation of the deconvolution effect, the same ranking procedure was applied to different binary classification tasks. In all those cases, the first positions in the ranking were systematically occupied by features deriving from deconvolved images, which proves the generality of the latter considerations.

5.4 CONCLUSION

Ultrasound tissue characterization has become a major research subject. This is motivated by the fact that from the backscattered echo signal important structural information on the insonified tissue can be inferred, such as acoustic impedance, scatterer density and distribution pattern. As similar information can discriminate between healthy and diseased tissue, this principle can be exploited within a clinical setting.

As the ultrasound echo suffers from the degradation effect due to the system PSF, deconvolution could be effectively employed as a pre-processing step to mitigate this effect and restore much representative information on the true tissue response. In this context, this chapter presents a novel deconvolution technique for medical ultrasound, which contains a major conceptual element of novelty over the existing related literature. Indeed, the goal here is not a mere visual quality improvement, but rather the restoration of a more informative tissue representation, exploitable for an improved characterization.

The performance of the algorithm was evaluated on tissue mimicking phantom studies, which revealed remarkable improvements in classification accuracy associated to the implementation of a deconvolution pre-processing step, quantifiable in a 20% misclassification error reduction when multiple classes were considered. The contribution of other popular deconvolution schemes was also measured for comparison, and the superiority of the proposed scheme was correspondingly attested.

When the GGD shape parameter was adopted as the only feature, different populations were observed to be clearly sorted by concentration in the feature space (Figure 5.8). This effect was particularly evident after the proposed deconvolution framework was applied (Figure 5.7). This interesting property could be ex-

ploited for ultrasound based scatterer concentration sensing and opens the way to further studies in that direction.

Obtained results encourage additional studies on the suitability of the proposed algorithm on clinical data. In this sense, few considerations have to be made. At first, the algorithm exploits a stationary model of tissue reflectivity, represented by the fact that constant shape parameter and variance are assumed on the whole image. As a similar assumption is well satisfied on the homogeneous phantoms realized for the validation, otherwise it is largely violated on images of biological tissues, where multiple tissues with different acoustical properties are simultaneously imaged. In order to deal with these general cases, a parametric image of the shape parameter could tentatively be considered instead of a constant value, to be iteratively updated within the optimization flow. Moreover, complementary studies on homogeneous tissue-mimicking phantoms containing polydisperse suspension of scatterers could be done. This will allow to better mimic situations of clinical interest, like cellular size variance during cell death [72], or the simultaneous presence of glandular acini (100 μm diameter) and cell nuclei (14 μm diameter) in breast tumor [73].

A second issue involves the computational complexity. In the current, unoptimized Matlab implementation, every frame is processed in roughly two minutes, which is clearly too much for assistance during real-time intervention. In this context substantial speed-up could be achieved with modern GPU computing techniques. In particular, the intrinsic parallelism of the algorithm, established by the piecewise constant approximation of the PSF, could be efficiently handled within GPU architectures, for a fastened execution.

III - Segmentation in Echocardiography

Summary

CARDIOVASCULAR disease (CVD) is the main cause of long-term disability and death for both genders within industrialized nations. According to the World Health Organization CVD causes over 4 million deaths in Europe and over 2 millions in the European Union (EU), corresponding to nearly half of all deaths, cf. Figure B. Besides heavily human, CVD involves major economic costs as well. In particular, the overall economic cost of CVD has been estimated to be over 192 euro billions in EU in 2006. These staggering numbers clearly motivate the intense research activity carried out worldwide dedicated to the prevention, early diagnosis and cure of this class of diseases.

Modern technology provides a number of tools to access myocardial function. Echocardiography, in particular, has evolved as a well-established imaging technique for the noninvasive evaluation of myocardium abnormalities. Besides the obvious aspects connected to safety, low cost and portability, echocardiography alone offers the possibility of real time observation of ventricular wall motion and deformation, which are fundamental for assessing the extent of possible myocardial ischemia and infarction. In clinical practice, the analysis mainly relies on visual inspection or manual measurements, off line performed by experienced cardiologists. Manual methods are tedious and time-consuming, and visual assessment leads to qualitative and subjective diagnoses that suffer from considerable inter and intra observer variability.

The main clinical application concerns the delineation (segmentation) and tracking of the myocardium muscle from US scans. Indeed from the position of the muscle boundaries throughout the cardiac cycle several indexes of its function can be inferred, as ejection fraction, ventricle volume and wall thickness. For these reasons automating the segmentation of echocardio-

Respiratory disease 7%
Injuries and poisoning 12%
Other cancer 11%
All other causes 17%
Lung cancer 6%
Colo–rectal cancer 2%
Stomach cancer 2%
Other CVD 11%
Coronary heart disease 21%
Stroke 11%

(a)

Respiratory disease 6%
Injuries and poisoning 5%
All other causes 18%
Other cancer 9%
Breast cancer 3%
Lung cancer 2%
Colo–rectal cancer 2%
Stomach cancer 1%
Coronary heart disease 22%
Other CVD 15%
Stroke 17%

(b)

Figure B: Deaths by cause for men (a) and women (b) in the last available year in Europe. Data form [74].

graphy images is highly desirable but also challenging. Several algorithms have been proposed in literature, documenting successful results for the detection of the *endocardium*, i.e. the innermost layer of tissue surrounding the ventricular cavity. Otherwise, for what concerns *epicardial* contour, no satisfactory solution has been provided yet, in particular when short-axis acquisitions are considered. The motivation is that in this kind of acquisition epicardial fibers happen to be collinear with the ultrasonic beam, and correspondingly the echo returned by these

regions is extremely week. The immediate consequence is that epicardial contours appear highly heterogeneous and discontinuous on ultrasound frames, which makes their automatic detection particularly challenging.

In this context, we developed a novel semi-automatic framework for myocardium segmentation, designed for allowing an efficient detection of endo- and epicardial contours of the left ventricle in short-axis acquisitions. Level-set technique is adopted. Statistics of the echo signal along with geometrical considerations on the myocardium expected shape are originally exploited for constraining the curve motion towards anatomically meaningful shapes. A localizing strategy recently proposed in literature was adopted, ensuring robustness against low contrast and missing contours.

This project has been developed during a one year internship within the ultrasound equipe of the Centre de recherche en imagerie médicale (CREATIS-LRMN) of Lyon - France, under the supervision of Prof. Olivier Bernard and Prof. Olivier Basset. This work has resulted in 4 publications up to the present time [75, 76, 77, 78]. Due to their general formulation, the developed tools have proved their suitability for a much wider class of images which exceeds medical ultrasound [76, 78].

This part of the thesis is structured as follows: in Chapter 6 the reader is supplied with the necessary theoretical background on level-set segmentation. Starting from standard region-based approaches, general ways for including localization and shape information are presented as well. In Chapter 7 the basics of heart morphology and functioning are described. In Chapter 8 the proposed segmentation framework is described and a validation on clinical data is presented.

Level-set segmentation

FROM THE pioneering work of Kass *et al.*[79], variational techniques have been widely exploited for image segmentation purposes. Due to their extreme generality and flexibility, these methods can be suitably specialized in order to deal with data of the most general nature: medical, biological and astronomical are some examples. In synthesis, active contours segmentation consists in the displacement of a curve on the image plane. This motion pursues the attachment to the boundaries of the interesting target and is normally governed by the minimization of an *a priori* defined cost functional. An elegant and efficient formalization of active contours segmentation is provided by level-sets [80, 81], which are the object of this chapter. The interested reader may find almost exhaustive reviews on level-sets in [19, 82].

The rest of the chapter is structured as follows. In §6.1 the general level-set formalism is introduced. In §6.2 the standard region-based approach is illustrated, while in §6.3 a more recent, localized solution is discussed. In §6.4 and §6.5 it is shown how statistical and geometrical information can be exploited in the segmentation process. In §6.6 some concluding considerations are made.

6.1 GENERAL

Let $I : \Omega \rightarrow \mathbb{R}^m$ be a given image, where $\Omega \subset \mathbb{R}^d$ and m is the cardinality of the feature space in which the image exists. For instance is $m = 1$ for gray-scale data and $m = 3$ for colored ones. In the level-set formalism, the evolving interface $\Gamma \subset \Omega$ is rep-
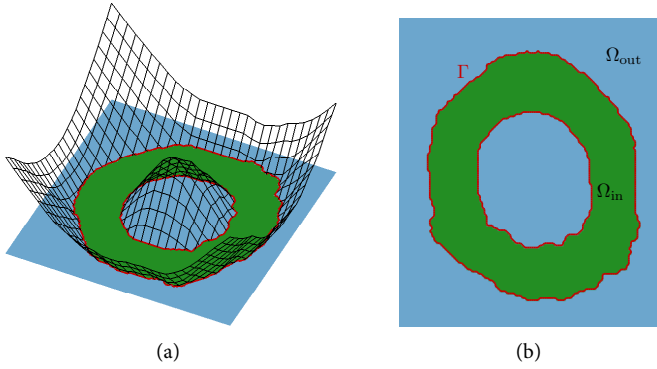
Figure 6.1: Example of level set function (a) and related partition of the image space Ω (b)

resented as the zero level-set of a Lipschitz-continuous function $\phi : \Omega \to \mathbb{R}$, called *level set function*, that satisfies:

$$
\begin{cases}
\phi < 0, & \forall \mathbf{x} \in \Omega_{\text{in}} \\
\phi > 0, & \forall \mathbf{x} \in \Omega_{\text{out}} \\
\phi = 0, & \forall \mathbf{x} \in \Gamma
\end{cases}
\tag{6.1}
$$

where $\mathbf{x}$ is a pixel coordinate and $\{\Omega_{\text{in}}, \Omega_{\text{out}}, \Gamma\}$ is a partition of Ω. An example of level set function is reported in Figure 6.1, along with the partition of its domain.

The classical problem of segmenting the target from the background is handled by the evolution of one level-set function driven by the minimization of a specific energy criterion E; its steady state partitions the image into two regions that delimit the boundaries of the object to be segmented:

$$
\hat{\phi} = \underset{\phi}{\text{arginf}}\{E(\phi)\}
\tag{6.2}
$$

The Euler-Lagrange equation associated to (6.2) is:

$$
\frac{\partial E}{\partial \phi} = 0
\tag{6.3}
$$

where $\partial/\partial\phi$ represents the first variation of the functional E with respect to the function ϕ, defined as:

$$
\frac{\partial E}{\partial \phi} = \lim_{t \to 0} \frac{1}{t}\{E(\phi + t\psi) - E(\phi)\}
\tag{6.4}
$$

where ψ is a test function of the same type of ϕ.

With these definitions, the segmentation flow associated to (6.2) writes as:

$$\frac{\partial \phi}{\partial \tau} = -\frac{\partial E}{\partial \phi} \tag{6.5}$$

where the artificial time parameter τ has been introduced. In level-set literature (6.5) is normally referred as the *level set equation*. Using first-order forward differencing for the numerical computation of $\partial/\partial\tau$, the n^{th} step of the segmentation flow is in practice implemented as:

$$\phi^{n+1} = \phi^n + \Delta\tau \cdot \left.\frac{\partial E}{\partial \phi}\right|_{\phi^n} \tag{6.6}$$

where the time step $\Delta\tau$ influences the curve displacement velocity, and is upper limited by the Courant-Friedrich-Levy (CFL) condition, that ensures the stability of the evolution [80, pp.30–31].

6.2   REGION BASED LEVEL-SETS

Many energy functionals have been proposed in literature, specifically designed for the particular application at hand. First generation level sets were basically based on the detection of edges [81]. These methods are intrinsically sensitive to the presence of noise and low image contrast, which can lead to poor segmentation results. In this context, region-based criteria have been shown to be a valuable alternative to edge-based ones. These proceed by partitioning the image into regions which are homogeneous with respect to some defined criterion.

The region based approach was first proposed by Zhu *et al.* [83] and became a consolidate segmentation modality thanks to the work of Chan and Vese [84]. The region based energy functional, in the most general form, can be written as:

$$E(\phi) = \mu_{in} \int_{\Omega_{in}} f_{in}(\phi, \mathbf{x}) d\mathbf{x} + \mu_{out} \int_{\Omega_{out}} f_{out}(\phi, \mathbf{x}) d\mathbf{x} + \\ + \mu_c \int_{\Gamma} f_c(\phi, \mathbf{x}) \|\nabla\phi\| d\mathbf{x} \tag{6.7}$$

where the first two terms, referred ad region terms, represent energy criteria associated to the inside and outside of the evolving

contour, while the third one, called contour term, is associated to the interface Γ. By introducing the Heaviside step function

$$\mathcal{H}(x) = \begin{cases} 1, & \text{if } x \geq 0 \\ 0, & \text{if } x < 0 \end{cases}. \tag{6.8}$$

and the Dirac delta function $\delta(\cdot)$, (6.7) becomes:

$$E(\phi) = \mu_{\text{in}} \int_{\Omega} f_{\text{in}}(\phi, \mathbf{x}) \mathcal{H}(-\phi) d\mathbf{x} +$$
$$+ \mu_{\text{out}} \int_{\Omega} f_{\text{out}}(\phi, \mathbf{x}) \mathcal{H}(\phi) d\mathbf{x} + \tag{6.9}$$
$$+ \mu_{\text{c}} \int_{\Omega} f_{\text{c}}(\phi, \mathbf{x}) \|\nabla \phi\| \delta(\phi) d\mathbf{x}.$$

Let's note that, in order to compute the Euler-Lagrange equation, a regularized version of $\mathcal{H}$ and δ must be considered. A possible choice is given by [17]:

$$\mathcal{H}_\tau(x) = \frac{1}{2}\left(1 + \frac{2}{\pi}\arctan\left(\frac{x}{\tau}\right)\right)$$
$$\delta_\tau(x) = \frac{d}{dx}\mathcal{H}_\tau(x) = \frac{1}{\pi} \cdot \frac{\boxtimes}{\tau^2 + x^2} \tag{6.10}$$

where $\tau \ll 1$ is the regularizing constant.

It is common to adopt the contour force $f_c(\cdot, \cdot)$ constantly equal to 1 [84]. With this substitution, the related energy term becomes a measure of the curve length and imposes a constraint on the regularity of its shape. The first variation of the corresponding contour term, referred as E_c, is expressed by:

$$\frac{\partial E_c}{\partial \phi} = \delta(\phi) \cdot \kappa, \quad \kappa = \nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|}\right) \tag{6.11}$$

where κ is the curvature of the interface [80, Chapter 1]. In standard level-set segmentation techniques this term is commonly introduced to preserve the stability of the evolution, and the corresponding weight μ_c is tuned so to obtain a contour of the desired smoothness.

6.2.1   Example: Chan-Vese functional

The framework introduced by Chan and Vese [84] fairly represents the most commonly adopted level set strategy, at least at

(a)   (b)

Figure 6.2: Result of Chan-Vese algorithm on images with
homogeneous (a) and heterogeneous intensity profiles (b).

a first attempt. This popular criterion will be adopted in the se-
quel in order to illustrate the developed concepts and is briefly
described in this section.

The proposed energy function writes as:

$$E = \int_{\Omega} |I(\mathbf{x}) - c_{\text{in}}|^2 \mathcal{H}(-\phi) d\mathbf{x} +$$
$$+ \int_{\Omega} |I(\mathbf{x}) - c_{\text{out}}|^2 \mathcal{H}(\phi) d\mathbf{x} + \qquad (6.12)$$
$$+ \mu \int_{\Omega} \|\nabla\phi\| \delta(\phi) d\mathbf{x}$$

which leads to the evolution law:

$$\frac{\partial \phi}{\partial \tau} = \delta(\phi) \left[-|I(\mathbf{x}) - c_{\text{in}}|^2 + |I(\mathbf{x}) - c_{\text{out}}|^2 + \mu \nabla \cdot \left(\frac{\nabla\phi}{|\nabla\phi|} \right) \right]$$

where c_{in} and c_{out} represent the mean values of pixel intensity
$I(\mathbf{x})$ measured inside and outside of the evolving contour re-
spectively. Correspondingly, the functional in (6.12) partitions
the image into regions which are maximally homogeneous with
respect to the average intensity.

6.3   Localized level sets

Region-based approaches present several advantages with respect
to edge-based ones, including robustness against initial curve
placement and insensitivity to noise. However, these techniques
are derived on the base of a strong hypothesis, consisting in the
fact that regions of interest can be differentiated in terms of their

Figure 6.3: Difference between region-based and localized approach. In the first case (a) the whole image content is used for interface evolution, in the second case (b) a smaller sub-image is considered instead.

global statistics. In many practical situations, the object to be segmented may be occluded, or its intensity profiles may be heterogeneous, for instance due to spatial variations in the illumination. In these situations traditional region based algorithms may return erroneous segmentation results, see Figure 6.2.

