Cardiac C-arm computed tomography

Cyril Mory

Thèse CIFRE

Tomographie cardiaque en angiographie rotationnelle

présentée devant

l'Université Claude Bernard Lyon 1

pour obtenir le grade de DOCTEUR

en Traitement du Signal et des Images

par

Cyril Mory

Ecole doctorale : Ecole Doctorale Interdisciplinaire Sciences Santé

Soutenue le 26 février 2014

Jury :

Rapporteurs :	Michel Defrise	Université Libre de Bruxelles
	Xiaochuan Pan	University of Chicago
Examinateurs :	Laurent Desbat	Université Joseph Fourier, Grenoble
	Françoise Peyrin	INSERM (CREATIS)
Directeur de thèse :	Philippe Douek	Université Claude Bernard Lyon 1 (CREATIS)

Encadrement

Cette thèse a été préparée au sein du laboratoire CREATIS, CNRS UMR 5220 – INSERM U1044 – Université Lyon 1 – INSA Lyon, dont le siège est situé au Bâtiment Blaise Pascal (4ème étage), 7 Avenue Jean Capelle, 69621 Villeurbanne Cedex.

Mes travaux se sont déroulés dans d'autres locaux du laboratoire CREATIS, au sein de l'hôpital cardiologique Louis Pradel, 28 avenue du Doyen Lépine, 69677 Bron cedex.

J'ai également été encadré par les membres du laboratoire Medisys de l'entreprise Philips, 33 Rue de Verdun, 92156 Suresnes, et par leurs collègues du laboratoire Philips Technologie GmbH Innovative Technologies, Research Laboratories, Röntgenstrasse 24-26; 22335 Hamburg, Germany.

Remerciements

Pour son expertise dans tant de domaines,
Qu’il partage volontiers et avec enthousiasme,
Pour son soutien indéfectible depuis tant d’années,
Avec une sincère humilité,
Avec une gratitude telle que seule l'exclusivité peut en rendre compte,
Je dédie le fruit de ces trois ans de travail à Jean-Baptiste Paquet.

Résumé

Un C-arm est un appareil d'imagerie médicale par rayons X utilisé en radiologie interventionnelle. La plupart des C-arms modernes sont capables de tourner autour du patient tout en acquérant des images radiographiques, à partir desquelles une reconstruction 3D peut être effectuée. Cette technique est appelée angiographie rotationnelle et est déjà utilisée dans certains centres hospitaliers pour l'imagerie des organes statiques. Cependant son extension à l'imagerie du cœur ou du thorax en respiration libre demeure un défi pour la recherche. Cette thèse a pour objet l'angiographie rotationnelle pour l'analyse du myocarde chez l'homme. Plusieurs méthodes nouvelles y sont proposées et comparées à l'état de l'art, sur des données synthétiques et des données réelles.

La première de ces méthodes, la déconvolution par FDK itérative synchronisée à l'ECG, consiste à effacer les artéfacts de stries dans une reconstruction FDK synchronisée à l'ECG par déconvolution. Elle permet d'obtenir de meilleurs résultats que les méthodes existantes basées sur la déconvolution, mais reste insuffisante pour l'angiographie rotationnelle cardiaque chez l'homme. Deux méthodes de reconstruction 3D basées sur l'échantillonnage compressé sont proposées : la reconstruction 3D régularisée par variation totale, et la reconstruction 3D régularisée par ondelettes. Elles sont comparées à la méthode qui constitue l'état de l'art actuel, appelée « Prior Image Constrained Compressed Sensing » (PICCS). Elles permettent d'obtenir des résultats similaires à ceux de PICCS.

Enfin, deux méthodes de reconstruction 3D+temps sont présentées. Leurs formulations mathématiques sont légèrement différentes l'une de l'autre, mais elles s'appuient sur les mêmes principes : utiliser un masque pour restreindre le mouvement à la région contenant le cœur et l'aorte, et imposer une solution régulière dans l'espace et dans le temps. L'une de ces méthodes génère des résultats meilleurs, c'est-à-dire à la fois plus nets et plus cohérents dans le temps, que ceux de PICCS.

Ces méthodes ont été implémentées avec soin en utilisant la bibliothèque RTK, une bibliothèque C++ de tomographie basée sur ITK. Les opérations les plus gourmandes en temps de calcul ont été portées en CUDA pour être exécutées sur carte graphique. Une reconstruction 3D+temps standard peut ainsi être calculée en environ une heure. Avec plus d'optimisation, et en utilisant du matériel plus performant, ce temps de calcul pourrait encore être réduit et devenir compatible avec les exigences de certaines procédures cliniques.

Résumé moins résumé

Un C-arm est un arc en forme de C fixé sur un bras robotique articulé, utilisé dans le cadre d'examens radiologiques. L'une des extrémités de cet arc est équipée d'un tube à rayons X, l'autre d'un détecteur. Grâce à la mobilité du bras robotisé, un C-arm peut acquérir des images radiologiques 2D depuis presque n'importe quel angle. C'est particulièrement utile en radiologie interventionnelle : les coronarographies et les angioplasties des coronaires, par exemple, nécessitent d'observer le cœur sous différents angles spécifiques pour éviter que les artères ne se superposent les unes aux autres, ou ne soient masquées par d'autres structures. Les C-arms sont utilisés pour des interventions très variées : angiographie, embolisation pulmonaire, remplacement de valve mitrale, exploration électrophysiologique, pose de stent, …

À la fin des années 1990, des systèmes C-arm équipés de fonctionnalités 3D sont apparus. Le principe est similaire à celui du scanner : le C-arm tourne autour du patient tout en acquérant des images (des « projections »), puis ces projections issues de différentes perspectives sont traitées pour obtenir un volume 3D. L'ensemble des méthodes permettant d'obtenir un volume 3D à partir d'une série de projections 2D est l'objet d'une branche des mathématiques : la « tomographie ».

En tomographie cardiaque, il faut composer avec le mouvement du cœur, qui se déforme à chaque battement et se déplace lors de la respiration. Un battement cardiaque peut, très grossièrement, être divisé en deux parties. Pendant la première, appelée « systole », le ventricule gauche du cœur se contracte pour expulser le sang dans l'aorte. Pendant la seconde, appelée « diastole », il se relâche et se gonfle du sang venant des veines pulmonaires. Le cœur se déforme beaucoup pendant la systole, et peu pendant la diastole. Pour l'instant, l'imagerie cardiaque à partir de projections acquises en respiration libre semble hors de portée. De plus, les médecins ne s'intéressent pas au mouvement du cœur induit par la respiration. Il est donc systématiquement demandé aux patients de retenir leur respiration pendant l'acquisition. La tomographie cardiaque a donc pour objet la reconstruction à partir de projections entachées du seul mouvement de battement cardiaque. L'objectif est soit d'obtenir un volume statique du cœur, en général une reconstruction de la diastole, soit d'obtenir un volume dynamique montrant le battement du cœur. Durant cette thèse, l'objectif principal a été d'obtenir une reconstruction précise du mouvement du cœur.

Dans les deux cas, on acquiert l'électrocardiogramme du patient en même temps que les projections. Pour reconstruire un instant donné du cycle cardiaque (une « phase »), on n'utilise que les projections acquises durant cette phase, et on ignore les autres. Par conséquent, pour chaque phase, un sous-ensemble incomplet du jeu de données est utilisé, et le reste est ignoré.

Dans de telles conditions, les méthodes de reconstruction classiques que sont la rétroprojection filtrée de Feldkamp, Davis et Kress (Feldkamp *et al* 1984) et la méthode algébrique SART d'Andersen et Kak (Andersen and Kak 1984) donnent des résultats décevants. Le manque de données génère des artéfacts qui masquent les structures anatomiques et empêchent l'interprétation. Des méthodes spécifiques permettant de reconstruire à partir de peu de données doivent être développées. Cette thèse présente et compare entre elles plusieurs méthodes spécifiquement conçues pour résoudre ce problème.

La première approche que nous avons développée est basée sur une déconvolution ; en effet, un volume reconstruit à partir de peu de projections, contenant des artéfacts dus au manque de données, peut être vu comme la convolution d'un volume « idéal », sans artéfacts, avec une réponse impulsionnelle responsable de l'apparition d'artéfacts. Inverser ce processus de convolution devrait donc permettre de supprimer les artéfacts. Cependant les limites de cette approche apparaissent rapidement. Tout d'abord, le volume reconstruit à partir de peu de projections ne peut être interprété comme le résultat d'une convolution que lorsque le faisceau de rayons X utilisé pour l'acquisition est parallèle, c'est-à-dire que le générateur de rayons X tire des rayons tous parallèles entre eux. Ce n'est malheureusement le cas d'aucun équipement clinique actuel : tous les C-arms utilisent des faisceaux de rayons X coniques. Une approximation est donc nécessaire. Ensuite, le spectre de Fourier de la réponse impulsionnelle contient beaucoup de coefficients nuls. Convoluer avec une telle réponse impulsionnelle implique donc la perte de beaucoup d'information. Dans ces conditions, la déconvolution n'est qu'une manière de remplacer l'information spectrale perdue lors de la convolution par une information spectrale estimée, mais les mécanismes de cette estimation sont peu explicites. La méthode de déconvolution présentée en 2011 par Badea et al. est décrite en détails, et une amélioration est proposée, qui permet entre autres d'exploiter des données acquises avec un faisceau de rayons X divergent. Malgré une amélioration significative de la qualité des résultats, cette nouvelle méthode de déconvolution reste insuffisante pour la reconstruction cardiaque à partir de données C-arm.

Il apparaît nécessaire, étant donné l'échec relatif des méthodes de déconvolution, de traiter explicitement le problème du manque de données. C'est précisément l'objet d'un domaine de recherche ouvert durant la dernière décennie par les travaux de Candès, Donoho et Tao : l'échantillonnage compressé. Fondé sur le constat que la plupart des signaux réels sont compressibles car parcimonieux, c'est-à-dire qu'il existe une base dans laquelle ils peuvent être exprimés avec peu de coefficients non-nuls, l'échantillonnage compressé inclut des travaux sur des méthodes de reconstruction adaptées à des données sous-échantillonnées au sens de Nyquist. Les deux chapitres suivants de cette thèse décrivent des méthodes de reconstruction basées sur l'échantillonnage compressé. Certaines sont adaptées de la littérature, d'autres introduites par nous.

Dans un premier temps, on s'intéresse à la reconstruction d'une seule phase cardiaque. Trois méthodes de reconstruction sont présentées : deux régularisent la solution en minimisant sa variation totale, l'autre s'appuie sur la minimisation de la norme L1 de sa transformée en ondelettes. Ces méthodes diffèrent peu de méthodes existantes (voire pas pour PICCS), mais leur implémentation, leur application à la reconstruction cardiaque de données acquises chez l'humain par un C-arm et la comparaison des résultats qu'elles permettent d'obtenir constituent des nouveautés. Les résultats obtenus montrent la supériorité de la régularisation par variation totale pour reconstruire les données synthétiques, mais sont beaucoup plus homogènes sur les cas réels.

Bien que ces méthodes soient conçues pour reconstruire une seule phase cardiaque, on peut obtenir une reconstruction de l'intégralité du cycle cardiaque en mettant bout-à-bout les reconstructions de chaque phase. Cet exercice fait apparaitre plusieurs problèmes : certaines régions censées être statiques ont une absorption variable au cours du cycle cardiaque, les frontières des organes en mouvement sont floues et l'amplitude des mouvements est parfois inférieure à celle attendue. Ces problèmes appellent le développement de méthodes plus performantes, permettant d'exploiter la cohérence temporelle entre phases cardiaques.

Dans un second temps, on s'intéresse donc à la reconstruction de l'intégralité du cycle cardiaque en une fois. Ceci permet d'ajouter une régularisation temporelle, qui impose que les volumes représentant des phases cardiaques successives diffèrent peu l'un de l'autre. Afin de contraindre au maximum le problème, une segmentation préalable du cœur est utilisée : les différences entre phases cardiaques successives sont acceptées à l'intérieur de la région contenant le cœur, et découragées (ou interdites, suivant la méthode) à l'extérieur. Une première méthode est présentée, mais elle s'avère extrêmement difficile à calibrer. Une seconde méthode, plus simple et plus intuitive, a donc été développée. Les résultats de cette dernière méthode, aussi bien sur données synthétiques que sur données réelles, sont meilleurs que ceux des méthodes de reconstruction 3D ne se focalisant que sur une phase à la fois.

Au début de cette thèse, l'application ciblée était l'imagerie de l'infarctus du myocarde en rehaussement tardif, c'est-à-dire quelques minutes après injection d'un agent de contraste dans les artères coronaires : l'agent de contraste est évacué en quelques secondes des tissus sains, mais stagne plusieurs minutes dans les tissus nécrosés. En effectuant l'acquisition au bon moment, on observe des régions hyper atténuantes dans le myocarde à cause d'une accumulation d'agent de contraste. Ces régions sont les parties infarcies du myocarde, et leur volume est un élément crucial du pronostic vital. Ce phénomène a été observé sur des images acquises par un scanner, et l'objectif de cette thèse était de tenter de l'observer avec un C-arm. Malheureusement, les résultats de reconstruction d'acquisition en rehaussement tardif sont encore décevants, et les méthodes développées durant cette thèse sont peu adaptées au problème. Des modifications du protocole d'acquisition sont suggérées.

Ce document s'achève sur une série de pistes de recherche futures, parmi lesquelles l'application de la méthode 4D ROOSTER développée durant cette thèse, ou de ses principes de base, à d'autres problèmes d'imagerie, et l'utilisation de méthodes de régularisation plus efficaces et mieux adaptées aux images réelles que celles proposées dans ce travail.

Mots clés

Tomographie dynamique, angiographie rotationnelle, cœur, cardiaque, électrocardiogramme, échantillonnage compressé, variation totale, ondelettes, régularisation, 3D, 4D

Cardiac C-arm computed tomography

List of abbreviations

ART: Algebraic Reconstruction Technique (Gordon *et al* 1970)

SART: Simultaneous Algebraic Reconstruction Technique (Andersen and Kak 1984)

SIRT: Simultaneous Iterative Reconstruction Technique (Kak and Slaney 1988)

FBP: Filtered Back Projection (Kak and Slaney 1988)

FDK: Feldkamp, Davis and Kress method (Feldkamp *et al* 1984)

bpm: beats per minute

CT: Computed Tomography

TV: Total Variation (Chan *et al* 2001)

PICCS: Prior Image Constrained Compressed Sensing (Chen *et al* 2008)

4D ROOSTER: Four Dimensions RecOnstructiOn using Spatial and TEmporal Regularization

ECG: ElectroCardioGram

IFDK: Iterative FDK (Medoff *et al* 1983, Sunnegaardh 2009)

IFBP: Iterative FBP

ADMM: Alternating Direction Method of Multipliers (Figueiredo *et al* 2009, Afonso *et al* 2010)

CG: Conjugate Gradient

GPU: Graphics Processing Unit

LVEF: Left Ventricular Ejection Fraction

Abstract

A C-arm is an X-ray imaging device used for minimally invasive interventional radiology procedures. Most modern C-arm systems are capable of rotating around the patient while acquiring radiographic images, from which a 3D reconstruction can be performed. This technique is called C-arm computed tomography (C-arm CT) and is used in clinical routine to image static organs. However, its extension to imaging of the beating heart or the free-breathing thorax is still a challenging research problem. This thesis is focused on human cardiac C-arm CT. It proposes several new reconstruction methods and compares them to the current state of the art, both on a digital phantom and on real data acquired on several patients.

The first method, ECG-gated Iterative FDK deconvolution, consists in filtering out the streak artifacts from an ECG-gated FDK reconstruction in an iterative deconvolution scheme. It performs better than existing deconvolution-based methods, but it is still insufficient for human cardiac C-arm CT. Two 3D reconstruction methods based on compressed sensing are proposed: total variation-regularized 3D reconstruction and wavelets-regularized 3D reconstruction. They are compared to the current state-of-the-art method, called prior image constrained compressed sensing (PICCS). They exhibit results that are similar to those of PICCS.

Finally, two 3D+time reconstruction methods are presented. They have slightly different mathematical formulations but are based on the same principles: using a motion mask to restrict the movement to the area containing the heart and the aorta, and enforcing smoothness of the solution in both space and time. One of these methods outperforms PICCS by producing results that are sharper and more consistent throughout the cardiac cycle.

Much care has been dedicated to the implementation of all these methods: they were all implemented in the RTK framework, a C++ tomography framework based on ITK, and the most demanding operations were ported to GPU. As a result, a typical 3D+time reconstruction can be performed in about one hour. With further optimization and using more powerful hardware, the processing time could be reduced again and become compatible with some existing clinical workflows.

Keywords

Dynamic tomography, C-arm CT, heart, cardiac, electrocardiogram, compressed sensing, total variation, wavelets, regularization, 3D, 4D

Table of contents

Chapter I : Introduction

At the end of the 19th century, several scientists, including William Crookes, Johann Wilhelm Hittorf and Nikola Tesla conducted research on discharge tubes and generated the first X-rays. Wilhelm Conrad Röntgen, a German physicist, was the first person to systematically study X-rays. He was awarded the first Nobel Prize in physics in 1901. The technology of X-ray tubes then evolved quickly: in 1913, William D. Coolidge invented the high vacuum tube, still in use nowadays, and X-ray radiography became a common clinical practice.

In 1917, the Austrian mathematician Johann Karl August Radon invented the Radon transform (Radon 1917): he showed that a two-dimensional function can be uniquely determined from its integrals along parallel lines. His work, combined with Stefan Kaczmarz's discoveries on solving large linear systems in 1937 (Kaczmarz 1937), and later with Alan Cormack's work (Cormack 1963, 1964), laid the mathematical foundations for computed tomography. The first medical CT scanner was developed by Godfrey Hounsfield, and first used for clinical practice in 1971.

Systems combining an X-ray tube and an X-ray detector mounted together on a C-shaped gantry, and therefore called C-arms, appeared in the 50s. C-arms are designed to be mobile, and perform 2D imaging for diagnostic and interventional radiology. In the 90s, C-arms capable of rotating around the patient while acquiring radiographic images were released, which opened the way to C-arm computed tomography, i.e. computed tomography using data acquired by a C-arm (Saint-Félix *et al* 1994, Moret *et al* 1998, Grass *et al* 1999).

This work focuses on cardiac C-arm computed tomography. The heart motion results in specific reconstruction problems. As a consequence, cardiac examinations using CT scanners have only been possible since 2004, with the availability of multi-slice detectors, and cardiac C-arm CT is not yet widely used in clinical practice. In this work, we shall explain why, and propose computed tomography algorithms to address some of the open problems.

I.1. Basics of X-Ray computed tomography

In classical X-ray radiographs, the image is a map of the attenuation of the X-rays along the path they followed from the source to the detector. Bones absorb the X-rays more than soft tissues, soft tissues more than air, and because of these differences in X-ray absorption, various anatomical structures can be identified in the radiographs. But radiographs only contain information about the cumulated attenuation of rays along their path: the contribution of the different organs crossed by the X-ray superimpose, yielding a 2D image sequence with overlapped structures. In contrast, the aim of computed tomography is to determine the local X-ray attenuation of each point on a 3D grid in the body, which gives access to the properties of each anatomical structure.

X-Ray computed tomography is the science of combining several radiographs taken from different perspectives in order to determine the X-ray attenuation of each voxel of the object. Depending on the shape of the X-ray beam and the trajectory of the X-ray source, the methods vary a lot. A short state of the art is presented here. Whenever possible, the methods will be presented using a convex optimization framework, as a means to highlight their similarities and differences.

I.1.a. Reconstruction of a 2D function from its projections

Parallel beam: Radon transform

The 2D Radon transform of a $\mathbb{R}^2 \rightarrow \mathbb{R}$ function is the set of line integrals of this function for angles ranging from 0° to 180°. It is similar to X-ray imaging with a parallel beam X-ray source (a source that generates X-rays that are parallel to each other) which performs a 180° rotation around the object.

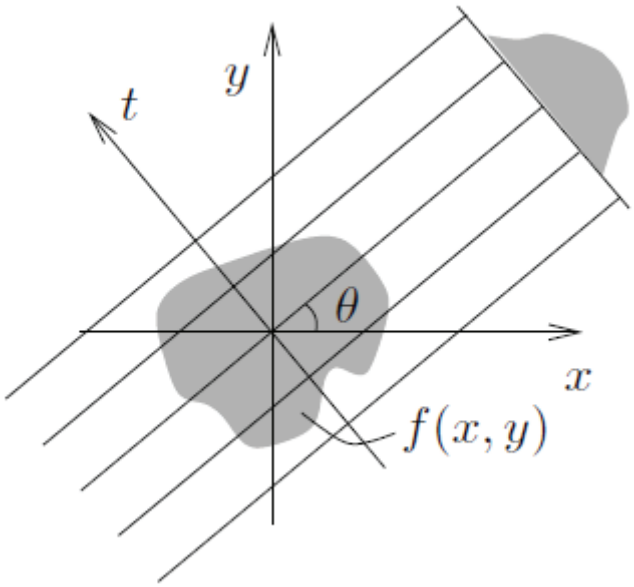

Figure 1. Parallel beam projection of the function f at angle θ [1]

Figure 2 shows a Shepp and Logan phantom (Shepp and Logan 1974) and its 2D Radon transform, also called "sinogram".

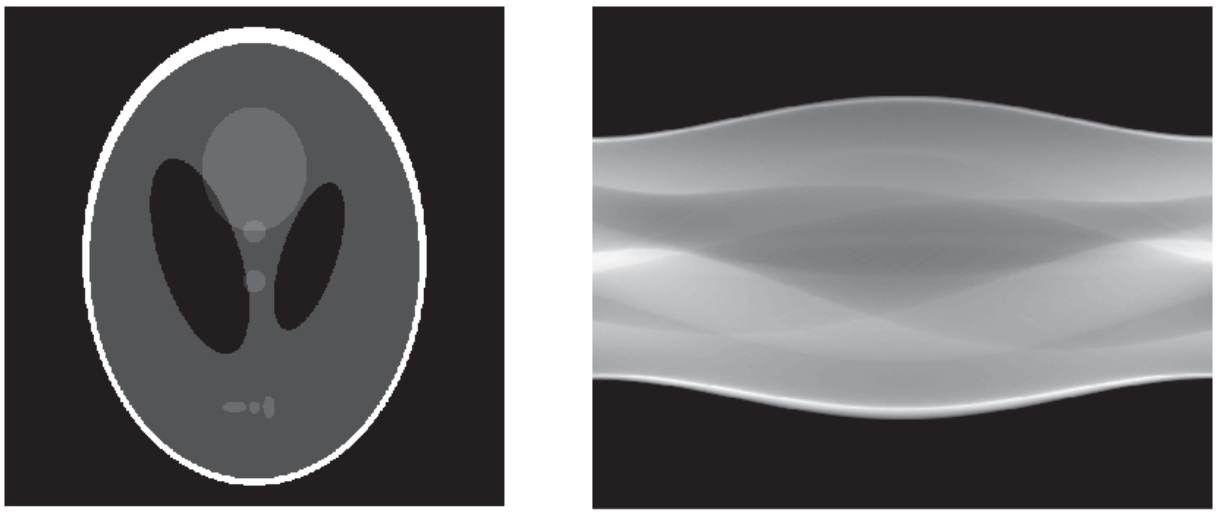

Figure 2. Shepp & Logan phantom (left) and its Radon transform, or "sinogram" (right)

[1] Figure taken from (Turbell, H. 2001)

Fan beam

In its original form, the Radon transform is of little practical interest in modern tomographic imaging: it is hard to design an X-ray source that generates rays parallel to each other. Instead, modern C-arms and CT scanners embed a divergent beam X-ray source: the rays all originate from a single point-shaped source[2], and form a fan (in 2D) or a cone (in 3D).

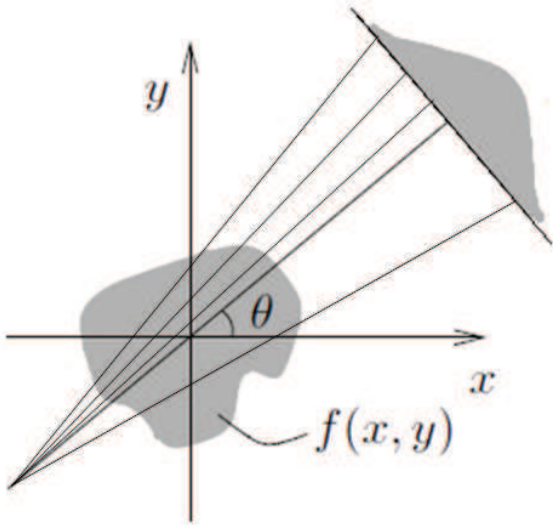

Figure 3. Fan beam projection of the function f at angle θ

Throughout this document, we will refer to the linear map of a 2D function to its parallel beam projections as the "Radon transform", and use the term "X-ray transform" for all the other cases (fan beam and cone beam projection operators). In all calculations, the letter R will denote the projection operator, no matter which projection geometry is used.

Reconstructing a 2D function from its parallel or fan-beam projections is a well studied problem, and essentially three classes of solutions exist:

- the direct Fourier methods (Henry Stark 1981, Gottlieb *et al* 2000), which will not be covered here, because they are not used in any of the methods presented in the next sections
- the Algebraic Reconstruction Technique and all the methods derived from it (Gordon *et al* 1970, Andersen and Kak 1984)
- the Filtered Back Projection and the methods building upon it (Feldkamp *et al* 1984, Mc Kinnon and Bates 1981, Medoff *et al* 1983)

[2] In practice, the source has a small radius, but can be approximated by a point

The Algebraic Reconstruction Technique (ART)

The Algebraic Reconstruction Technique (Gordon *et al* 1970) has been the first iterative reconstruction method applied to X-ray computed tomography. It can be derived from a convex optimization approach.

Let f be the 2D object we are trying to reconstruct, p the set of projections of f we have measured, and R the linear operator performing a discrete approximation of the Radon transform. f is a column vector with as many components as the number of pixels, p is also a column vector, with as many components as the number of samples per projection times the number of projections, and R is a non-square matrix. Each line of R is the set of coefficients used to compute the projection through f along a single ray.

We suppose that $Rf = p + \varepsilon$, where ε is the noise on the projections, and we want to recover f given R and p. This problem can be solved in the least squares sense by defining a cost function $(f) = \|Rf - p\|_2^2$, and finding $\hat{f} = \mathrm{argmin}_f \ J(f)$

$J(f)$ is convex, so it reaches its global minimum when its gradient is **0**, i.e. when

$$R^T(Rf - p) = 0$$

The operator R^T is the adjoint of R, and is called the back projection operator. Back projection consists in setting the values of all points located on the path of a ray to the projection's value. In simple words, the projection is smeared out along the path of the rays. This is illustrated in Figure 4.

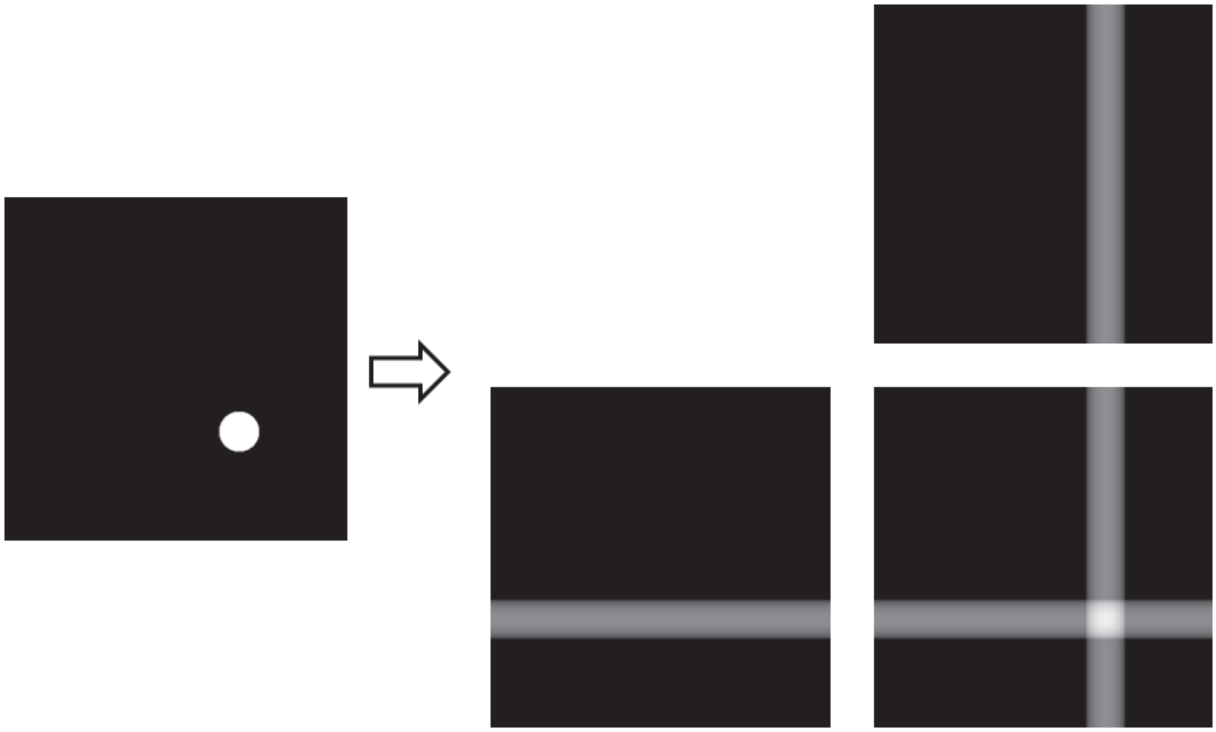

Figure 4. Projections of a circle are back projected along the direction with which they have been acquired. Their sum is shown in the bottom right corner.

The most straightforward approach to minimize $J(f)$ is to perform a gradient descent. The algorithm would start from an initial image f_0, and iteratively update it using the following update step:

$$f_{k+1} = f_k - \alpha_k R^T(Rf_k - p)$$

where α_k is the gradient descent step size. In words, this update step consists in the following operations:

- Projections are calculated through the image f_k: it yields Rf_k

- These projections are compared with the ones actually measured: it yields $Rf_k - p$

- The difference between the computed projections Rf_k and the measured projections p is back projected: it yields $R^T(Rf_k - p)$

- The volume yielded by back projection is subtracted to the current volume, with a weighting coefficient α_k: it yields $f_k - \alpha_k R^T(Rf_k - p)$

This method is called Simultaneous Iterative Reconstruction Technique, short SIRT (Kak and Slaney 1988). It converges to a volume f that minimizes $J(f)$, but the convergence is slow, because every iteration step involves a computationally demanding forward and back projection along all rays. The ART minimizes $J(f)$ by the Kaczmarz method (Kaczmarz 1937), which consists in forward projecting and back projecting one ray at a time. The Simultaneous ART, short SART (Andersen and Kak 1984), is a compromise between SIRT and ART in the sense that it forward projects and back projects simultaneously all the rays acquired with a given angle, or alternatively with a small set of angles, at each iteration. The SIRT, SART and ART methods have different convergence properties (in particular, ART converges only with consistent, noise-free data), but are very similar from a convex optimization standpoint: they all minimize the same cost function $J(f)$, and differ only by the minimization method employed.

Filtered Back Projection (FBP)

Filtered Back Projection is a reconstruction method that implements the analytical inversion of the Radon transform. It can also be adapted to reconstruct from fan beam projections. It is described extensively in (Kak and Slaney 1988, Turbell, H. 2001), in which all theoretical foundations can be found.

As the name suggests, Filtered Back Projection consists in two steps: filtering, and back projecting. The projections are first convolved with a high-pass filter, called “ramp filter” because of its V-shaped profile in the Fourier domain, and then back projected. Using linear operators, FBP can be expressed as follows:

$$\hat{f} = R^T K p$$

where K is the ramp filtering operator. FBP requires only a single back projection, and can be performed in a few seconds on modern hardware. This is therefore the most widespread reconstruction method in commercial CT scanners (Pan *et al* 2009).

I.1.b. Generalization to 3D

Cone beam geometry

Modern X-ray imaging devices rely on a source that generates a cone-shaped beam of X-rays. It rotates around the patient following either a helical trajectory, for CT scanners, or a circular trajectory, for CT scanners and C-arm systems.

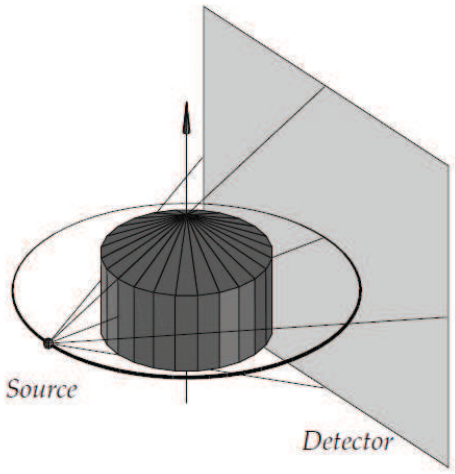

Figure 5. Cone beam source and planar detector following a circular acquisition trajectory[3]

Practical cone beam algorithm

The practical cone beam algorithm (Feldkamp *et al* 1984), also called FDK from the names of its inventors Feldkamp, Davis and Kress, is a method adapted from the 2D fan beam filtered back projection in order to reconstruct an object from cone beam projections acquired on a circular trajectory. It consists in weighting the filtered projection data before back projecting it: the further the considered ray is away from the cone's axis, the lower the weight.

In 1983, Tuy identified a criterion for exact reconstruction to be possible: every plane intersecting the object must also intersect the trajectory of the source. A circular trajectory does not meet this criterion for exact reconstruction (Tuy 1983), therefore FDK can only be an approximate method. Only the central slice (the intersection of the volume and the plane containing the source trajectory) can be reconstructed exactly. The further away from the central slice, the more the volume is affected by so-called "cone artifacts" (Sunnegaardh 2009, Valton 2007).

I.1.c. Managing motion in cardiac CT

Cardiac imaging requires specific acquisition and processing methods because of the observed motion. In a human being at rest, this motion has two main causes: breathing and heart beating.

Breathing motion

The heart undergoes a movement close to a translation during inhalation, and the opposite movement during exhalation. The amplitude of this movement is typically larger than what is induced by heart beating. It is difficult to compensate for it, especially when it is mixed with

[3] Figure taken from (Turbell, H. 2001)

heartbeat motion. In addition, the displacement of the heart induced by breathing is irrelevant in most clinical situations. As a consequence, patients are asked to hold their breath during acquisitions, so as to avoid breathing motion. Most patients are old (Go *et al* 2013) and suffer from various cardiac and vascular diseases, which affects their breathing and breath holding capabilities. Physicians estimate that patients should not be required to hold their breath for more than 10 to 15 seconds.

Breathing motion is periodic, but hard to reproduce. This means that if a patient holds his breath for a while, then breathes normally, and then starts holding his breath again, there is little chance that his chest will be in same mechanical state as during the first breath hold. In cardiac C-arm CT, this means that acquisitions must be performed during a single breath hold.

Heartbeat motion

There are several ways to deal with the motion induced by heart beating. One can just choose not to take it into account, and reconstruct the cardiac volume as if the heart had been static. The result is called an "ungated" reconstruction: the static structures like the rib cage and spine are well reconstructed, but the contours of the heart are not sharp, and the location of fine and highly contrasted objects like iodine-injected coronary arteries is imprecise. Alternatively, one can also acquire the ECG signal of the patient simultaneously to the projections, select the cardiac phase (the moment in the cardiac cycle) one wants to reconstruct based on the ECG, and perform the reconstruction using only the projections acquired around this phase. This approach is called retrospective gating and is illustrated in Figure 6. Mathematically, it amounts to multiplying all projections by a temporal weighting function shaped like a gate, thus the term "gating"(Kachelriess and Kalender 1998). A third approach is to acquire only the data that corresponds to a certain cardiac phase, and switch off the X-ray tube during the other phases. This last approach is called prospective gating, and is commonly referred to as "step and shoot" (Hsieh *et al* 2006). It delivers a lower dose to the patient, but requires real-time ECG R-peaks detection and prevents from performing an ungated reconstruction. Step and shoot acquisitions require fast acceleration and braking of the gantry, which is currently possible on CT scanners but not on C-arms. All the experiments presented in this work have been carried out with retrospectively gated acquisitions.

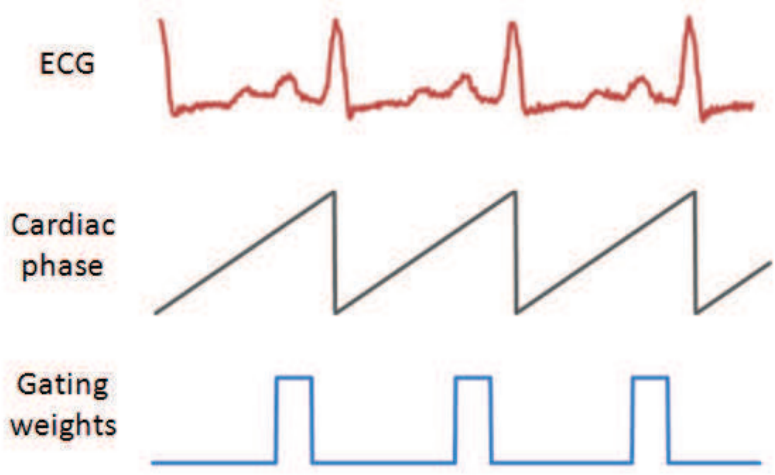

Figure 6. Retrospective ECG-gating: the target cardiac phase is selected (here 80% of the cardiac cycle), and only the projections acquired during a cardiac phase close to the target are kept. The other ones are discarded. The projections kept form clusters of consecutive projections

I.2. C-arm computed tomography

I.2.a. Presentation

Figure 7. Three modern C-arm systems attached to the ceiling of the catheter lab

A C-arm is an X-ray imaging device made of an X-Ray tube and a detector mounted on a C-shaped gantry, as can be seen on Figure 7. Because the gantry can be translated and rotated around all three axes, the C-arm is a convenient tool to image the vascular system, in particular during an intervention. It is designed to acquire and display 2D images (also called 'projections' throughout this thesis) in real time.

In the late 90's, C-arm systems capable of automatically and precisely rotating around the patient were released, and C-arm computed tomography became possible.

I.2.b. Clinical interest

The initial objective of this PhD thesis was to develop a specific reconstruction method for late enhancement myocardium imaging in C-arm CT. The principle is the following: acute myocardial infarction occurs when a patient's artery gets suddenly obstructed (e.g. by a clot or by the remains of an atherosclerosis plaque). In a few minutes, the muscle cells that were irrigated by the culprit artery die. In most cases, the patient is treated in the catheter lab, and the culprit artery is re-opened. Some of the contrast medium injected during the treatment in the catheter lab accumulates in the infarcted region, because it washes out slower from dead tissues than from healthy ones. It has been proved in (Boussel *et al* 2008) that an 80kV CT scanner acquisition on the patient's heart performed a few minutes after the intervention presents a brighter region in the infarction area. This technique allows evaluating the infarction size and location, both of which are key prognosis parameters. If this reconstruction could be performed in the catheter lab, instead of requiring moving the patient to the CT scanner, it would save both the patient and the physician a lot of time. Developing the tomography algorithms for this application was the initial aim of this thesis, which over the course of these three years took a slightly different direction. The reasons for this change, and the results obtained on late enhancement datasets, are discussed in section IX.2.

Late enhancement imaging, however, is not the only potential application of C-arm CT. There are many situations in which 3D imaging during a C-arm guided intervention is desirable. In vascular surgery, for example, physicians could start their intervention by a 3D acquisition with contrast medium, which would then serve as a roadmap over the fluoroscopy to guide the intervention (Glöckler *et al* 2013). The system would merge the *live* fluoroscopic images, which show the current position of the surgical tools in the patient's body, with a 2D image calculated from the 3D, which shows the location of the vessels. Such a procedure would require less contrast medium than the

current clinical practice, which consists in injecting a bolus of iodine to make the vessels visible in fluoroscopy every time an ambiguity arises.

A comprehensive review of the various ways C-arm computed tomography is used or could be used can be found in (Wallace *et al* 2008, Grangeat 2009).

I.2.c. Challenges specific to cardiac C-arm CT

Cardiac reconstruction from C-arm CT implies specific challenges that do not exist on CT scanners. Therefore, CT reconstruction algorithms cannot be exploited without major modifications on C-arm CT data.

Modern CT scanners can perform a 360° rotation in 250ms, while C-arm devices are slower and limited to short-scan acquisitions. The fastest C-arms currently available require around 3 seconds for a 220° rotation. The cardiac rhythm at rest, for humans, is on average 60 beats per minute. As a consequence, a cardiac C-arm CT acquisition is never free of cardiac motion and will capture at least 3 cardiac cycles.

Some C-arm devices perform several short scans around the patient, by going back and forth several times (Lauritsch *et al* 2006). Assuming that the patient's heart beats regularly, and each sweep can be started by ECG-triggering, this kind of acquisition procedure can fill some of the gaps in the angular sampling for a particular heart phase. However, multi-sweep acquisitions require a long breath hold and a long contrast medium injection.

In this work, the focus is on developing tomography algorithms for a particular acquisition procedure, which has the following properties:

- Only one sweep, in order to be compatible with most patients' breath hold capabilities
- 308 projections
- 10.3 seconds
- 210° rotation on a circular arc

The first clinical application targeted is the measurement of the left ventricular ejection fraction (short LVEF), i.e. the ratio between the amount of blood that the heart pumps out at each cycle and the amount of blood the left ventricle stores when fully dilated. Measuring the LVEF requires an easy-to-segment reconstruction of the left ventricle over the whole cardiac cycle.

Impact of the angular distribution

In tomography with few views, the reconstruction quality depends not only on the number of projections used, but also on the way they cover the angular range around the object (Schwartz *et al* 2011, Chen *et al* 2012). In ECG-gated cardiac C-arm CT, the angular distribution of the projections is determined by the periodic motion of the patient's heart, and results in a few clusters of consecutive projections (one cluster per heart beat), separated by empty angular regions (where other cardiac phases are acquired), as shown in Figure 6.

Projections from the same cluster are very similar, and therefore each cluster brings only marginally more information than its central projection alone. Consequently, a dataset with many clusters of few projections is preferable to a dataset with few clusters of many projections. This effect is

highlighted in Figure 8, which shows SART reconstructions of a Shepp & Logan phantom from 60 projections grouped in 10, 20, 30 and 60 equally spaced clusters of 6, 3, 2 and 1 projections respectively.

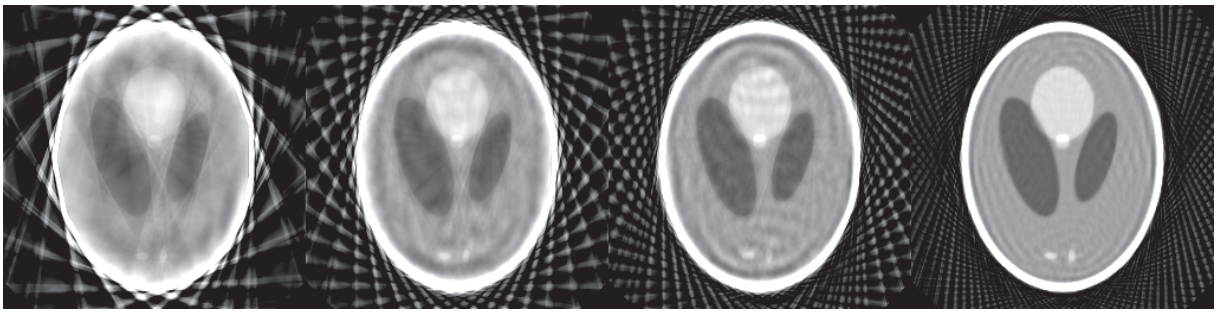

Figure 8. SART reconstructions of the Shepp and Logan phantom from 60 projections, grouped, from left to right, into 10, 20, 30 and 60 equally spaced clusters of 6, 3, 2 and 1 projections respectively.

The number of heart cycles captured during the acquisition is therefore a crucial parameter in the final reconstruction quality of ECG-gated cardiac C-arm CT data. Unfortunately, little can be done to influence it since it only depends on the patient's heart rhythm (the higher, the better) and on the acquisition time (the longer, the better) – and the latter has to remain short to fit in a single breath hold and limit the amount of contrast medium.

A consequence is that the number of heart cycles during the acquisition must be considered with great attention when assessing the quality of a reconstruction

Evaluation of the results

The goal of this work is to design a 3D + time reconstruction method for human cardiac C-arm CT. Ideally, the resulting 3D+t reconstruction would have the same textures' structure as the ungated FDK reconstruction, be free of artifacts, and animated by a realistic cardiac movement. It seems however illusory to target for a perfect dynamic reconstruction of the heart in the difficult conditions detailed above.

In order to compare the reconstruction results between different methods, or between different sets of parameters within the same method, we must evaluate both the quality of each volume representing a given cardiac phase, and the quality of the motion pattern of the whole sequence. Several metrics are proposed to quantify these aspects on the simulated data, where a ground truth is available, and real data reconstructions are assessed by visual evaluation and Contrast to Noise Ratio (CNR) measurements. Visual evaluation was performed by myself and by trained radiologists, with the following criteria in mind:

- Size, shape, contrast with the background and sharpness of the edges of the beating ellipse, for the Shepp and Logan phantom
- Sharpness of the contours of the left ventricle, for real data
- Intensity of the streak artifacts and temporal consistency, for both kinds of datasets

Chapter II : State of the art in Cardiac C-arm Computed Tomography

A number of methods have been developed to reconstruct 3D or 3D+time volumes from cardiac C-arm data. As C-arm systems are mostly used to diagnose and treat vascular pathologies, it is not surprising that many of these methods aim to reconstruct the coronary tree (Langet *et al* 2011, 2012, Blondel *et al* 2006, 2004, Hansis *et al* 2010, 2008). Reconstructing the surface of the left ventricle, in order to measure functional parameters such as the left ventricular ejection fraction, has also been a topic of research (Müller *et al* 2013). In this work, we focus on whole heart reconstruction, and will therefore not describe in details the methods designed specifically for coronary arteries or left ventricle reconstruction.

II.1. Methods of reference

ECG-gated FDK

ECG-gated FDK consists in performing an FDK reconstruction with ECG-gated data. When the gated projections are sufficiently well distributed in the angular space, it yields satisfactory images. When they are not, the reconstructed image contains line-shaped artifacts called “streak artifacts” or “streaks”, which hamper the medical interpretation. In our case, the gaps in the angular sampling are large, and ECG-gated FDK is not suitable. Figure 9 shows a simulated ECG-gated reconstruction of a phantom with the angular sampling obtained in practice on clinical data.

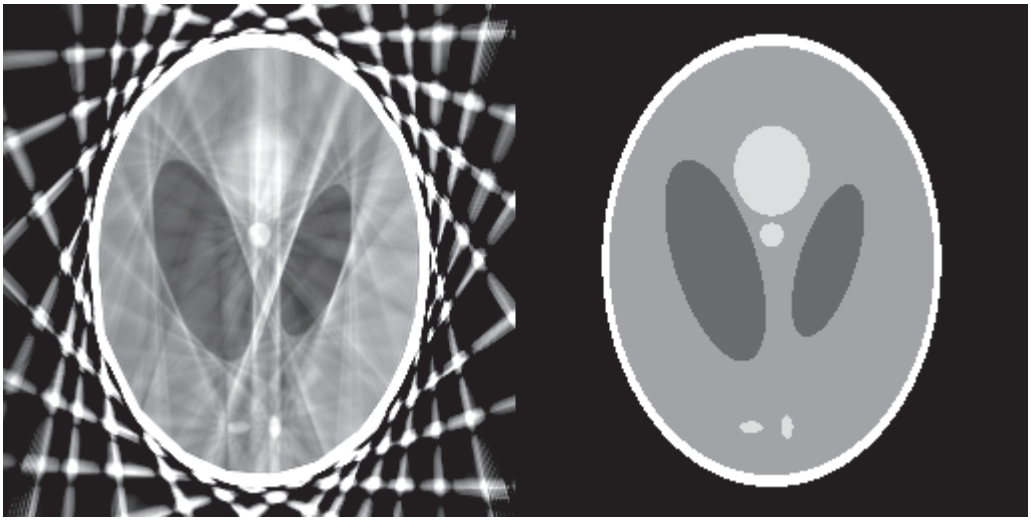

Figure 9. ECG-gated FDK reconstructions of the Shepp and Logan phantom from 300 projections (left) and ground truth (right). The streaks are too intense for such a reconstruction to be usable in clinical practice

ECG-gated ART (convex optimization approach)

The same ECG-gating approach can be applied to ART. It is interesting to formalize it in a convex optimization framework since most of the novel algorithms introduced in this thesis can be viewed as more complex convex optimization-based algorithms. ECG-gated ART aims at minimizing the following cost function:

$$J(f) = \|G(Rf - p)\|_2^2$$

where G is the linear operator performing the gating, and the rest of the notations are identical to those used in section I.1.a. G is a diagonal matrix, and its diagonal coefficients are the square root of the gating weights. The minimum of this cost function is reached when the gradient of J is null, i.e. when

$$R^T G^T G(Rf - p) = 0$$

$$\Leftrightarrow R^T G^2(Rf - p) = 0$$

because G is diagonal. Thus in any gradient-based method to minimize $J(f)$, the update step involves multiplying the difference between artificial projections through the current volume, and the measured projections, by G^2. Note that in the case of binary gating (when a projection is either selected or rejected, so the weights are either 0 or 1), $G^2 = G$. The gating matrix G will be used extensively in chapter Chapter I.

McKinnon Bates(Mc Kinnon and Bates 1981)

In 1984, McKinnon and Bates proposed the following method that builds upon ECG-gated FDK:

- Start with a standard ungated FDK. The regions undergoing movement (in practice, the heart) are blurred, but the static regions are well defined
- Forward project the resulting volume
- Subtract these projections from the measured ones: only the dynamic region containing the heart should be non-zero
- Compute an *ECG-gated* FDK of these subtracted projections
- Add the ECG-gated FDK from difference projections to the ungated FDK

This method is fast, it significantly reduces the amount of streak artifacts with respect to ECG-gated FDK, but neither the reduction in streaks nor the temporal resolution obtained are sufficient in our case (see section V.3.b, as the iterative method described therein boils down to McKinnon-Bates' method if only one iteration is performed).

II.2. Deconvolution

A second class of methods performs an ECG-gated reconstruction (FBP or ART), and then attempts to filter out the streak artefacts using a deconvolution scheme. Such a method has first been proposed in (Dhawan *et al* 1985). The deconvolution was performed using the Wiener filter. It has been recently reintroduced by Badea et al. in the context of 4D micro-CT (Badea et al 2011) with a modified Wiener filtering. Their method is simple to implement, fast, has only one parameter, and gives satisfying results on 4D cardiac micro CT in the mouse. Unfortunately, the results on human

data are disappointing because of the difference in cardiac rhythms between humans and small animals (mice have a cardiac rhythm approximately ten times higher than humans) and of the difference in rotation speed between C-arms and CT scanners. Section Chapter I is dedicated to a deconvolution method using iterative FBP and its comparison with Badea's method.

II.3. Motion compensation

An intuitive approach to reconstruct from data containing movement is to estimate the motion and compensate for it. Every motion-compensated reconstruction method relies on a first step of motion estimation, either from the projections, from the reconstructed volumes, or from both. Unfortunately, estimating the movement in cardiac C-arm CT data is difficult, and none of the motion compensation methods published at the moment has proven to be efficient in the situation considered here (single sweep human cardiac C-arm CT for whole heart reconstruction). Since I have not explored this kind of methods in our work, the following is only a quick overview of the state of the art in the domain.

Prümmer has worked on whole heart motion-compensated reconstruction from multi sweep C-arm CT data in his PhD thesis (Prümmer 2009) and in various articles (Prümmer *et al* 2009, Rohkohl *et al* 2010). The methods proposed in these articles are unfortunately not suitable for single sweep acquisitions: in single sweep acquisitions, there are large gaps in the angular sampling, while in multi sweep acquisitions the gaps left in the angular sampling during one sweep can be partially filled by subsequent sweeps. For this reason the motion-compensated methods underlying Prümmer's methods cannot be applied with success in our case.

In 2012, Müller *et al* proposed a reconstruction method for whole heart cardiac C-arm CT from single sweep acquisitions, called Combined Multiple Heart Phase Registration (CMHPR) (Müller *et al* 2012). CMPHR consists in computing a series of ECG-gated reconstructions (based on compressed sensing, using the PICCS method, see 0), then estimating a 4D movement vector field from it, and finally performing a motion-compensated reconstruction. The results are demonstrated on two acquisitions on porcine models, containing 25 and 32 heart beats respectively. Müller *et al* also proposed a second method using surface-based motion compensation in order to reconstruct the left ventricle: a moving endocardium is first reconstructed from the projections, then its movement is estimated, and then extended to a dense movement vector field. Finally, the dense movement vector field is used for the motion-compensated reconstruction (Müller *et al* 2013). These methods rely on the possibility to estimate the movement accurately, which in turn depends on the efficiency of the compressed sensing-based method used for the initial ECG-gated reconstructions.

II.4. Compressed sensing

Compressed sensing is a recent development of the signal processing theory. It has drawn considerable attention during the last years, mostly following the work of Candès *et al*. (Candes *et al* 2006b, Candes and Wakin 2008, Candes *et al* 2006a). Very roughly, this research field aims to investigate the links between signal sampling conditions, reconstruction methods and the concept of sparsity. A signal is said to be N-sparse in a certain basis if it has at most N non-zero coefficients in this basis. Numerical methods have been developed to recover a signal from far fewer (adequately performed) measurements than what is dictated by the Nyquist criterion, under the assumption that the signal has a correct sparsity in a certain basis.

Real signal are never exactly sparse in bases we use, but they may have only a few high energy coefficients and many very low energy coefficients if the basis is well chosen. Even with relaxed hypotheses on the signal's sparsity and no control on the signal sampling process, compressed sensing-based methods remain efficient for solving ill-posed inverse problems. Constraining the solution to be sparse in a certain basis enforces a certain form of regularity (which depends on the basis), and rules out irregular solutions. Thus it often referred to as "regularization". The choice of a basis in which the signal is supposed to be sparse, or equivalently of a sparsifying transform, is a critical question.

Sparsity in space

Several authors have proposed 3D reconstruction methods based on the spatial sparsity of the coronary tree (Li *et al* 2002, Hansis *et al* 2008, Langet *et al* 2012). Langet later proposed an improvement using both spatial and temporal regularization (Langet *et al* 2012). While coronary arteries filled with contrast medium are thin and very bright objects (they are sparse in space), the myocardium is not. These methods cannot be straightforwardly adapted to whole heart reconstruction. However, the idea of combining spatial and temporal regularization, proposed in (Langet *et al* 2012) is developed in the methods presented in sections 0 and 0.

Sparsity of the gradient

A widespread application of the compressed sensing theory is the reconstruction using total variation (short TV) minimization. The total variation of an image f is defined as the L1 norm of the gradient's magnitude. Images with low TV are typically piecewise constant, therefore imposing a low TV during a reconstruction forces the output to be spatially regular.

ASD-POCS (Sidky and Pan 2008, Bergner *et al* 2010) is a compressed sensing reconstruction method based on such a regularization. It iteratively minimizes the current volume's TV, under the constraint that the difference between the measured projections and the estimated ones lies under a given threshold.

Another recent method is the Prior Image Constrained Compressed Sensing (Chen *et al* 2008, 2012, Lauzier *et al* 2012a), which also enforces a low TV, both on the reconstructed image itself and on the difference between the reconstructed image and a "prior image". The prior image is supposed to be close to the image to reconstruct. In practice the prior image used is the ungated FDK reconstruction. PICCS is currently the only method available in the literature to reconstruct from single sweep C-arm CT data with a good temporal resolution. At the moment, it has only been demonstrated on animals (several pigs and a dog), not yet on human C-arm CT data. Section 0 of this document is dedicated to PICCS.

Sparsity on a wavelets basis

Real images, including C-arm CT reconstructions, contain textured regions, which typically have a high total variation, and therefore cannot be accurately reconstructed by enforcing a low TV constraint. Several reconstruction methods using a wavelet transform as sparsifying transform have been proposed (Daubechies *et al* 2004, Jia *et al* 2011, Dong *et al* 2013), which supposedly circumvent this inadequacy of TV. A method using the Daubechies wavelet transform as sparsifying basis is presented in details in section 0.

Dictionary-based sparsifying transform

A number of methods propose to learn an efficient sparsifying transform from a close prior image (or a set of close prior images) (Babacan *et al* 2011). A dictionary is learnt from the prior images, and is used for the reconstruction.

Non-local means regularization (Jia *et al* 2010, Hao *et al* 2012) can also be considered a dictionary-based compressed sensing method in which the basis is learnt from the reconstructed image itself.

Chapter III : Data

Throughout this document, three kinds of datasets will be used to evaluate the presented methods: a Shepp & Logan phantom modified to mimic a beating heart (one of the ellipsoids periodically dilates and shrinks over time), several injected human C-arm CT scans, and a late enhancement human C-arm CT scan. The injected scans were acquired at the department of medicine, division of cardiology, University of Colorado Denver, Aurora, USA. The late enhancement dataset was acquired at the radiology department, Hôpital Cardiologique Louis Pradel, Bron, France. These three kinds of datasets share the same geometry and number of projections, but differ by many aspects. In this chapter, each kind of dataset is briefly described.

III.1. Shepp and Logan phantom

The Shepp and Logan phantom (Shepp and Logan 1974) is a numerical phantom built out of ellipsoids with different shapes and absorptions. The line-integrals through an ellipsoid can be computed exactly; therefore it is possible to calculate exact projections through this phantom, which made it very popular. It was initially designed to simulate a human head, with a low contrast similar to that existing between white and grey matter in the brain. For the purpose of this study, a modified Shepp & Logan phantom with higher contrast has been used. One of the ellipsoids was used to simulate a beating heart: it shrinks and dilates periodically, its diameter following a sinusoidal function. Only the volume of the ellipsoid changes over time; its absorption remains constant. Figure 10 shows cuts of this Shepp & Logan phantom in systole and diastole. The beating ellipsoid is pointed out by a red arrow.

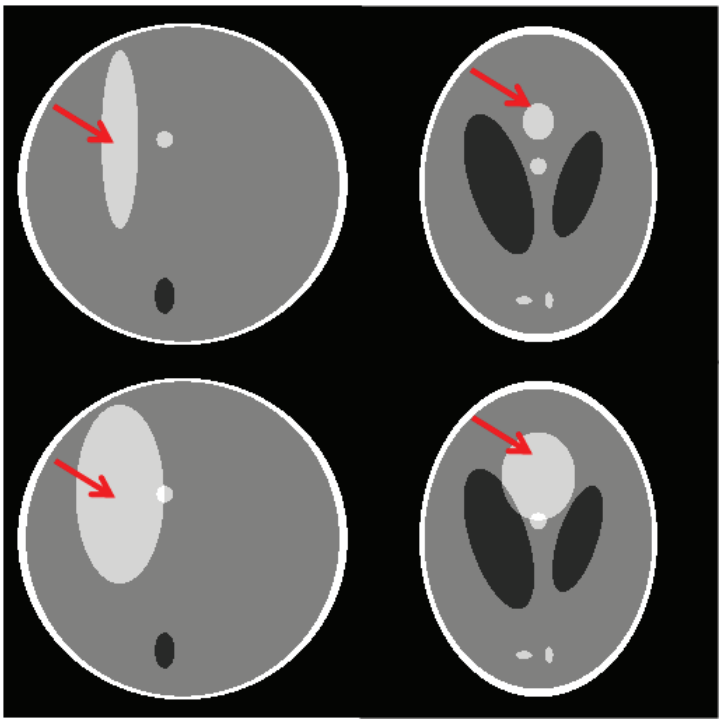

Figure 10. Ground truth of the beating Shepp and Logan phantom used throughout this thesis. Systole is displayed on the top row, diastole on the bottom row. The left column contains cuts along a plane containing the rotation axis, and the right column a cut along a plane normal the rotation axis.