To accurately segment these objects, a new class of active contours energies has been presented in [85], which utilizes local information instead of global one. The initiative is that, if the homogeneity hypothesis is not satisfied globally, there's a greater probability for it to be satisfied within smaller sub-regions. Correspondingly, a variational formulation is derived, which generalizes the region-based formalism by including such a localizing principle. This is realized by the binary mask:

$$W(\mathbf{x}, \mathbf{y}) = \begin{cases} 1, & \text{if } |\mathbf{x} - \mathbf{y}| < \rho \\ 0, & \text{otherwise} \end{cases} \tag{6.13}$$

selecting the local regions (spheres of radius ρ in this case) in which the driving force is computed. In (6.13) $\mathbf{x}$ and $\mathbf{y}$ represent the coordinates of two points in Ω. The localized energy function, in its most general form, writes as:

$$E(\phi) = \int_\Omega \delta(\phi(\mathbf{x})) \left\{ \int_\Omega W(\mathbf{x}, \mathbf{y}) F(I(\mathbf{y}), \phi(\mathbf{y})) d\mathbf{y} \right\} d\mathbf{x} +$$
$$+ \mu \int_\Omega \delta(\phi) \|\nabla \phi\| d\mathbf{x} \tag{6.14}$$

Figure 6.4: Standard region-based Chan-Vese (a), localized Chan-Vese (b).

where F is a generic internal energy measure, used to represent local adherence to a given model at each point along the contour. A visual representation of the localizing approach is given in Figure 6.3. Instead of considering the entire image, curve velocity is computed in the local subsets $\Omega_{\text{in},\mathbf{x}}$ and $\Omega_{\text{out},\mathbf{x}}$, defined as:

$$\Omega_{\text{in},\mathbf{x}} = \{\mathbf{x} | \phi(\mathbf{x}) < 0, \mathbf{x} \in N(\mathbf{x})\}$$
$$\Omega_{\text{out},\mathbf{x}} = \{\mathbf{x} | \phi(\mathbf{x}) > 0, \mathbf{x} \in N(\mathbf{x})\} \tag{6.15}$$

where $N(\mathbf{x}) = \{\mathbf{y} | W(\mathbf{x}, \mathbf{y}) = 1\}$.

The level set equation then is derived as:

$$\frac{\partial \phi}{\partial \tau} = \delta(\phi(\mathbf{x})) \int_{\Omega} W(\mathbf{x}, \mathbf{y}) \nabla_\phi F(I(\mathbf{y}), \phi(\mathbf{y})) d\mathbf{y} +$$
$$+ \mu \nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \tag{6.16}$$

Equation (6.16) provides a simple way to convert virtually every region-based criterion to its localized equivalent. In Figure 6.4 a comparison is presented between the tradition Chan-Vese scheme and its localized version. Due to the presence of an heterogeneous illumination and poorly contrasted boundaries the global approach does not succeed in producing a correct segmentation, which is obtained when the localized version is employed instead.

This important capability provided by localized schemes will be profitably exploited in the design of a left-ventricle segmentation algorithm. Indeed, due to the complex interactions between

the tissue and the ultrasonic beam, left ventricle contours appear highly discontinuous and heterogeneous in the image. Localized level sets represent therefore a powerful candidate to deal with these problematic image characteristics.

6.4   STATISTICAL LEVEL SETS

The extremely general level sets formalism allows to embed several sources of information in the segmentation process. Signal statistics, textures, colors, expected target shape are examples, see [86] for a review.

In the context of ultrasound images, methods based on speckle statistics are commonly developed, indeed in such images the high noise makes standard approaches based on intensity profiles and edges unusable. Hereto, the focus of this section will be on statistically inspired level-set segmentation. Two approaches are reviewed, the first, more common, Maximum-likelihood (ML) criterion by Sarti *et al.*[17], along with a more recent one, based on the Bhattacharyya distance by Michailovich *et al.*[87].

6.4.1   *Maximum Likelihood Segmentation*

Let I be the image to be segmented and Γ the segmenting contour, the ML criterion seeks the contour which maximizes the a-posteriori probability $P(I|\Gamma)$ [88]. In a level set formalism this is equivalent to finding the level set function $\hat{\phi}$ maximizing:

$$E(\phi) = P(I|\phi) = P_{\text{in}}(I|\phi) \cdot P_{\text{out}}(I|\phi) \qquad (6.17)$$

where statistical independence between the regions defined by Ω_{in} and Ω_{out} has been assumed. Assuming as well independence of the pixels within each region, then it is:

$$P_{\text{in}} = \prod_{\mathbf{x} \in \Omega_{\text{in}}} p_{\text{in}}(I(\mathbf{x})), \quad P_{\text{out}} = \prod_{\mathbf{x} \in \Omega_{\text{out}}} p_{\text{out}}(I(\mathbf{x})), \qquad (6.18)$$

where $p_{\text{in}}(\cdot)$ and $p_{\text{out}}(\cdot)$ represent the probability density function (pdf) associated to the pixel intensity I in the inside and outside regions.

Due to the complexity related to the product operations, it is common to maximize the logarithm of (6.17) instead of (6.17) itself. Since the logarithm function is monotonically increasing these two operations are indeed equivalent. By making use of

the Heaviside step function, then the energy function becomes:

$$E(\phi) = \int_\Omega \log p_{\text{in}}(I)\mathcal{H}(-\phi)d\mathbf{x}+$$
$$\int_\Omega \log p_{\text{out}}(I)\mathcal{H}(\phi)d\mathbf{x}. \tag{6.19}$$

In [17] a ML framework is presented for the segmentation of medical ultrasound images. In that work a standard Rayleigh distribution is adopted for the uncompressed envelope signal. Interesting results are reported in echocardiography and obstetrics.

6.4.2 Bhattacharyya Distance Segmentation

In [87] a level-set segmentation criterion is proposed based on the notion of statistical distance. Specifically, it proceeds by partitioning the image into region which are maximally distant with respect to the given metric. In particular the Bhattacharyya distance is there adopted. Given two pdfs p_{in} and p_{out}, this is defined as $D = -\log B$, where:

$$B(p_{\text{in}}|p_{\text{out}}) = \int_{\mathbb{R}} \sqrt{p_{\text{in}}(I)p_{\text{out}}(I)}dI \tag{6.20}$$

is the Bhattacharyya coefficient. In particular, the segmentation flow in [87] is driven by the minimization of B, which, due to the to fact that the logarithm function is monotonically increasing, is perfectly equivalent maximization of D. This avoids computational issues associated to the evaluation of the logarithm function. Non parametric kernel-based estimates of the two pdfs are used in [87], which makes the algorithm independent on any prior assumption on the intensity distribution.

According to the authors, their framework presents computational advantages over the ML one. This is motivated by the fact that the logarithm function used in (6.19) is known to be very sensitive to variations of its argument in vicinity of relatively small values of the latter.

6.5 Shape prior constraints

It has been shown how level sets provide suitable tools for addressing extremely general segmentation tasks, allowing to easily embed several kind of image information, like intensity, edges,

texture, color and many others. In many practical cases occlusions, low contrast or low SNR make image data insufficient for producing meaningful results. In these cases higher-level prior knowledge about the shape of the expected object, if available, can be exploited in the segmentation process, see e.g. [89, 90, 91, 86, 92].

In level-set literature, the shape information is often taken into account by considering the linear combination of two terms [89, 90]. The first one, referred as *data attachment term*, drives the segmentation to particular image features. The second, referred as *shape prior term*, preserves the similarity between the moving interface and a reference shape model. The choice of the prior model is a fundamental issue: typically it is derived from the statistical analysis of a set of training examples [89, 90, 93]. Since orientation and scale of the target are a priori unknown, a transformation is performed at each iteration to align the segmented contour with the prior shape. This step is mainly addressed in a steepest descent framework [90, 89].

Let's denote data attachment and shape prior terms as E_{data} and E_{shape} respectively. Here E_{data} can be any of the energy criterion previously introduced. The global energy function then writes as:

$$E(\phi, \lambda) = E_{\text{data}}(\phi) + \alpha \cdot E_{\text{shape}}(\phi, \lambda) \qquad (6.21)$$

where α corresponds to a positive hyper-parameter that balances the influence between the two terms, and λ represents the parameters of the adopted model (for instance, λ could denote scale, position and orientation of the reference shape [90]). The shape prior term is commonly implemented as a distance measure between the active contour and the expected shape, and its effect is to penalize those solutions which are not consistent with the model itself.

Shape prior segmentation is then defined as:

$$(\hat{\phi}, \hat{\lambda}) = \operatorname*{argmin}_{\phi, \lambda} \{ E(\phi, \lambda) \} \qquad (6.22)$$

which can be solved iteratively with the two step scheme:

$$\phi^{n+1} = \phi^n + \Delta\tau \cdot \left. \frac{\partial E}{\partial \phi} \right|_{(\phi^n, \lambda^n)} \qquad (6.23)$$

$$\lambda^{n+1} = \operatorname*{argmin}_{\lambda} \{ E_{\text{shape}}(\phi^{n+1}, \lambda) \} \qquad (6.24)$$

The first operation corresponds to the displacement of the active contour, while the second step seeks the affine transformation mapping the prior shape onto the active contour itself. For arbitrary shapes E_{shape} must be computed numerically and the corresponding optimization (6.24) must be solved iteratively, mainly with steepest descent schemes [90]. In these situations the computational burden associated to the shape prior may become considerable.

6.6 Conclusion

In this chapter the basic elements of level-set theory have been provided. Traditional region based schemes have been described in §6.2, with some detail on the renown Chan-Vese functional, which represents the pioneering work in this field. The recent framework by Lankton *et al.*has been presented in §6.3, which improves robustness against heterogeneous or missing boundaries by introducing a localization concept. Since edges and intensity profiles are not a trustful source of information in highly noisy situations, a common way to deal with these data is to exploit noise statistics within the segmentation flow. Hereto, in §6.4 two different approaches have been reviewed, based on the concepts of Maximum likelihood estimation and statistical distance. Finally, §6.5 has been dedicated to the formalization of the shape prior segmentation framework. Guesses on the expected shape of the desired target are extremely useful when image data are not sufficient for an accurate segmentation, like in the case of occlusions, low contrast or low SNR.

Myocardium anatomy and functioning

T HE HEART is one of the most vital organs in the human body. It is located between the lungs in the middle of the chest, behind and slightly to the left of the breastbone (sternum). Its fundamental role consists in preserving the activity of the circulatory systems by pumping blood through the entire human body. In particular it cyclically receives deoxygenated blood and pumps it to the lungs to enrich it with fresh oxygen.

In this chapter the basics of heart anatomy and functioning are described in §7.1 and §7.2. In §7.3 an emphasis is put on the left ventricular chamber, which exploits the most onerous mechanical effort in the heart cycle and is the region maximally subject to possible diseases. The importance of left ventricle segmentation as a diagnostic tool to observe and quantify heart deficiencies is then stressed in §7.4, where a literature overview of left ventricle segmentation algorithms is presented. In §7.5 some conclusions are drawn.

7.1 HEART ANATOMY

The heart is the muscular organ responsible for pumping oxygenated blood throughout the vessels by repeated, rhythmic contractions. It is composed of an involuntary striated muscle tissue, called *myocardium*. The innermost and outermost layers of tissues are called *endocardium* and *epicardium* respectively. The endocardium is manly made of epithelial and connective tissues. It lines the inner cavities of the heart, covers heart valves and is continuous with the inner lining of blood vessels, moreover it

Figure 7.1: Anatomy and functioning of human heart muscle.
http://www.childrenshospital.org/az/Site509/Images/normal_large.gif

hosts the Purkinje fibers, which participate in the contraction of the heart muscle. Differently, the epicardium is mainly constituted of connective tissue and functions as a protective layer. The heart muscle is contained within a sac called *pericardium*, which is attached to the central part of the diaphragm and the back of the breastbone. The pericardium is filled with *pericardial fluid*, which acts as a lubricant to allow normal heart movement within the chest.

The heart is composed of four chambers: the two upper atria (singular *atrium*) and the two lower *ventricles*. Left ventricle is separated from the right one by the *ventricular septum*, while left atrium is separated from the right one by the *atrial septum*. In an healthy heart blood flow is unidirectional, the actual path is indicated in Figure 7.1 by means of white arrows. This is made possible by a system of valves. In particular, the right atrium communicates with the right ventricle by means of the *tricuspid valve*, while the *mitral valve* links left atrium to the left ventricle. The *aortic valve* connects the left ventricle to the *aorta*, the major artery of the human circulatory system, while the *pulmonary valve* connects the right ventricle to the *pulmonary artery*, cf. Figure 7.1.

7.2 Heart functioning

The cardiac muscle acts fundamentally as a pump: blood is received by the atria, moved to the ventricles, and then returned to the circulatory system. Blood flow within the heart is regulated by coordinated contractions (*systole*) and relaxations (*diastole*) of the four chambers. In particular both atria contract simultaneously as the ventricles are relaxed and vice versa. This allows blood to flow from the atria into the ventricles, and then again into the circulatory system. Refluxes are prevented by the action of the valves. Valves functioning is purely mechanical: they open when gradient of blood pressure has a certain sign, and close when the opposite situation happens.

A schematic of heart functioning is illustrated in Figure 7.1. Oxygen-depleted blood is received in the right atrium from the *superior vena cava* during atrial diastole. Then right atrium contracts and blood flows through the tricuspid into the right ventricle. During ventricular systole the tricuspid closes and blood in the right ventricle is pumped into the *pulmonary artery* and then to the lungs, where it is oxygenated. Oxigenated blood flows into the left half of the heart: form the left atrium to the left ventricle, and then into the aorta. The portion of the circulatory system which starts in the aorta and ends in the vena cava is called *systemic circulation*, and distributes oxygen and nutrients to the entire body. The portion which starts in the pulmonary artery, passes through the lungs, and ends in the pulmonary veins, is called *pulmonary circulation*, and is responsible for blood re-oxygenation.

This operation is repeated cyclically and is referred as *cardiac cycle*. In an adult healthy heart, it repeats with a frequency of roughly 72 beats per minute (bpm). The sequence of contractions and relaxations is regulated by the propagation of electric pulses, released periodically by the *sinoatrial node*, or peacemaker, located in the right atrium.

7.3 Left ventricle

The left ventricle is the heart chamber to which the most intense mechanical effort is demanded, indeed it is responsible for pumping blood throughout the entire human body. For this reason the large majority of cardiopathies affect this region.

Efficiency of the left ventricle is mainly evaluated by measuring its capability to feed the human physic with satisfactory

Measure	Value
End Diastolic Volume (EDV)	120 ml
End Systolic Volume (ESV)	50 ml
Stroke Volume (SV)	70 ml
Heart Rate (HR)	72 bpm
Cardiac Output (CO)	4.9 L/minute
Wall Thickness	1.1 mm

Table 7.1: Function indexes of an healthy heart. Note that these values are subject to variations depending on age, gender and physical activity.

quantities of blood. In this context *ejection fraction* and *cardiac output* are the principal indexes, commonly adopted as predictors for prognosis. Ejection fraction is the fraction of blood pumped out of left ventricle with each heart beat. By introducing the end-diastolic (EDV) and end-systolic (ESV) volumes, ejection fraction E_f is defined as:

$$E_f = \frac{EDV - ESV}{EDV} = \frac{SV}{EDV} \tag{7.1}$$

where $SV = EDV - ESV$ is the *stroke volume*. Damage to the muscle of the myocardium, such as that sustained during myocardial infarction or in cardiomyopathy, impairs the heart's ability to eject blood and therefore reduces ejection fraction. This reduction in the ejection fraction can manifest itself clinically as heart failure. Physiological values for these quantities for an healthy adult male of 75 Kg are reported in Table 7.3, where the *cardiac output* is indicated as well, defined as $CO = SV \times HR$.

Ejection fraction can be evaluated from echocardiographic images. This is done by tracing the endocardial contours at end-diastole and end-systole followed by extrapolating estimates of EDV and ESV from these contours [94]. Ejection fraction is then computed by means of (7.1).

Besides ejection fraction, left ventricular hypertrophy (LVH) is a further index of fundamental significance. This consists in the muscle thickening and consequent increment of mass, and can be connected with pathological situations implying a muscle overload, which is the case of aortic stenosis, aortic insufficiency, and hypertension. LVH as well can be evaluated in echocardiography by jointly tracing endo- and epicardial contours, and evaluating the distance between the two.

Figure 7.2: Parasternal Short Axis View.
http://www.yale.edu/imaging/echo_atlas/views/graphics/short_axis_lv.gif

In echocardiography, different views of the heart can be privileged according to the particular condition to be investigated. Parasternal Short Axis (SAx), Parasternal Long Axis, Apical 4 Chambers and Apical 2 Chambers are the four most common ones. In particular the focus is put here on the short axis acquisitions. In Figure 7.2 a typical short-axis ultrasound frame is illustrated along with a drawing of the imaged slice. Short axis views are mainly adopted for observing the contractile motion of the myocardium and evaluating suspicious hypertrophies.

7.4   LEFT VENTRICLE SEGMENTATION

As mentioned in the previous section, segmentation of the left ventricle is the standard procedure for assessing clinically relevant measurements of heart function. Due to the intrinsic low quality of ultrasound frames, this task may suffer from a subjective bias, besides being tedious and time consuming [95]. Hence, an automated segmentation procedure is highly desirable both to reduce inter and intra operator variability in borders detection and to speed up the segmentation process. An exhaustive review on echocardiography and, more in general, ultrasound image segmentation is given in Noble *et al.*[96].

In this context, a great attention has been given to the segmentation of the endocardium [17, 96]. This is indeed fundamental for retrieving an important index as the ejection-fraction. Otherwise, very limited literature addresses the detection of the

epicardium. The motivation is that in this kind of acquisitions epicardial fibers happen to be collinear with the ultrasonic beam and the echo returned by these regions is extremely week, with the immediate consequence that epicardial contours appear highly heterogeneous and discontinuous on the resulting image, and difficult to be detected automatically. Nevertheless a trustful detection of both boundaries has a high clinical relevance, since it would allow the computation of fundamental parameters as the ventricular mass, which has been proven to be an important precursor for a variety of conditions such as cardiomyopathy, hypertension and valvular disease [97]. Among the very few works addressing the segmentation of the whole myocardium from SAx acquisitions, the most popular ones are [98, 99, 100].

In [98] Dias *et al.* proposed an algorithm searching for endo and epicardial contours along privileged radial directions starting from a guess of myocardium center. The estimation problem was modeled in a Bayesian framework as a maximum a posteriori estimation problem with Rayleigh statistics for the data term and a spatio-temporal Markov Random Field for the regularization term. An iterative multigrid dynamic programming algorithm was used to solve the optimization problem.

In [99] Chalana *et al.* developed a multiple active contour method for segmentation of 2D+T sequences. They used the image intensity gradient as the attracting force and invoked temporal continuity via an external energy term that constrained the motion between consecutive frames. The different contributions to the curve evolution were balanced according to a set of 14 empirically set parameters. The algorithm was initialized with a manual delineation of the epicardium and the epicardium was correspondingly searched in an automatic way. The two contours were then segmented in sequence throughout the entire acquisition.

In [100] Setarehdan *et al.* developed a fuzzy multiscale edge detection (FMED) method that uses a wavelet transform to define the various levels of resolution of image content. In this work, an edge was defined as a point with maximal membership of the edge fuzzy set. Temporal information was included in edge detection by defining a moving edge fuzzy membership function, as well as an edge fuzzy membership function. The combined edge and motion membership function was then maximized to define edge points.

In this scenario we propose a novel algorithm for combined segmentation of endo and epicardium from SAx acquisitions

which is believed to overcome several shortcomings associated with existing methods. In particular:

- Results of [99] rely on empirically tuning a set of 14 parameters, which highly reduces the clinical usability of the segmentation tool. On the contrary, the proposed algorithm requires 3 parameters only to be set, all with a very specific meaning. As a consequence they can be intuitively adjusted by the user in order to obtain the best results.

- In [99] and [100] image gradient is adopted as edge indicator employed for location of endo and epicardium. It is well established that image gradient alone is not a reliable source of information for heavily speckle corrupted images as ultrasound ones are. For similar data, to employ speckle statistics as driving criterion has become the consolidated practice instead. In particular, in the proposed algorithm a Rayleigh distribution is assumed for the speckle and myocardium is segmented by finding the maximum statistical separation from the background.

- The work in [98] formalizes the segmentation problem as the optimization of a non convex cost function, which is solved by a computationally cumbersome and hardly implementable multigrid dynamic programming algorithm. Since the advent of level sets, they became one of the most popular alternatives for formalizing segmentation problems, indeed they allow both simple and efficient implementation. For this reason we formalize our segmentation problem on the base of the level set theory.

A complete description of the proposed algorithm will be addressed in the next chapter.

7.5   Conclusion

In this chapter the basis of heart morphology and functioning have been provided and several indexes of cardiac function have been presented. The importance of left ventricle segmentation for assessing those performance metrics has been justified and several algorithms addressing this task in an automatic or semi-automatic way have been briefly reviewed. In such a scenario the proposed myocardium segmentation tool has been contextualized. This will be the subject of the next chapter, where an evaluation on clinical data will be presented.

Myocardium segmentation

I N THIS CHAPTER we present an original semi-automatic tool for combined segmentation of endo- and epicardium in 2D echocardiography. The algorithm is based on the level set theory and exploits all the concepts of localization, statistical segmentation and shape priors introduced in Chapter 6.

The rest of the chapter is structured as follows. In §8.1 we recall the general form of the shape prior level-set framework, describe our shape prior term and derive the minimization of the corresponding energy functional. A strategy for preventing collisions between the two contours is described as well. In §8.2, implementation issues are addressed. We explain in particular how the level-set is initialized for each experiment. The experimental setup and the adopted performance metrics are illustrated in §8.3 and §8.4 respectively. In §8.5 we then provide an evaluation of the method on cardiac images. Though the main emphasis will be put on US images a validation on natural images and MRI will be presented as well in §8.6, showing the flexibility of the framework. The main conclusions of this work are given in §8.7.

8.1 PROPOSED SEGMENTATION FRAMEWORK

Many energy functionals have been proposed in literature to incorporate shape priors into level-set formulation. As in [93, 89, 86] we adopt the general expression for the energy function of (6.21), reported here for the sake of completeness:

$$E(\phi, \lambda) = E_{\text{data}}(\phi) + \alpha \cdot E_{\text{shape}}(\phi, \lambda) \qquad (8.1)$$
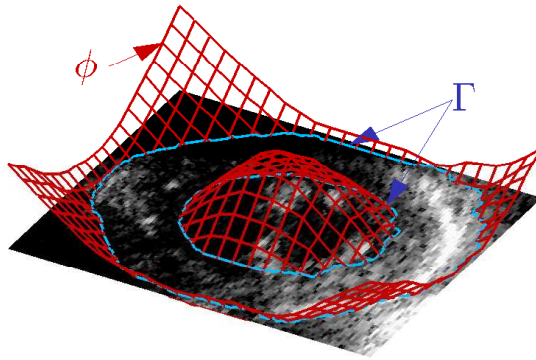
Figure 8.1: Example of level set function used in the presented
algorithm. The zero level set Γ is made out of two disjoint
connected components, representing endo- and epicardium.

where E_{data} represents the chosen data attachment term and E_{shape}
embeds the shape prior. The weight α corresponds to a positive
hyper-parameter that balances the influence between the two
terms. The expression for the two energy terms adopted for the
segmentation task at hand is derived in the following, along with
the variational framework to address the optimization problem.

In order to detect simultaneously endo- and epicardial con-
tours, a level set function with a disjoint zero-level-set is adop-
ted, as the one represented in Figure 8.1.

8.1.1   Data attachment term

As mentioned, the peculiar nature of ultrasound frames makes
standard image processing techniques, developed for natural im-
ages, ineffective [101]. This is mainly due to the presence of the
characteristic speckle noise. Hereto, the standard approach is
to adopt speckle statistics as the driving criterion to address the
segmentation task [96].

In particular, the standard Rayleigh model is assumed for the
first-order statistics of the uncompressed envelope signal. Al-
though this distribution is known to be ineffective when non dif-
fusive scattering conditions are encountered, it has been widely
used for segmentation purposes with satisfactory results [17].
This is further motivated by the fact that its analytical simplicity

allows the closed form solution for the integrals involved in the segmentation flow, which is desirable for a fast computation.

The Bhattacharyya coefficient (6.20) is adopted as data attachment term E_{data}. By embedding the Rayleigh model the two discriminating distributions write as:

$$p_{\text{in}}(I) = \frac{I}{\sigma_{\text{in}}^2} \exp\left\{-\frac{1}{2\sigma_{\text{in}}^2}\right\}, \quad p_{\text{out}}(I) = \frac{I}{\sigma_{\text{out}}^2} \exp\left\{-\frac{1}{2\sigma_{\text{out}}^2}\right\} \tag{8.2}$$

where σ_{in} and σ_{out} represent the Rayleigh parameter estimated inside and outside of the moving interface. In the context of level-sets, this parameter can be estimated in a ML sense inside and outside the evolving interface from the following two expressions:

$$\begin{cases} \sigma_{\text{in}}^2(\phi) = \dfrac{1}{2} \cdot \dfrac{\int_\Omega I^2(\mathbf{x})\mathcal{H}(-\phi(\mathbf{x}))d\mathbf{x}}{\int_\Omega \mathcal{H}(-\phi(\mathbf{x}))d\mathbf{x}} \\[4mm] \sigma_{\text{out}}^2(\phi) = \dfrac{1}{2} \cdot \dfrac{\int_\Omega I^2(\mathbf{x})\mathcal{H}(\phi(\mathbf{x}))d\mathbf{x}}{\int_\Omega \mathcal{H}(\phi(\mathbf{x}))d\mathbf{x}} \end{cases} \tag{8.3}$$

By substituting the Rayleigh model into the expression of the Bhattacharyya coefficient, it becomes:

$$B_R(\phi) = 2 \cdot \frac{\sigma_{\text{in}}(\phi) \cdot \sigma_{\text{out}}(\phi)}{\sigma_{\text{in}}^2(\phi) + \sigma_{\text{out}}^2(\phi)} \tag{8.4}$$

with σ_{in} and σ_{out} defined as in (8.3).

While region-based active contours have been successfully used for endocardium segmentation [17], they result otherwise ineffective in providing meaningful segmentations of the epicardium. This is due to the fact that, when a short axis acquisition is made, myocardial fibers and the ultrasonic beam happen to be collinear at the epicardium. Correspondingly, it returns an extremely weak echo and its boundaries appear poorly defined in the image and are therefore difficult to be located in an automatic way. It was shown in the Chapter 6 that a localizing approach can be exploited in order to deal with similar problematic image features. Hereto a localized version of the Bhattacharyya distance is adopted.

The localized energy then writes as [75]:

$$B_{R,\mathbf{x}}(\phi) = 2 \cdot \frac{\sigma_{\text{in},\mathbf{x}}(\phi) \cdot \sigma_{\text{out},\mathbf{x}}(\phi)}{\sigma_{\text{in},\mathbf{x}}^2(\phi) + \sigma_{\text{out},\mathbf{x}}^2(\phi)} \tag{8.5}$$

Figure 8.2: Comparison between local and global approach. In red the global method in [17] and in yellow the local one in [75]. In green the manual reference by the physician.

where $\sigma_{\text{in},\mathbf{x}}$ and $\sigma_{\text{out},\mathbf{x}}$ correspond to the ML estimates of the Rayleigh parameters computed on the two localized sets $\Omega_{\text{in},\mathbf{x}}$ and $\Omega_{\text{out},\mathbf{x}}$, defined in (6.15). Using the general framework of (6.16), the level set equation can be derived:

$$\frac{\partial \phi}{\partial \tau}(\mathbf{x}) = \delta(\phi(\mathbf{x}))f(\mathbf{x}) \tag{8.6}$$

where

$$f(\mathbf{x}) = \int_{\Omega_y} \left\{ W(\mathbf{x},\mathbf{y})\delta(\phi(\mathbf{y}))\frac{B_{R,\mathbf{x}}}{2}\left[\frac{\sigma^2_{\text{out},\mathbf{x}} - \sigma^2_{\text{in},\mathbf{x}}}{\sigma^2_{\text{in},\mathbf{x}} + \sigma^2_{\text{out},\mathbf{x}}}\right] \times \right.$$
$$\left. \times \left[\frac{1}{A_{\text{out},\mathbf{x}}}\left(\frac{I^2(\mathbf{x})}{2\sigma^2_{\text{out},\mathbf{x}}} - 1\right) + \frac{1}{A_{\text{in},\mathbf{x}}}\left(\frac{I^2(\mathbf{x})}{2\sigma^2_{\text{in},\mathbf{x}}} - 1\right)\right]\right\}d\mathbf{y} +$$
$$+\lambda\kappa\delta(\phi(\mathbf{x})) \tag{8.7}$$

where $A_{\text{in},\mathbf{x}}$ and $A_{\text{out},\mathbf{x}}$ represent the areas of the two sets $\Omega_{\text{in},\mathbf{x}}$ and $\Omega_{\text{out},\mathbf{x}}$. The derivation of (8.7) is addressed in Appendix A.

A visual evaluation of the contribution due to the localized approach is given in Figure 8.2, where a comparison with the standard region-based method in [17] is shown. It is evident how the global approach completely fails in locating the epicardial contours, so producing nonsense results. Differently the local approach manages to retrieve meaningful contours, showing a good correspondence with the manual reference by the cardiologist. These preliminary results have been published in [75].

Although good performances are observed, the algorithm still suffers from undesired irregularities in the final solution, due to the attachment of the active contour to small structures, like papillary muscles, and speckle noise. In order to deal with these effects, an higher level knowledge is embedded in the segmentation framework, consisting in the assumption of a prior model of myocardium shape. Such a model is described in the next section, along with an effective way to deal with it.

8.1.2 Shape Prior Term

It is assumed that the myocardial shape to be segmented may be approximated by two ellipses, an inner and an outer one (for brevity sake, we will refer in the sequel to the shape comprised between these two elliptic contours as annular). This assumption is indeed well supported by observation when SAx acquisitions are considered [102].

Inspired by the framework proposed in [90], we introduce an annular shape constraint into our level-set framework by minimizing the following energy criterion:

$$E_{\text{shape}} = \int_\Omega \Psi^2(\mathbf{x}, \boldsymbol{\lambda}) \|\nabla\phi(\mathbf{x})\| \delta(\phi(\mathbf{x})) \qquad (8.8)$$

where $\Psi(\mathbf{x}, \boldsymbol{\lambda})$ is the implicit function representing the distance of a point $\mathbf{x}$ to the annular shape defined by the parameters $\boldsymbol{\lambda}$. Equation (8.8) reads as a measure of the distance between the active contour and the prior shape, and therefore imposes a similarity between the segmentation result and the prior itself. The following parametric expression for Ψ is proposed:

$$\Psi(\mathbf{x}, \boldsymbol{\lambda}) = \max\{\mathcal{E}(\mathbf{x}, \boldsymbol{\lambda}_{\text{out}}), -\mathcal{E}(\mathbf{x}, \boldsymbol{\lambda}_{\text{in}})\} \qquad (8.9)$$

where $\boldsymbol{\lambda}_{\text{in}}$ and $\boldsymbol{\lambda}_{\text{out}}$ represent the parameters of the inner and outer ellipses, and $\mathcal{E}$ is the algebraic distance of a point $\mathbf{x} = (x, y)$ to the ellipse, represented by the standard quadratic equation for conic sections:

$$\mathcal{E}(\mathbf{x}, \boldsymbol{\lambda}) = \lambda_{i,5} x^2 + \lambda_{i,4} xy + \lambda_{i,3} y^2 + \lambda_{i,2} x + \lambda_{i,1} y + \lambda_{i,0} \\ \text{with } \lambda_{i,4}^2 - 4\lambda_{i,5}\lambda_{i,3} < 0 \qquad (8.10)$$

with $i \in \{\text{in}, \text{out}\}$. Consider that, although six parameters appear in (8.10), the inequality constraint reduces by one the number of degrees of freedom, which, for an ellipse, are indeed 5. An example of Ψ function is given in Figure 8.3.

(a)

(b)

Figure 8.3: Example of Ψ function (a) along with a slice (b).

We note here that another, more trivial, choice for Ψ could be adopted, consisting in the pointwise product between the two ellipse distance functions. However, that choice would no longer represent a distance function [103] and, further, it would not allow an optimized solution to the least-squares fitting problem, as the one we are presenting in the sequel. It is also to be noted that (8.10) corresponds to a representation of the ellipse through the algebraic distance. While the Euclidean distance could also be used, the algebraic distance has the advantages of yielding an analytical evaluation of (8.8) and allowing the use of the fast elliptic fitting algorithm of [104], as described in the next section. On the opposite, the Euclidean distance implies numerically evaluation of (8.8) and thus heavier iterative techniques [105, 106].

The minimization of energy (8.8) leads to finding a geodesic

zero level-set (in our case two contours) in a Riemannian space derived from the shape prior content. As compared to the initial work of [107] and derived approaches [90], the arc length function of (8.8) is no longer weighted by an image based information but only by our shape prior term. From the observation that the minimum of this expression is reached when the zero level of ϕ perfectly fits the zero level of the parametric implicit function Ψ, one can anticipate that the minimization of this energy criterion will make the level-set evolve toward an annular shape. The numerical minimization of (8.8) is addressed using a two phase scheme [84, 90]. Specifically, keeping λ fixed, the minimization of E_{shape} with respect to ϕ leads to the following equation:

$$\frac{\partial \phi}{\partial \tau} = \delta(\phi(\mathbf{x}))g(\mathbf{x}, \lambda) \qquad (8.11)$$

where

$$g(\mathbf{x}, \lambda) = \left\{ 2\Psi \frac{\nabla \Psi \cdot \nabla \phi}{\|\nabla \phi\|} + \Psi^2 \text{div}\left(\frac{\nabla \phi}{\|\nabla \phi\|} \right) \right\} \qquad (8.12)$$

where the dependence of ϕ and Ψ on $\mathbf{x}$ has been omitted for compactness of notation. Then, keeping ϕ fixed, λ is updated according to the following least-squares fitting problem:

$$\lambda = \operatorname*{argmin}_{\lambda'} \oint_{\Gamma} \Psi^2(s, \lambda')ds = \operatorname*{argmin}_{\lambda'} \sum_{\mathbf{x}_i \in \Gamma} \Psi^2(\mathbf{x}_i, \lambda') \qquad (8.13)$$

By noting that $\|\nabla \phi\| = 1$ because of the signed distance property [80], then (8.13) corresponds to the exact minimization of (8.8) w.r.t. λ. The rightmost of (8.13) is justified by the fact that the image space is in practice discrete. To the best of author's knowledge, no method for the fitting of annular shapes has been proposed yet. We derive such a solution in the following.