The absorption coefficient of this Shepp and Logan phantom are shown in Figure 11.

Figure 11. Attenuation coefficients of the Shepp & Logan phantom used in this thesis

The projections of this phantom were computed using exact line-integral calculations. 308 projections evenly distributed over a 205 degrees circular arc trajectory were used. A 10.3 seconds acquisition was simulated, with the beating ellipsoid performing 64 beats per minute. 10 phases were reconstructed, each one as a 256 * 256 * 256 voxels volume.

Figure 12 shows the ungated FDK reconstruction of this phantom, in which the motion blurring can be observed in both slices.

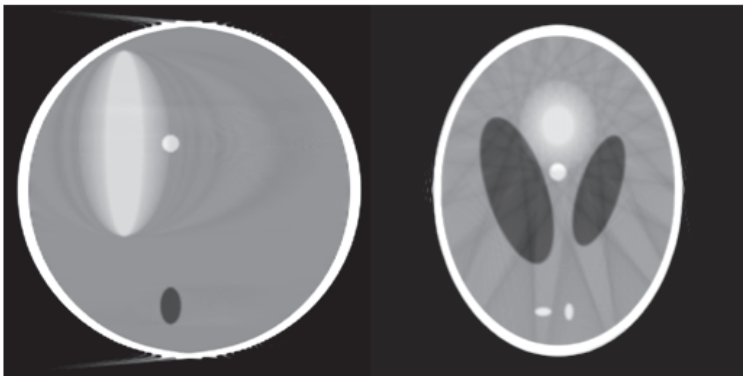

Figure 12. Ungated FDK reconstruction of the beating Shepp & Logan phantom

III.2. Injected C-arm CT scans

The X-rays absorption of blood is close to that of water, and consequently to that of soft tissues. Therefore in a cardiac CT reconstructed volume, there is no noticeable difference between the myocardium and the blood inside the ventricle. In order to create a distinction between blood and muscle, an iodine-based contrast medium, which is opaque to X-rays, is injected into the patient's vascular system. It mixes up with the blood and travels with it. The injection site, the volume of contrast medium and the rate at which it is injected are important parameters in order to obtain opaque blood at the proper position and time for the acquisition.

The datasets used throughout this thesis have been acquired following the protocol described in (Schwartz *et al* 2011): on average, 68 mL of contrast medium were injected to each patient, starting

with a high injection rate (10 mL/s), then lowering the rate (4-5 mL/s) and then injecting saline in order to push the contrast medium outside the heart's left ventricle. 308 projections evenly distributed on a 205° circular arc trajectory were acquired in 10.3 seconds. Patient DENVER1 had a 60 bpm cardiac rhythm, and patient DENVER2 88 bpm. Projections 50 and 250 for both patients are displayed on Figure 13.

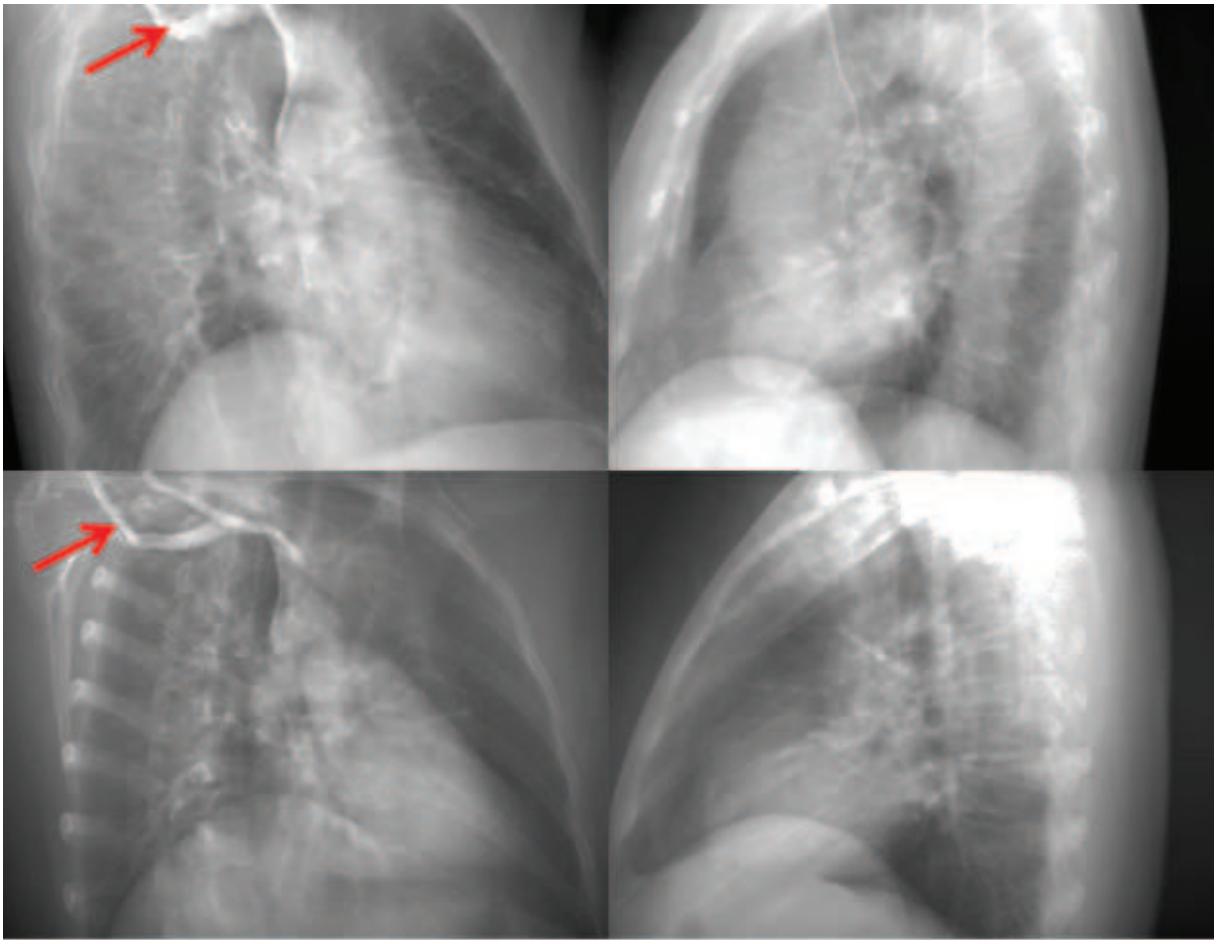

Figure 13. Projection 50 out of 308 (left column) and 250 out of 308 (right column) of the injected C-arm CT scan performed on patient DENVER1 (top row) and patient DENVER2 (bottom row). The red arrow points out the highly concentrated contrast agent in the superior vena cava in projection 50.

As can be observed in projection 50, at its initial concentration, before it gets diluted into blood, the contrast medium has a high absorption. This creates hyper-attenuation artifacts in the reconstructed volumes, just like metallic objects would do. These artifacts are shown in Figure 14.

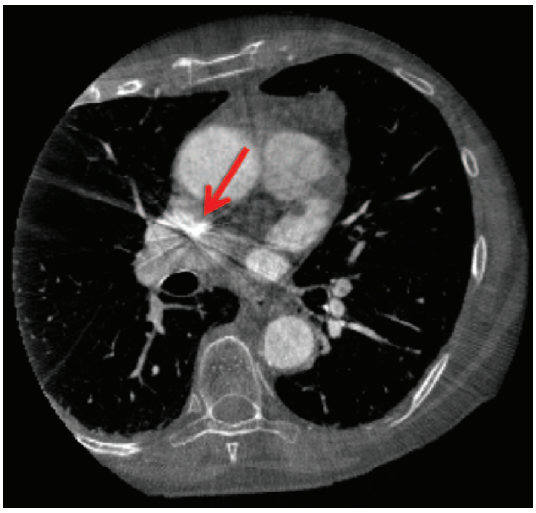

Figure 14. Hyper attenuation streak artifacts caused by highly concentrated contrast agent in the superior vena cava in patient DENVER1. The location of the contrast agent, from which streak artifacts originate, is pointed out by a red arrow

Moreover, the contrast medium flows in the vascular system. Its distribution in space can be considered periodic inside the heart if the injection is performed carefully, but is never periodic outside the heart. Therefore, projections acquired at the same ECG-phase can represent different volumes: the flow of the contrast medium makes the projection data inconsistent. Cardiac reconstruction algorithms are nevertheless built on the assumption that the attenuation of a voxel depends only on its position in space and on the ECG-phase considered, or in other words, that the imaged object's attenuation varies periodically like the ECG.

Another source of error is the truncation of data: patients are typically larger than the field of view, so that certain parts of their chest are visible on some projections and not on the others. This also makes the projection dataset inconsistent. In order to mitigate the artifacts caused by truncation, the reconstructed volume has been slightly extended, as recommended in (Zhang and Zeng 2007). 10 phases were reconstructed, each one as a 284 * 216 * 284 voxels volume.

III.3. Late enhancement scan

The principle of late enhancement acquisition has already been described in section I.2.b. Contrary to injected scans, no contrast medium is injected specifically for this acquisition. The one used during diagnosis and treatment is considered sufficient. The acquisition is performed a few minutes after the end of the intervention, once the catheter has been removed from the patient's body. Therefore the contrast medium has long been diluted into the blood, and does not cause the inconsistencies in the projections mentioned in section III.2. A few projections extracted from this dataset are shown in Figure 50.

Though late enhancement scans were performed on several patients, only the best case is used in this work, because even on that case the results are not yet satisfying. 308 projections evenly distributed on a 205° circular arc trajectory were acquired in 10.3 seconds. The patients' cardiac rhythm was 102 bpm. 10 phases were reconstructed, each one as a 284 * 216 * 284 voxels volume.

III.4. Relevance of phantom studies

Compressed sensing reconstruction methods, which constitute the core of this work, assume that the image to reconstruct is sparse in a certain basis, and therefore has some kind of regularity. Their efficiency depends on whether the image to reconstruct indeed presents the kind of regularity that the method enforces.

Unfortunately, most synthetic phantoms are schematic and do not contain the same textures, edge structure and attenuation levels distribution as real images. Phantoms are typically piecewise constant images with very sharp edges and no texture, while real cardiac CT images contain textured regions separated by less sharp transitions.

Methods using TV regularization will typically perform extremely well on phantoms, while being less efficient on real data. On the contrary, methods based on wavelet regularization are expected to perform better on real data than on piecewise phantoms. Methods based on deconvolution or on motion compensation usually do not make assumptions on the image regularity, so that their performances on phantoms are usually consistent with those on real data.

In this work, an evaluation of the different algorithms on phantom data will be provided, because their available ground truth allows to study their behavior in detail, and to quantify their efficiency. However, the definitive assessment of their performance will be carried out on real data.

Chapter IV : Results with a method of reference: ECG-gated SART

SART is a widespread reconstruction method, which can be used as a baseline to evaluate the efficiency of other methods. It has already been described in this document in section II.1, therefore only results are presented here.

These results have been obtained using a relaxation parameter $\lambda = 0.5$, and performing 10 iterations. The forward projection operator is a ray cast method with trilinear interpolation (similar to Joseph's forward projector, but with an arbitrary step length), and the back projection operator is a voxel-based back projector (in their RTK implementation on GPU).

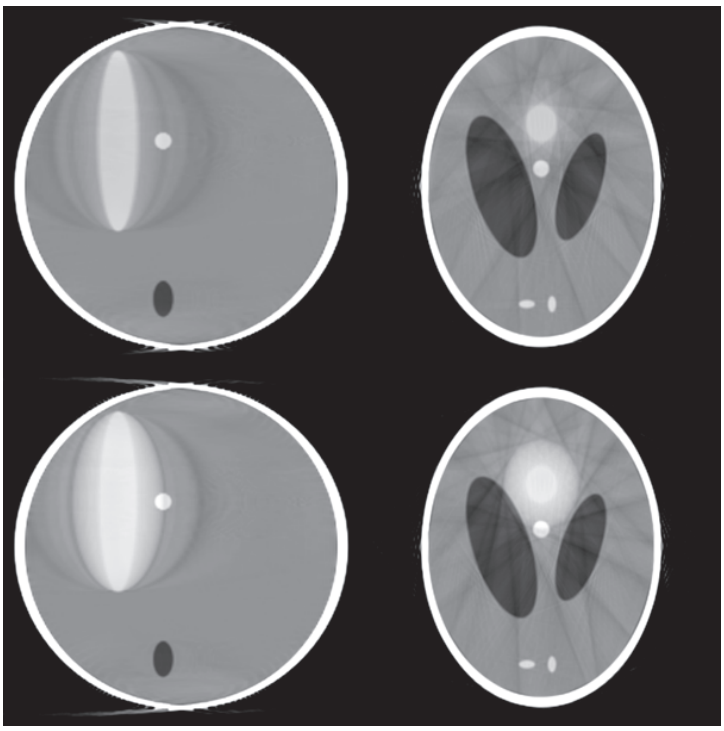

Figure 15. Reconstructions of the moving Shepp & Logan phantom by the SART method, in both systole (top row) and diastole (bottom row)

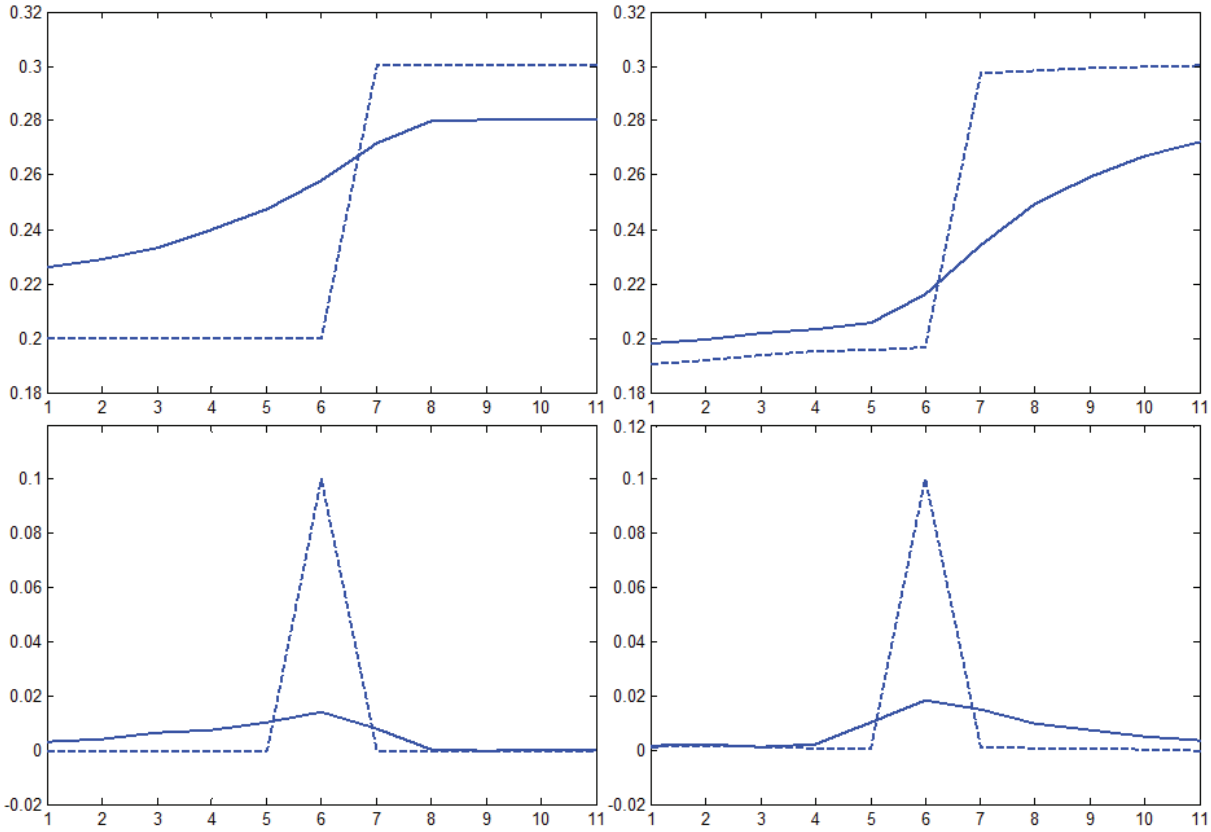

Figure 16. Profiles through the border of the beating ellipsoid (top row) and their first order derivative (bottom row) for systole (left column) and diastole (right column). The dashed line is the ground truth, and the solid line is the SART method.

On the Shepp & Logan phantom in Figure 15, there is a clear difference between systole and diastole results. However, the borders of the beating ellipse in diastole are not sharp, and a lot of streak artifacts corrupt the reconstructions.

In oblique cuts like the ones presented in Figure 17 and Figure 18, the streak artifacts exist but are hard to identify: the regions that are on the path of a streak artifact appear brighter than they should, and those close to the path of a streak appear darker than they should. This effect is hard to detect on static images, but appears clearly on animated sequences of the whole cardiac cycle. The reader in encouraged to visit my personal page on the CREATIS website (http://www.creatis.insa-lyon.fr/~mory/), where animated gif sequences are available.

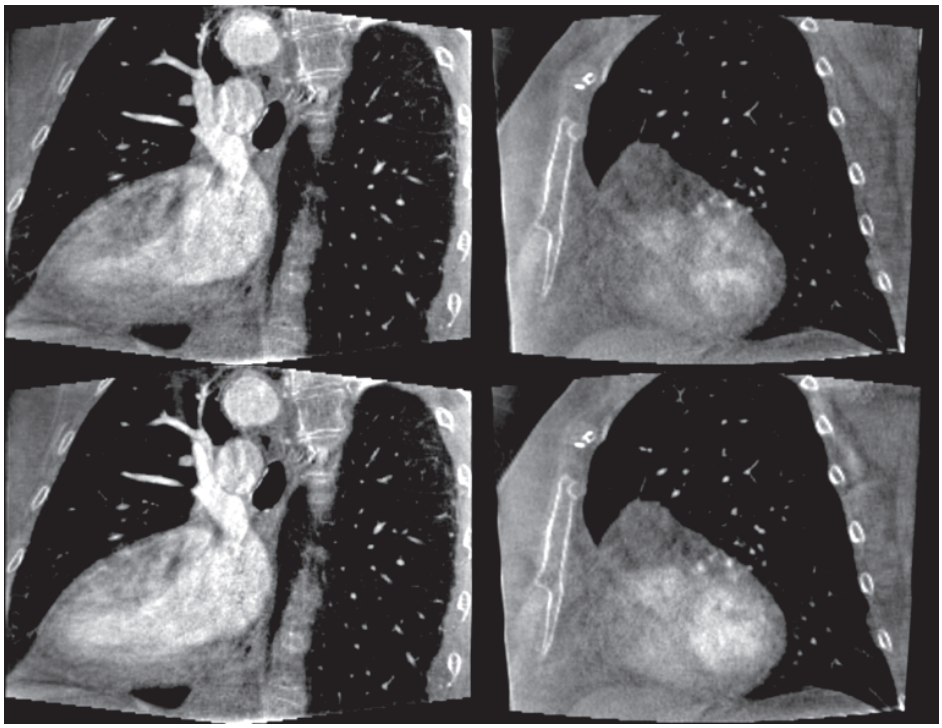

Figure 17. Reconstructions of patient DENVER1 by the SART method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

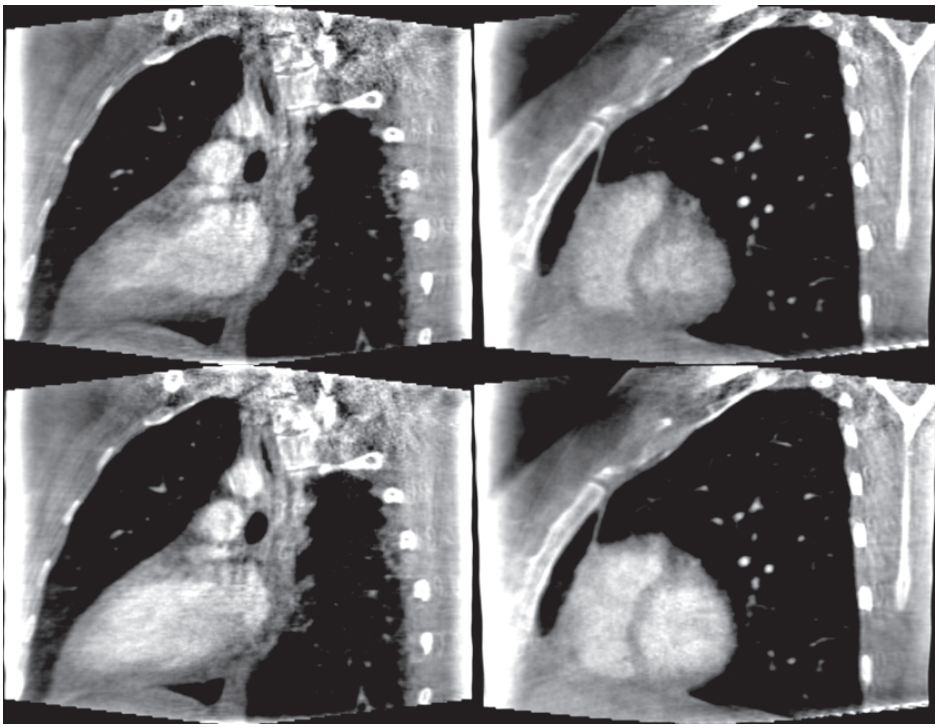

Figure 18. Reconstructions of patient DENVER2 by the SART method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

Chapter V : Deconvolution for streak artifacts removal

In the 1980's, a number of limited-view computed tomography methods based on deconvolution have been proposed (Dhawan *et al* 1985, Gordon and Rangayyan 1983). Limited-view reconstruction is a convolution process (see section V.1), and therefore these methods propose to first perform a gated reconstruction, and then to filter out the streak artifacts of the gated reconstruction by deconvolution. Deconvolution was mainly performed by Wiener filtering and was only marginally efficient.

In a recent paper (Badea *et al* 2011), Badea et al. introduced a new deconvolution-based reconstruction algorithm for 4D cardiac micro CT. It uses inverse filtering and fills the gaps in the Fourier data with the corresponding Fourier coefficients of the ungated reconstruction. This new algorithm is simple to implement, fast, has only one parameter, and gives satisfying results on 4D cardiac micro CT.

The aim of this chapter is to compare the method proposed in (Badea *et al* 2011) with one based on convex optimization, and equivalent to ECG-gated Iterative Filtered Back Projection (IFBP) (Sunnegardh and Danielsson 2008, Sunnegaardh 2009, Medoff *et al* 1983, Nassi *et al* 1982). They are evaluated on cardiac micro CT and human cardiac C-Arm CT phantom simulations. ECG-gated IFBP has the advantage of being theoretically suited not only to parallel beam geometry (like (Badea *et al* 2011) is), but also to fan beam and cone beam, which are the geometries used in practice.

V.1. Gated FBP is a convolution

Before we describe the deconvolution-based methods in details, we prove that, in the continuous and infinite extent case, and *in parallel beam geometry*, a gated reconstruction can indeed be expressed as a convolution. The proof is only valid in this context. In the next subsection, and throughout the paper, we will use the following notations, or their obvious extension to 3D:

$$f: \mathbb{R}^2 \to \mathbb{R}$$
$$(x, y) \to \text{Attenuation}(x, y)$$

f is the object function that we aim to reconstruct from its projections.. $\mathcal{F}_{1D}$ is the one-dimensional Fourier transform, $\mathcal{F}_{2D}$ the two-dimensional Fourier transform. The symbol × denotes the pointwise product, and * the convolution.

Let us define Θ the set of angles for which we keep the projections. The gating function g is:

$$g: \mathbb{R} \times [0; \pi[\to \{0; 1\}$$
$$(s, \theta) \to \begin{cases} 1 \text{ if } s = 0 \\ 1 \text{ if } \theta \in \Theta \\ 0 \text{ otherwise} \end{cases}$$

g is a 2D function, but for any fixed value of θ, $g(s, \theta)$ is constant and can be replaced by g_θ. The gated reconstruction can be written as $f_{\text{gated}} = R^{-1}(R(f) \times g)$, or equivalently $R(f_{\text{gated}}) = R(f) \times$

g. Let $\theta\epsilon[0;\pi[$, and let $R_\theta(f)$ be the projection of f along direction θ. The Fourier Slice Theorem (Kak and Slaney 1988) states that for any $r\epsilon\mathbb{R}$

$$\mathcal{F}_{2D}\left(f_{\text{gated}}\right)(r,\theta) = \mathcal{F}_{1D}(R_\theta(f) \times g_\theta)(r) \tag{V.1}$$

g_θ being a constant, it can be taken out of the Fourier transform:

$$\mathcal{F}_{1D}(R_\theta(f) \times g_\theta)(r) = \mathcal{F}_{1D}\left(R_\theta(f)\right)(r) \times g_\theta \tag{V.2}$$

On the other hand, from the Fourier Slice Theorem,

$$\mathcal{F}_{1D}\left(R_\theta(f)\right)(r) = \mathcal{F}_{2D}(f)(r,\theta) \tag{V.3}$$

And so, by combining all three equations,

$$\mathcal{F}_{2D}\left(f_{\text{gated}}\right)(r,\theta) = \mathcal{F}_{2D}(f)(r,\theta) \times g(r,\theta) \tag{V.4}$$

This holds for any θ and any r, thus

$$\mathcal{F}_{2D}\left(f_{\text{gated}}\right) = \mathcal{F}_{2D}(f)g \tag{V.5}$$

Which is equivalent, in the spatial domain, to

$$f_{\text{gated}} = f * {\mathcal{F}_{2D}}^{-1}(g) \tag{V.6}$$

This demonstrates that the gated reconstruction is a convolution. The Point Spread Function (PSF) is the inverse 2D Fourier transform of the gating function. If we substitute f with a Dirac function, we obtain $\delta_{\text{gated}} = \delta * {\mathcal{F}_{2D}}^{-1}(g) = {\mathcal{F}_{2D}}^{-1}(g)$, which gives a simple way to construct this PSF: it is the gated reconstruction of a Dirac function. This function is plotted on the center of Figure 19.

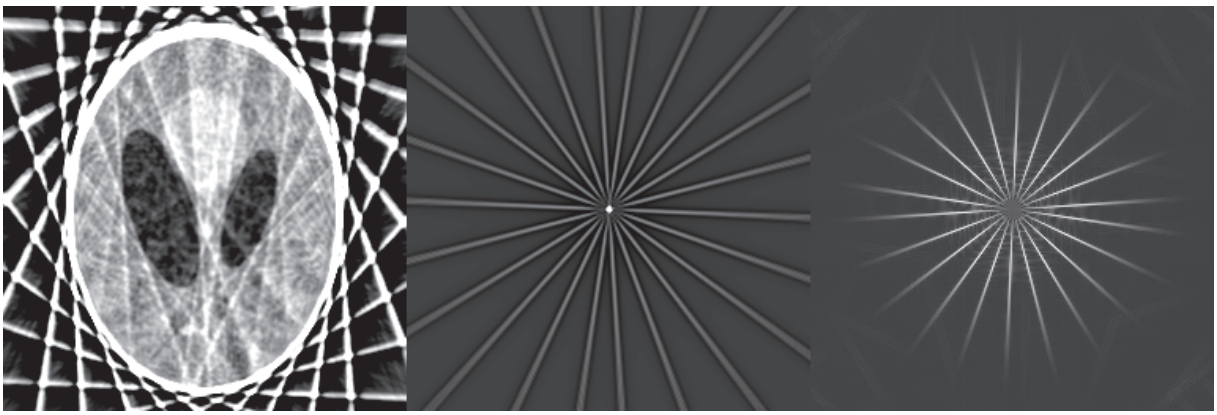

Figure 19. From left to right, gated reconstruction of a phantom, gated reconstruction of a Dirac function, and shifted FFT of the Dirac peak's gated reconstruction. The streak pattern is the same on the Dirac and the phantom

We have shown that in the 2D parallel beam and continuous case, a gated reconstruction is a convolution between the object function and the gated reconstruction of a Dirac function. Recovering f from f_{gated} from (V.6) thus boils down to performing a deconvolution. However some Fourier domain information is lost during the gated reconstruction process, when $g(r,\theta) = 0$ in (V.6). As shown on the right image of Figure 19, which contains many zero (or almost zero)

coefficients, the amount of information lost can be large. This information will have to be extrapolated, explicitly or implicitly, in the deconvolution process.

In divergent beam geometry, the gated reconstruction process is linear, but not shift-invariant. To be more precise, in fan beam it is shift-invariant for very special cases of Θ, for example when $\Theta = \emptyset$ or $\Theta = [0; 2\pi[$, but not for most subsets of $[0; 2\pi[$ (Gullberg 1979). And in particular, $R_\theta^T K R_\theta$ (with K the ramp filter) is not shift-invariant, as is illustrated on Figure 20, which shows the single projection reconstructions of two Dirac functions, one centered, and the other one off-center. It is obvious that the reconstructions are not shifted version of one another. As a result, in the case of fan and cone beam projections, the gated FBP reconstruction process cannot be interpreted as a convolution.