Least-squares fitting of annular shapes

Considering (8.9), the sum in (8.13) can be rewritten as:

$$J(\mathbf{x}, \lambda) = \sum_{\mathbf{x} \in \Gamma_A} \mathcal{E}^2(\mathbf{x}, \lambda_{\text{out}}) + \sum_{\mathbf{x} \in \Gamma_B} \mathcal{E}^2(\mathbf{x}, \lambda_{\text{in}}) \qquad (8.14)$$

where the partition $\Gamma = \{\Gamma_A, \Gamma_B\}$ has been introduced

$$\begin{aligned} \Gamma_A(\lambda_{\text{in}}, \lambda_{\text{out}}) &= \{\mathbf{x} \in \Gamma | \mathcal{E}(\mathbf{x}, \lambda_{\text{out}}) \geq -\mathcal{E}(\mathbf{x}, \lambda_{\text{in}})\} \\ \Gamma_B(\lambda_{\text{in}}, \lambda_{\text{out}}) &= \{\mathbf{x} \in \Gamma | \mathcal{E}(\mathbf{x}, \lambda_{\text{out}}) < -\mathcal{E}(\mathbf{x}, \lambda_{\text{in}})\} \end{aligned} \qquad (8.15)$$

Input data: $\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}$, $k = 1$, tol = 1e-2

Initialization:

$$E^{(0)} = J\left[\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}\right];$$

$$\Gamma_{\text{A}}^{(0)} = \Gamma_{\text{A}}\left[\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}\right]; \quad \Gamma_{\text{B}}^{(0)} = \Gamma_{\text{B}}\left[\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}\right];$$

---

while $e > $ tol **do**:

$$\hat{\boldsymbol{\lambda}}_{\text{in}}^{(k)} = \text{fitLS}\left[\Gamma_{\text{B}}^{(k-1)}\right]; \quad \hat{\boldsymbol{\lambda}}_{\text{out}}^{(k)} = \text{fitLS}\left[\Gamma_{\text{A}}^{(k-1)}\right];$$

$$\Gamma_{\text{A}}^{(k)} = \Gamma_{\text{A}}\left[\hat{\boldsymbol{\lambda}}_{\text{in}}^{(k)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(k)}\right]; \quad \Gamma_{\text{B}}^{(k)} = \Gamma_{\text{B}}\left[\hat{\boldsymbol{\lambda}}_{\text{in}}^{(k)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(k)}\right];$$

$$E^{(k)} = J\left[\hat{\boldsymbol{\lambda}}_{\text{in}}^{(k)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(k)}\right];$$

$$e = \|E^{(k)} - E^{(k-1)}\| / \|E^{(k-1)}\|; \quad k = k + 1;$$

end while

Table 8.1: Least Squares Fitting Algorithm

From this formulation, we observe that (8.14) can be minimized by fitting two separate ellipses on Γ_{A} and Γ_{B}, for which fast direct solvers exist. Consequently we propose to minimize J by alternatively fitting the two ellipses and updating Γ_{A} and Γ_{B} according to (8.15). The resulting algorithm is summarized in Table 8.1. By doing so, the energy J is ensured to decrease at each step. In Table 8.1 we call fitLS the function performing the direct least-squares ellipse fitting described in [104].

In Figure 8.4, two examples are reported. In both cases, the two initial ellipses defined by $\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}$ and $\hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}$ are coincident. We adopt at first three different initializations from the same dataset (blue contours in Figure 8.4(a)). In each case, the algorithm converges to the same results after only 3 iterations, as shown in Figure 8.4(b). This shows the robustness of our method with respect to the initialization. We then tested our fitting algorithm on a more challenging case, where some of the points that describe both the inner and outer contours are superimposed (Figure 8.4(c)). Figure 8.4(d) shows the good behavior of our algorithm in such situation. In this case, convergence has been reached after only 4 iterations.

The computational complexity of Table 8.1 is dominated by

(a)   (b)

(c)   (d)
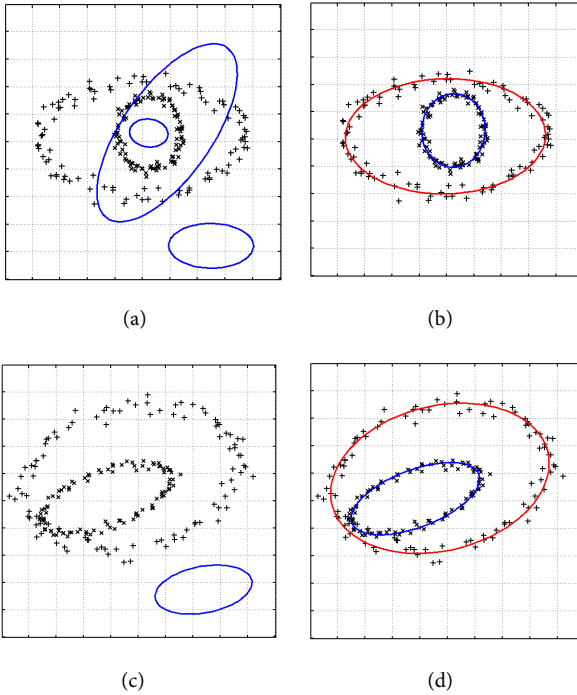
Figure 8.4: Illustration of the use of annular least square fitting. In both cases, the two initial ellipses are coincident. First row: illustration on a simple case. a) Initialization: the three ellipses corresponds to three different initializations. b) Results obtained at convergence: the final annular contours are identical. Second row: illustration on a more challenging case. c) Initialization. d) Results obtained at convergence after only 4 iterations.

the few single ellipse least-squares fitting operations involved, indicated by the function fitLS. As shown in [104], such a problem is equivalent to an eigenvalue problem for a 6×6 matrix, which can be efficiently solved with $\mathcal{O}(n^3)$ floating point operations [108]. As an example, the solution of Table 8.4 is obtained in 13 ms in a MATLAB (R2010b,The Math Works) implementation of the algorithm, executed on a 2.27 GHz Intel Core i5 laptop equipped with 4 GB of RAM and running Windows 7 64-bit.

As detailed in the sequel, the segmentation method consists in a two stage procedure, i.e.alternating the fit of the annular shape to the level-set described in Table 8.1 and the constrained

evolution of the level-set through (8.6). It is then to be noted that the fitting stage can be made very fast by initializing the algorithm in Table 8.1 with the fitting result computed at the previous iteration. In practice, the experiments showed that this strategy allows performing the fitting stage in one single iteration

As a further remark, it should be noted that, although the framework has been specialized for annular shapes, the proposed formalism is general. In particular it can be extended to geometric primitives for which an efficient solution to the least squares fitting problem can be individuated, to be used in place of FitLS in Table 8.1. In this context, an extension of the algorithm to arbitrary views of the heart based on hyperquadrics is currently under our investigation [109].

We conclude this section by pointing out that the penalty term in (8.8) is not new in literature: Chen *et al.* introduced an analogous one in [90], in order to embed a shape constraint in the geodesic active contour flow [107]. Nevertheless, the difference between the presented framework and the one in [90] is substantial: in [90] Ψ represents the distance, computed numerically, of a point to an arbitrary shape, and consequently a similarity transformation between the segmenting contour and the prior shape has to be searched in a steepest descent scheme. In our work, we develop an analytic representation for Ψ, expressed by (8.9), allowing thereby to express the parameters update step as a direct least squares fitting. As a consequence, this stage of the algorithm avoids slow steepest descent iterative procedures. These aspects will be further stressed in the results section, where a comparison of the proposed algorithm with the one by Chen is presented, both in terms of segmentation accuracy and execution time.

8.1.3   Collision preventing term

Up to this point the two contours are let to move independently one from the other. For this reason there's the risk for them to collide by merging into a single contour. It could happen for instance in the case the structures of interest are close to each other. This fact would preclude a meaningful segmentation result.

We stress that this situation was never encountered as long as ultrasound images were considered. This was due to the fact that image size was high enough to preserve a sufficient distance between endo and epicardium. As a consequence, the term we are describing in this section was not considered when dealing

with ultrasound images. Anyway, when testing our algorithm on images of different nature, in particular on cardiac MRI, where endo and epicardium were only few pixels distant, we experienced several collision events.

To avoid these situations, an energy term $E_{\text{collision}}$, referred as *collision preventing term*, can be added to the total energy function. This strategy is not new in literature, where $E_{\text{collision}}$ typically implements an incompressibility constraint [110], imposing the area (or volume) enclosed by the two contours to be similar to a certain reference value, obtained on the base of a training process. Besides that this learning phase highly limits the generality such an approach, we note that this criterion, due to its global nature, does not really prevent the two contours from touching. In this sense, an alternative would be to introduce a constraint based on the minimum distance between the two contours. Unfortunately, this solution is not compatible with our formalism, as the two contours are defined by a single level set function and indistinguishable one from the other.

In this context, we developed an original $E_{\text{collision}}$ term which exploits the advantages of both these approaches by remaining consistent with the adopted formalism. Specifically, as the first solution, it makes use of area information, which can be immediately deduced from the sign of ϕ, while, as the second, it looks locally, instead of globally, for collision warning situations.

The rationale behind the proposed term is illustrated in Fig.8.5. The example is relative to a cardiac MRI image. The basic idea is that, when looking within a surrounding B_R of a contour pixel $\mathbf{x}$, if the two contours are close in that point, then B_R will be crossed by a thin strip of pixels of Ω_{in} and will be prevalently filled by pixels of Ω_{out} (cf. Figure 8.5(e)). Instead, the two populations will be approximately equal if the two contours are distant enough (cf. Figure 8.5(d)). The warning distance is determined by the extent of B_R.

Following the above considerations and inspired by the localized framework in [85], the proposed energy criterion writes as:

$$E_{\text{collision}} = \int_\Omega \delta(\phi(\mathbf{x})) \int_\Omega B_R(\mathbf{x}, \mathbf{y}) F_{\text{collision}}(\phi(\mathbf{y})) d\mathbf{x} d\mathbf{y}$$

$$(8.16)$$

where $B_R(\mathbf{x}, \mathbf{y})$ is the disk of radius R analogous to (6.13), represented as a white circle in Figure 8.5 and.

$$F_{\text{collision}} = (A_{\text{in},\mathbf{x}} - A_{\text{out},\mathbf{x}})^2 \qquad (8.17)$$
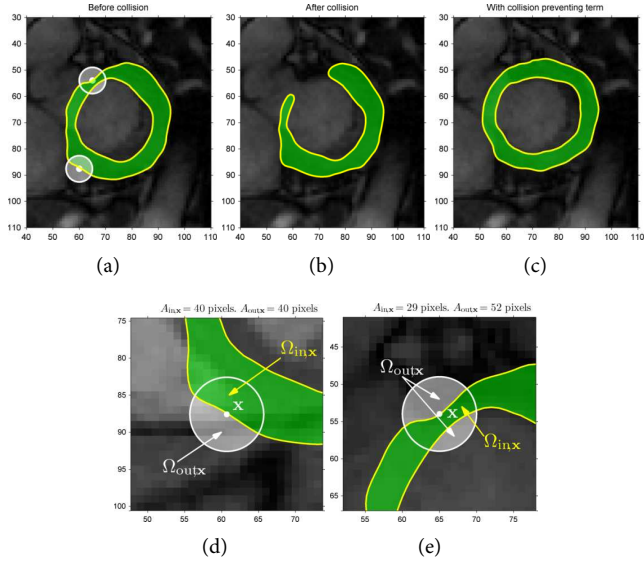
(a)   (b)   (c)

(d)   (e)

Figure 8.5: In (a) the two contours soon before the collision. Two regions, corresponding to a safe distant-contours situation and a warning close-contours situation are represented as white disks. In (b) the active contour after the collision. In (c) the result obtained thanks to the collision preventing term. In (d) and (e) a zoom of the two regions corresponding to distant and close contours respectively. The image is taken from a cardiac MRI exam.

where $A_{\text{in},\mathbf{x}}$ and $A_{\text{out},\mathbf{x}}$ represent the areas of the two sets $\Omega_{\text{in},\mathbf{x}}$ and $\Omega_{\text{out},\mathbf{x}}$ in Figure 8.5, and can be expressed as:

$$
\begin{aligned}
A_{\text{in},\mathbf{x}} &= \int_\Omega B_R(\mathbf{x},\mathbf{y}) H(-\phi(\mathbf{y})) d\mathbf{y}, \\
A_{\text{out},\mathbf{x}} &= \int_\Omega B_R(\mathbf{x},\mathbf{y}) H(\phi(\mathbf{y})) d\mathbf{y}.
\end{aligned}
\tag{8.18}
$$

The effect of (8.16) to keep the two areas $A_{\text{in},\mathbf{x}}$ and $A_{\text{out},\mathbf{x}}$ similar, so preventing the close contours situations of Figure 8.5(e). Using the general framework in [85] the minimization of $E_{\text{collision}}$ with respect to ϕ leads to:

$$
\frac{\partial \phi}{\partial \tau} = h(\mathbf{x}) \delta(\phi(\mathbf{x}))
\tag{8.19}
$$

where

$$h(\mathbf{x}) = 2 \int_{\Omega} \delta(\phi(\mathbf{y})) B_R(\mathbf{x}, \mathbf{y}) (A_{\text{in},\mathbf{x}} - A_{\text{out},\mathbf{x}}). \qquad (8.20)$$

An example of the effect of this term is illustrated in Figure 8.5(c). Note that no collision happens thanks to the proposed term and a meaningful result is obtained. The value of ρ has been fixed equal to 5 pixels in that case.

We conclude this section with a remark. The term in (8.17) does not distinguish between close contours and peaky corners, indeed in this second circumstance as well the number of inside-pixels may be substantially smaller than the outside-ones. Though this fact may represent a shortcoming in general, since it hinders the detection of high curvature details, otherwise this is not the case for the application at hand, in which the interesting patters can be represented by ellipses and therefore present an intrinsic smoothness.

8.2   Implementation issues

Level-set evolution equation is implemented using a standard finite difference scheme [80], where the implicit function ϕ is represented by a signed distance function. In order to improve efficiency, values of ϕ are only computed in a narrow band around the zero level set. Consequently, ϕ is re initialized every few iterations using a fast marching scheme [80]. The final, shape-prior level set equation, writes as:

$$\frac{\partial \phi}{\partial \tau}(\mathbf{x}) = \delta(\phi(\mathbf{x})) \left[f(\mathbf{x}) + \alpha \cdot g(\mathbf{x}, \lambda) \right] \qquad (8.21)$$

where f and g are defined as in (8.6) and (8.12) respectively. The choice of the weight term α is dependent on the specific application. For instance, when the quality of the image is poor, it is recommended to use higher value of α in order to put more importance on the shape prior term. When presenting results in the experiments section, we will adopt the α value which produces the best segmentation result. In particular a value $\alpha = 1.2$ was found to produce the best result on available images.

Concerning the initialization, let us note the localized framework of [85] requires that this is made not too far from the desired solution. Indeed while such a local data attachment term allows making the algorithm more robust to the variation of image properties, it makes it more sensitive to initialization. In the experiments described in the next sections, the following simple procedure is thus considered. The user is asked to position

six points, as this is the minimum number of points needed to define 2 concentric ellipses. The five first points are used to set an ellipse (one ellipse being uniquely defined by 5 points) and the last point is used to obtain the second concentric ellipse. For each experiments, the initial points used to achieve the given results are displayed in green. Finally, the localizing approach is implemented by assuming as mask W a disk whose radius was experimentally set to 10 pixels.

8.3   MATERIALS AND METHODS

The algorithm was validated on a set of 59 B-mode images acquired on 5 different patients. Data were acquired using a Toshiba Powervision 6000 (Toshiba Medical Systems Europe, Zoetermeer, the Netherlands) equipped with a 3.75 MHz-probe. Image size was equal to 249×168 pixels^2. The algorithm was initialized by a non-expert user as described in the previous section. Manual segmentation drawn by an expert cardiologists was used as reference.

8.4   PERFORMANCE METRICS

Performance of the algorithm was assessed by measuring Root Mean Square Distance (RMSD), Hausdorff Distance (HD) and Dice coefficient (D) between automatic and reference contours. These metrics are standardly adopted for evaluating medical images segmentation algorithms [111].

Considering two contours $\mathcal{C}_1$ and $\mathcal{C}_2$, the MAD between the two is defined as:

$$RMSD(\mathcal{C}_1, \mathcal{C}_2) = \frac{1}{2}\left\{ \sqrt{\frac{1}{L(\mathcal{C}_1)} \int_{\mathbf{x}\in\mathcal{C}_1} \left[\inf_{\mathbf{y}} d(\mathbf{x},\mathbf{y})\right]^2 d\mathbf{x}} + \right.$$
$$\left. + \sqrt{\frac{1}{L(\mathcal{C}_2)} \int_{\mathbf{y}\in\mathcal{C}_2} \left[\inf_{\mathbf{x}} d(\mathbf{x},\mathbf{y})\right]^2 d\mathbf{y}} \right\} \quad (8.22)$$

where $d(\mathbf{x},\mathbf{y})$ is the Euclidean distance between $\mathbf{x}$ and $\mathbf{y}$, while $L(\mathcal{C})$ denotes the length of $\mathcal{C}$. The HD is instead defined as:

$$HD(\mathcal{C}_1, \mathcal{C}_2) = \max\left\{ \sup_{bx\in\mathcal{C}_1} \inf_{\mathbf{y}\in\mathcal{C}_2} d(\mathbf{x},\mathbf{y}), \sup_{\mathbf{y}\in\mathcal{C}_2} \inf_{\mathbf{x}\in\mathcal{C}_1} d(\mathbf{x},\mathbf{y}) \right\}. \quad (8.23)$$

	Dice	HD	RMSD
WITH SHAPE PRIOR	**0.89±0.02**	**5.57±1.18**	**2.22±0.46**
WITHOUT SHAPE PRIOR	0.87±0.03	7.78±2.12	2.55±0.49

Table 8.2: Performance metrics value (MEAN ± STANDARD DEVIATION) for the proposed algorithm on 59 echocardiographic images, respectively with (α=1.2) and without (α=0) shape information.

Finally, defining Ω_1 and Ω_2 as the sets of pixels enclosed by the closed contours $\mathcal{C}_1$ and $\mathcal{C}_2$, the Dice coefficient writes as:

$$D(\Omega_1, \Omega_2) = \frac{2\,\#(\Omega_S \cap \Omega_R)}{\#(\Omega_1) + \#(\Omega_2)} \tag{8.24}$$

where $\#(\cdot)$ extracts the set cardinality.

The RMSD measures the average distance between the two contours, and is thus suited for evaluating the global performance of the algorithm. On the opposite, the Haussdorf metric represents the maximum distance between the contours, and is therefore a local index of performance [112]. More specifically, the HD keeps trace of the localized discrepancies which are largely hidden in the average operation involved in the RMSD computation. While RMSD and HD depend on the image size, the Dice coefficient returns instead an intrinsic index of performance, since it measures the superposition in percentage between the two enclosed areas. These three metrics are adopted to provide an exhaustive evaluation of the behavior of the proposed algorithm.

8.5   Results

8.5.1   Evaluation of shape prior contribution

In order to evaluate the effect of the shape prior information we ran our algorithm with and without using the shape prior term. When the shape information was taken into account a value α=1.2 was used. The complete set of results is reported in Table 8.2. This table clearly shows that the introduction of the shape prior yields an improvement of the segmentation, yielding a slightly higher Dice and smaller HD and RMSD.

Figure 8.6 provides segmentation examples allowing a finer interpretation of these results. It may be observed that the shape prior allows to avoid the local irregularities, linked to the attraction of the active contours to small scale noisy structures. This
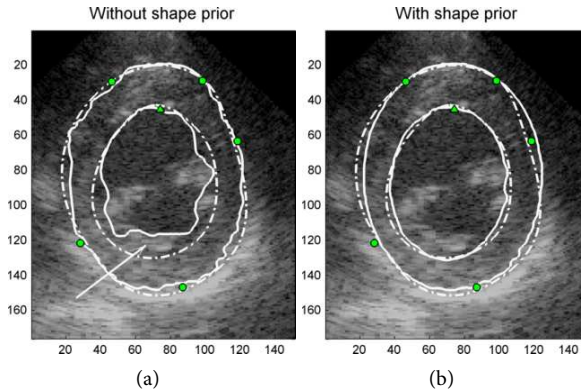
Figure 8.6: Manual reference and segmentation result as a thin and continuous line respectively. Note how the shape constraint prevents the contour from adhering to irregularities in the endocardium (white arrow in (a)), principally ascribable to the presence of the papillary muscles. In general much regular results are obtained with the shape prior approach, with an augmented correspondence with the cardiologist reference. Error metrics for those images are: (a) HD=14.3, RMSD=5.35, D=0.83; (b) HD=5.83, RMSD=2.52, D=0.89.

explains in particular the large decrease of HD associated with the shape information.

8.5.2   Comparison with Chen algorithm

We compare in this section the performance of our algorithm with the approach described in [90]. The basic equations for Chen algorithm are reported in Appendix D. The use of Chen's approach implies the learning of the shape prior through a training phase performed on a set of manually segmented contours. In this experiment, the prior contour was obtained for each modality from a training image set, built by randomly selecting reference contours. Specifically, 10 images out of 59 were used. The prior contour was obtained by averaging the reference contours drawn by the physician. Due to potential misalignment among the references, a registration step must precede the averaging operation, as detailed in [90]. The so obtained prior contour is illustrated in Figure 8.7(b). From this figure it can also be clearly seen that ellipses represent a reliable approximation of myocar-

Figure 8.7: Prior contour computation for (a) MRI and (b) ultrasound images. The cluster of reference curves is represented as thin contours and the corresponding average shape as a bold red contour.

	DICE	HD	RMSD
PROPOSED ALGORITHM	**0.89±0.02**	**5.57±1.18**	**2.22±0.46**
CHEN'S ALGORITHM	0.86±0.03	8.05±1.55	2.50±0.72

Table 8.3: Performance metrics value (MEAN ± STANDARD DEVIATION) for the proposed algorithm and Chen's one for US images. The training images are eliminated from the data set before evaluation. The total number of testing images is the equal to 49.

dial contours.

The complete set of results is reported in Table 8.3. These results clearly show that the proposed algorithm outperforms Chen's approach for the present application. Let us note that the difference in performance is particularly pronounced for the Hausdorff distance. This behavior is linked to the fact that in Chen's approach, the only non-rigid degree of freedom applied to the shape prior in the course of the segmentation is a uniform scaling. On the opposite, our approach performs a complete fit of the prior at each iteration, thus yielding a better adjustment to the myocardial shape variations.

Another important issue concerns the computational speed. The average cpu time needed for segmenting one image was 13.5 seconds for our method and 36.7 seconds for Chen's method. This speed-up is a consequence of the availability of a parametric representation for the prior, combined with the existence of a fast fitting procedure for such a geometry. In particular, the

most costly step in Chen's algorithm is the initial alignment of
the prior shape onto the initial contour (cf. Appendix D). For the
considered images, this step can take up to several thousands of
iterations. On the opposite, in our approach the initial registra-
tion corresponds to fit the initial contour with an annular shape,
which is typically done in few iterations.

8.6   Beyond ultrasound

Due to its extreme flexibility, the shape prior segmentation frame-
work described in this chapter, can be suitably applied to every
kind of image containing annular-like patterns. Indeed similar
patterns are recurrent in images from several fields: medicine,
biology and astronomy are some examples. In this section we
present some results obtained on various natural images and on
cardiac MRI.

As all those images are less challenging than US, the sophisti-
cated statistical framework described in §8.1.1 is no longer nec-
essary and the standard Chan-Vese framework, in its localized
version, has been proved to be largely sufficient.

8.6.1   Segmentation of natural images

Figure 8.8 illustrates the application of the proposed method for
the segmentation of the pupil and iris in human eye, of the zona
pellucida in human embryo, and the rings of Saturn. We used
for these experiments the same α value equal to 0.8. We ob-
serve that in all those situations gray scale information alone is
not sufficient for obtaining correct results. In the eye image (Fig-
ure 8.8(a)), the eyelash hinders a correct iris detection, moreover
the variations in the iris color make pupil segmentation impre-
cise. A similar effect can be observed in the embryo image (Fig-
ure 8.8(b)) where the segmentation is inaccurate due to low con-
trast profiles and to proximity of the cell to the internal bound-
ary of the zona pellucida. In the Saturn image, the ring portion
hidden by the planet body cannot be followed without a shape
prior approach (Fig.8.8(c)).

The results obtained using the shape prior term are given in
the second row. In the 3 cases, they are in good agreement with
the expected shapes. They illustrate the usefulness of the pro-
posed annular shape prior for such images. Regarding the com-
putational cost, Table 8.4 provides the cpu times corresponding
to the segmentations performed with our approach and shown

(a)   (b)   (c)

(d)   (e)   (f)

Figure 8.8: Segmentation results obtained on natural scene images. The green points correspond to the initialization. Images (a) to (c) show the result obtained when the proposed approach is applied without the shape constraint. Images (d) to (f) show the result obtained when the shape prior is applied.

	EYE IMAGE (Fig.8.8(d))	EMBRYO IMAGE (Fig.8.8(d))	SATURN IMAGE (Fig.8.8(f))
OVERALL	8.31	17.93	13.86
SHAPE FITTING	0.12	0.33	0.36

Table 8.4: Cpu time (s) corresponding to the segmentation with the proposed approach on natural images.

in Figure 8.8(d) to Figure 8.8(f). Note that the 3 images share the same size 200x200 pixels. In order to give an insight on the impact of the shape prior term on the overall computational burden, Table 8.4 also gives the cpu time corresponding to the fitting step of the algorithm in §8.1.2. These results confirm that the improvements in segmentation accuracy ascribable to the shape prior term come at a substantially negligible price in terms of computational complexity.

	DICE	HD	RMSD
WITH SHAPE PRIOR	**0.82±0.05**	**3.63±0.73**	**1.46±0.21**
WITHOUT SHAPE PRIOR	0.80±0.05	4.62±0.82	1.72±0.27

Table 8.5: Performance metrics value (MEAN ± STANDARD DEVIATION) for the proposed algorithm on 56 cardiac MRI images, respectively with (α=0.8) and without (α=0) shape information.

8.6.2   Segmentation of cardiac MRI

The algorithm has been evaluated on a set of 52 cardiac MRI images. A value α =0.8 was used to balance the shape prior influence. We note that this is smaller than the one adopted for ultrasound as MRI has an intrinsically higher image quality. As mentioned, due to the reduced size of MRI images, endo and epicardium happened to be too close to prevent collision between the two contours. Consequently the collision preventing energy term presented in §8.1.3 was added to the energy function, weighted by an opportune hyper parameter β. A value β=0.4 was experimentally found to avoid collision events on all the available images. A value R=4 pixels was adopted for the corresponding collision indicator $\mathcal{B}_R$. The same set of experiments performed for US were made for MRI and are reported in this section.

From the evaluation of the shape prior contribution we obtained the results reported in Table 8.5. Note that both HD and RMSD have smaller values for MRI than for US due to the smaller size of the former. A visual comparison between the shape prior and the shape free approach is reported in Figure 8.9.

A comparison with Chen algorithm is presented as well. In this case the prior shape was built from a training set composed of 10 out of the 56 images. The so obtained reference contour is illustrated in Figure 8.7(a). The obtained results are reported in Table 8.6. An example image is illustrated in Figure 8.10.

8.6.3   Dependence on the initialization

As previously mentioned, the adoption of the localized framework of [85] imposes a good initialization in order to guarantee adequate results. In practice, this implies that the distance between the initial contour and the object to be detected should be on the order of the extent of the mask $W(\cdot)$. In this section, we thus propose to study the influence of initialization by applying a
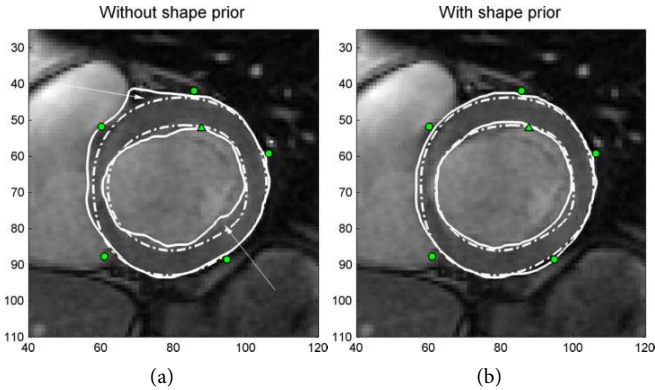
Figure 8.9: Segmentation results on cardiac MRI. Manual reference and segmentation result as a dashed and continuous line respectively. Note how the shape constraint prevents the contour from adhering to the right ventricle (upper arrow in (a)) and allows to cope with the low contrast of the endocardium (lower arrow in (a)). Error metrics for those images are: (a) HD=5, RMSD=1.52, D=0.83; (c) HD=2.23, RMSD=0.99, D=0.89;(b) HD=2.23, RMSD=0.99, D=0.89.

	Dice	HD	RMDS
PROPOSED ALGORITHM	**0.82±0.05**	**3.63±0.73**	**1.46±0.21**
CHEN'S ALGORITHM	0.76±0.18	4.89±1.06	1.74±0.45

Table 8.6: Performance metrics value (MEAN ± STANDARD DEVIATION) for the proposed algorithm and Chen's one for MRI images. The training images are eliminated from the data set before evaluation. The total number of testing images is the equal to 46.

random variation to the location of the 6 points used in our initialization procedure. Although the study is presented here for MRI images only, the considerations that follow are absolutely general.

This experiment is illustrated in Figure 8.11 and is done as follows. We delineate from a specific MRI image six regions in which initial points are randomly positioned (red disk in Figure 8.11(a)). Each region corresponds to a disk with a radius of 3 pixels. This dimension has been chosen to simulate the uncertainty in the positioning of the points by the user. In order to show the influence of such random scheme in the initial-

Figure 8.10: Comparison between the proposed (a) and Chen (b) algorithm applied to an MRI image. The prior for Chen algorithm is not appropriate for this particular frame and compensates for the lack of data information in a way which is not consistent with the real heart morphology, as shown by the arrows in (b) and (c). Differently, the adopted parametric formulation easily allows to second myocardial shape. Error measures are: (a) D=0.87, HD=3.60 and RMSD=1.42; (b) D=0.76, HD=5.88 and RMSD=2.09.

ization, we display in Figure 8.11(a) two particular initial contours obtained from the proposed random initializations and Figure 8.11(b) shows the final contours obtained from 10 different initializations. In order to evaluate the variability of the resulting segmentations, we compute the similarity measures by considering all the possible combinations of two final contours. We obtained mean values of D=0.98, HD=1.1 pixels and RMSD=0.3 pixels. These values illustrate the robustness of the proposed method with respect to the initialization phase.

As a conclusion, although the local data attachment term implies an initialization close to the desired contour, within this limit, this initialization thus needs not to be very accurate.

8.7   CONCLUSION

in this chapter a novel framework based on level-sets has been introduced and described. The algorithm contains several elements of novelty over the related literature, specifically:

- a parametric implicit function defining the distance from two ellipses is introduced;

- an efficient solution to the least- squares fitting problem of annular shapes is developed;

<table>
<tr><td>(a)</td><td>(b)</td></tr>
</table>

Figure 8.11: Evaluation of the influence of the initialization points.
a) Regions where the initial points are randomly positioned with
two particular derived contours. b) Segmentation results obtained
from 10 different initializations.

- thanks to its parametric representation, the shape prior
  may be adjusted to the evolving contour through an ef-
  ficient least square strategy. Correspondingly, the com-
  putational cost associated to the introduction of the prior
  shape information is made essentially negligible.

This framework provides two additional benefits. First, since
a parametric representation is adopted, a training phase for model
selection is avoided, for which the availability of training sam-
ples and their number are difficult issues. Secondly, the detec-
tion of both contours is addressed with a single level-set func-
tion, which is desirable in terms of memory consumption and
computation.

The algorithm was evaluated on a set of clinical data where
good performance were observed. These promising results en-
courage further efforts in order to make the algorithm usable in
a realistic clinical setting. In particular, we believe they must go
in two main directions:

- Reduce user interaction. This required the proposal of a
  less demanding initialization procedure;

- Reduce computation time. The biggest computation here
  is requested by the data attachment term, in which all lo-
  calized region must be processed in sequence. This step
  might be made substantially lighter with the employment

of GPU programming techniques, which would allow to exploit the parallelism offered by modern GPUs architectures.

Conclusion

Tʜɪꜱ ᴍᴀɴᴜꜱᴄʀɪᴘᴛ described the research activity carried out by the author between January 2008 and December 2010 in the context of his PhD at the Advanced Research Center on Electronic Systems and Information Technology of the University of Bologna, Italy. Part of the work was made during a one year internship within the ultrasound team of the the Centre de Recherche et d'Applications en Traitement de l'Image et du Signal (CREATIS) in Lyon, France.

Two major topics inherent with medical ultrasound images were addressed during the PhD research. The first project had its principal goal in the development of an effective pre-processing tool for reducing system dependencies in the performance of ultrasound based tissue characterization algorithms. This was made possible thanks to the proposal of an original deconvolution framework allowing to restore statistically consistent maximum *a posteriori* estimates of the tissue reflectivity. A performance evaluation was documented in the thesis, obtained from phantom data mimicking biological tissues with different particle concentrations. The results showed a relevant decrease in the misclassification error when processed data were considered instead of the raw signal, along with the superiority of the proposed paradigm over standard deconvolution tools. Future work must include the validation of the proposed scheme on clinical data. In this sense the main interest relies in the detection of potentially cancerous areas on trans rectal ultrasound images of the prostate for guiding biopsy procedures.

The second activity stemmed instead from the purpose to provide the cardiology community with software tools for easing

the reading and analysis of echocardiographic sequences. In this context a semi-automatic segmentation tool for joint detection of endo- and epicardial contours from short axis acquisition was designed. The algorithm was derived on the base of the level-set formalism and effectively incorporated information on signal statistics and myocardium expected geometry for constraining the results to anatomically meaningful shapes. The obtained results have shown that the proposed algorithm can reduce the high inter- and intra-operator variability and is feasible and accurate even for realistic clinical data. A natural continuation of this activity includes the generalization of the proposed framework to the case of arbitrary probe orientations, in this sense some preliminary results have been recently obtained [P1][P6], and the use of the detected contours for the calculation of parameters commonly adopted in the clinical routine for the diagnosis of heart disease as ejection fraction and wall thickening.