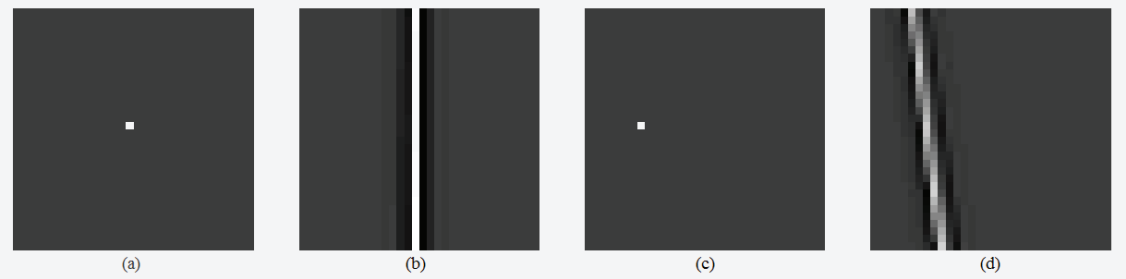

Figure 20. (a): Centered Dirac, (b): Reconstruction of (a) using only one projection, (c): Off-center Dirac, (d): Reconstruction of (c) using only one projection. For both reconstructions, the source is at the same location (under the image). (b) and (d) are not shifted versions of one another.

V.2. Badea's method

A straightforward deconvolution method is called inverse filtering, and consists in dividing the Fourier transform of the convolved image f_c by the Fourier transform of the PSF, and applying an inverse Fourier transform on the result (Bracewell 2003). With $\mathcal{F}_{2D}$ the 2D Fourier transform and $\hat{f}$ the deconvolution result, inverse filtering is simply expressed as:

$$\mathcal{F}_{2D}(\hat{f})(u,v) = \frac{\mathcal{F}_{2D}(f_c)(u,v)}{\mathcal{F}_{2D}(\mathrm{PSF})(u,v)} \tag{V.7}$$

The major problem is to handle the divisions by 0, in particular when $\mathcal{F}_{2D}(\mathrm{PSF})(u,v)$ has many zeros, which is the case with ECG-gating.

In (Badea *et al* 2011), an important adaptation of inverse filtering was proposed to handle the zeros of the spectrum of the PSF. The voxel-by-voxel division in the Fourier domain is performed only when the magnitude of the Fourier coefficient of the PSF is above a certain threshold. In the other voxels, where $\left|\mathcal{F}_{2D}(\delta_{\text{gated}})(u,v)\right|$ is too close to zero, the division is considered unreliable and the corresponding Fourier coefficients are copied from the ungated FBP reconstruction, which is the reconstruction obtained by taking into account the projections corresponding to all cardiac phases. This can be summarized as:

$$\mathcal{F}_{2D}(f)(u,v) = \begin{cases} \dfrac{\mathcal{F}_{2D}(f_{\text{gated}})(u,v)}{\mathcal{F}_{2D}(\delta_{\text{gated}})(u,v)} & \text{if } \left|\mathcal{F}_{2D}(\delta_{\text{gated}})(u,v)\right| > \text{threshold} \\ \mathcal{F}_{2D}(f_{\text{ungated}})(u,v) & \text{otherwise} \end{cases} \tag{V.8}$$

In practice, inverse filtering deconvolution causes border effects. This is because by multiplying the discrete Fourier transforms of two images, one obtains the discrete Fourier transform of their

circular convolution, not of their zero-padded convolution (DFT-based convolution requires that the images be infinite and periodic, and they usually aren't). The reverse process, inverse filtering, therefore assumes that the image to be deconvolved is the result of a circular convolution between the target and the PSF, which is usually wrong, and leads to border effects.

It turns out that, with sufficiently large zero-padding of the images, a circular convolution can be equivalent to a zero-padded one. Here the images are reconstructed in a field of view twice as large as the object and multiplied by a 2D cosine window, which is close to zero-padded deconvolution, and helps mitigate the border effects. The threshold is set to 15% of the maximum value of $\left|\mathcal{F}_{2D}\left(\delta_{\text{gated}}\right)\right|$, as recommended in (Badea *et al* 2011).

V.3. ECG-gated Iterative FBP

V.3.a. Van Cittert deconvolution method

A straightforward iterative deconvolution method has been proposed by Van Cittert (Van Cittert 1931). It can be applied to streak removal and leads naturally to iterative filtered back-projection.

In an attempt to iteratively deconvolve an image f_c, the Van Cittert method uses the following update step:

$$f_{k+1} = f_k + \alpha(f_c - \text{PSF} * f_k) \tag{V.9}$$

where f_k is the deconvolved image at the k-th iteration, f_c is the observed image (here the gated reconstruction), and α a relaxation weight. Section V.3.c gives some insight on how to determine α.

This scheme is of particular interest, since it does not require the explicit estimation of the PSF. One only needs to compute PSF $* f_k$, i.e. the forward projection of f_k, followed by a gated reconstruction. As an iterative method, it requires an initialization: the ungated reconstruction provides an excellent starting point in practice.

V.3.b. Iterative FDK

With the notations introduced in sections I.1.a and II.1, the update step of the Van Cittert method, modified as described in the previous subsection, is as follows:

$$f_{k+1} = f_k + \alpha\left(f_{gated} - R^T K G R f_k\right) \tag{V.10}$$

$$f_{k+1} = f_k + \alpha(R^T K G R f - R^T K G R f_k) \tag{V.11}$$

$$f_{k+1} = f_k + \alpha R^T K G (R f - R f_k) \tag{V.12}$$

Note that here, R is not the exact Radon transform: it has a finite number of rays and projections, and is not limited to parallel projection. It is referred to as *X-ray transform* in the following.

Equation (V.12) can be translated into the following steps:

- Start from the ungated reconstruction
- Until a stopping criterion is met
  - Apply ECG-gating, which selects a subset of projection angles
  - Forward project the current volume along the selected angles

- Subtract the estimated projections to the measured ones
- Apply an FBP or FDK to these “difference projections”
- Add the result to the current volume (with a weight α)

There are many ways to set the stopping criterion. The algorithm for example can be stopped:

- When the difference between the current result and that of the previous iteration is below a certain threshold
- When a given metric reaches a certain threshold
- After a given number of iterations

In this work, the stopping criterion used was the number of iterations (set to 100).

It turns out that this scheme has already been studied and is referred to as “Iterative Filtered Back Projection” (short IFBP). However, to the best of our knowledge, only ungated IFBP has been the topic of recent research, in order to remove cone beam or metal artifacts (Sunnegaardh 2009, Sunnegardh and Danielsson 2008). Limited view IFBP has been studied, but never exactly on the same problem: in (Medoff *et al* 1983), the authors study the so-called “bagel problem”, and in (Nassi *et al* 1982) the study is restricted to fan beam and uses iterative reconstruction-reprojection to estimate the missing views. Its application to cardiac micro CT and cardiac C-Arm CT simulations, as well as human cardiac C-Arm CT data, is novel.

V.3.c. Convex optimization interpretation

Since the iterative scheme (V.12) does not model the streak removal process as a strict deconvolution problem, the stationarity assumption can be relaxed. Thus, the method can be adapted to fan and cone-beam projections. This is justified by the following convex optimization interpretation of IFBP. The optimization perspective also helps understand how to set the parameter α.

The ECG-gated reconstruction problem can be formulated as follows. The image f to be determined minimizes the following energy:

$$E(f) = \frac{1}{2}\left\| K^{\frac{1}{2}} G(Rf - p_{\text{measured}}) \right\|_2^2 \tag{V.13}$$

Note that without the K operator, $E(f)$ would be the energy minimized by ECG-gated ART. It would also be a valid approach for ECG-gated reconstruction. The aim here, though, is to give a convex optimization interpretation of IFBP, so ramp filtering needs to be performed. The gradient of this energy reads:

$$\nabla E(f) = R^T G^T K G(Rf - p_{\text{measured}}) \tag{V.14}$$

$$\nabla E(f) = R^T K G(Rf - p_{\text{measured}}) \tag{V.15}$$

because the ECG-gating and the ramp filtering commute (they can be performed in any order), and $G^T G = G$ because G is diagonal and binary. From equation (V.15), it is clear that a gradient descent procedure to find the minimum of this energy would have the following update step:

$$f_{k+1} = f_k + \alpha_k R^T K G(Rf_k - p_{\text{measured}}) \tag{V.16}$$

with α_k the gradient descent step at iteration k. Setting $\alpha_k = \alpha$ exactly boils down to the same update step as in equation (V.12). This shows that IFBP can be seen as a steepest descent to find the minimum of the energy defined in equation (V.13), with α the step length of the steepest descent. The optimal value of the parameter α_k can be determined at each iteration: the problem was solved analytically by Lalush and Tsui in (Lalush and Tsui 1994). Other suboptimal, yet easier choices can be made. In this work, we used $\alpha_k = \alpha$, with a small fixed value.

V.4. Results

This section contains reconstruction results for both Badea's deconvolution method and the ECG-gated iterative FDK method, on the Shepp & Logan phantom and the injected clinical datasets described in Chapter III. Badea's deconvolution method is presented on the left and ECG-gated IFDK on the right. In all figures, the top row contains slices extracted from the systole reconstruction, and the bottom row slices from the diastole one.

Figure 21 shows reconstructions of the moving Shepp & Logan phantom by Badea's method and iterative FDK. In Badea's results, the beating ellipsoid has a better-defined shape, closer to that of the ground truth, than in ECG-gated IFDK results. On the other hand, high frequency artefacts, probably caused by errors on the high frequency components in the Fourier domain, appear in Badea's results.

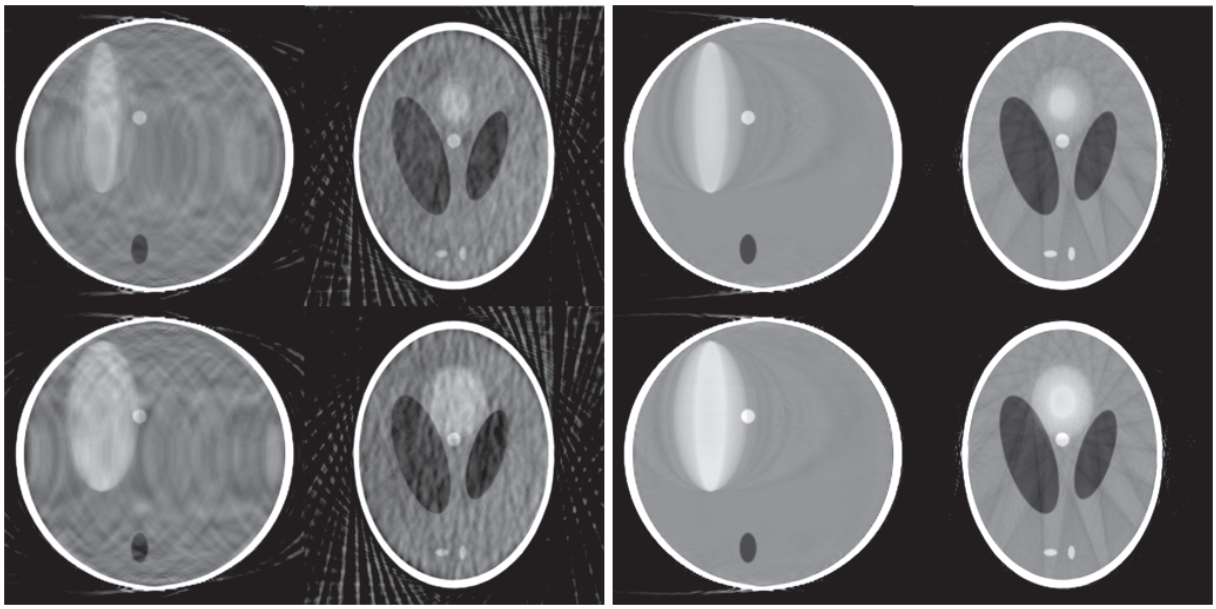

Figure 21. Reconstructions of the moving Shepp & Logan phantom by Badea's method (left panel) and iterative FDK (right panel), in both systole (top row) and diastole (bottom row).

Extracting the derivative of profiles through the edges of the beating ellipsoid allows measuring how sharp they are. We averaged several profiles in order to deal with noise and streak artifact.

More precisely, we proceed as follows: the beating ellipsoid taken from the ground truth is dilated and its surface is extracted, then the process is repeated several times. This generates several ellipsoid surfaces of different sizes, and each value of the profile is computed by sampling the reconstructed volume on one of these surfaces and averaging the samples. The same process is applied using shrinking instead of dilation, which generates profile data for the area inside the beating ellipsoid. Figure 22 shows profiles obtained by this method and their first order derivative, on

both Badea and IFDK reconstruction, in systole and diastole. As expected from visual evaluation, the beating ellipsoid in Badea's results appears marginally sharper than in IFDK results.

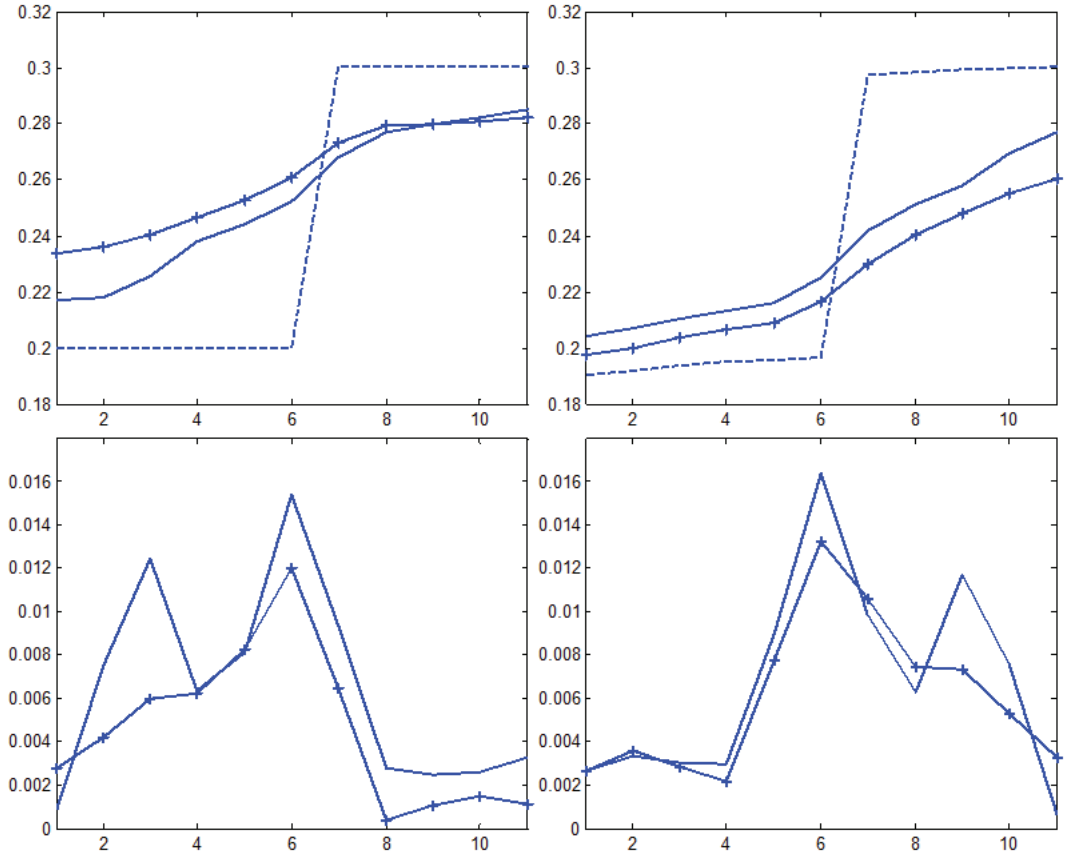

Figure 22. Profiles through the border of the beating ellipsoid (top row) and their first order derivative (bottom row), in systole (left column) and diastole (right column). The dashed line is the ground truth (not represented in the bottow row), the solid line is Badea's method, and the solid line with crosses is the ECG-gated iterative FDK. Badea's results are slightly sharper than ECG-gated IFDK results.

The methods are also compared in terms of Root Mean Squared Error (RMSE) and RMSE in a ROI containing the beating ellipsoid. The results are listed in Table 1.

		Badea's method	ECG-gated IFDK
RMSE	Systole	0.0834	0.0443
	Diastole	0.0821	0.0454
RMSE in ROI	Systole	0.0250	0.0238
	Diastole	0.0301	0.0302

Table 1. Root Mean Squared Errors for Badea's method and ECG-gated IFDK on the beating Shepp and Logan phantom

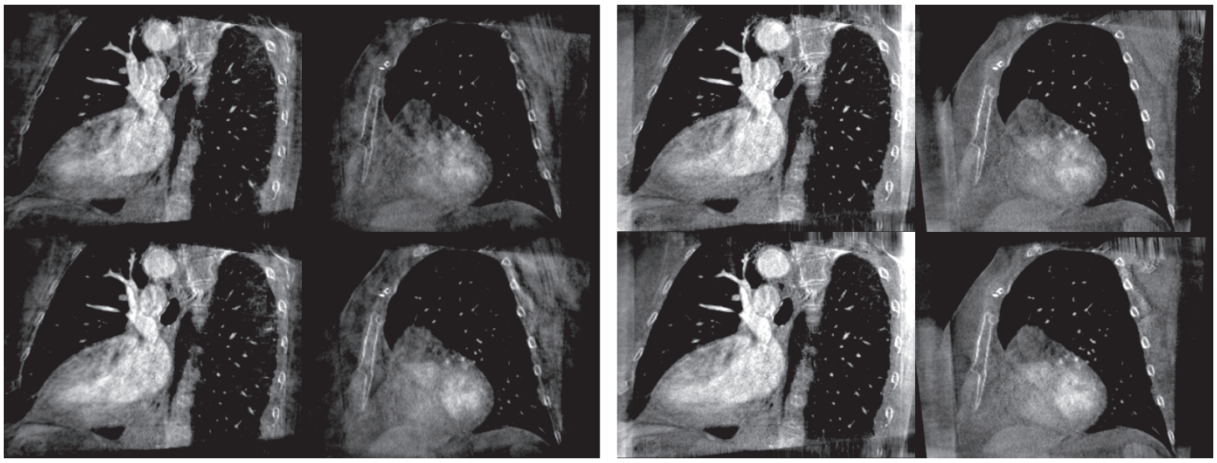

Figure 23. Reconstructions of patient DENVER1 by Badea's method (left of the white column) and iterative FDK (right of the white column), in both systole (top row) and diastole (bottom row). The visualization window is [0.006; 0.058].

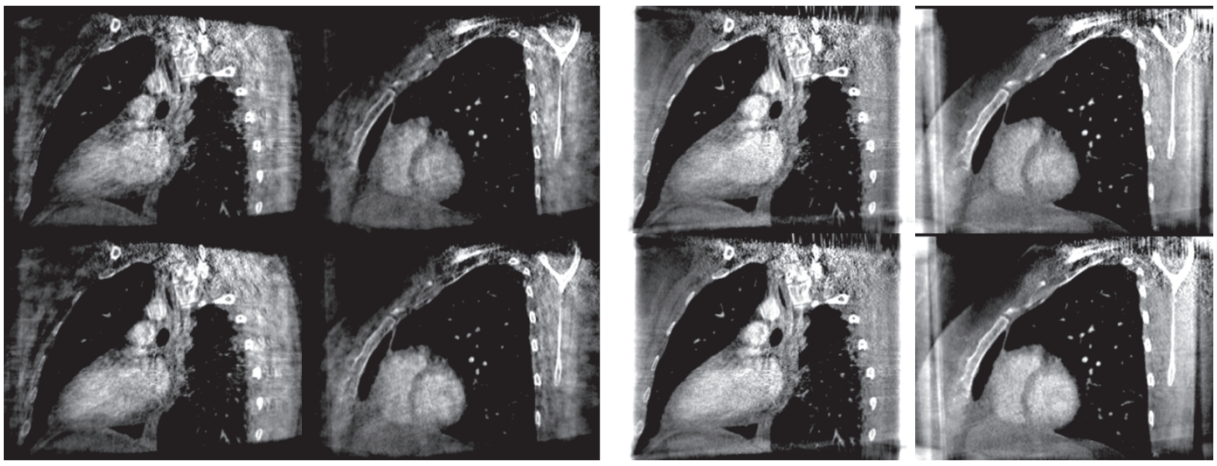

Figure 24. Reconstructions of patient DENVER2 by Badea's method (left of the white column) and iterative FDK (right of the white column), in both systole (top row) and diastole (bottom row). The visualization window is [0.006; 0.054].

Figure 23 and Figure 24 show reconstructions of patients DENVER1 and DENVER2 respectively. Except for the hyper-attenuation artefacts on the borders of the volume, caused by truncation, the Iterative FDK results are less noisy than those reconstructed with Badea's method. The difference between systole and diastole is stronger in Iterative FDK results than in Badea's, although in both cases it is smaller than clinically expected.

The RMSE and local RMSE results on the Shepp and Logan phantom are also in favor of iterative FDK.

V.5. Conclusion on deconvolution

Two deconvolution methods were presented, and the reconstruction results shown and commented. Their respective advantages and drawbacks are listed in Table 2.

Badea's method	ECG-gated Iterative FDK
+ Fast + Slightly sharper + Only one parameter	- Slow (several iterations) - Slightly less sharp
- Rigorous in parallel beam only - High frequency Fourier noise	+ Suited to fan and cone beam

Table 2. Respective advantages and drawbacks of Badea's method and ECG-gated IFDK

Although the measurements indicate that iterative FDK outperforms Badea's method, neither of both seems well suited to human cardiac C-Arm CT reconstruction from a single sweep acquisition: the results are either too noisy or too close to an ungated reconstruction, i.e. affected by beating movement blurring. No acceptable compromise has been found on real clinical data.

Two things are worth noting: first, the sharpness of the ECG-gated Iterative FDK results is partly controlled by the cutoff frequency of the ramp filter (in order to avoid amplify high-frequency noise, the ramp filter is often multiplied by an apodization window in the Fourier domain, which implies a cutoff frequency). Therefore, the conclusions on the compared sharpness of the results of Badea's method and Iterative FDK depend on the parameters used, and should be considered less robust than the other comparison items. Second, the tradeoff between motion blurring and artifact level exists in all methods presented throughout this document, as it is inherent to the problem of cardiac C-Arm CT reconstruction itself. The next section contains a mathematical interpretation of the problem, details the reasons of this tradeoff, and suggests research directions for the rest of this thesis.

V.6. Mathematical considerations

When expressed as a convex optimization problem, reconstruction can be interpreted in terms of kernel and image of the data-attachment term's linear map. In most ECG-gated approaches, the data attachment term is $\|G(Rf-p)\|_2^2$, and the linear map GR. In the 2D parallel beam case, R is the Radon transform, and from the Fourier Slice Theorem it can be deduced that the kernel of GR, denoted $Ker(GR)$, is the set of functions that have null Fourier coefficients along the angles selected by G. Its orthogonal complement, denoted $Ker(GR)^{\perp}$ is the set of functions that have null Fourier coefficients outside those angles. f can be expressed as the sum of its projection f_{Ker} on $Ker(GR)$ and its projection f_{Ortho} on the orthogonal complement $Ker(GR)^{\perp}$:

$$f = f_{Ker} + f_{Ortho}$$

Iterations of ECG-gated forward and back projections of f only update f_{Ortho}, and leave f_{Ker} unchanged (a proof is given at the end of this section). If the gating window width is 20% of the cardiac cycle, the dimension of $Ker(GR)$ is 4 times higher than the dimension of $Ker(GR)^{\perp}$. Therefore it is critical to find a consistent f_{Ker}. This can be done in two ways:

- Regularization: iterative algorithms can enforce constraints like limited support or non-negativity. Compressed sensing methods constrain the solution to be sparse in a certain basis. Both kinds of regularization modify f_{Ker}
- Initialization: if no regularization is enforced or if it has little impact on f_{Ker}, initialization is critical. Most of the time, either f is initialized with zero, which means that f_{Ker} is initially zero, which leads to streak artifacts, or f is initialized with the ungated FBP reconstruction, which fills f_{Ker} with data from the wrong cardiac phases, which causes blurring. Neither is really satisfying

As a consequence, an optimal f_{Ker} can only be obtained with strong regularization, which allows, to some extent, to progressively derive f_{Ker} from f_{Ortho}. Compressed sensing methods therefore appear as a reasonable research direction.

In the fan beam and cone beam cases, the Fourier Slice Theorem cannot be applied in a straightforward manner, therefore $Ker(GR)$ does not have a simple expression in the Fourier domain, nor is it directly linked to gaps in the angular sampling. However the conclusions drawn from the parallel beam case can still guide us, and lead us to investigate compressed sensing reconstruction methods.

Proof that $Ker(GR)$ is not modified during steepest descent:

Let A be a linear map $A: \mathbb{R}^K \to \mathbb{R}^L$, $f \in \mathbb{R}^K$ and $b \in \mathbb{R}^L$.

We want to minimize $\|Af - b\|_2^2$. The steepest descent procedure consists in iteratively applying the following update step:

$$f_{k+1} = f_k + \alpha_k A^T(Af_k - b)$$

We want to prove that the kernel of A is not updated. To this end, we must prove that the orthogonal projections of f_{k+1} and of f_k on $Ker(A)$ are identical, which amounts to proving that the orthogonal projection of $A^T(Af_k - b)$ on $Ker(A)$ is **0**. This is obvious because in finite dimension, $Im(A^T) = Ker(A)^\perp$.

Chapter VI : Regularized 3D reconstruction

This chapter is meant as an in-depth study of three reference reconstruction methods introducing spatial regularization. More specifically, we have implemented in a common framework the following 3D reconstruction methods:

- Method 1: Total variation (TV)-regularized reconstruction based on the Augmented Lagrangian (AL) method and the Alternating Direction Method of Multipliers (ADMM)
- Method 2: Daubechies wavelets-regularized reconstruction, again with AL and ADMM
- Method 3: TV-regularized Prior Image Constrained Compressed Sensing (PICCS), implemented by alternating ECG-gated SART and TV minimization by steepest descent

Some of these methods are directly taken from the literature, and some have been modified during this PhD:

- Method 1 is novel, but the way it states the reconstruction problem is almost identical to ASD-POCS, and is a sub-problem of Method 3
- Method 2 minimizes a cost function mentioned in (Daubechies *et al* 2004). Its application to cardiac C-arm CT has not been published before
- Method 3 was published in (Chen *et al* 2008), and applied to cardiac C-arm CT. At this point in time, it is considered the best method for reconstructing cardiac anatomies in C-arm CT.

An abuse in notations is used repeatedly in this chapter and the next one: the notation

$$\hat{f} = \arg\min_{f} \quad Cost(f)$$

is used even when $\arg\min_f \ Cost(f)$ is not a singleton, but a set containing several elements. In these cases, it means "the $\hat{f}$ obtained by minimizing $Cost(f)$". It usually results in $\hat{f}$ being the projection of f_0, the f with which the iterative search is initialized, onto the set of minimizers of $Cost$, namely $\arg\min_f \ Cost(f)$.

VI.1. 3D reconstruction with TV regularization

VI.1.a. Cost function

The cost function minimized by this reconstruction method is the following:

$$J(f) = \|G(Rf - p)\|_2^2 + \alpha TV(f)$$

with α a fixed parameter and TV defined, as in section II.4, by

$$TV(f) = \sum_{v=1}^{V} \sqrt{[\nabla_x f(v)]^2 + \left[\nabla_y f(v)\right]^2 + [\nabla_z f(v)]^2}$$

where v is the voxel position in the volume and V the number of voxels. J is a convex cost function.

VI.1.b. Minimization algorithm

Many approaches could be used to find an f that minimizes J. We have chosen the augmented lagrangian method combined with the Alternating Direction Method of Multipliers (ADMM) (Afonso *et al* 2010, 2011, Figueiredo *et al* 2009) because it allows to use wavelet-based regularization instead of total variation with minimal modifications.

Prerequisite: soft thresholding on a vector

Let $\boldsymbol{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \epsilon \mathbb{R}^n$, and let us consider the following problem:

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{x}} \|\boldsymbol{x} - \boldsymbol{y}\|_2^2 + \lambda \|\boldsymbol{x}\|_2$$

The solution to this problem(Tibshirani 1994) is given by the vectorial soft thresholding operator ST of threshold $\frac{\lambda}{2}$, defined as

$$\boldsymbol{x}^* = ST_{\frac{\lambda}{2}}(\boldsymbol{y}) = \max\left(\|\boldsymbol{y}\|_2 - \frac{\lambda}{2}, 0\right) \frac{\boldsymbol{y}}{\|\boldsymbol{y}\|_2}$$

Augmented lagrangian and ADMM

The unconstrained problem of finding $\arg\min_f J(f)$ is equivalent to the constrained problem

$$\begin{cases} (\hat{f}, \hat{g}) = \arg\min_{f,g} \|G(Rf - p)\|_2^2 + \alpha \sum_{v=1}^{V} \|g(v)\|_2 \\ \text{subject to } g = \nabla f \end{cases}$$

The Augmented Lagrangian method consists in iteratively

- minimizing $\|G(Rf - p)\|_2^2 + \alpha \sum_{v=1}^{V} \|g(v)\|_2 + \beta \|\nabla f - g - d_k\|_2^2$ over f and g, without constraints,

- and updating d_k at each iteration.

In the AL method, β is a constant, and $\|\nabla f - g - d_k\|_2^2$ is a term added to control the difference between ∇f and g. Initially, d_k is set to $d_0 = \mathbf{0}$.