Derivation of the Bhattacharyya level-set function

The Bathhacharyya distance for two Rayleigh pdfs is given by (8.4):

$$B_R(\phi) = 2 \cdot \frac{\sigma_{\text{in}}(\phi) \cdot \sigma_{\text{out}}(\phi)}{\sigma_{\text{in}}^2(\phi) + \sigma_{\text{out}}^2(\phi)} \qquad (A.1)$$

where the two Rayleigh parameters are defined as (8.3):

$$\begin{cases} \sigma_{\text{in}}^2(\phi) = \dfrac{1}{2} \cdot \dfrac{\int_\Omega I^2(\mathbf{x})\mathcal{H}(-\phi(\mathbf{x}))d\mathbf{x}}{\int_\Omega \mathcal{H}(-\phi(\mathbf{x}))d\mathbf{x}} \\[3mm] \sigma_{\text{out}}^2(\phi) = \dfrac{1}{2} \cdot \dfrac{\int_\Omega I^2(\mathbf{x})\mathcal{H}(\phi(\mathbf{x}))d\mathbf{x}}{\int_\Omega \mathcal{H}(\phi(\mathbf{x}))d\mathbf{x}} \end{cases} \qquad (A.2)$$

After some algebra the first variation of B_R w.r.t. ϕ can be written as:

$$\frac{\partial B_R}{\partial \phi} = 2\frac{\sigma_{\text{in}}^2 - \sigma_{\text{out}}^2}{\left(\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2\right)^2}\left\{\sigma_{\text{in}}\frac{\partial \sigma_{\text{out}}}{\partial \phi} - \sigma_{\text{out}}\frac{\partial \sigma_{\text{in}}}{\partial \phi}\right\} \qquad (A.3)$$

where the terms $\partial \sigma_{\text{in,out}}/\partial \phi$ have to be computed. It is useful to use:

$$\frac{\partial \sigma_{\text{in}}}{\partial \phi} = \frac{1}{2\sigma_{\text{in}}}\frac{\partial \sigma_{\text{in}}^2}{\partial \phi}, \qquad \frac{\partial \sigma_{\text{out}}}{\partial \phi} = \frac{1}{2\sigma_{\text{out}}}\frac{\partial \sigma_{\text{out}}^2}{\partial \phi} \qquad (A.4)$$

and then compute the first variation of the two variances instead of the standard deviations. Let's introduce:

$$\begin{cases} I_{\text{in}}^m = \int_\Omega I^2\mathcal{H}(-\phi)d\mathbf{x} \\ I_{\text{out}}^m = \int_\Omega I^2\mathcal{H}(\phi)d\mathbf{x} \end{cases} \qquad \begin{cases} A_{\text{in}} = \int_\Omega \mathcal{H}(-\phi)d\mathbf{x} \\ A_{\text{out}} = \int_\Omega \mathcal{H}(\phi)d\mathbf{x} \end{cases} \qquad (A.5)$$

where the dependency of I and ϕ on $\mathbf{x}$ has been omitted. These allow to write:

$$\frac{\partial \sigma_{in}^2}{\partial \phi} = \frac{1}{A_{in}^2} \left\{ \frac{\partial I_{in}^m}{\partial \phi} A_{in} - \frac{\partial A_{in}}{\partial \phi} I_{in}^m \right\} \qquad (A.6)$$

By using the theory of variations, it is:

$$\frac{\partial I_{in}^m}{\partial \phi} = -I^2 \delta(\phi), \qquad \frac{\partial A_{in}}{\partial \phi} = -\delta(\phi) \qquad (A.7)$$

and, after some calculation:

$$\frac{\partial \sigma_{in}}{\partial \phi} = \frac{\sigma_{in}}{A_{in}} \left\{ -\frac{I^2}{2\sigma_{in}^2} + 1 \right\} \delta(\phi). \qquad (A.8)$$

Analogously one can derive the same expression for σ_{out}:

$$\frac{\partial \sigma_{out}}{\partial \phi} = \frac{\sigma_{out}}{A_{out}} \left\{ \frac{I^2}{2\sigma_{out}^2} - 1 \right\} \delta(\phi). \qquad (A.9)$$

By substituting (A.8) and (A.9) into (A.3), the final expression can be obtained:

$$\frac{\partial B_R}{\partial \phi} = \frac{B_R}{2} \left[\frac{\sigma_{out}^2 - \sigma_{in}^2}{\sigma_{in}^2 + \sigma_{out}^2} \right] \left[\frac{1}{A_{out}} \left(\frac{I^2(\mathbf{x})}{2\sigma_{out}^2} - 1 \right) + \right.$$
$$\left. + \frac{1}{A_{in}} \left(\frac{I^2(\mathbf{x})}{2\sigma_{in}^2} - 1 \right) \right] \delta(\phi). \qquad (A.10)$$

Computation of spatially-variant blurring operators

B.1 Spatially invariant PSF

Consider the image formation process:

$$Y = P * X \tag{B.1}$$

where $X \in \mathcal{M}_{N \times M}$ is the real data and $Y \in \mathcal{M}_{N \times M}$ is the observation, obtained with the kernel $P \in \mathcal{M}_{n \times m}$.

It's common to represent the observation process in matrix vector notation:

$$\mathbf{y} = H\mathbf{x}$$

where $\mathbf{x}$ and $\mathbf{y}$ are the column vectors obtained via lexicographical ordering of X and Y. Here we want to show how $H \in \mathcal{M}_{NM \times NM}$ looks like for case $N = 9$, $M = 5$, $n = 3$, $m = 3$.

When performing the convolution operation, boundary conditions for X have to be specified. Specifically, this is done by padding of $n - 1$ elements in the axial direction, and $m - 1$ in the lateral one. For computational reasons we will consider here

circular padding. The padded matrix becomes:

$$
X_{\mathrm{p}} =
\begin{bmatrix}
\boxed{x_{8,4}} & \boxed{x_{8,5}} & \boxed{x_{8,1}} & x_{8,2} & x_{8,3} & x_{8,4} & \boxed{x_{8,5}} \\
\boxed{x_{9,4}} & \boxed{x_{9,5}} & \boxed{x_{9,1}} & x_{9,2} & x_{9,3} & x_{9,4} & \boxed{x_{9,5}} \\
\boxed{x_{1,4}} & \boxed{x_{1,5}} & x_{1,1} & x_{1,2} & x_{1,3} & x_{1,4} & x_{1,5} \\
\boxed{x_{2,4}} & \boxed{x_{2,5}} & x_{2,1} & x_{2,2} & x_{2,3} & x_{2,4} & x_{2,5} \\
\boxed{x_{3,4}} & \boxed{x_{3,5}} & x_{3,1} & x_{3,2} & x_{3,3} & x_{3,4} & x_{3,5} \\
\boxed{x_{4,4}} & \boxed{x_{4,5}} & x_{4,1} & x_{4,2} & x_{4,3} & x_{4,4} & x_{4,5} \\
\boxed{x_{5,4}} & \boxed{x_{5,5}} & x_{5,1} & x_{5,2} & x_{5,3} & x_{5,4} & x_{5,5} \\
\boxed{x_{6,4}} & \boxed{x_{6,5}} & x_{6,1} & x_{6,2} & x_{6,3} & x_{6,4} & x_{6,5} \\
\boxed{x_{7,4}} & \boxed{x_{7,5}} & x_{7,1} & x_{7,2} & x_{7,3} & x_{7,4} & x_{7,5} \\
\boxed{x_{8,4}} & \boxed{x_{8,5}} & x_{8,1} & x_{8,2} & x_{8,3} & x_{8,4} & x_{8,5} \\
\boxed{x_{9,4}} & \boxed{x_{9,5}} & x_{9,1} & x_{9,2} & x_{9,3} & x_{9,4} & x_{9,5}
\end{bmatrix}
$$

where padded elements are included within boxes for clarity purposes.

Let's call p_{ij} the generic element of P and introduce the Circulant matrix $\mathbb{P}_k \in \mathcal{M}_{N \times N}$

$$
\mathbb{P}_k =
\begin{bmatrix}
p_{1k} & 0 & 0 & 0 & 0 & 0 & 0 & p_{2k} & p_{3k} \\
p_{2k} & p_{1k} & 0 & 0 & 0 & 0 & 0 & 0 & p_{3k} \\
p_{3k} & p_{2k} & p_{1k} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & p_{3k} & p_{2k} & p_{1k} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & p_{3k} & p_{2k} & p_{1k} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & p_{3k} & p_{2k} & p_{1k} & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & p_{3k} & p_{2k} & p_{1k} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & p_{3k} & p_{2k} & p_{1k} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & p_{3k} & p_{2k} & p_{1k}
\end{bmatrix}
\qquad (\text{B.2})
$$

Let's note that the matrix corresponds to the 1D convolution of the $k - th$ column of P with a vector of N entries.

The H matrix then writes as

$$
H =
\begin{bmatrix}
\mathbb{P}_1 & \bigcirc & \bigcirc & \mathbb{P}_3 & \mathbb{P}_2 \\
\mathbb{P}_2 & \mathbb{P}_1 & \bigcirc & \bigcirc & \mathbb{P}_3 \\
\mathbb{P}_3 & \mathbb{P}_2 & \mathbb{P}_1 & \bigcirc & \bigcirc \\
\bigcirc & \mathbb{P}_3 & \mathbb{P}_2 & \mathbb{P}_1 & \bigcirc \\
\bigcirc & \bigcirc & \mathbb{P}_3 & \mathbb{P}_2 & \mathbb{P}_1
\end{bmatrix}
\qquad (\text{B.3})
$$

where $\bigcirc$ is a $N \times N$ matrix with all entries equal to 0. Matrix H is so made out of $M \times M$ blocks of $N \times N$ circular matrices, and is therefore referred as block circulant with circulant blocks (BCCB). This property is desirable since it makes H diagonalizable via 2D Fourier transform, i.e.:

$$
Y = \mathtt{ifft2D}(\mathtt{fft2D}(P) \cdot \mathtt{fft2D}(X)) \qquad (\text{B.4})
$$

If zero padding is chosen instead then Y can be obtained from:

$$
Y_{\mathrm{padded}} = \mathtt{ifft2D}(\mathtt{fft2D}(P)\mathtt{fft2D}(X_{\mathrm{padded}})) \qquad (\text{B.5})
$$

after discarding rows and columns corresponding to the padded elements.

B.2   Spatially variant PSF

Suppose now that X is blurred with different kernels at different depths. We represent these regions with different capital letters A, B and C and represent an entry of X belonging to one of these regions with a convenient apex.

$$
X = \left[\begin{array}{ccccc}
x_{11}^A & x_{12}^A & x_{13}^A & x_{14}^A & x_{15}^A \\
x_{21}^A & x_{22}^A & x_{23}^A & x_{24}^A & x_{25}^A \\
x_{31}^A & x_{32}^A & x_{33}^A & x_{34}^A & x_{35}^A \\
\hline
x_{11}^B & x_{12}^B & x_{13}^B & x_{14}^B & x_{15}^B \\
x_{21}^B & x_{22}^B & x_{23}^B & x_{24}^B & x_{25}^B \\
x_{31}^B & x_{32}^B & x_{33}^B & x_{34}^B & x_{35}^B \\
\hline
x_{11}^C & x_{12}^C & x_{13}^C & x_{14}^C & x_{15}^C \\
x_{21}^C & x_{22}^C & x_{23}^C & x_{24}^C & x_{25}^C \\
x_{31}^C & x_{32}^C & x_{33}^C & x_{34}^C & x_{35}^C
\end{array} \right] = \left[\begin{array}{c}
X^A \\
\hline
X^B \\
\hline
X^C
\end{array} \right] \tag{B.6}
$$

We introduce the three blurring operators corresponding to the three regions as:

$$
P^L = \left[\begin{array}{ccc}
p_{11}^L & p_{12}^L & p_{13}^L \\
p_{21}^L & p_{22}^L & p_{23}^L \\
p_{31}^L & p_{32}^L & p_{33}^L
\end{array} \right] \tag{B.7}
$$

with $L \in \{A, B, C\}$. If we assume a piecewise constant interpolation between the different kernels, then the global blurring

matrix equivalent to (B.2) becomes:

$$
\mathbb{P}_k =
\left[
\begin{array}{ccc|ccc|ccc}
p_{1k}^A & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
p_{2k}^A & p_{1k}^A & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
p_{3k}^A & p_{2k}^A & p_{1k}^A & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & p_{3k}^B & p_{2k}^B & p_{1k}^B & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & p_{3k}^B & p_{2k}^B & p_{1k}^B & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & p_{3k}^B & p_{2k}^B & p_{1k}^B & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & p_{3k}^C & p_{2k}^C & p_{1k}^C & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & p_{3k}^C & p_{2k}^C & p_{1k}^C & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & p_{3k}^C & p_{2k}^C & p_{1k}^C
\end{array}
\right]
\tag{B.8}
$$

where the sub-matrices corresponding to the single, space in-variant kernels have been enclosed with black contours. The associated H is then build as:

$$
H =
\begin{bmatrix}
\mathbb{P}_1 & \bigcirc & \bigcirc & \bigcirc & \bigcirc \\
\mathbb{P}_2 & \mathbb{P}_1 & \bigcirc & \bigcirc & \bigcirc \\
\mathbb{P}_3 & \mathbb{P}_2 & \mathbb{P}_1 & \bigcirc & \bigcirc \\
\bigcirc & \mathbb{P}_3 & \mathbb{P}_2 & \mathbb{P}_1 & \bigcirc \\
\bigcirc & \bigcirc & \mathbb{P}_3 & \mathbb{P}_2 & \mathbb{P}_1
\end{bmatrix}
\tag{B.9}
$$

Note that linear convolution is considered instead of circular one. As noticed, this does not prevent fast computation since it just needs a padding operation on the input. By introducing the matrices corresponding to the single kernels:

$$
\mathbb{P}_k^L =
\begin{bmatrix}
p_{1k}^L & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
p_{2k}^L & p_{1k}^L & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
p_{3k}^L & p_{2k}^L & p_{1k}^L & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & p_{3k}^L & p_{2k}^L & p_{1k}^L & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & p_{3k}^L & p_{2k}^L & p_{1k}^L & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & p_{3k}^L & p_{2k}^L & p_{1k}^L & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & p_{3k}^L & p_{2k}^L & p_{1k}^L & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & p_{3k}^L & p_{2k}^L & p_{1k}^L & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & p_{3k}^L & p_{2k}^L & p_{1k}^L
\end{bmatrix}
\tag{B.10}
$$

and the corresponding H^L built as in (B.3), then the interpolated H can be written as

$$
H = \sum_L D^L \cdot H^L
\tag{B.11}
$$

Where the diagonal matrices D^L have their entry (j, j) equal to one if the $j-th$ pixel in $\mathbf{x}$ belongs to the region L, otherwise this

is null. We note that the product $D^L \cdot H^L$ corresponds to selecting the entries of $\mathbb{P}$ in (B.8) denoted by an apex L and discarding the others.

The above discussions can be generalized to an arbitrary number of regions N_R as:

$$H = \sum_{i=1}^{N_R} D^i H^i \qquad \text{(B.12)}$$

where H^i is the BTTB matrix corresponding to the blurring i-th kernel and D^i is the diagonal interpolation matrix. Consider that other type of interpolation could be considered a part of linear. They maybe useful to smooth border effects due to substantial changes in the blurring kernel. In these cases the D matrices are still diagonal but have no longer binary entries, cf. [36].

B.3 Preconditioning of spatially variant kernels

In several applications as image restoration it is useful to find suitable preconditioners of the matrix $H^T H$. A preconditioner P of a matrix A is a matrix for which the condition number of the product $P \cdot A$ is smaller than the one of A. A prerequisite of a suitable preconditioner is that matrix vector products of the kind $P\mathbf{x}$ are more easily computable than products $A\mathbf{x}$.

The simplest preconditioner is called *diagonal preconditioner*, or *Jacobi preconditioner*, and consists in the diagonal matrix whose entries are equal to the inverse of the diagonal entries of A [42]. Since the exact form of the Jacobi preconditioner associated to H involves complex products between the different kernels entries which is difficult to formalize and implement, we introduce here a simplified expression of the same. This is derived on the base of the considerations in [113].