The minimization is performed first by finding the f that minimizes the cost function for a fixed value of g, then finding the g that minimizes the cost function for a fixed value of f (this is the Alternating Direction Method of Multipliers, short ADMM). In total, each iteration consists in the following steps:

$$\begin{cases} f_{k+1} = \arg\min_f \ \|G(Rf - p)\|_2^2 + \beta\|\nabla f - g_k - d_k\|_2^2 \\ g_{k+1} = \arg\min_g \alpha \sum_{v=1}^{V} \|g(v)\|_2 + \beta\|\nabla f_{k+1} - g - d_k\|_2^2 \\ d_{k+1} = d_k - \nabla f_{k+1} + g_{k+1} \end{cases}$$

The calculations required to compute f_{k+1} and g_{k+1} are detailed hereinafter:

f_{k+1} is the zero of the gradient of $\|G(Rf - p)\|_2^2 + \beta\|\nabla f - g_k - d_k\|_2^2$, that is to say the value of f such that:

$$R^T G^2 R f - R^T G p + \beta \nabla^T (\nabla f - g_k - d_k) = 0$$

$$\Leftrightarrow \underbrace{(R^T G^2 R + \beta \nabla^T \nabla)}_{A} f = \underbrace{R^T G p + \beta \nabla^T (g_k + d_k)}_{b}$$

f_{k+1} can thus be computed by a conjugate gradient algorithm aimed at solving $Af = b$. Note that $\nabla^T = -\mathrm{div}$, if both ∇ and ∇^T are computed with circular padding on the borders.

g_{k+1} can be computed voxel-by-voxel, and for each voxel the problem boils down to the soft thresholding on a vector presented earlier in this section. Therefore, for any v,

$$g_{k+1}(v) = ST_{\frac{\alpha}{2\beta}}(\nabla f_{k+1} - d_k)$$

With ST the soft thresholding operator with threshold $\frac{\alpha}{2\beta}$.

VI.1.c. Implementation

In the practical implementation, the parameter β is not set to a fixed value, but rather changes over the course of the iterations according to the following law: $\beta_k = \beta \times \frac{k}{k_{\max}}$. This way β_k is low during the first iterations, which implies a strong regularization (the threshold in the computation of g_{k+1} is high). It then gets larger over the course of the iterations, which gradually attenuates the impact of the regularization, in a “coarse to fine” fashion.

VI.1.d. Results

The proposed method is evaluated on the Shepp & Logan phantom and the real data with contrast medium described in Chapter III. The performance on phantom data is measured with the metrics and methods introduced in section V.4 (slices, profiles for sharpness measurement, RMSE and RMSE in the beating region). Figure 25 shows reconstructions of the Shepp & Logan phantom with the

proposed method. Some streak artifacts remain, especially in the right column images (cut plane normal to the rotation axis) but the beating ellipsoid borders seem sharp.

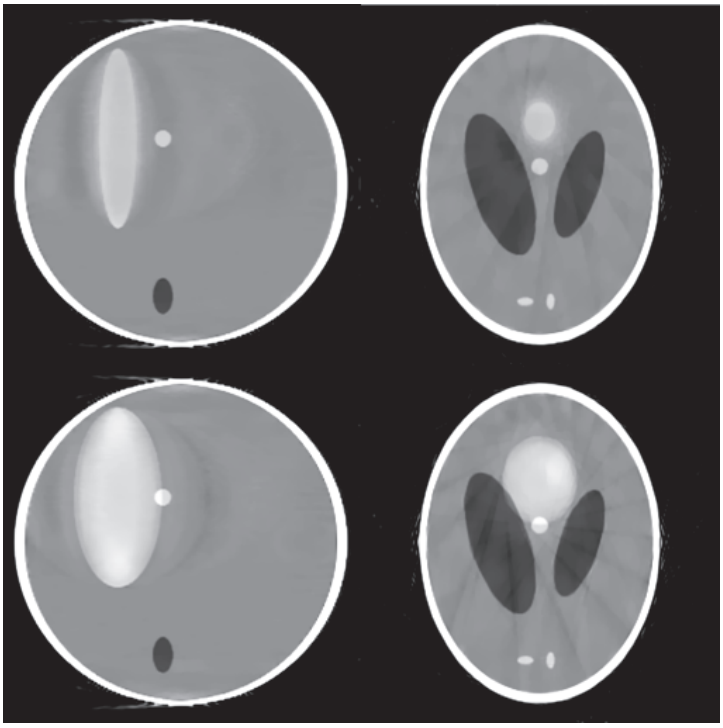

Figure 25. Reconstructions of the moving Shepp & Logan phantom by the TV-regularized 3D method, in both systole (top row) and diastole (bottom row).

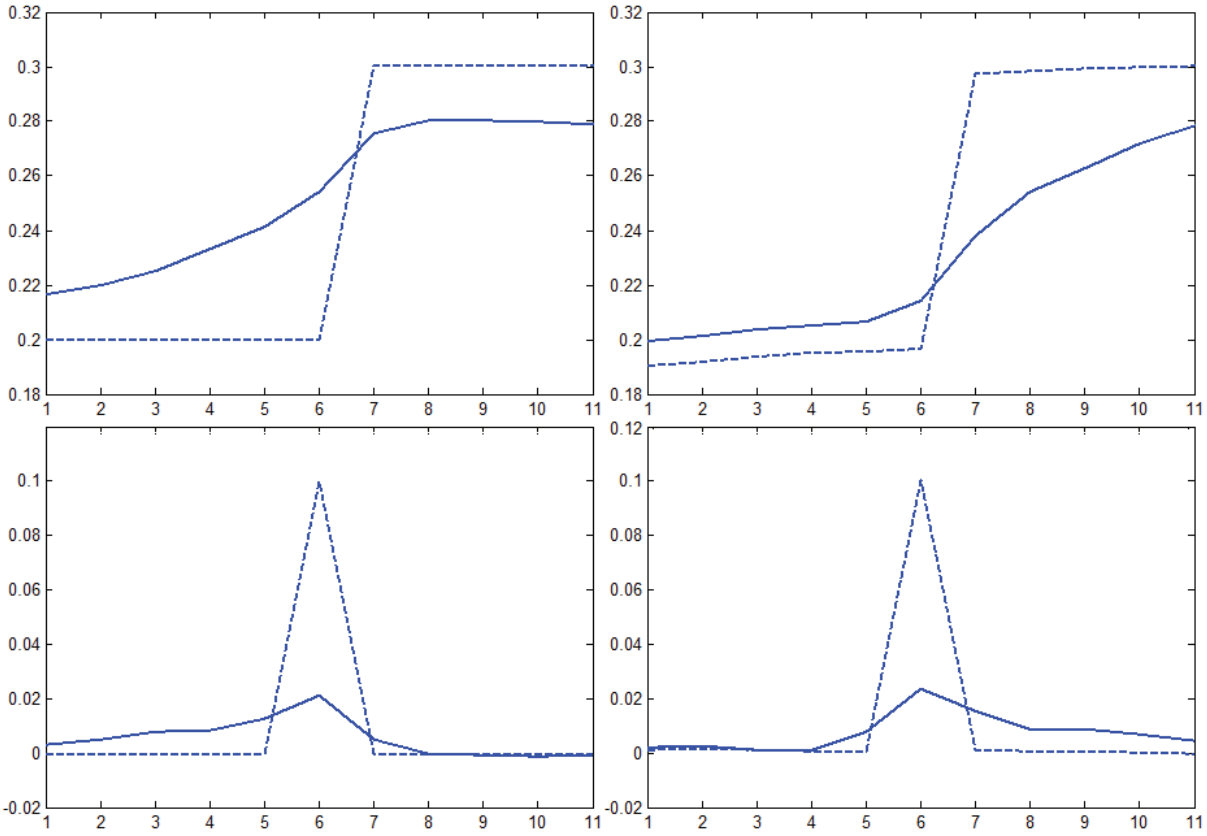

Figure 26. Profiles through the border of the beating ellipsoid (top row) and their first order derivative (bottom row) for systole (left column) and diastole (right column). The dashed line is the ground truth, and the solid line is the TV regularized 3D method.

The first derivative of the profiles confirms that the edges of the beating ellipsoid are sharper than in the reconstructions obtained by the deconvolution methods: the maximum of the first derivative is slightly above 0.02 in both systole and diastole, while it is around 0.016 for deconvolution methods.

In addition to this increase in sharpness, the global RMSE is similar to that of IFDK reconstructions and the RMSE in the beating region is improved with respect to deconvolution results, as shown in Table 3.

		TV-regularized 3D
RMSE	Systole	0.0452
	Diastole	0.0463
RMSE in ROI	Systole	0.0175
	Diastole	0.0233

Table 3. Root Mean Squared Errors for TV-regularized 3D reconstruction on the beating Shepp and Logan phantom

Minimizing total variation leads to piecewise constant images, therefore it is not surprising that TV regularization performs well on the piecewise constant Shepp & Logan phantom. Whether or not TV regularization is efficient on clinical data, which contains textures, must also be evaluated.

Figure 27 and Figure 28 show the reconstruction results on patients DENVER1 and DENVER2 respectively.

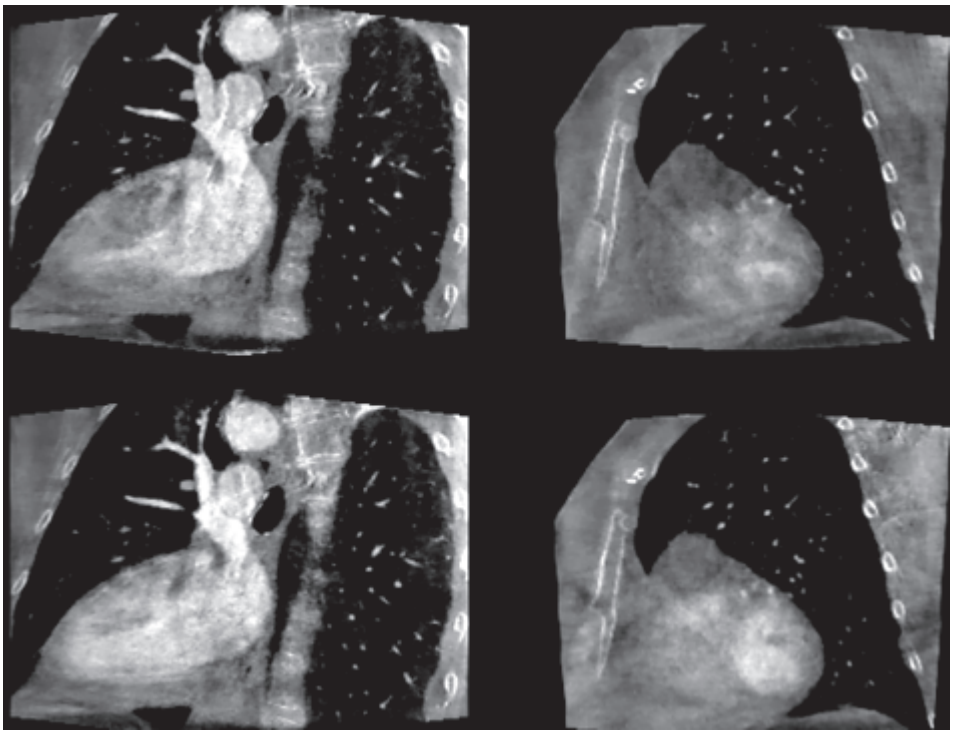

Figure 27. Reconstructions of patient DENVER1 by the TV-regularized 3D reconstruction method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

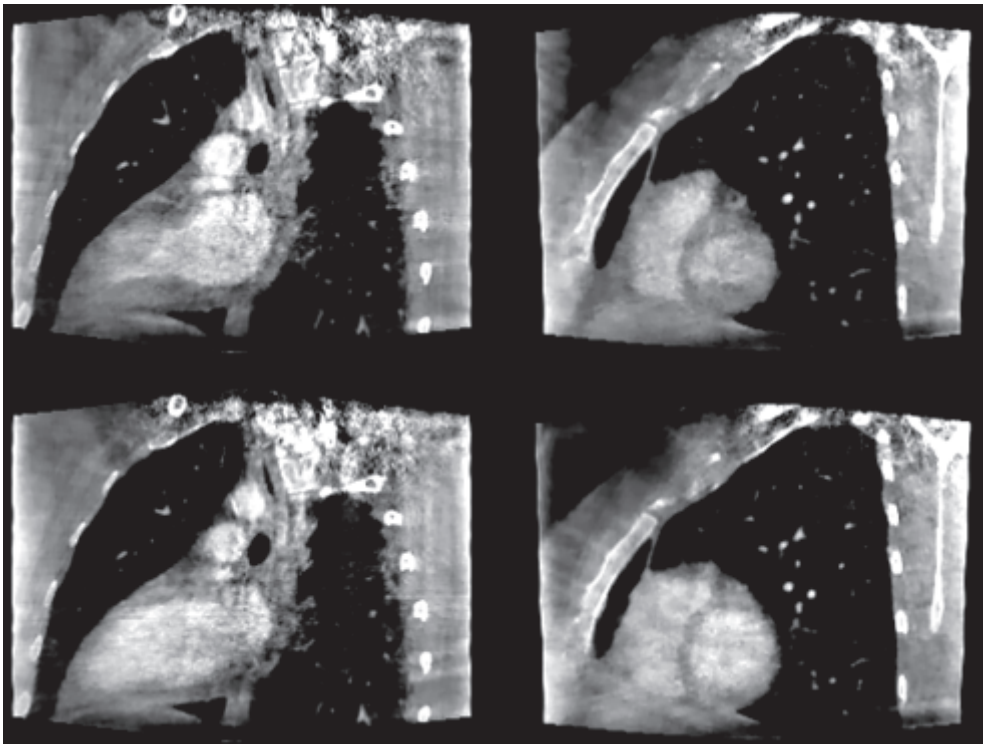

Figure 28. Reconstructions of patient DENVER2 by the TV-regularized 3D reconstruction method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

In Figure 27, the border between muscle and blood containing contrast medium is well defined in the long axis cut of the systole, but blurry and hard to identify in the other images. In Figure 28, this

border is well defined in all images. While this difference can be caused by a number of factors, it is consistent with the observation made in section I.2.c on the impact of heart rate on reconstruction quality: patient DENVER2 had a higher heart rate than patient DENVER1 (88 bpm vs. 60 bpm).

The results presented here are the best tradeoff I achieved between texture preservation (which requires low TV-regularization) and accuracy of the movement reconstruction (which requires high TV-regularization, as explained in section V.6).

This method seems sufficient for a patient with a high heart rate, but cannot be used on a patient with a standard heart rate, which makes it of little use for clinical practice. Other methods have to be investigated. Finding a sparsifying transform better suited to cardiac CT images than total variation would allow applying a stronger regularization (which is necessary to reconstruct the movement more accurately) while still preserving textures. Wavelets have been shown to be efficient in compressing natural images, and therefore seem a rather natural choice. The next section describes a reconstruction method using Daubechies wavelets regularization.

VI.2. 3D reconstruction with Daubechies wavelets regularization

VI.2.a. Cost function

The cost function we try to minimize in this section is the following

$$J(f) = \|G(Rf - p)\|_2^2 + \alpha\|Wf\|_1$$

where W is a Daubechies wavelet transform. The order of the wavelet filter and the number of decomposition levels are parameters of the reconstruction method.

VI.2.b. Minimization algorithm

We employ the same method as in section 0 to find a minimum of J. The calculations, though, differ a bit because the operator $W^T = W^{-1}$, which allows to simplify several equations. As previously, the problem is first converted to a constrained optimization problem using variable splitting:

$$\begin{cases} (\hat{f}, \hat{g}) = \underset{f,g}{\arg\min} \|G(Rf - p)\|_2^2 + \alpha\|Wf\|_1 \\ \text{subject to } g = Wf \end{cases}$$

The corresponding ADMM iterations are as follows:

$$\begin{cases} f_{k+1} = \underset{f}{\text{argmin}} \quad \|G(Rf - p)\|_2^2 + \beta\|Wf - g_k - d_k\|_2^2 \\ g_{k+1} = \underset{g}{\text{argmin}} \quad \alpha\|Wf\|_1 + \beta\|Wf_{k+1} - g - d_k\|_2^2 \\ d_{k+1} = d_k - Wf_{k+1} + g_{k+1} \end{cases}$$

As in section VI.1.b, f_{k+1} is the zero of the gradient of $\|G(Rf - p)\|_2^2 + \beta\|Wf - g_k - d_k\|_2^2$:

$$R^T G^2 Rf - R^T Gp + \beta W^T(Wf - g_k - d_k) = 0$$

$$\Leftrightarrow \underbrace{(R^T G^2 R + \beta I)}_{A} f = \underbrace{R^T Gp + \beta W^{-1}(g_k + d_k)}_{b}$$

where I denotes the identity operator. f_{k+1} is obtained by conjugate gradient algorithm aimed at solving $Af = b$.

g_{k+1} is, again, obtained by soft thresholding (Daubechies *et al* 2004)

$$g_{k+1}(v) = ST_{\frac{\alpha}{2\beta}}(Wf_{k+1} - d_k)$$

VI.2.c. Implementation

In the RTK framework, it was easier to store and handle a volume than its wavelets decomposition. Since W is invertible, both options are mathematically equivalent. In order to ease the implementation, we modified the algorithm described in section VI.2.b to store g'_k and d'_k instead of g_k and d_k, where $g'_k = Wg_k$ and $d'_k = Wd_k$.

The soft thresholding operation, in practice, is performed on all wavelet bands but the lowest frequency one. It gives better results in practice, as the low frequency coefficients are critical.

VI.2.d. Results

The proposed method is evaluated on the datasets described in Chapter IIIChapter I. The performance on phantom data is measured with the metrics and methods described in section V.4 (slices, profiles for sharpness measurement, RMSE and RMSE in the beating region).

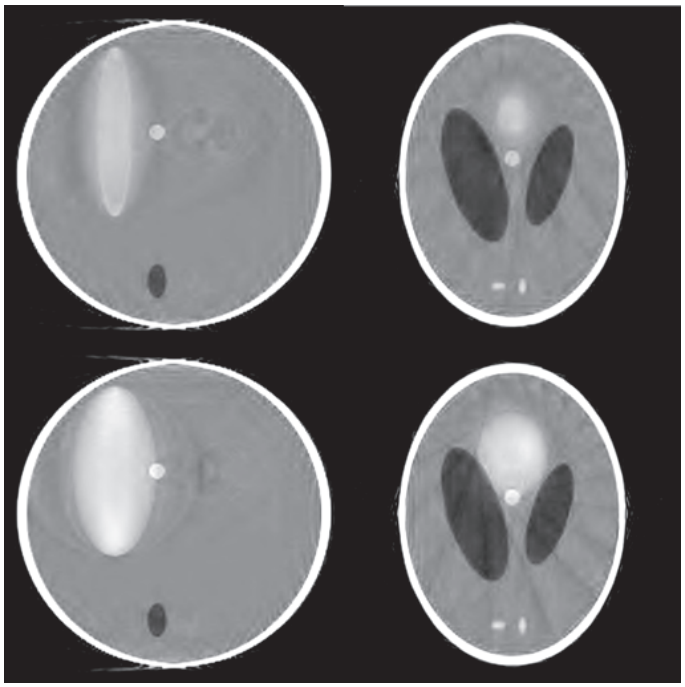

Figure 29. Reconstructions of the moving Shepp & Logan phantom by the wavelets-regularized 3D method, in both systole (top row) and diastole (bottom row).

The beating ellipsoid is well defined, and systole and diastole can be recognized, but artifacts appear on every edge: since sharp edges are not sparse in the wavelet domain used here, the strong regularization required to obtain a satisfying difference between systole and diastole has generated artifacts.

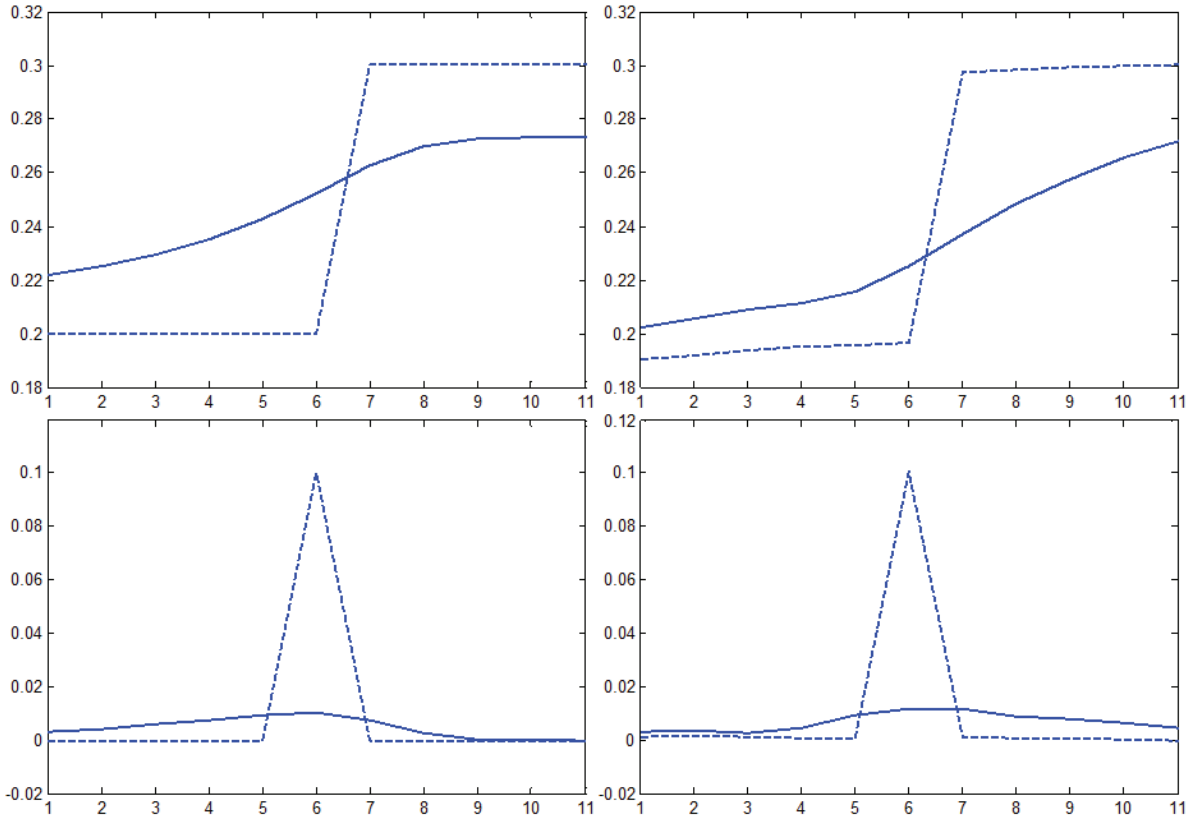

Figure 30. Profiles through the border of the beating ellipsoid (top row) and their first order derivative (bottom row) for systole (left column) and diastole (right column). The dashed line is the ground truth, and the solid line is the wavelets regularized 3D method.

The profiles and their first derivative confirm that the edges of the beating ellipsoid are blurrier than in any of the previously shown reconstructions. The global RMSE and the RMSE in the beating region also prove the poor performance of this method on the Shepp & Logan phantom, as shown in Table 4.

		Wavelets-regularized 3D
RMSE	Systole	0.0478
	Diastole	0.0493
RMSE in ROI	Systole	0.0205
	Diastole	0.0270

Table 4. Root Mean Squared Errors for wavelets-regularized 3D reconstruction on the beating Shepp and Logan phantom

As already mentioned in section III.4, the efficiency of a compressed sensing method depends on whether the regularization employed matches the actual properties of the image to be reconstructed. Minimizing the L1-norm of the Daubechies wavelets decomposition was expected to be a poor regularization choice in the case of a piecewise constant phantom. On the other hand, it should be well suited to the reconstruction of textured images. Figure 31 and Figure 32 show the results of this method on patient DENVER1 and patient DENVER2 respectively.

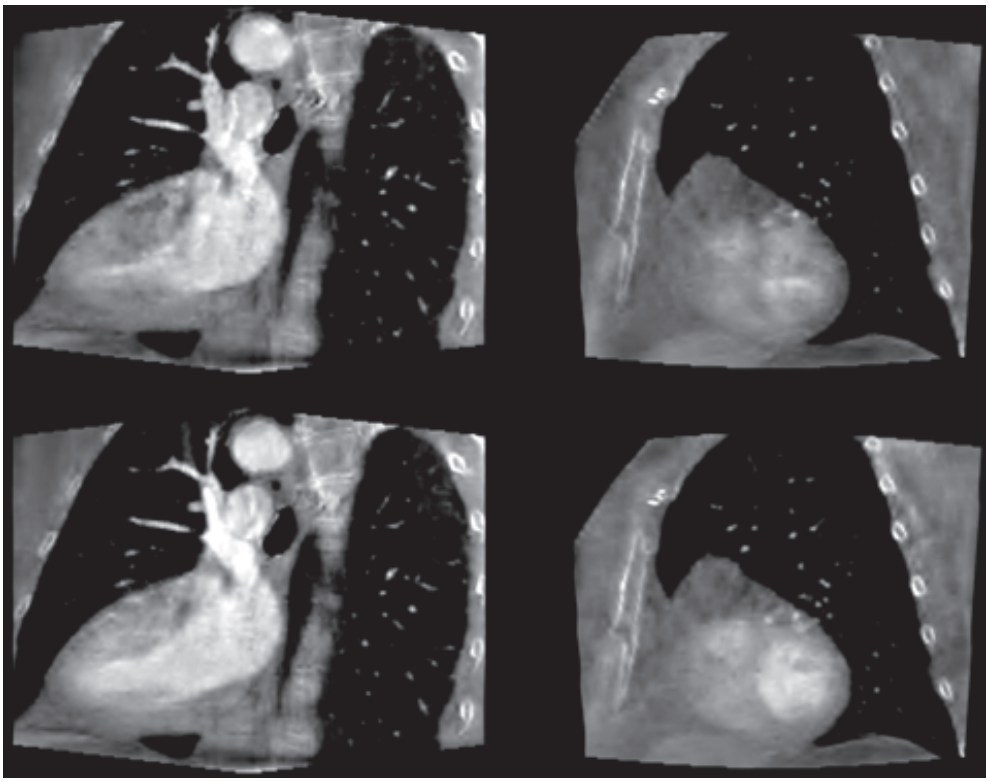

Figure 31. Reconstructions of patient DENVER1 by the wavelets-regularized 3D reconstruction method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

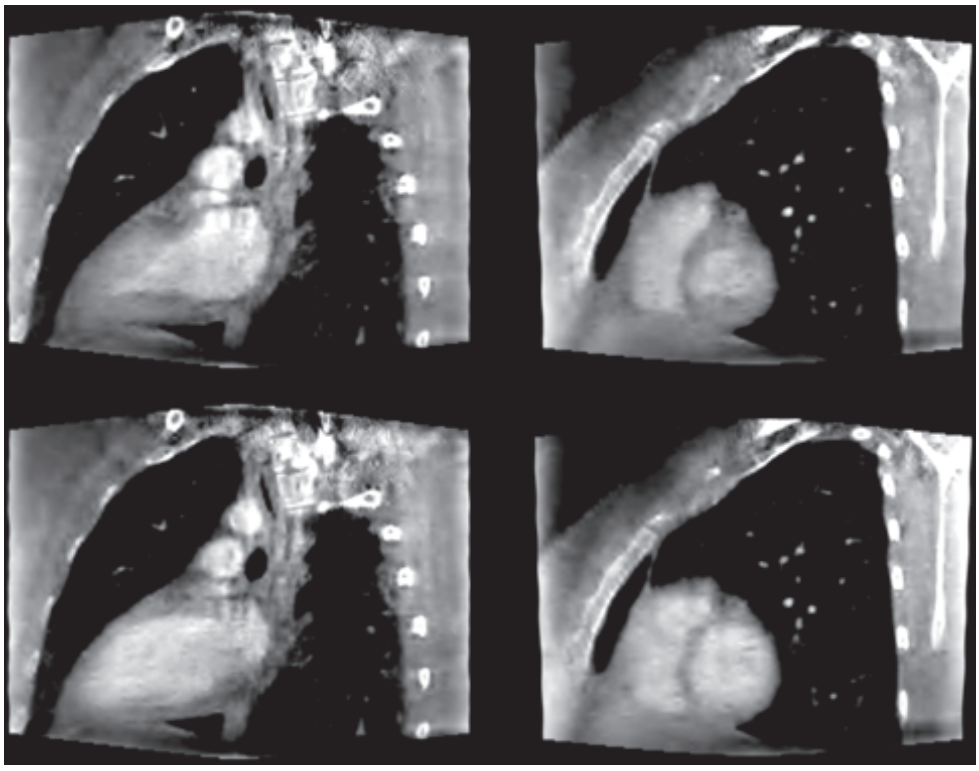

Figure 32. Reconstructions of patient DENVER2 by the wavelets-regularized 3D reconstruction method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

On real data, the results are more convincing than on the Shepp & Logan phantom: the regularization has removed much of the high-frequency noise that was still present in the TV-regularized images (red arrows in Figure 31 and Figure 32), and led to well differentiated systole and diastole reconstructions. The change in regularization improved the overall reconstruction quality. It should now be compared with a state of the art method.

The Prior Image Constrained Compressed Sensing method, short PICCS, is a generalization of the TV-regularized 3D reconstruction method described in section 0. It has been used several times in the literature to perform cardiac C-arm CT reconstruction, and is therefore the *de facto* gold standard for cardiac C-arm CT.

VI.3. PICCS with TV regularization

Prior Image Constrained Compressed Sensing has been developed to reduce the texture-erasing effect of TV regularization. TV is balanced with another regularization term: the attachment to a "prior" image, known before the reconstruction and chosen close to the expected result. In practice, this second regularization term forces the reconstruction to have a texture similar to that of the prior image.

There are numerous versions of the PICCS algorithm available in the literature (Chen *et al* 2008, 2012, 2009, Lauzier *et al* 2012b, Bergner *et al* 2010, Nett *et al* 2008). The definition of the cost function is stable throughout these publications, but the algorithm used for minimization varies a lot from one paper to another. In all of them, the regularization used to compute the results is TV, although the cost function is usually defined with a more general sparsifying transform.

In this section, we describe in details the version of the PICCS algorithm with TV regularization we used.

VI.3.a. Cost function

PICCS aims to find the minimum of the following cost function:

$$J(f) = \mu\|G(Rf - p)\|_2^2 + (1-\alpha)TV(f) + \alpha TV(f - f^*)$$

where f^* is the prior image (in our case, the ungated FDK reconstruction), μ is a parameter controlling the relative weights of data attachment and regularization terms, and $\alpha \in [0;1]$. α controls the relative weights of the two regularization terms. Note that when α is set to 0, the PICCS cost function boils down to that of TV-regularized 3D reconstruction method presented in section 0.

VI.3.b. Minimization algorithm

Where TV is differentiable, its gradient can be calculated analytically, therefore steepest descent and non-linear conjugate gradient can be employed to minimize $J(f)$ (Lauzier *et al* 2012a). Where it is not differentiable, i.e. when $\nabla_x f(v) = \nabla_y f(v) = \nabla_z f(v) = 0$, we avoid the division by zero by adding a small ε to the denominator. It is a bold way of dealing with the problem, and a more elegant one would be to use Chambolle's method to compute the proximal operator of TV (Chambolle 2004), but it is computationally efficient and gives good results in practice. The augmented lagrangian and ADMM with 3 terms could also be used, but it would require a third level of nested iterations.

We have chosen a simple approach, proposed in (Chen *et al* 2008), which consists in minimizing separately the data attachment term $\|G(Rf - p)\|_2^2$ and the sum of the TV terms $(1-\alpha)TV(f) + \alpha TV(f - f^*)$ in an alternative fashion. Let us denote by $SART(R, p, f_0)$ the result of SART using the forward-projection operator R, the set of measured projections p and the volume f_0 as initialization. Let us similarly denote by $SD(\alpha, f^*, f_0)$ the minimization by steepest descent of the sum of TV terms using the parameter α, the prior image f^* and the volume f_0 as initialization. It results in the following update steps (Chen *et al* 2008):

$$\begin{cases} f_{k+1}^{SART} = SART(GR, Gp, f_k) \\ f_{k+1} = SD\left(\alpha, f^*, f_{k+1}^{SART}\right) \end{cases}$$

This algorithm is simpler than the various conjugate gradient approaches studied in (Lauzier *et al* 2012a) and easier to implement. However, it behaves badly when α is close to zero: with enough iterations, $SD(0, f^*, f_0)$ returns a uniform volume, because the steepest descent is not constrained to preserve any structure from f_0 (when α is large enough, typically larger than 0.1, the attachment to f^* is sufficient to avoid this effect). Therefore we have slightly modified the steepest descent part, adding a data-attachment term to the initial volume, as suggested in (Chan *et al* 2001):

$$\text{Modified } SD(0, f^*, f_0) = \arg\min_f \ \lambda\|f - f_0\|_2^2 + (1-\alpha)TV(f) + \alpha TV(f - f^*)$$

and the algorithm we obtain is

$$\begin{cases} f_{k+1}^{SART} = SART(GR, Gp, f_k) \\ f_{k+1} = \text{Modified } SD\left(\alpha, f^*, f_{k+1}^{SART}\right) \end{cases}$$

There is no trivial proof that this algorithm indeed minimizes the PICCS cost function introduced in section VI.3.a. As a consequence, proving that it actually converges is hard. Section VII.2.c contains a partial proof of convergence of the 4D ROOSTER method (presented later in this document, in section 0) using the theory of non-expansive mappings. It can be adapted to PICCS in order to prove that if this algorithm has at least one fixed point, it converges to one of its fixed points.

VI.3.c. Results

Figure 33 shows reconstructions of the moving Shepp and Logan phantom using PICCS with TV regularization, and Figure 34 displays profiles through the border of the beating ellipsoid. As already stated in section III.4, it would make little sense to analyze these results in detail: PICCS was designed to work best on real data, and the best results on piecewise constant phantoms are obtained by setting α to zero or almost zero, which amounts to using the TV-regularized 3D reconstruction method presented in section 0.

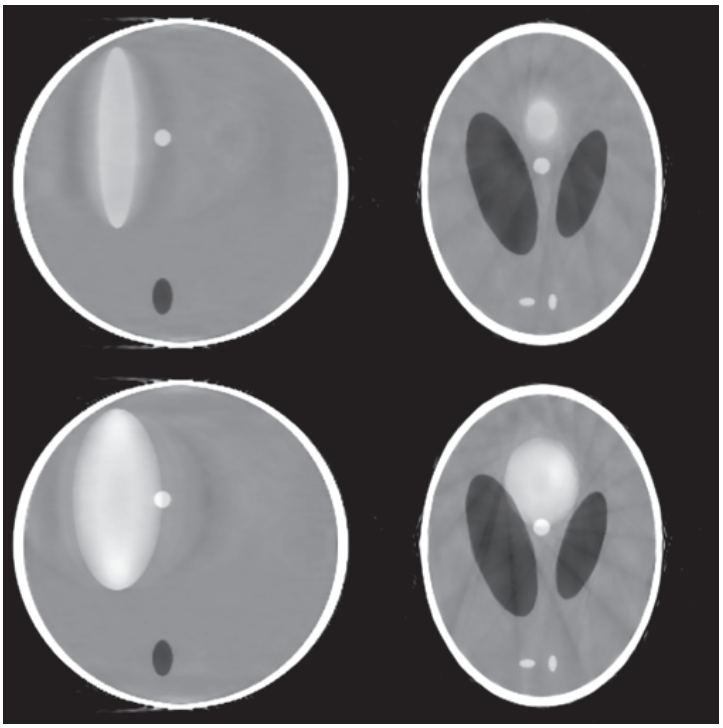

Figure 33. Reconstructions of the moving Shepp & Logan phantom by the PICCS TV method, in both systole (top row) and diastole (bottom row).

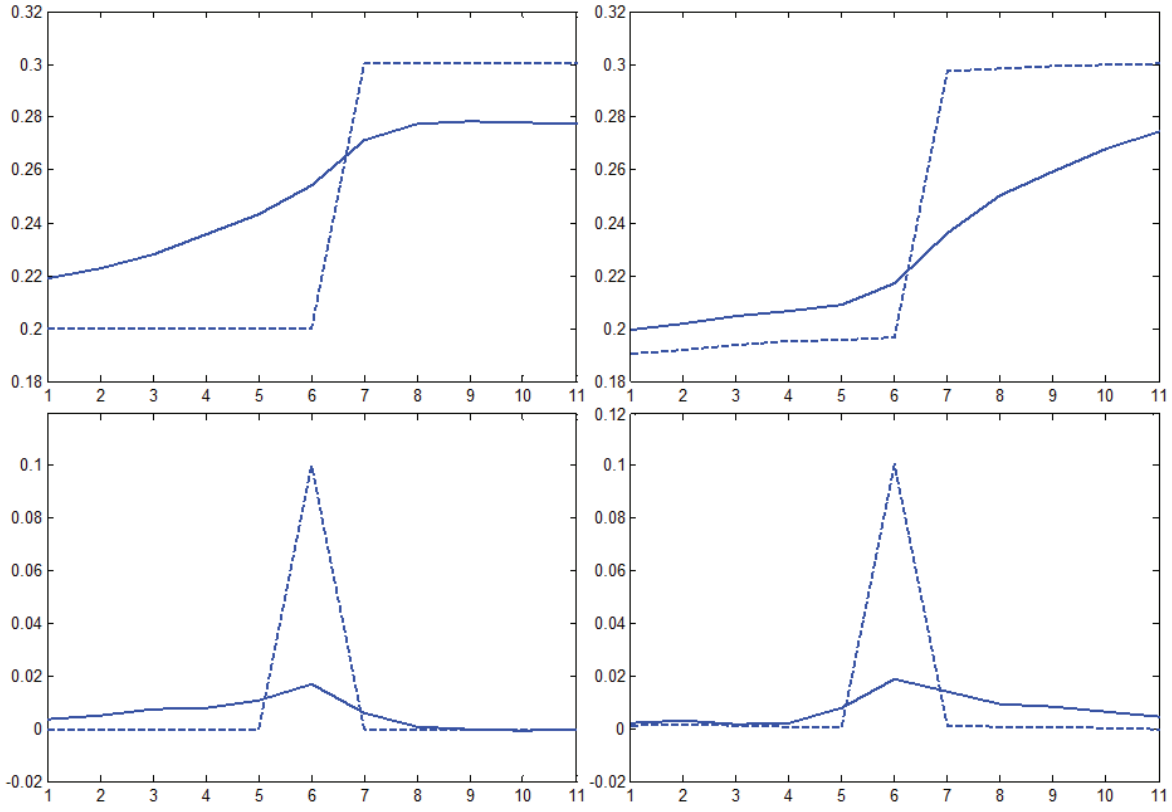

Figure 34. Profiles through the border of the beating ellipsoid (top row) and their first order derivative (bottom row) for systole (left column) and diastole (right column). The dashed line is the ground truth, and the solid line is the PICCS TV method.

		PICCS
RMSE	Systole	0.0441
	Diastole	0.0454
RMSE in ROI	Systole	0.0189
	Diastole	0.0252

Table 5. Root Mean Squared Errors for PICCS reconstruction on the beating Shepp and Logan phantom

Figure 35 and Figure 36 contain the reconstruction results using PICCS on patient DENVER1 and patient DENVER2 respectively. The results are visually similar to those of the previously presented compressed sensing-based reconstruction methods.

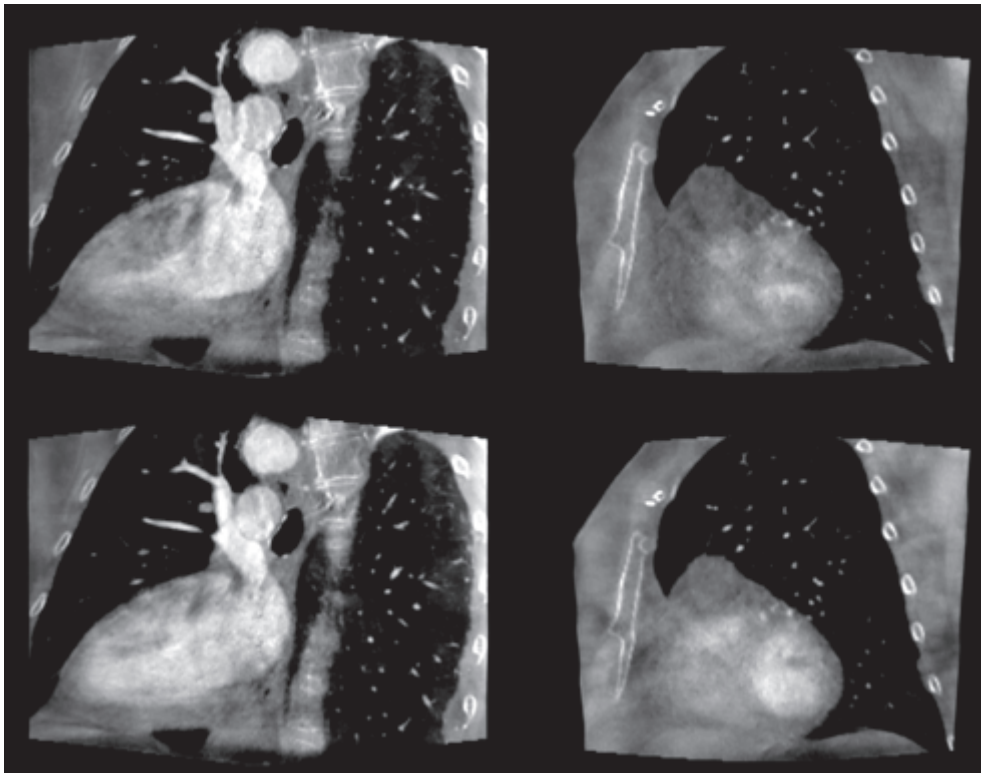

Figure 35. Reconstructions of patient DENVER1 by the PICCS TV method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

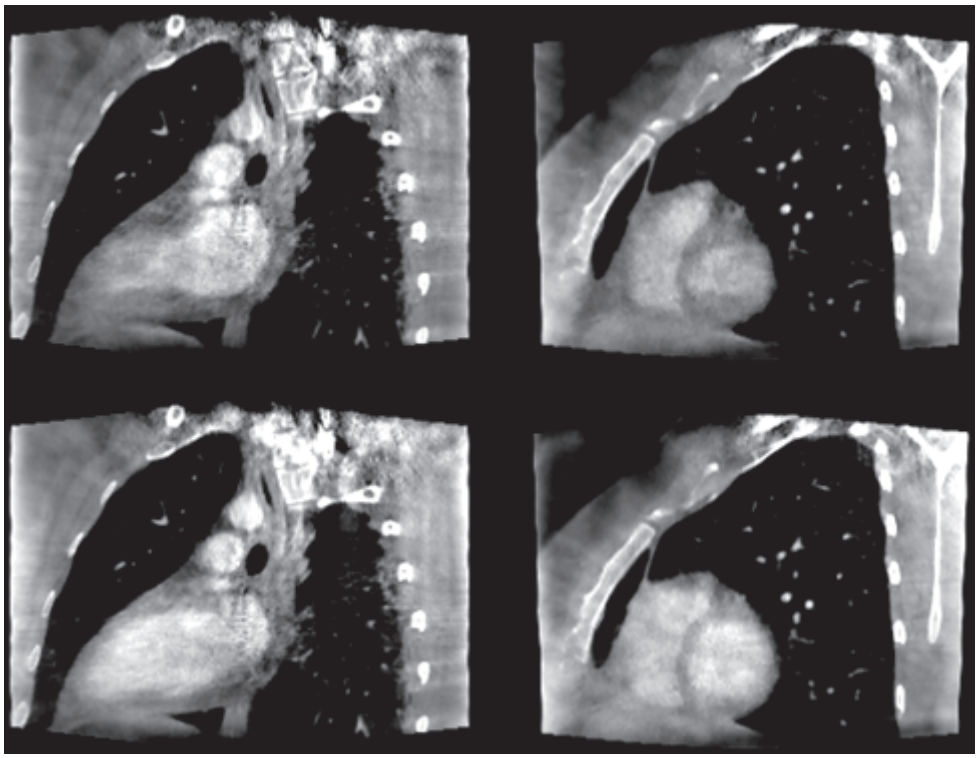

Figure 36. Reconstructions of patient DENVER2 by the PICCS TV method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

VI.4. Conclusions on 3D reconstruction of a single cardiac phase

VI.4.a. Choice of the sparsifying transform

In section V.6, we have shown that a strong regularization is essential to obtain a sharp reconstruction of the moving structures in each phase. Whether the regularization can be set strong enough to reach good reconstruction results without generating artifacts or completely smoothing out textures, depends on how sparse the solution indeed is in the chosen basis. We have seen that TV minimization is well suited to piecewise constant images with sharp edges, and performs reasonably well on real data, while Daubechies wavelets are efficient on real data and suboptimal on piecewise constant images like the Shepp and Logan phantom.

VI.4.b. Consistency of the 3D+time sequences

The methods described in this chapter reconstruct each cardiac phase independently. The volumes can then be assembled to form a 3D+time reconstruction.

It turns out that they often yield solutions with strong attenuation variations over the cardiac cycle, even in regions where no movement is expected. This is illustrated on Figure 37, which shows the difference between two consecutive phases (separated by 10% of the cardiac cycle).

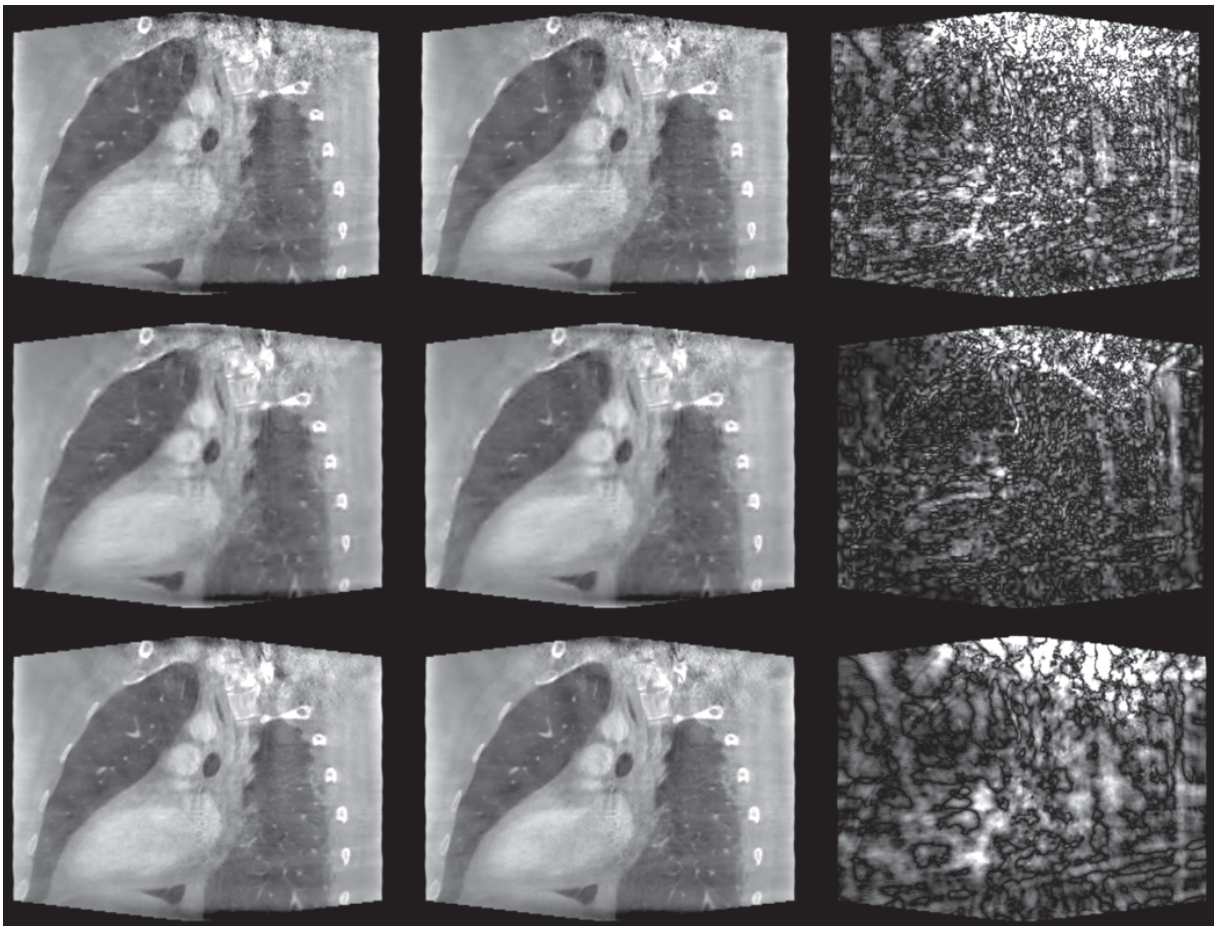

Figure 37. Phases 0% (left column) and 10% (middle column) of the cardiac cycle obtained by TV-regularized reconstruction (top row), Wavelets-regularized reconstruction (middle row) and PICCS (bottom row). The right column contains the absolute value of the difference between these two phases, with a 10 times thinner visualization window.

This limitation of 3D reconstruction algorithms can only be overcome by reconstructing the whole cardiac cycle at once and enforcing some consistency between different instants of the cardiac cycle. Moreover, such an approach will exploit more information for the reconstruction of each specific cardiac phase, so that we can expect a global improvement of the reconstruction quality. It is the purpose of the algorithms proposed in the next chapter.

Chapter VII : Regularized 3D+t reconstruction

In this chapter, we introduce two 3D+time reconstruction methods, in order to address the problems of pure 3D methods pointed out in section VI.4.b. Both require a segmentation of the regions undergoing movement for the reconstruction. For the cases presented in this document, the segmentation has been performed on the ungated FDK reconstruction (see Figure 38) using the semi-automatic segmentation tool LiveMorph (Mory *et al* 2012). It could be replaced by a fully automatic method like the one used in (Schwartz *et al* 2011), especially because the segmentation used by 4D ROOSTER can be very rough. In this work, the segmentation obtained from the ungated FDK reconstruction was artificially expanded to include a margin of error.

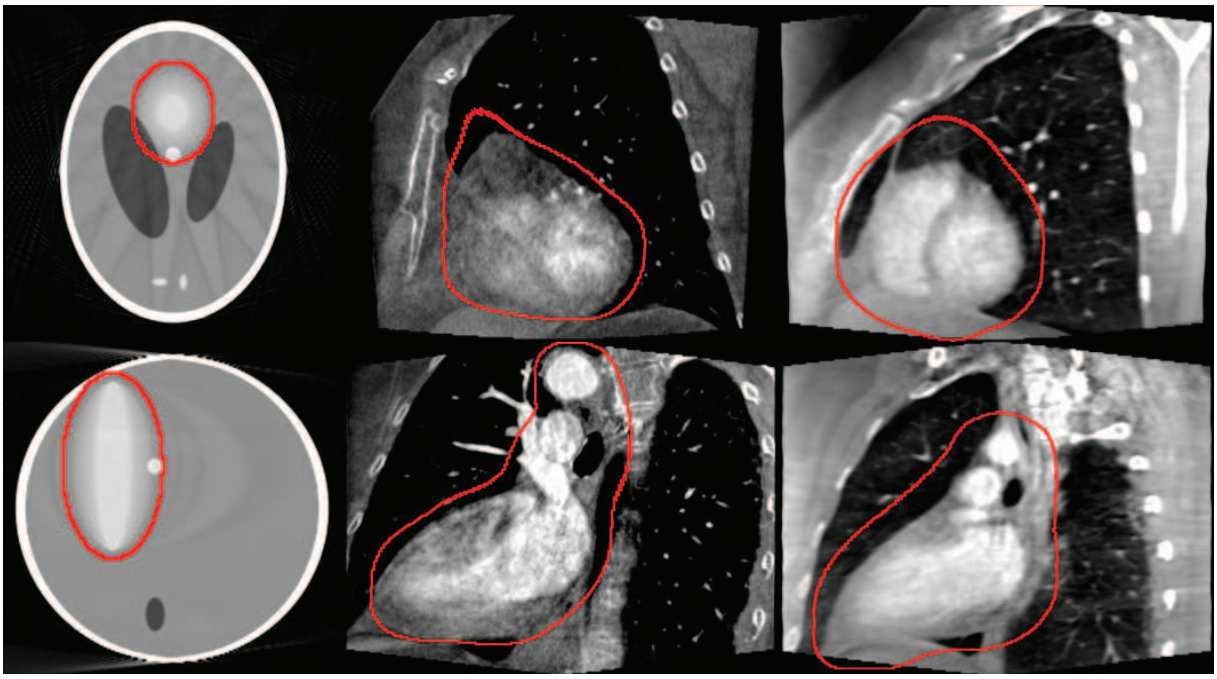

Figure 38. Motion masks displayed on top of ungated FDK reconstructions of the moving Shepp & Logan phantom (left column), patient DENVER1 (middle column) and patient DENVER2 (right column). The top row contains short axis cuts, and the bottom row long axis cuts. The contours of the motion mask appear in red.

VII.1. 3D+t reconstruction with ROI-based 4D TV regularization

VII.1.a. Cost function

The notations used in the previous chapter must be slightly adapted to fit in a 3D+t reconstruction context:

- $f_1, f_2, \dots f_{\mathrm{N}}$ are 3D volumes, each one representing a single cardiac phase, and $f = \begin{pmatrix} f_1 \\ \vdots \\ f_N \end{pmatrix}$ is a 4D sequence of volumes. $f_1, f_2, \dots f_{\mathrm{N}}$ are column vectors of size M, where M is the number of voxels in a 3D volume. Thus f is a column vector of size MN
- R_θ is the X-ray transform for angle θ
- p_θ is the projection measured at angle θ. The p_θ are column vectors of size P, where P is the number of pixels in a projection, and the R_θ are matrices with P lines and M columns.
- S_θ is an interpolation operator which, from the 3D + time sequence, estimates the 3D volume through which projection θ has been acquired.

S_θ is defined as follows: from the ECG, we know in which phase $t(\theta)$ the patient's heart was when projection p_θ was acquired. The forward projection at angle θ, which will be compared to the measured projection p_θ, should be computed through an estimate of the patient's thorax at $t(\theta)$, which by a slight abuse in notation we will denote $f_{t(\theta)}$, and which is obtained by interpolating between some of the volumes f_i. If N=10, f_1 is the volume at 10% of the cardiac cycle, f_2 the volume at 20%, and so on. With $t(\theta) = 87\%$, the interpolated volume would be $S_\theta f = f_{t(\theta)} = 0.3 f_8 + 0.7 f_9$. In this example, and in the implementation we used, S_θ is a linear interpolator. However, other interpolation methods could be tested, like nearest neighbors or splines. Note that S_θ is a huge matrix with M lines and MN columns, but is not explicitly computed in practice.

This method consists in minimizing the following cost function $J(f)$:

$$J(f) = \sum_\theta \|R_\theta S_\theta f - p_\theta\|_2^2 + \alpha ROI\_TV(f)$$

- α is a parameters that determines the relative weight between the data attachment and regularization terms
- ROI_TV is a modified 4D total variation regularization function defined as follows:

$$ROI\_TV(f) = \sum_{m=1}^{M} \sqrt{[\nabla_{\mathrm{x}} f(m)]^2 + \left[\nabla_{\mathrm{y}} f(m)\right]^2 + [\nabla_{\mathrm{z}} f(m)]^2 + [\omega(m) \nabla_t f(m)]^2}$$

where $\omega(m)$ weighs the temporal component of the gradient with respect to the spatial ones, and depends on whether the voxel v is inside or outside the ROI containing the regions undergoing movement. ω is a vector of size M. If m is outside the ROI, its intensity should not vary much between consecutive phases. Therefore $\omega(m)$ should be high. On the other hand, if m is inside the ROI, it is normal that its intensity varies over the cardiac cycle, and $\omega(m)$ should be low in order to allow this variation. In practice we use only two values of $\omega(m)$, one for the voxels inside the ROI and one for the voxels outside the ROI, denoted ω_{in} and ω_{out} respectively, with $\omega_{in} < \omega_{out}$.

From these weights $\omega(m)$, we define a diagonal matrix H of size 4M, which performs the voxel-by-voxel multiplication of the gradient of f by either 1 (for x, y and z components of the gradient), ω_{in} (for voxels inside the ROI) or ω_{out} (for voxels outside the ROI).

In order to have a finer temporal resolution during systole, I have tried using a denser temporal sampling during systole and a looser one during diastole, instead of having a sequence f of 3D volumes regularly spaced in the cardiac cycle (for example, every 10%). Noticing no improvement, I reverted to regularly spaced volumes for the sake of code simplicity.

VII.1.b. Minimization algorithm

Using the matrix H, and similarly to section VI.2.b, we transform the unconstrained problem of finding the minimum of J into a constrained problem by variable splitting:

$$\begin{cases} (\hat{f}, \hat{g}) = \underset{f,g}{\arg\min} \sum_{\theta} \|R_\theta S_\theta f - p_\theta\|_2^2 + \alpha \sum_{v=1}^{V} \|g(v)\|_2 \\ \text{subject to } g = H\nabla f \end{cases}$$

Using the Augmented Lagrangian and the Alternating Direction Method of Multipliers, this problem is solved by the following iterative procedure:

$$\begin{cases} f_{k+1} = \underset{f}{\arg\min} \quad \sum_{\theta} \|R_\theta S_\theta f - p_\theta\|_2^2 + \beta \|H\nabla f - g_k - d_k\|_2^2 \\ g_{k+1} = \underset{g}{\arg\min} \quad \alpha \sum_{v=1}^{V} \|g(v)\|_2 + \beta \|H\nabla f_{k+1} - g - d_k\|_2^2 \\ d_{k+1} = d_k - H\nabla f_{k+1} + g_{k+1} \end{cases} \quad \text{(VII.1)}$$

f_{k+1} is the zero of the gradient of $\sum_\theta \|R_\theta S_\theta f - p_\theta\|_2^2 + \beta \|H\nabla f - g_k - d_k\|_2^2$, that is to say the value of f such that:

$$\sum_{\theta} {S_\theta}^T {R_\theta}^T R_\theta S_\theta f - {S_\theta}^T {R_\theta}^T p_\theta + \beta \nabla^T H^T (H\nabla f - g_k - d_k) = 0$$

$$\Leftrightarrow \underbrace{\left(\sum_{\theta} {S_\theta}^T {R_\theta}^T R_\theta S_\theta + \beta \nabla^T H^T H \nabla \right)}_{A} f = \underbrace{\sum_{\theta} {S_\theta}^T {R_\theta}^T p_\theta + \beta \nabla^T H^T (g_k + d_k)}_{b}$$

f_{k+1} can thus be computed by a conjugate gradient algorithm aimed at solving $Af = b$. Note that $H^T = H$ (H is diagonal) and therefore $H^T H = H^2$.

As in section VI.1.b, g_{k+1} can be computed voxel-by-voxel, and for each voxel the problem boils down the soft thresholding on a vector. Therefore, for any v,

$$g_{k+1}(v) = ST_{\frac{\alpha}{2\beta}}(H\nabla f_{k+1} - d_k) \quad \text{(VII.2)}$$

VII.1.c. Results

The proposed method is evaluated on the datasets described in Chapter IIIChapter I. Performance on phantom data is measured with the metrics and methods introduced in section V.4 (slices, profiles for sharpness measurement, RMSE and RMSE in the beating region).