Consider the approximation matrix $\tilde{H}$ of H built as:

$$\tilde{H} = \sum_L D^L H^L D^L \qquad \text{(B.13)}$$

It can be shown that, after appropriately reordering of the entries so to make pixels of the same regions contiguous after lexicographical ordering, H can be rewritten as:

$$\tilde{H} = \begin{bmatrix} \tilde{H}^A & \bigcirc & \bigcirc \\ \bigcirc & \tilde{H}^B & \bigcirc \\ \bigcirc & \bigcirc & \tilde{H}^C \end{bmatrix} \qquad \text{(B.14)}$$

where each $\tilde{H}^L$ is an $(3M) \times (3M)$ BTTB matrix built as:

$$
\tilde{H}^L = \begin{bmatrix}
\tilde{P}_1^L & \bigcirc & \bigcirc & \bigcirc & \bigcirc \\
\tilde{P}_2^L & \tilde{P}_1^L & \bigcirc & \bigcirc & \bigcirc \\
\tilde{P}_3^L & \tilde{P}_2^L & \tilde{P}_1^L & \bigcirc & \bigcirc \\
\bigcirc & \tilde{P}_3^L & \tilde{P}_2^L & \tilde{P}_1^L & \bigcirc \\
\bigcirc & \bigcirc & \tilde{P}_3^L & \tilde{P}_2^L & \tilde{P}_1^L
\end{bmatrix}
\tag{B.15}
$$

and

$$
\tilde{P}_k^L = \begin{bmatrix}
p_{1k}^L & 0 & 0 \\
p_{2k}^L & p_{1k}^L & 0 \\
p_{3k}^L & p_{2k}^L & p_{1k}^L
\end{bmatrix}
\tag{B.16}
$$

By consider the BCCB approximation C^L of each BTTB matrix $\tilde{H}^L$ as:

$$
C^L = \begin{bmatrix}
\tilde{P}_1^L & \bigcirc & \bigcirc & \tilde{P}_3^L & \tilde{P}_2^L \\
\tilde{P}_2^L & \tilde{P}_1^L & \bigcirc & \bigcirc & \tilde{P}_3^L \\
\tilde{P}_3^L & \tilde{P}_2^L & \tilde{P}_1^L & \bigcirc & \bigcirc \\
\bigcirc & \tilde{P}_3^L & \tilde{P}_2^L & \tilde{P}_1^L & \bigcirc \\
\bigcirc & \bigcirc & \tilde{P}_3^L & \tilde{P}_2^L & \tilde{P}_1^L
\end{bmatrix}
\tag{B.17}
$$

then the circulant approximation of $\tilde{H}$ can be built as:

$$
C = \begin{bmatrix}
\tilde{C}^A & \bigcirc & \bigcirc \\
\bigcirc & \tilde{C}^B & \bigcirc \\
\bigcirc & \bigcirc & \tilde{C}^C
\end{bmatrix}
\tag{B.18}
$$

Matrix $H^T H$ is then approximated by $C^T C$

$$
C^T C = \begin{bmatrix}
(\tilde{C}^A)^T \tilde{C}^A & \bigcirc & \bigcirc \\
\bigcirc & (\tilde{C}^B)^T \tilde{C}^B & \bigcirc \\
\bigcirc & \bigcirc & (\tilde{C}^C)^T \tilde{C}^C
\end{bmatrix}
\tag{B.19}
$$

It is then straightforward to verify from (B.17) and (B.16) that diagonal entries of (B.19) coincide with the energy of the blur of the corresponding region. This is indeed the way the Jacobi preconditioner is implemented in our algorithm.

Complex Generalized Gaussian Distribution

Consider the complex random variable $z = x_r + jx_i$, where both x_r and x_i obey a GGD with zero mean, variance σ^2 and shape parameter ξ. Assuming mutual independence of x_r and x_i, then $p(z) = p(x_r)p(x_i)$ [16], where $p(x_r)$ and $p(x_i)$ are defined as in (5.1). Hence it is:

$$p(z) = a^2 \exp\left(-\left|\frac{x_r}{b}\right|^\xi - \left|\frac{x_i}{b}\right|^\xi \right) \qquad (C.1)$$

and correspondingly $\ln p(z) \propto (|x_r|^\xi + |x_i|^\xi)$, which, unless $\xi = 2$, is different from $|z|^\xi$. This fact prevents the restoration problem to be formalized as an l^p-norm optimization task.

Hereto, in order to preserve the desired formalism, we define here a variation of the GGD for complex variables, which writes as:

$$p(z) = a^2 \exp\left(-\left|\frac{z}{d}\right|^\xi \right) \qquad (C.2)$$

where d is determined so to satisfy the normalization condition $\int \int p(z)dx_i dx_r = 1$, which leads to

$$d = b \cdot \frac{\sqrt{2/\xi} \cdot \Gamma(2/\xi)}{\sqrt{\Gamma(2/\xi)} \cdot \sqrt{\pi}} \qquad (C.3)$$

The requested integral can be easily computed by substituting $z = \rho \exp(j\theta)$, so that $dx_r dx_i = \rho d\rho d\theta$.

Let's note that in (C.2) the real and imaginary part of z are no longer independent. In particular, the pdf (C.3) is an approximation of the one in (C.1), so that the phaseo z is uniformly

Figure C.1: In (a) the pdf in (C.1) is displayed, while in (b) the pdf (C.2). The two pdfs are displayed for unit variance and $p = 1$.

distributed in the $[0, 2\pi]$ interval (cfr. Figure C.1). Note otherwise that the two pdfs coincide when $p = 2$.

Chen algorithm

The basic of Chen's algorithm are here reported. The interested reader is addressed to the original paper [90] for a detailed explanation. The segmentation process is defined by:

$$\min_{\phi,\mu,R,T} \int_\Omega \delta(\phi(\mathbf{x}))\left\{e(|\nabla I|) + \frac{\lambda}{2}\Psi^2(\mu R(\theta)\mathbf{x} + T)\right\}|\nabla\phi(\mathbf{x})|\mathbf{x}$$

(D.1)

where Ψ is the distance from a reference shape, obtained from a training phase, $e(\cdot)$ is an edge indicator function, while μ, θ and T represent scale, rotation and translation for the optimal registration of the prior shape onto the active contour. R is the rotation matrix associated to the angle θ. Ψ can be computed with numerical methods as fast marching. The optimization is handled using the gradient descent method:

$$\frac{\partial\phi}{\partial\tau} = \delta(\phi)div\left\{\left(e(|\nabla I|) + \frac{\lambda}{2}\Psi^2\right)\frac{\nabla\phi}{|\nabla\phi|}\right\}$$

(D.2)

$$\frac{\partial\mu}{\partial\tau} = -\lambda\int_\Omega \delta(\phi)\Psi\nabla\Psi(\mathbf{x})|\nabla\phi|d\mathbf{x}$$

(D.3)

$$\frac{\partial\theta}{\partial\tau} = -\lambda\int_\Omega \delta(\phi)\mu\Psi\nabla\Psi\left(\frac{dR}{d\theta}\mathbf{x}\right)|\nabla\phi|d\mathbf{x}$$

(D.4)

$$\frac{\partial T}{\partial\tau} = -\lambda\int_\Omega \delta(\phi)\Psi\nabla\Psi|\nabla\phi|d\mathbf{x}$$

(D.5)

where Ψ must be evaluated at $\mu R\mathbf{x} + T$.

Since at time zero the reference shape and the initial contour may be substantially misaligned, a registration step must

precede the segmentation flow. This is implemented by iterating equations (D.3) thru (D.5) until Ψ^2 reaches a minimum bound.

We observe that image gradient, exploited for edges location, is not a reliable indicator in low SNR situations, as in ultrasound images. For this reason, and for comparison purpose, we use in all the experiments the localized framework described in Section 8.1.1 as data attachment term. This implies that the left hand side of (D.2) is in practice replaced by equation (8.7).

Publications list

Journal papers

[P1] T. Dietenbeck, M. Alessandrini, D. Barbosa, J. D'hooge, D. Friboulet and O. Bernard. Detection of the whole myocardium in echocardiography for multiple orientations using a geometrically constrained level-set. *Elsevier - Medical Image Analysis* (submitted for publication).

[P2] M. Alessandrini, S. Maggio, J. Poré e, L. De Marchi, N. Speciale, E. Franceschini, O. Bernard and O. Basset. A restoration framework for ultrasound tissue classification. *IEEE Trans. on Ultrasonics Ferroelectrics and Frequency Control* (submitted for publication).

[P3] M. Alessandrini, T. Dietenbeck, O. Basset, D. Friboulet and O. Bernard. Using a geometric formulation of annular-like shape priors for constraining variational level-sets. *Elsevier - Pattern Recognition Letters* (submitted for publication).

[P4] S. Maggio, A. Palladini, L. De Marchi, M. Alessandrini, N. Speciale, and G. Masetti. Predictive deconvolution and hybrid feature selection for computer-aided detection of prostate cancer. *Medical Imaging, IEEE Transactions on*, 29(2):455 –464, 2010.

Conference papers

[P5] S. Maggio, M. Alessandrini, N. Speciale, O. Bernard, D. Vray, O. Basset, and M. Unser. Continuous domain ARMA modeling for ultrasound tissue characterization. In *Proc. IEEE Trans. on International Symposium on Biomedical Imaging* (ISBI 2011), (accepted for publication).

[P6] T. Dietenbeck, M. Alessandrini, D. Barbosa, J. D'hooge, D. Friboulet, and O. Bernard. Multiview myocardial segmentation in echocardiographic images using a piecewise parametric shape prior. In *Proc. IEEE Trans. on International Symposium on Biomedical Imaging* (ISBI 2011), (accepted for publication)

[P7] M. Alessandrini, S. Maggio, J. Poreé, L. De Marchi, N. Speciale, E. Franceschini, O. Bernard, and O. Basset. An expectation maximization framework for an improved tissue characterization using ultrasounds. In *Proc. of SPIE Medical Imaging (2011)*, 2011, Orlando, USA. (accepted for publication).

[P8] M. Alessandrini, T. Dietenbeck, D. Barbosa, J. D'hooge, O. Basset, N. Speciale, D. Friboulet, and O. Bernard. Segmentation of the full myocardium in echocardiography using constrained level-sets. In *Proc. of Computing in Cardiology Conference (CinC2010)*, pages 409–412, 26–29 Sept. 2010, Belfast, Northern Ireland.

[P9] M. Alessandrini, T. Dietenbeck, O. Basset, D. Friboulet, and O. Bernard. Using a geometric formulation of annular-like shape priors for constraining variational level-sets. In *Image Processing (ICIP), 2010 17th IEEE International Conference on*, pages 669–672, 26-29 Sept. 2010, Hong Kong, China.

[P10] T. Dietenbeck, M. Alessandrini, D. Friboulet, and O. Bernard. Creaseg: a free software for the evaluation of image segmentation algorithms based on level-sets. In *Image Processing (ICIP), 2010 17th IEEE International Conference on*, pages 665-668, 26-29 Sept. 2010, Hong Kong, China.

[P11] M. Alessandrini, D. Friboulet, O. Basset, J. D'hooge, and O. Bernard. Level-set segmentation of myocardium and epicardium in ultrasound images using localized bhattacharyya distance. In *Ultrasonics Symposium (IUS), 2009 IEEE International*, pages 2468 –2471, 2009, Rome, Italy.

[P12] M. Alessandrini, A. Palladini, L.De Marchi, and N. Speciale. Expectation maximization for joint deconvolution and statistics estimation. In *Proc. of International Acoustical Imaging Symposium (AI2009)*, Monterey USA, 2009. (accepted for publication).

[P13] M. Alessandrini, L.De Marchi, and N. Speciale. Recursive Least Squares adaptive filters for ultrasonic signals deconvolution. In *Proc. IEEE International Symposium on Circuits and Systems (ISCAS2008)*, pages 2937–2940, May 2008, Seattle, USA.

[P14] S. Maggio, L.De Marchi, M. Alessandrini, and N .Speciale. Computer aided detection of prostate cancer based on gda and predictive deconvolution. In *Proc. IEEE International Ultrasonics Symposium (IUS2008)*, pages 28–31, Beijing, China, October 2008.

Bilbiography

[1] T. L. Szabo. *Diagnostic Ultrasound Imaging: Inside Out*. 2004.

[2] Alessandro Palladini. *Statistical methods for biomedical signal analysis and processing*. PhD thesis, University of Bologna, 2009.

[3] G. D. Ludwig. The velocity of sound through tissues and the acoustic impedance of tissues. *The Journal of the Acoustical Society of America*, 22:862–866, 1950.

[4] Martin E. Anderson and Gregg E. Trahey. A seminar on k-space applied to medical ultrasound, 2000.

[5] J. Jensen. Linear description of ultrasound imaging systems, 2001. Notes for the International Summer School on Advanced Ultrasound Imaging Technical University of Denmark.

[6] J.A. Jensen and N.B. Svendsen. Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 39(2):262 –267, March 1992.

[7] J. Ng, R. Prager, N. Kingsbury, G. Treece, and A. Gee. Modeling ultrasound imaging as a linear, shift-variant system. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 53(3):549 –563, 2006.

[8] Alan V. Oppenheim, Ronald W. Schafer, and John. R. Buck. *Discrete Time Digital Signal Processing*. Prentice Hall, second edition, 1998.

[9] J. W. Goodman. *Statistical Optics*. Wiley Classic Library Edition, 2000.

[10] P. M. Shankar, R. Molthen, V. M. Narayanan, J. M. Reid, V. Genis, F. Forsberg, C. W. Piccoli, A. E. Lindenmayer, and B. B. Goldberg. Studies on the use of non-rayleigh statistics for ultrasonic tissue characterization. *Ultrasound in Medicine & Biology*, 22(7):873 – 882, 1996.

[11] V. Dutt and J. F. Greenleaf. Ultrasound echo envelope analysis using a homodyned k distribution signal model. *Ultrasonic Imaging*, 16(4):265 – 287, 1994.

[12] P. Mohana Shankar. A general statistical model for ultrasonic backscattering from tissues. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 47(3):727 –736, may. 2000.

[13] O. Bernard, J. D'hooge, and D. Friboulet. Statistics of the radiofrequency signal based on k distribution with application to echocardiography. *IEEE Trans. Ultrason., Ferroelect., Freq. Control*, 53:1689–1694, 2006.

[14] O. Bernard, B. Touil, J. D'hooge, and D. Friboulet. Statistical modeling of the radio-frequency signal for partially- and fully-developed speckle based on a generalized gaussian model with application to echocardiography. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 54(10):2189–2194, oct. 2007.

[15] R.F. Wagner, S.W. Smith, J.M. Sandrik, and H. Lopez. Statistics of speckle in ultrasound b-scans. *Sonics and Ultrasonics, IEEE Transactions on*, 30(3):156 – 163, may. 1983.

[16] Steven M. Kay. *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory (v. 1)*. Prentice Hall, 1 edition, April 1993.

[17] A. Sarti, C. Corsi, E. Mazzini, and C. Lamberti. Maximum likelihood segmentation of ultrasound images with rayleigh distribution. *IEEE Trans. on Ultrasonics Ferroelectrics and Frequency Control*, 52(6):947–960, June 2005.

[18] P. M. Shankar, V. A. Dumane, C. W. Piccoli, J. M. Reid, F. Forsberg, and B. B. Goldberg. Classification of breast masses in ultrasonic b-mode images using a compounding technique in the nakagami distribution domain. *Ultrasound in Medicine & Biology*, 28(10):1295–1300, 2002.

[19]   Olivier Bernard. *Segmentation in echocardiographic imaging using parametric level set model driven by the statistics of the radiofrequency signal.* PhD thesis, University of Lyon, 2006.

[20]   M. R. Spiegel and J. M. Liu. *Mathematical Handbook of Formulas and Tables.* 2nd edition, 1968.

[21]   M.N. Do and M. Vetterli. Wavelet-based texture retrieval using generalized gaussian density and kullback-leibler distance. *Image Processing, IEEE Transactions on*, 11(2):146 –158, February 2002.

[22]   Patrizio Campisi & Karen Egiazarian, editor. *Blind Image Deconvolution: Theory and Applications.* CRC Press, 2007.

[23]   Simon Haykin, editor. *Blind Deconvolution.* Prentice Hall Information and System Science Series. Prentice Hall, 2002.

[24]   E. A. Robinson. *Predictive Decomposition of Time Series with applications to seismic exploration.* PhD thesis, Massachusetts Institute of Technology, 1954.

[25]   J. K. Perloff. Estimation of in vivo pulses in medical ultrasound. *Ultrasonic Imaging*, 16:190–203, 1994.

[26]   S. Maggio, A. Palladini, L. De Marchi, M. Alessandrini, N. Speciale, and G. Masetti. Predictive deconvolution and hybrid feature selection for computer-aided detection of prostate cancer. *Medical Imaging, IEEE Transactions on*, 29(2):455 –464, 2010.

[27]   Nicola Testoni. *Adaptive Multiscale Biological Signal Processing.* PhD thesis, University of Bologna, 2008.

[28]   Simon Haykin. *Adaptive filter theory (3rd ed.).* Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1996.

[29]   T. Taxt and R. Jirik. Superresolution of ultrasound images using the first and second harmonic signal. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 51(2):163 –175, feb. 2004.

[30]   T. Taxt and J. Strand. Two-dimensional noise-robust blind deconvolution of ultrasound images. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 48(4):861 –866, jul. 2001.

[31]   O.V. Michailovich and D. Adam. A novel approach to the
2-d blind deconvolution problem in medical ultrasound.
Medical Imaging, IEEE Transactions on, 24(1):86 –104, jan.
2005.

[32]   T. Taxt. Three-dimensional blind deconvolution of ultra-
sound images. *Ultrasonics, Ferroelectrics and Frequency
Control, IEEE Transactions on*, 48(4):867 –871, 2001.

[33]   J. Ng, H. Gomersall, R. Prager, N. Kingsbury, G. Treece,
and A. Gee. Wavelet restoration of three-dimensional
medical pulse-echo ultrasound datasets in an em frame-
work. In Iwaki Akiyama, editor, *Acoustical Imaging*, vol-
ume 29 of *Acoustical Imaging*, pages 403–408. Springer
Netherlands, 2009.

[34]   Harry C. Andrews and B. R. Hunt. *Digital Image Restora-
tion*. Prentice Hall Professional Technical Reference, 1977.

[35]   J. Ng, R. Prager, N. Kingsbury, G. Treece, and A. Gee.
Wavelet restoration of medical pulse-echo ultrasound im-
ages in an em framework. *Ultrasonics, Ferroelectrics and
Frequency Control, IEEE Transactions on*, 54(3):550 –568,
2007.

[36]   James Nagy and Dianne P. O'leary. Fast iterative image
restoration with a spatially-varying psf. Technical report,
1997.

[37]   M.A.T. Figueiredo and R.D. Nowak. An em algorithm for
wavelet-based image restoration. *Image Processing, IEEE
Transactions on*, 12(8):906 – 916, aug. 2003.

[38]   J.M. Bioucas-Dias. Bayesian wavelet-based image decon-
volution: a gem algorithm exploiting a class of heavy-
tailed priors. *Image Processing, IEEE Transactions on*,
15(4):937 –951, apr. 2006.

[39]   Jonathan Richard Shewchuk. An introduction to the con-
jugate gradient method without the agonizing pain. Tech-
nical report, 1994.

[40]   T. Taxt and G.V. Frolova. Noise robust one-dimensional
blind deconvolution of medical ultrasound images. *Ultra-
sonics, Ferroelectrics and Frequency Control, IEEE Trans-
actions on*, 46(2):291 –299, mar. 1999.

[41] James Kee Huat Ng. Restoration of medical pulse echo ultrasound images. Technical report, 2006.

[42] S. Boyd and L. Vandenberghe. *Convex Optimization*. 2004.

[43] O. Michailovich and A. Tannenbaum. Blind deconvolution of medical ultrasound images: A parametric inverse filtering approach. *Image Processing, IEEE Transactions on*, 16(12):3005 –3019, dec. 2007.

[44] U.R. Abeyratne, A.P. Petropulu, and J.M. Reid. Higher order spectra based deconvolution of ultrasound images. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 42(6):1064 –1075, November 1995.

[45] O. Michailovich and D. Adam. Robust estimation of ultrasound pulses using outlier-resistant de-noising. *Medical Imaging, IEEE Transactions on*, 22(3):368 –381, mar. 2003.

[46] R. Jiric and T. Taxt. High-resolution ultrasonic imaging using fast two-dimensional homomorphic filtering. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 53(8):1440–1448, 2006.

[47] Dennis C. Ghiglia and Mark D. Pritt. *Two-Dimensional Phase Unwrapping: Theory, Algorithms, and Software*. Wiley, 1998.

[48] C. L. Nikias and A. P. Petropulu. *Higher Order Spectral Analysis: A Nonlinear Signal Processing Framework*. Prentice Hall, 1993.

[49] O.V. Michailovich and D. Adam. Phase unwrapping for 2d blind deconvolution of ultrasound images. *Medical Imaging, IEEE Transactions on*, 23(1):7 –25, jan. 2004.

[50] T. Taxt. Restoration of medical ultrasound images using two-dimensional homomorphic deconvolution. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 42(4):543 –554, July 1995.

[51] T. Taxt. Comparison of cepstrum-based methods for radial blind deconvolution of. ultrasound images. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 44(3):666 –674, May 1997.

[52] J. Strand and T. Taxt. Two-dimensional phase unwrapping using robust derivative estimation and adaptive integration. *Image Processing, IEEE Transactions on*, 11(10):1192 – 1200, October 2002.

[53] Jarle Strand and Torfinn Taxt. Performance evaluation of two-dimensional phase unwrapping algorithms. *Appl. Opt.*, 38(20):4333–4344, Jul 1999.

[54] Oleg Michailovich and Allen Tannenbaum. On approximation of smooth functions from samples of partial derivatives with application to phase unwrapping. *Signal Processing*, 88(2):358 – 374, 2008.

[55] Dennis C. Ghiglia and Louis A. Romero. Robust two-dimensional weighted and unweighted phase unwrapping that uses fast transforms and iterative methods. *J. Opt. Soc. Am. A*, 11(1):107–117, Jan 1994.

[56] S. Maggio, M. Alessandrini, L. De Marchi, and N. Speciale. Computer aided detection of prostate cancer based on gda and predictive deconvolution. In *Ultrasonics Symposium, 2008. IUS 2008. IEEE*, pages 28 –31, 2008.

[57] Mehdi Moradi, Parvin Mousavi, and Purang Abolmaesumi. Computer-aided diagnosis of prostate cancer with emphasis on ultrasound-based approaches: A review. *Ultrasound in Medicine & Biology*, 33(7):1010 – 1028, 2007.

[58] D. Fukuoka, Y. Ikedo, T. Hara, H. Fujita, E. Takada, T. Endo, and T. Morita. Development of breast ultrasound cad system for screening. *Lecture Notes in Computer Science*, pages 392–?398, September 2006.

[59] J.M. Thijssen, A. Starke, G. Weijers, A. Haudum, K. Herzog, P. Wohlsein, J. Rehage, and C.L. De Korte. Computer-aided b-mode ultrasound diagnosis of hepatic steatosis: a feasibility study. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 55(6):1343 –1354, jun. 2008.

[60] Frederic L. Lizzi, Michael Greenebaum, Ernest J. Feleppa, Marek Elbaum, and D. Jackson Coleman. Theoretical framework for spectrum analysis in ultrasonic tissue characterization. *The Journal of the Acoustical Society of America*, 73(4):1366–1373, 1983.

[61] T. Liu, F. L. Lizzi, J. A. Ketterling, R. H. Silverman, and G. J. Kutcher. Ultrasonic tissue characterization via 2-d spectrum analysis: Theory and in vitro measurements. *Medical Physics*, 3(34):1037–1046, March 2007.

[62] Frederic L. Lizzi, Ernest J. Feleppa, S. Kaisar Alam, and Cheri X. Deng. Ultrasonic spectrum analysis for tissue evaluation. *Pattern Recognition Letters*, 24(4-5):637 – 658, 2003.

[63] O. Basset, Z. Sun, JL. Mestas, and G. Gimenez. Texture analysis of ultrasonic images of the prostate by means of co-occurrence matrices. *Ultrason Imaging*, 15(3):218–237, July 1993.

[64] Chung-Ming Wu, Yung-Chang Chen, and Kai-Sheng Hsieh. Texture features for classification of ultrasonic liver images. *Medical Imaging, IEEE Transactions on*, 11(2):141 –152, jun. 1992.

[65] F. Rakebrandt, D. C. Crawford, D. Havard, D. Coleman, and J. P. Woodcock. Relationship between ultrasound texture classification images and histology of atherosclerotic plaque. *Ultrasound in Medicine & Biology*, 26(9):1393–1402, 2000.

[66] P. H. Tsui, C. C. Huang, and S. H. Wang. Use of nakagami distribution and logarithmic compression in ultrasonic tissue characterization. *Journal of Medical and Biological Engineering*, 26(2):69–73, May 2006.

[67] D. Donoho. On minimum entropy deconvolution. Technical report, 1981.

[68] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.

[69] Timothy E. Doyle, Adam T. Tew, Keith H. Warnick, and Brent L. Carruth. Simulation of elastic wave scattering in cells and tissues at the microscopic level. *The Journal of the Acoustical Society of America*, 125(3):1751–1767, 2009.

[70] A. R. Webb. *Statistical Pattern Recognition*. 2nd edition.

[71] Statistical pattern recognition toolbox (stprtool). online available at: http://cmp.felk.cvut.cz/cmp/cmp_software.html.

[72] Roxana M. Vlad, Ratan K. Saha, Nehad M. Alajez, Shawn Ranieri, Gregory J. Czarnota, and Michael C. Kolios. An increase in cellular size variance contributes to the increase in ultrasound backscatter during cell death. *Ultrasound in Medicine & Biology*, 36(9):1546 – 1558, 2010.

[73] M.L. Oelze, Jr. O'Brien, W.D., J.P. Blue, and J.F. Zachary. Differentiation and characterization of rat mammary fibroadenomas and 4t1 mouse carcinomas using quantitative ultrasound imaging. *Medical Imaging, IEEE Transactions on*, 23(6):764 –771, jun. 2004.

[74] Department of Public Health. European cardiovascular disease statistics, 2008 edition. online available at: http://www.heartstats.org/datapage.asp?id=7683.

[75] M. Alessandrini, D. Friboulet, O. Basset, J. D'hooge, and O. Bernard. Level-set segmentation of myocardium and epicardium in ultrasound images using localized bhattacharyya distance. In *Ultrasonics Symposium (IUS), 2009 IEEE International*, pages 2468 –2471, 2009.

[76] M. Alessandrini, T. Dietenbeck, O. Basset, D. Friboulet, and O. Bernard. Using a geometric formulation of annular-like shape priors for constraining variational level-sets. In *Proc. IEEE International Conference on Image Processing (ICIP2010)*, pages 669–672, Hong Kong, 2010.

[77] M. Alessandrini, T. Dietenbeck, D. Barbosa, J. D'hooge, O. Basset, N. Speciale, D. Friboulet, and O. Bernard. Segmentation of the full myocardium in echocardiography using constrained level-sets. In *Proc. of Computing in Cardiology Conference (CinC2010)*, Belfast, Northern Ireland, September 2010. accepted for publication.

[78] M. Alessandrini, T. Dietenbeck, O. Basset, D. Friboulet, and O. Bernard. Using a geometric formulation of annular-like shape priors for constraining variational level-sets. *Pattern Recognition Letters*. submitted for publication.

[79] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.

[80] Stanley J. Osher and Ronald P. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*. Springer, 2002.

[81] R. Malladi, J.A. Sethian, and B.C. Vemuri. Shape modeling with front propagation: a level set approach. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 17(2):158 –175, feb. 1995.

[82] Xavier Bresson. *Image Segmentation with Variational Active Contours*. PhD thesis, Ecole Politechnique Federale de Lausanne, 2005.

[83] S.C. Zhu, T.S. Lee, and A.L. Yuille. Region competition: unifying snakes, region growing, energy/bayes/mdl for multi-band image segmentation. pages 416 –423, jun. 1995.

[84] T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Trans. on Image Processing*, 10(2):266–277, February 2001.

[85] S. Lankton and A. Tannenbaum. Localizing region-based active contours. *IEEE Trans. on IP*, 17(11):2029–2039, 2008.

[86] D. Cremers, M. Rousson, and R. Deriche. A review of statistical approaches to level set segmentation: Integrating color, texture, motion and shape. *International Journal on Computer Vision*, 2(72):195–215, 2007.

[87] O. Michailovich, Y. Rathi, and A. Tannenbaum. Image segmentation using active contours driven by the bhattacharyya gradient flow. *IEEE Trans. on Image Processing*, 16(11):2787–2801, November 2007.

[88] Steven M. Kay. *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall, 1993.

[89] T. Chan and Wei Zhu. Level set based shape prior segmentation. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 1164 – 1170 vol. 2, 2005.

[90] Yunmei Chen, Hemant D. Tagare, Sheshadri Thiru-
      venkadam, Feng Huang, David Wilson, Kaundinya S.
      Gopinath, Richard W. Briggs, and Edward A. Geiser. Us-
      ing prior shapes in geometric active contours in a varia-
      tional framework. *Int. J. Comput. Vision*, 50:315–328, De-
      cember 2002.

[91] Mikael Rousson and Nikos Paragios. Shape priors for
      level set representations. In *In Proc. ECCV*, pages 78–92.
      Springer, 2002.

[92] X. M. Pardo, V. Leborán, and R. Dosil. Integrating prior
      shape models into level-set approaches. *Pattern Recogni-
      tion Letters*, 25(6):631–639, 2004.

[93] M. E. Leventon, W. Eric, L. Grimson, and O. Faugeras.
      Statistical shape influence in geodesic active contours. In
      *Computer Vision and Pattern Recognition, 2000. CVPR
      2000. IEEE Computer Society Conference on*, volume 1,
      page 1316, 2000.

[94] E. D. Folland, A. F. Parisi, P. F. Moyniham, D. R. Jones,
      C. L. Feldman, and D. E. Tow.   Assessment of left-
      ventricular ejection fraction and volumes by real-time
      two-dimensional echocardiography: A comparison of
      cineangiographic and radionuclide techniques. *Circula-
      tions*, 66:760 –766, 1979.

[95] D. A. Conetta, E. A. Geiser, L. H. Olivier, A. B. Miller, and
      R.Conti. Reproducibility of left-ventricular area and vol-
      ume measurements using a computer endocardial edge
      detection algorithm in normal subjects. *Amer. J. Cardiol*,
      56:947–952, 1985.

[96] J.A. Noble and D. Boukerroui. Ultrasound image segmen-
      tation: a survey. *Medical Imaging, IEEE Transactions on*,
      25(8):987 –1010, aug. 2006.

[97] J. K. Perloff. Development and regression of increased
      ventricular mass. *Am J Cardiol.*, 50(3):605–611, 1982.

[98] J.M.B. Dias and J.M.N. Leitao. Wall position and thick-
      ness estimation from sequences of echocardiographic im-
      ages. *Medical Imaging, IEEE Transactions on*, 15(1):25 –38,
      feb. 1996.

[99]   V. Chalana, D.T. Linker, D.R. Haynor, and Yongmin Kim. A multiple active contour model for cardiac boundary detection on echocardiographic sequences. *Medical Imaging, IEEE Transactions on*, 15(3):290 –298, jun. 1996.

[100]   S. Kamaledin Setarehdan and John J. Soraghan. Automatic cardiac lv boundary detection and tracking using hybrid fuzzy temporal and fuzzy multiscale edge detection. *Biomedical Engineering, IEEE Transactions on*, 46(11):1364 –1378, nov. 1999.

[101]   C.H. Chu, E.J. Delp, and A.J. Buda. Detecting left ventricular endocardial and epicardial boundaries by digital two-dimensional echocardiography. *Medical Imaging, IEEE Transactions on*, 7(2):81 –90, jun. 1988.

[102]   Maxime Taron, Nikos Paragios, and Marie-Pierre Jolly. Border detection on short axis echocardiographic views using a region based ellipse-driven framework. In *Medical Image Computing and Computer-Assisted Intervention MICCAI 2004*, volume 3216 of *Lecture Notes in Computer Science*, pages 443–450. Springer Berlin / Heidelberg, 2004.

[103]   J. M. Berg. On parameter estimation using level sets. *SIAM Journal on Control and Optimization*, 37:1372–1393, 1998.

[104]   A. Fitzgibbon, M. Pilu, and R.B. Fisher. Direct least square fitting of ellipses. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 21(5):476 –480, may. 1999.

[105]   Petko Faber and R.B. Fisher. Euclidean fitting revisited. In Carlo Arcelli, Luigi Cordella, and Gabriella Di Baja, editors, *Visual Form 2001*, volume 2059 of *Lecture Notes in Computer Science*, pages 165–175. Springer Berlin / Heidelberg, 2001.

[106]   Sung Joon Ahn, Wolfgang Rauh, and Hans-J?rgen Warnecke. Least-squares orthogonal distances fitting of circle, sphere, ellipse, hyperbola, and parabola. *Pattern Recognition*, 34(12):2283 – 2303, 2001.

[107]   V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal on Computer Vision*, 22:61–79, 1995.

[108] Gene H. Golub and Charles F. Van Loan. *Matrix computations (3rd ed.)*. Johns Hopkins University Press, Baltimore, MD, USA, 1996.

[109] T. Dietenbeck, M. Alessandrini, D. Barbosa, J. D'hooge, D. Friboulet, and O. Bernard. Multiview myocardial segmentation in echocardiographic images using a piecewise parametric shape prior. In *In* Proc. IEEE Trans. on International Symposium on Biomedical Imaging *(ISBI 2011)*, 2011. submitted for publication.

[110] Yun Zhu, Xenophon Papademetris, Albert J. Sinusas, and James S. Duncan. A coupled deformable model for tracking myocardial borders from real-time echocardiography using an incompressibility constraint. *Medical Image Analysis*, 14(3):429 – 448, 2010.

[111] V. Chalana and Y. Kim. A methodology for evaluation of boundary detection algorithms on medical images. *Medical Imaging, IEEE Transactions on*, 16(5):642 – 652, October 1997.

[112] D.P. Huttenlocher, G.A. Klanderman, and W.A. Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15:850–863, 1993.

[113] James G. Nagy and Dianne P. O'Leary. Restoring images degraded by spatially variant blur. *SIAM J. Sci. Comput.*, 19:1063–1082, July 1998.

Martino Alessandrini was born in Bologna, Italy, in 1983. He obtained a B.Sc. degree in Electronic Engineering in 2005 and a M.Sc. degree in Electronic Engineering for Biomedical Applications in 2007, both at the Università di Bologna. He was awarded the Ph.D. degree from the same university in 2011, with a thesis concerning deconvolution and segmentation of medical ultrasound images. His main research interests are in the filed of image and signal processing and include solution of ill conditioned problems, variational techniques for image segmentation, restoration and denoising, multi-resolution analysis, statistical estimation theory and adaptive filtering.