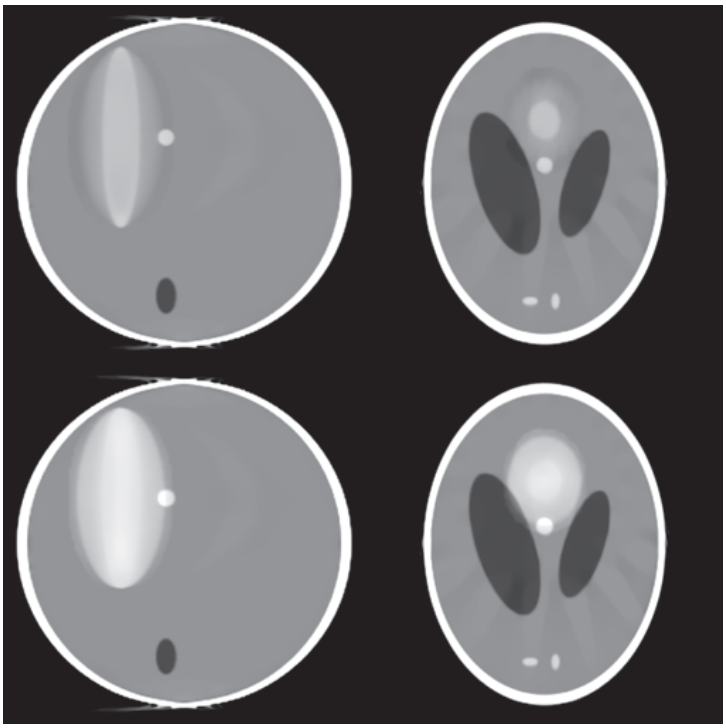

Figure 39. Reconstructions of the moving Shepp & Logan phantom by the ADMM 4D TV method, in both systole (top row) and diastole (bottom row)

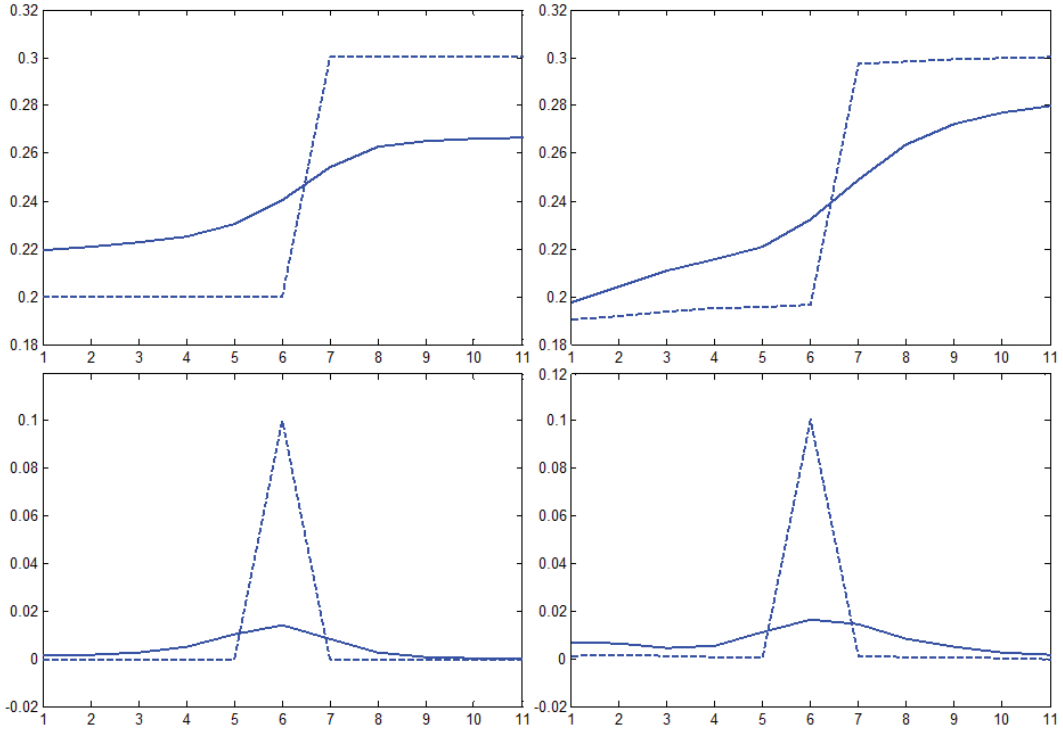

Figure 40. Profiles through the border of the beating ellipsoid (top row) and their first order derivative (bottom row) for systole (left column) and diastole (right column). The dashed line represents the ground truth, and the solid line represents the ADMM 4D TV method.

		ADMM 4D TV
RMSE	Systole	0.0430
	Diastole	0.0438
RMSE in ROI	Systole	0.0190
	Diastole	0.0227

Table 6. Root Mean Squared Errors for ADMM 4D TV reconstruction on the beating Shepp and Logan phantom

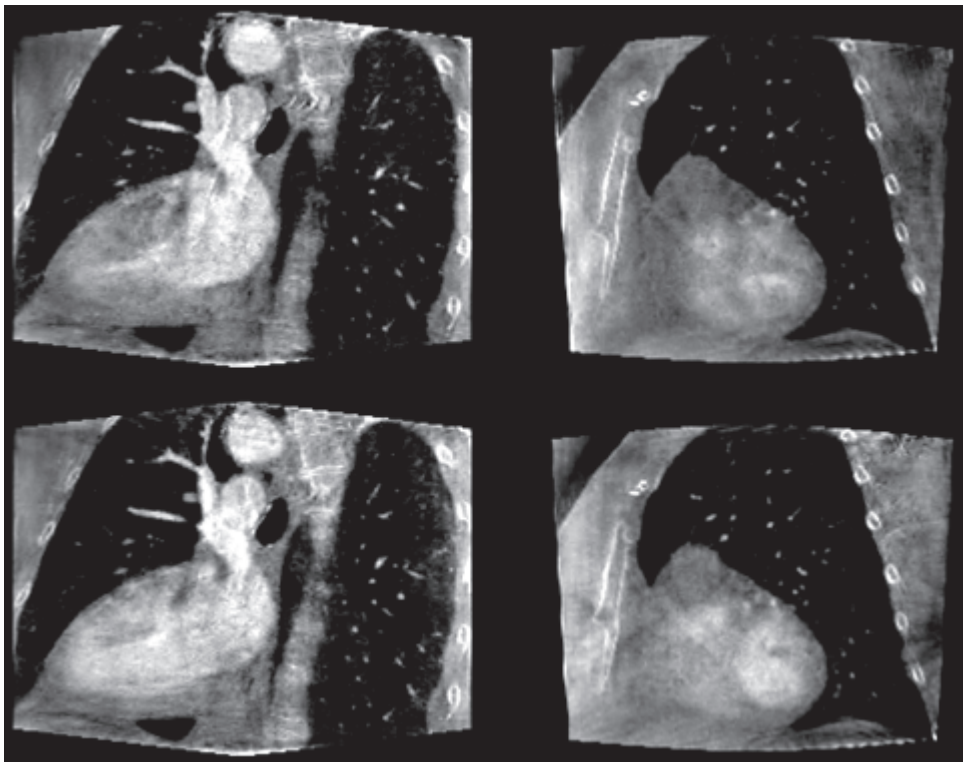

Figure 41. Reconstructions of patient DENVER1 by the ADMM 4D TV method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

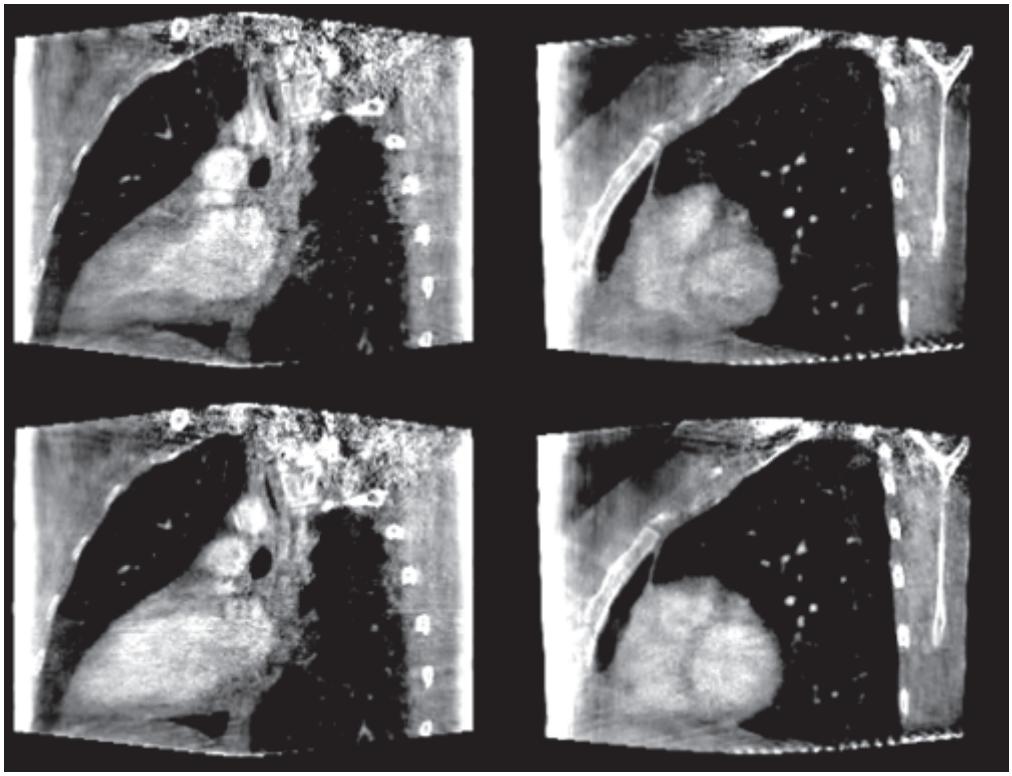

Figure 42. Reconstructions of patient DENVER2 by the ADMM 4D TV method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

The results obtained on the Shepp & Logan phantom have fewer streaks than those of PICCS or ADMM 3D TV, which reduces the global RMSE, but the beating ellipsoid has blurrier edges.

The results on real data are not convincing either: they are very similar to those of the TV regularized 3D reconstruction method presented in section Chapter I.

Considering the significant increase in the complexity of the algorithm, those are quite disappointing results.

VII.1.d. Practical difficulties

In theory, with a sufficiently large number of iterations, any value of the parameter β could be used. If the number of iterations of the main loop is kept relatively low, the choice of β is important. Even with an optimized implementation and a limited volume resolution, each iteration of the main loop takes about two minutes. Performing hundreds of iterations would make the execution time unacceptable for clinical practice, therefore a suitable β must be determined.

Equations (VII.1) and (VII.2) show that β is involved in at least two processes: it controls the tolerance on the difference between $H\nabla f$ and g, and determines the amount of soft thresholding in the computation of g_{k+1}. This one parameter controlling two very different processes, it is hard to predict which effect a change in the value of β can have in the reconstruction, and fine-tuning is delicate.

The choice of ω_{in} and ω_{out} is also difficult: $ROI\_TV(f)$ is a sum over all voxels, therefore setting ω_{out} too high makes the contribution of the voxels located in the ROI negligible, and the algorithm only regularizes outside the ROI. Choosing ω_{out} too low, on the other hand, allows motion to occur in regions that should be mostly static.

These practical fine-tuning difficulties have led me to develop another 3D + time reconstruction method, based on the same ideas (ROI, TV regularization, and 4D reconstruction) but with a different mathematical formulation. It is described in the next section.

VII.2. 4D ROOSTER

In this section, we introduce a new 3D + time reconstruction method, which we name 4D ROOSTER (for 4D RecOnstructiOn using Spatial and TEmporal Regularization). This method builds upon un-regularized 4D reconstruction by conjugate gradient, adding a series of regularization constraints one by one: first positivity, then restriction of motion to a ROI containing the heart and the vessels, and finally spatial and temporal total variation minimization.

VII.2.a. Expectations

The 4D ROOSTER method derives from the desired properties of the solution to the problem:

- (P1) The forward projection through the reconstructed 4D volume should match the measured projection data as accurately as possible
- (P2) All voxels should have non-negative attenuation
- (P3) No movement should occur outside the heart and the vascular system
- (P4) Each volume of the sequence should have some kind of spatial regularity, i.e. be spatially smooth except on the edges of the organs, and be free of streak artifacts
- (P5) Consecutive volumes in the sequence should be similar

The notations we will use in the next section to describe the method are similar to those of section 0. Only the matrix H has a slightly different definition: H is here a binary diagonal matrix with M lines, which "selects" the voxels located outside the heart. As a result, for any i,

$$\begin{cases} Hf_i(x,y,z) = f_i(x,y,z) \text{ if } (x,y,z) \text{ is outside the heart} \\ \quad Hf_i(x,y,z) = 0 \text{ if } (x,y,z) \text{ is inside the heart} \end{cases}$$

The multiplication by the matrix H is the linear algebra equivalent of multiplying by a motion mask.

Now, let us formalize our five requirements on the solution:

- (P_1), taken in the least squares sense, means that $\sum_\theta \|(R_\theta S_\theta f - p_\theta)\|_2^2$ should be small.
- (P_2) is equivalent to imposing that $\forall i$ and $\forall (x,y,z)$, $f_i(x,y,z) \geq 0$.
- (P_3) means that for any i and j, $Hf_i = Hf_j$.
- (P_4) can be enforced by numerous regularization methods. We chose to express it as a constraint on 3D total variation on each volume f_i. The total variation is defined by

  $$TV_{space}(f) = \left\| \sqrt{(\nabla_x f_i)^2 + \left(\nabla_y f_i\right)^2 + (\nabla_z f_i)^2} \right\|_1$$

  The ∇ operator accounts for the spacing between voxels, which can vary from one direction to another.
- (P_5) can also be expressed as a constraint on total variation, although this time a one dimension total variation has to be used. The following constraint is derived:

  $$TV_{time}(f) = \|\nabla_t f\|_1$$

  should be small. The discrete gradient along time is computed with a circular boundary condition because the sequence of volumes is meant to be cyclic.

VII.2.b. Method

The algorithm we propose in order to derive a 3D + time sequence f that has the aforementioned properties consists in the following steps:

- Start from $k = 0$ and $f^{(k)} = 0$
- Until $k = k_{max}$
    - Compute $\hat{f} = \underset{f}{\operatorname{argmin}} \sum_\theta \|(R_\theta S_\theta f - p_\theta)\|_2^2$ by a conjugate gradient descent initialized with $f^{(k)}$
    - Set all negative voxels to zero
    - In each $\hat{f}_i$ replace the area outside the heart by its temporal mean $\frac{1}{N}\sum_j H\hat{f}_j$
    - For each $\hat{f}_i$, compute $\hat{g}_i = \underset{g}{\operatorname{argmin}} \lambda_{space}\left\|g - \hat{f}_i\right\|_2^2 + TV_{space}(g)$ using a gradient descent, as described in (Chan *et al* 2001), and concatenate the $\hat{g}_i$ into $\hat{g}$.
    - Compute $f^{(k+1)} = \underset{f}{\operatorname{argmin}} \lambda_{time}\|f - \hat{g}\|_2^2 + TV_{time}(f)$ by the same method
    - k = k+1

The parameters of the algorithm are the number of main loop iterations k_{max}, the number of nested conjugate gradient iterations, the number of iterations used in the TV minimization, the parameters λ_{space} and λ_{time}, which are data attachment coefficients for TV minimization, and the steps of the gradient descents.

VII.2.c. Convergence

This section provides the theoretical background to prove that, if the main loop of 4D ROOSTER has at least one fixed point, 4D ROOSTER converges to one of these fixed points. However, we cannot prove that such a fixed point does exist. We introduce a few definitions from the theory on non-expansive mappings:

- A mapping T : $\mathbb{R}^{MN} \longrightarrow \mathbb{R}^{MN}$ is non-expansive if $\forall x, y \in \mathbb{R}^{MN}, \|Tx - Ty\|_2 \leq \|x - y\|_2$
- T is strongly non-expansive if T is non-expansive and whenever $(x_n)_{n\in\mathbb{N}}$ and $(y_n)_{n\in\mathbb{N}}$ are sequences in $\mathbb{R}^{MN}$ such that $(x_n - y_n)_{n\in\mathbb{N}}$ is bounded and $\|x_n - y_n\|_2 - \|Tx_n - Ty_n\|_2 \longrightarrow 0$, it follows that $(x_n - y_n) - (Tx_n - Ty_n) \longrightarrow 0$
- T is firmly non-expansive if $\forall x, y \in \mathbb{R}^{MN}, \|Tx - Ty\|_2^2 \leq \langle Tx - Ty, x - y \rangle$

All operators involved in 4D ROOSTER are strongly non-expansive:

- conjugate gradient is convergent. With enough iterations, it is strongly non-expansive
- positivity enforcement is a projector onto a non-empty closed convex set, therefore it is firmly non-expansive, and therefore strongly non-expansive (see fact 4.2 of (Bauschke *et al* 2012))
- averaging along time outside the motion mask is firmly non-expansive, and therefore strongly non-expansive. The proof is given in appendix of this document

- the spatial and temporal TV regularization operators are proximal mappings, therefore they are firmly non-expansive (Bauschke and Combettes 2009), and therefore strongly non-expansive

As stated in fact 4.2 of (Bauschke *et al* 2012), the composition of a finite number of strongly non-expansive mappings is a strongly non-expansive mapping. Thus, each iteration of the main loop of 4D ROOSTER amounts to applying a strongly non-expansive mapping T on the current 3D + time sequence f_k, such that $f_{k+1} = T(f_k)$. Applying the theorem reminded in fact 4.3 of (Bauschke *et al* 2012), if T has at least one fixed point, 4D ROOSTER converges to one of its fixed points.

VII.2.d. Results

The proposed method is evaluated on the datasets described in Chapter III. Performance on phantom data is measured with the metrics and methods introduced in section V.4 (slices, profiles for sharpness measurement, RMSE and RMSE in the beating region).

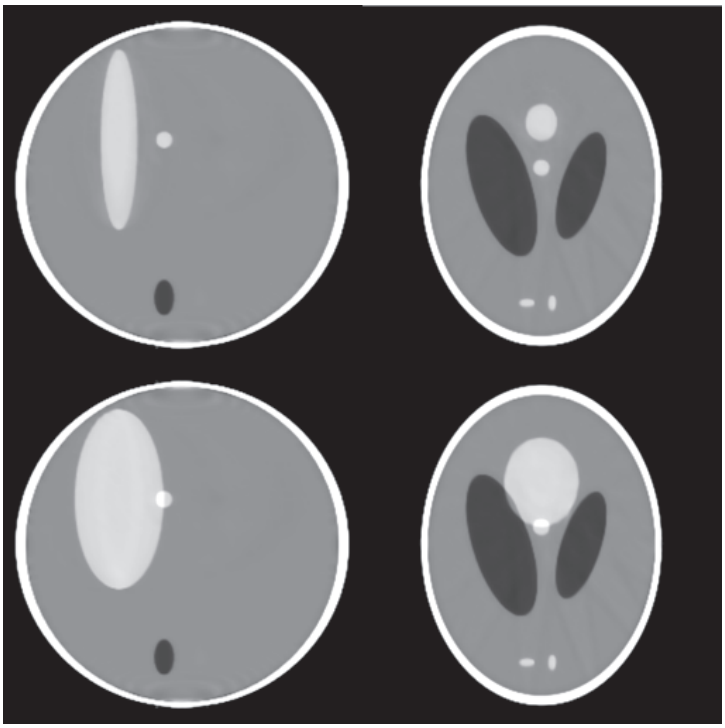

Figure 43. Reconstructions of the moving Shepp & Logan phantom by the 4D ROOSTER method, in both systole (top row) and diastole (bottom row)

The reconstruction results are visually close to the ground truth. Some slight streak artefacts remain (they are visible mainly on the right column of Figure 43), and some cone beam artefacts are also noticeable far away from the central slice (at the top and bottom of the slices displayed on the left column of Figure 43), but apart from these defects the method seems efficient.

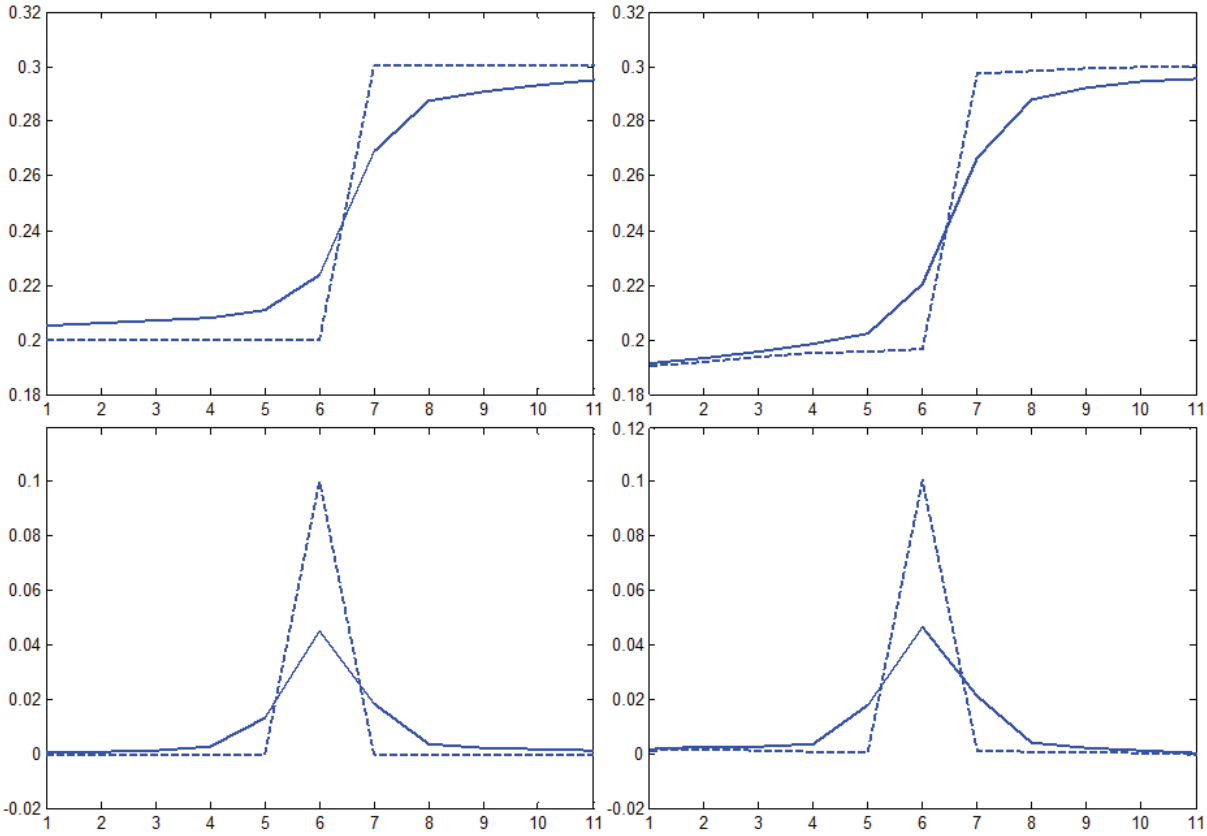

Figure 44. Profiles through the border of the beating ellipsoid (top row) and their first order derivative (bottom row) for systole (left column) and diastole (right column). The solid line represents the ground truth, and the solid line with crosses represents the 4D ROOSTER method.

Profiles through the beating ellipsoid prove that the edges are much sharper in these results than with any of the aforementioned methods. The global and local RMSE, shown in Table 7, are also lower than with the other methods.

		4D ROOSTER
RMSE	Systole	0.0346
	Diastole	0.0347
RMSE in ROI	Systole	0.0093
	Diastole	0.0118

Table 7. Root Mean Squared Errors for 4D ROOSTER reconstruction on the beating Shepp and Logan phantom

In Figure 45 and Figure 46, results of the 4D ROOSTER method are presented for patient DENVER1 and patient DENVER2.

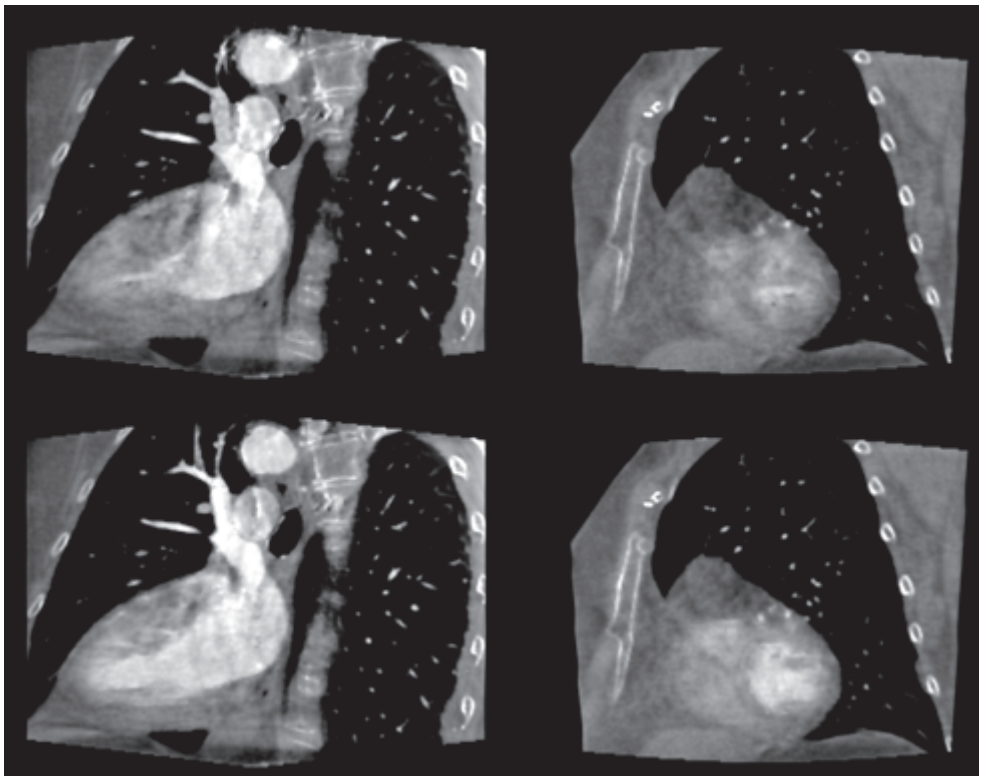

Figure 45. Reconstructions of patient DENVER1 by the 4D ROOSTER method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

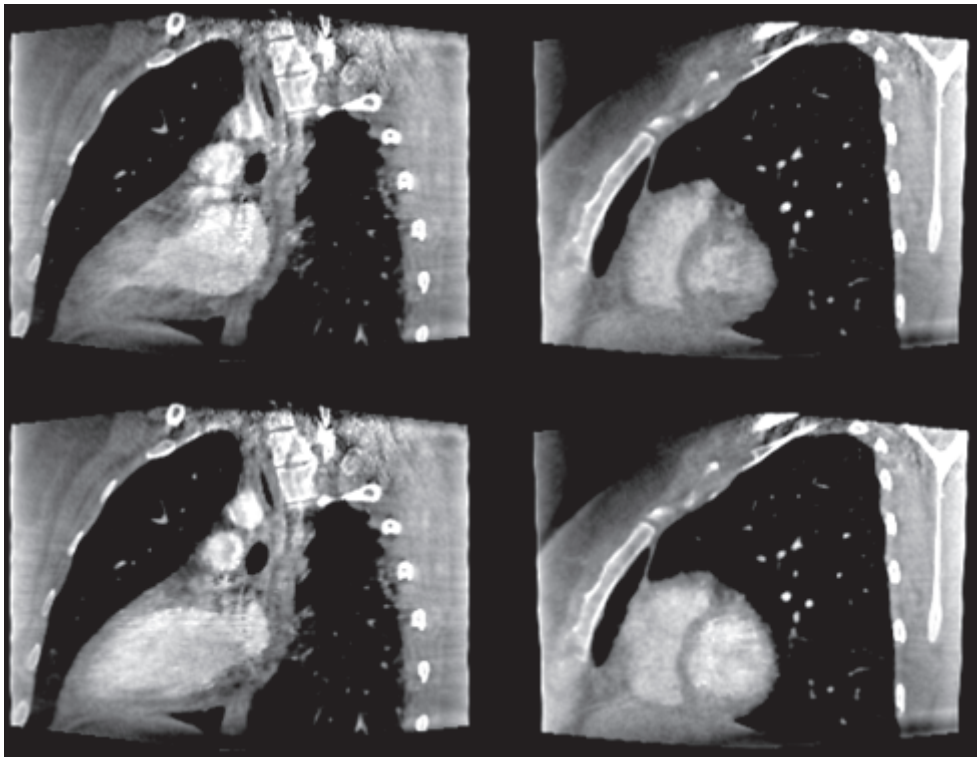

Figure 46. Reconstructions of patient DENVER2 by the 4D ROOSTER method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

In both cases, the border between blood and muscle is well visible (except in the short axis systole slice of patient DENVER1).

Chapter VIII : Conclusions on cardiac C-arm computed tomography methods

It could be argued that, with different parameters, it would have been possible to obtain better images with one of the presented methods. There are indeed many parameters in all methods, and I cannot guarantee that the results presented here are the best ones that can be obtained, only that they are the best ones that were obtained during this study, based on visual evaluation with the criteria listed in section I.2.c. The numerical measurements came later and confirmed the ranking between results obtained by visual evaluation.

The parameters used to obtain the results presented in this work are listed in Table 8. They have been chosen empirically, by running numerous tests and adjusting the parameters after each test. Unfortunately, the "optimal" parameters vary with the nature of the data, the size of the reconstructed volume, the sampling rate, … and I have not found a way to estimate how they vary.

	Shepp & Logan phantom	*Real data*
Gating window width (in % of cardiac cycle)	20	20
Volume size (in voxels)	256 * 256 * 256	284 * 216 * 284
ECG-gated SART		
Number of iterations	5	5
Relaxation parameter λ	0.5	0.5
Badea's method		
Threshold (multiplied by the maximum of the MTF)	0.15	0.15
ECG-gated IFDK		
Number of iterations	100	100
Steepest descent step α	0.001	0.001
ADMM 3D TV		
Main loop iterations	100	30
Nested conjugate gradient iterations	4	4
Regularization parameter α	0.5	0.1
ADMM parameter β	1000	1000
ADMM 3D Wavelets		
Main loop iterations	30	30
Nested conjugate gradient iterations	4	4
Regularization parameter α	0.1	2
ADMM parameter β	1000	10000
PICCS		
Main loop iterations	30	30
Nested ECG-gated SART iterations	4	4
SART relaxation parameter λ	0.5	0.5
Nested TV steepest descent iterations	20	20
TV steepest descent data attachment weight	500	500
TV steepest descent step size	0.01	0.002

PICCS balance TV / Prior parameter α	0.1	0.5
ADMM 4D TV		
Main loop iterations	30	30
Nested conjugate gradient iterations	10	4
Regularization parameter α	0.1	0.1
ADMM parameter β	1000	1000
ω outside the heart ROI	2	2
ω inside the heart ROI	10	10
Number of time points	10	10
4D ROOSTER		
Main loop iterations	30	30
Nested conjugate gradient iterations	4	4
Nested spatial TV steepest descent iterations	5	5
Spatial TV steepest descent data attachment weight	100	10000
Spatial TV steepest descent step size	0.001	0.002
Nested temporal TV steepest descent iterations	5	5
Temporal TV steepest descent data attachment weight	100	1000
Temporal TV steepest descent step size	0.001	0.002
Number of time points	10	10

Table 8. Parameters used with each method to obtain the results presented in this document

VIII.1. Quantitative evaluation of real data reconstructions

As a quantitative evaluation of the reconstructions of clinical datasets, the Contrast to Noise Ratio (CNR) on the long axis end-systolic slices has been measured, using one ROI in the blood and one in the muscle.

CNR gives an indication of how distinguishable from one another two adjacent regions are. It was computed as

$$CNR = \frac{|M_{muscle} - M_{blood}|}{\sigma_{muscle}}$$

where M_{blood} and M_{muscle} are the mean attenuations in the blood and muscle ROIs respectively, and σ_{muscle} is the standard deviation in the muscle ROI. The ROIs are shown in Figure 47 and the results listed in Table 9.

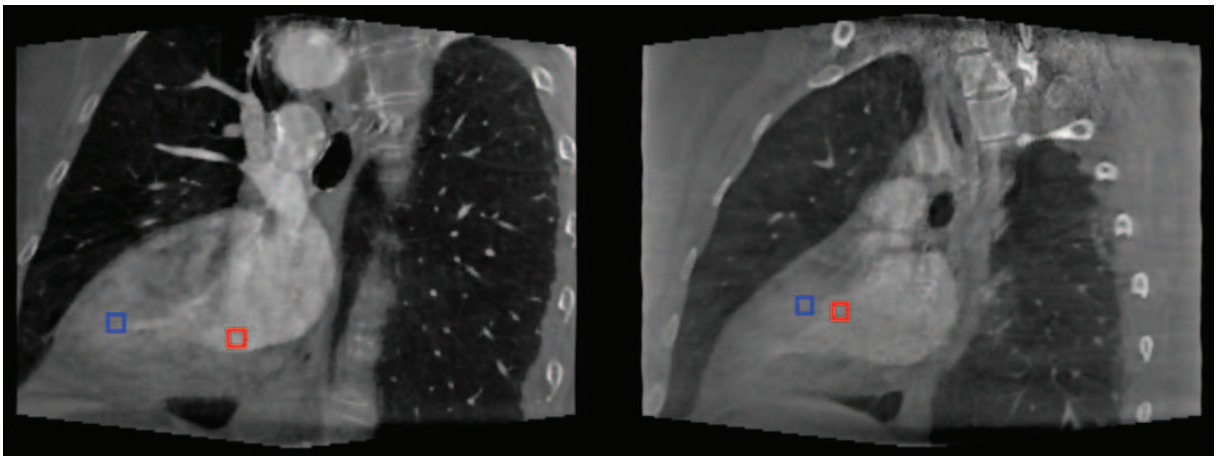

Figure 47. Regions of interest used for CNR computation. The red rectangle is the blood ROI, the blue one the muscle ROI

	Contrast to Noise Ratio	
	Patient DENVER1	Patient DENVER2
SART	3.4536	2.0273
Badea	2.2294	0.4332
Iterative FDK	4.0144	0.9303
ADMM 3D TV	5.0809	1.8846
ADMM 3D Wavelets	4.5681	5.5014
PICCS	5.4482	3.3044
ADMM 4D TV	4.8053	1.9992
4D ROOSTER	6.4691	4.4689

Table 9. CNR measured in the long axis end-systolic reconstruction results of all methods, on both patients

Like every image quality criterion, CNR should be considered with care when trying to rank methods, for it has biases:

- it is high in images consisting of homogeneous regions separated by sharp edges, which favors strongly-regularized reconstruction methods
- it depends on the ROIs chosen to measure it

Therefore, only large and consistent differences in CNR should be used to rank methods. With these precautions, a few comments on the CNR results can be made:

- Regularized methods provide a better CNR than deconvolution ones
- PICCS provides a better CNR than TV-regularized ADMM
- 4D ROOSTER provides a better CNR than PICCS

VIII.2. Comparison between methods

Figure 48 and Figure 49 show some reconstruction results with all the methods presented in this work. On these figures, however, it is hard to determine which method performs best. The motion is a crucial aspect of the image quality for cardiac imaging, and it cannot be rendered on paper. The problem is the same for all the results presented in this document.

The reader is advised to visit my personal page on the CREATIS website (http://www.creatis.insa-lyon.fr/~mory/), which contains animated gif sequences. Unlike static images, these animated

sequences allow to rank the reconstruction methods. They show that 4D ROOSTER outperforms the other methods on all three cases.

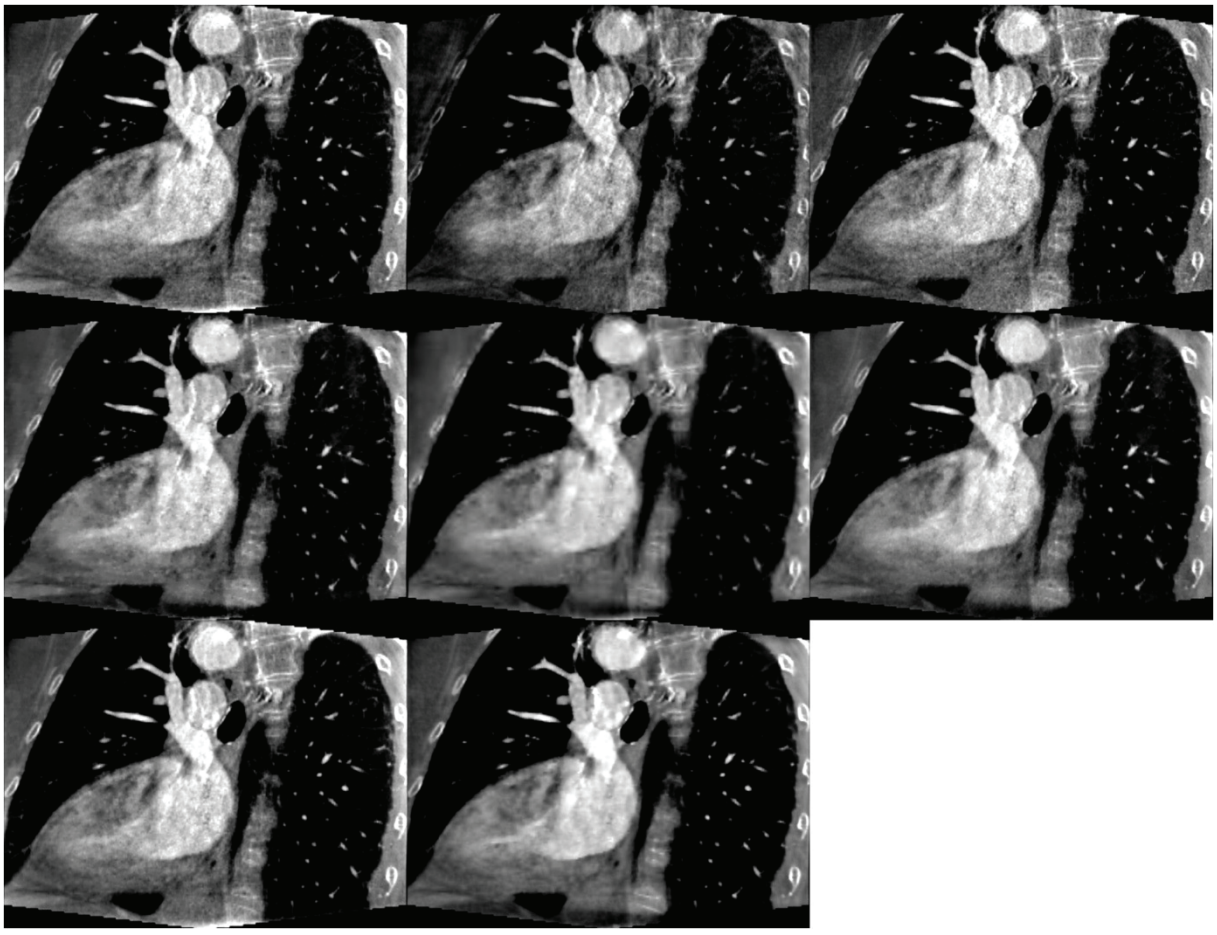

Figure 48. Long axis cuts of systole in patient DENVER1 with all methods presented in this paper. From left to right: on the top row, ECG-gated SART, Badea, ECG-gated IFDK; on the middle row, ADMM 3D TV, ADMM 3D Wavelets, PICCS; on the bottom row, ADMM 4D TV and 4D ROOSTER

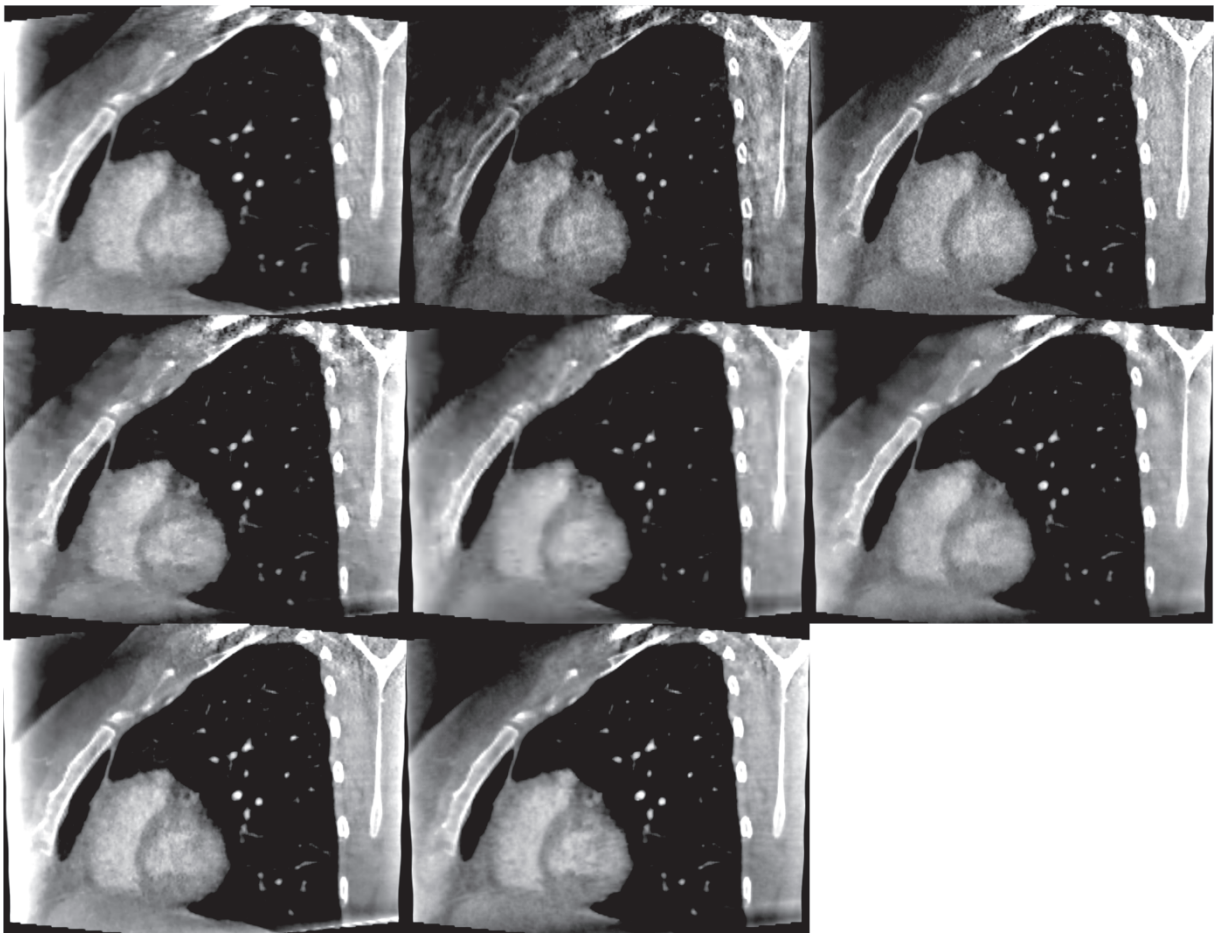

Figure 49. Short axis cuts of systole in patient DENVER2 with all methods presented in this paper. From left to right: on the top row, ECG-gated SART, Badea, ECG-gated IFDK; on the middle row, ADMM 3D TV, ADMM 3D Wavelets, PICCS; on the bottom row, ADMM 4D TV and 4D ROOSTER

Chapter IX : Clinical applications

The clinical applications of cardiac C-arm CT can roughly be grouped into two classes:

- Those requiring precise information on the heart's movement, e.g. left ventricular ejection fraction measurement. In this case, the whole cardiac cycle must be reconstructed and a good contrast between blood and soft tissue is essential, but soft tissue textures and precise recovery of Hounsfield attenuation values are unnecessary
- Those requiring precise information on soft tissue, e.g. late enhancement imaging for myocardial infarction detection. Soft tissue analysis requires only the reconstruction of the end-diastolic phase, which must be as precise as possible

This section describes what has been accomplished for both kinds of applications during this PhD thesis.

IX.1. Reconstructing the heart's movement from an injected scan

The protocol used to acquire the injected scans presented in this thesis, described in (Schwartz *et al* 2011), is well suited to the reconstruction of the whole cardiac cycle: it gathers the same quantity of data on each cardiac phase, thus allowing to reconstruct them all at once to visualize the heart's beating motion. On the other hand, it provides only little data on a given phase, making it impossible to precisely reconstruct soft tissues in end-diastole.

IX.2. Late enhancement

IX.2.a. Rationale

When the infarcted region is small, or when little contrast agent was injected during the intervention, the artifacts and the blur caused by heartbeat motion make it very hard to distinguish between healthy and infarcted myocardium.

The late enhancement CT scanner data available in Lyon (Boussel *et al* 2008) allowed the doctors to conclude that very little information could be obtained from ungated reconstructions of late enhancement scans, and that ECG-gating was required. Therefore I started investigating the field of ECG-gated late enhancement cardiac C-arm CT.

IX.2.b. Data and first results

A few projections extracted from a late enhancement scan performed on a patient in Lyon, at the cardiology hospital, are displayed on Figure 50, and a coronal slice of the FDK reconstruction obtained from this dataset is displayed in Figure 51.

An important problem is that the X-ray beam used for this acquisition is almost completely absorbed when passing through the liver. This causes high noise in the projections (see Figure 50), because only a few photons reach the detector, and hyper attenuation artifacts in the reconstruction (see Figure 51). As shown in Figure 51, it can be hard to distinguish the infarcted myocardium from these artifacts in late enhancement C-arm CT. This problem can be mitigated if the patients manage to

inhale deeply and hold their breath at end-inhale, but the conditions in which the acquisition is performed (emergency, usually at night, only a few minutes after the percutaneous coronary intervention) make it difficult to train the patients to do it correctly.

Patient LYON1 had a late enhancement MRI (currently the gold standard for infarction imaging) within a few days after the C-arm scan, which confirmed the location and size of the infarction. Therefore, in this case, we are certain that the patient indeed had a myocardial infarction, and that the hyper attenuated region pointed by the red arrows in Figure 51 is not caused by artifacts.

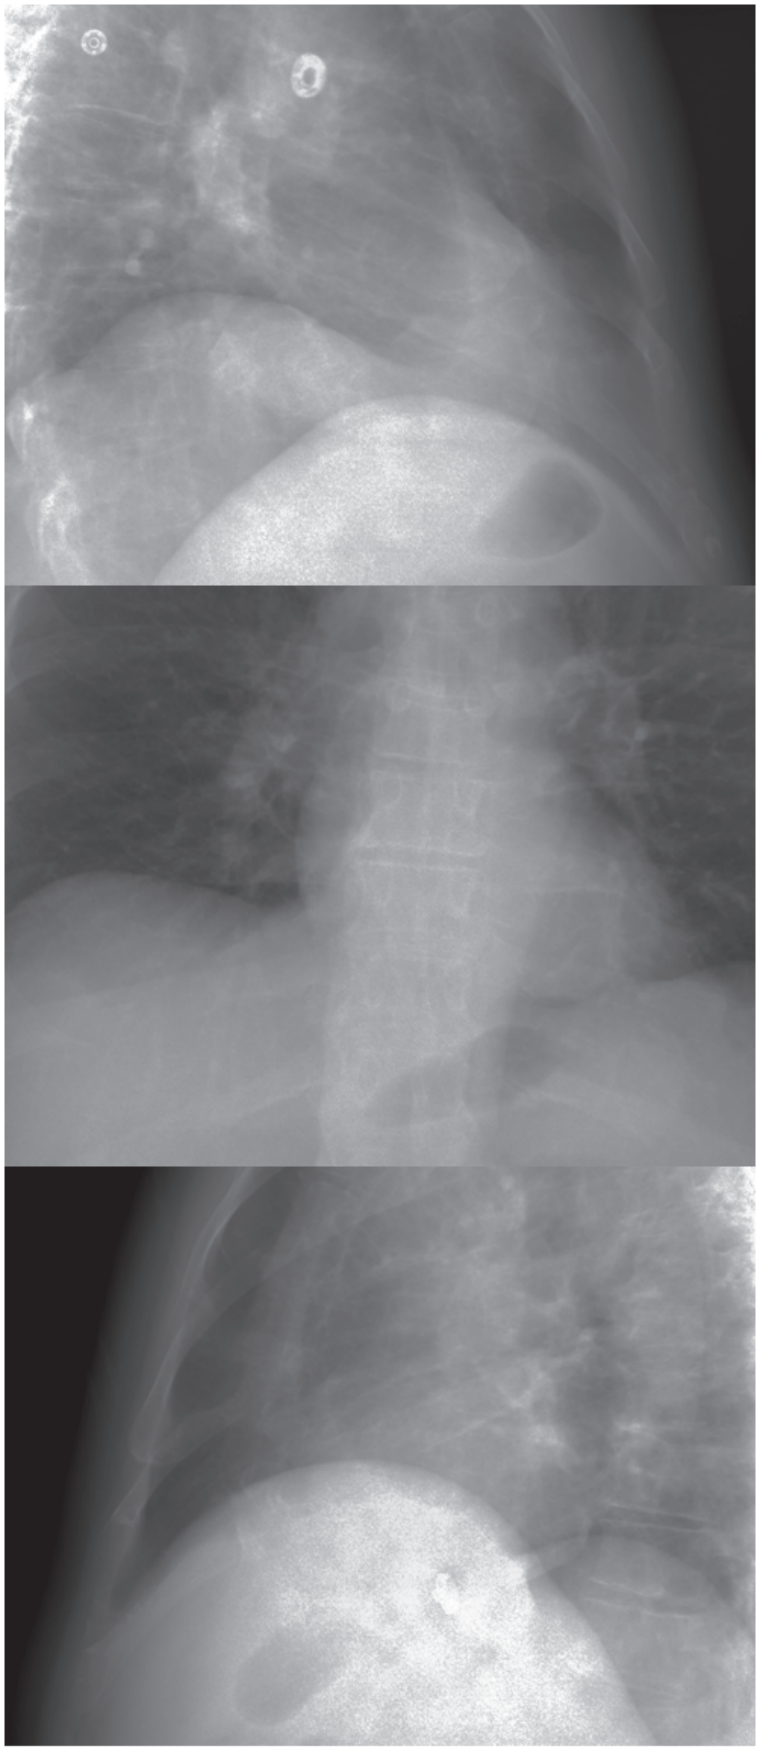

Figure 50. Projections 50 (top row), 150 (middle row) and 250 (bottom row) extracted from the late enhancement scan performed on patient LYON1

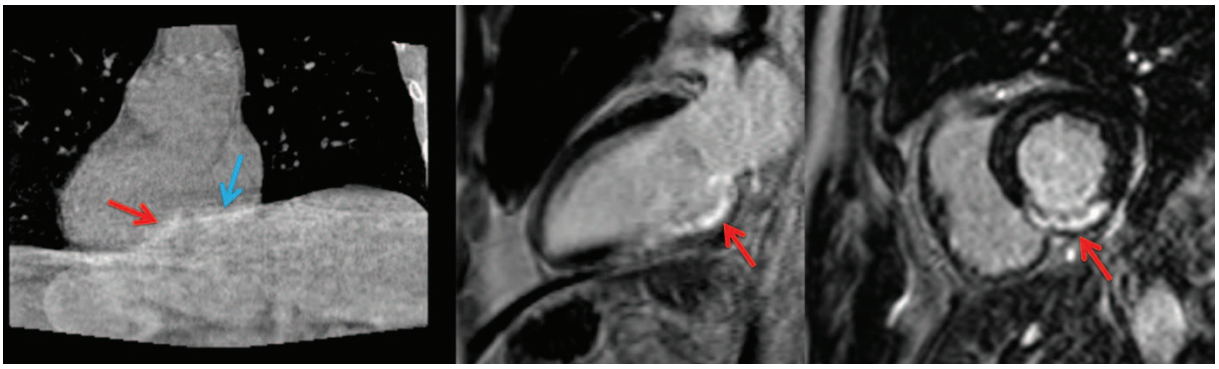

Figure 51. Heart of the patient LYON1, in late enhancement cardiac C-arm CT (on the left) and late enhancement MRI (in the middle with a long axis cut, and on the right with a short axis cut). The red arrows points to the infarcted region, while the blue arrow shows hyper attenuation streak artifacts caused by the liver.

As shown in Figure 52, even with the 4D ROOSTER method, ECG-gated reconstruction of late enhancement data is still far from being of clinically acceptable quality.

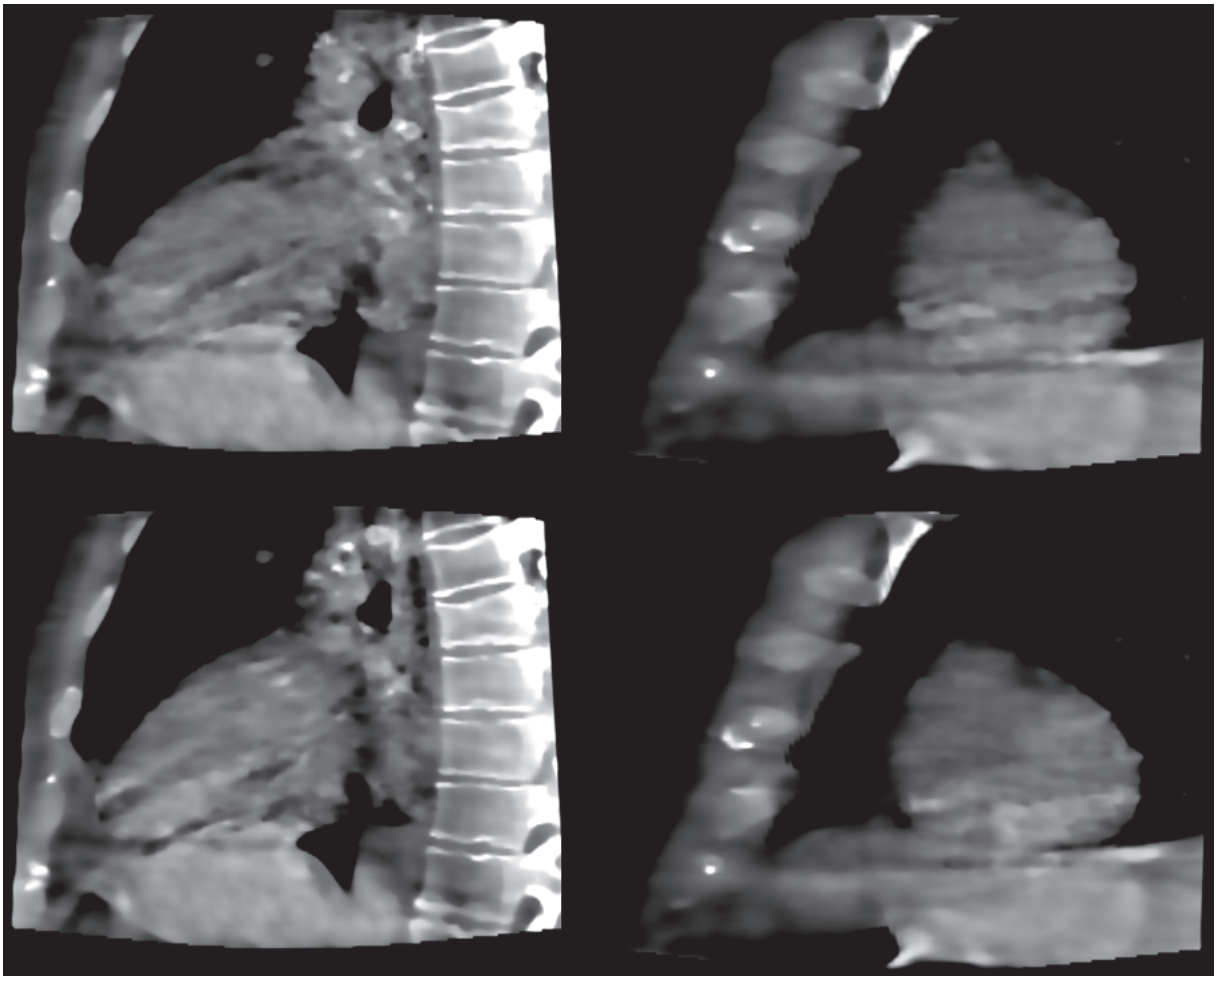

Figure 52. Reconstructions of patient LYON1 by the 4D ROOSTER method, in both systole (top row) and diastole (bottom row). The slices show the left ventricle, in long-axis cut in the left column and in short axis cut in the right column

IX.2.c. Perspectives

Soft tissue imaging requires as much information as possible on a single cardiac phase, usually end-diastole. An ECG-triggered multi sweep acquisition like the one described in (Lauritsch *et al* 2006) seems the only way, using currently available C-arms, to obtain a decent angular distribution of the projections to reconstruct a single phase. As multi sweep acquisitions are designed to acquire as much information as possible on a single phase and as little as possible on the other ones (to avoid unnecessary X-ray dose), 4D ROOSTER would not be a suitable reconstruction method for this kind of data. PICCS, or motion compensation-based methods, should be used instead.

Multi-sweep acquisitions are currently rarely used because they imply a long breath hold and the injection of large quantities of contrast medium. In the case of late enhancement imaging, only the breath hold duration would remain an issue, since no injection is performed during the acquisition.

Chapter X : Perspectives

X.1. Clinical use of 4D ROOSTER

4D ROOSTER has only been tested on a few cases. The method could be tested on a larger number of patients for clinical validation.

X.2. Improvements on the method

X.2.a. Other sparsifying transforms

The 4D ROOSTER method can easily be modified in order to use other regularization methods. A straightforward change would be to replace spatial TV minimization by some wavelets regularization, using Daubechies wavelets like in the method presented in section 0 or other kinds of wavelets or curvelets (Starck *et al* 2001). Another lead that could be investigated is the opportunity to perform non-local means regularization, as used in (Jia *et al* 2010).

X.2.b. Taking into account the flow of contrast

Instead of modeling the volume to reconstruct as a function of space and ECG-phase, we could investigate the possibility of segmenting the vessels and cavities into which the contrast flows, and modeling the attenuation of the object as the sum of two 4D functions: the first one, representing all the soft tissues and bones, would be periodic with the ECG, and the second one would be a-periodic but have a very limited support. Such an approach has already been proposed by Langet for angiographic C-arm CT (Langet *et al* 2012). However, is not certain that a cardiac C-arm CT acquisition contains enough data to perform this kind of reconstruction for the whole heart.

X.3. Application to other problems

Cardiac C-arm computed tomography is not the only application in which the 4D ROOSTER method, or a similar one, can be used. Thorax imaging during free breathing and standard cardiac CT from a CT scanner acquisition are both problems in which either the method itself or some of its basic principles could bring improvements. Replacing the forward projection operator with the Fourier transform operator, the method could also be tested on compressed sensing cardiac MRI.

Acknowledgment

I owe many thanks to Vincent Auvray, with whom I have spent a lot of time either on the phone or in meetings at Philips Medisys. Vincent has always provided reasonable advice on scientific matters and reliable support in the interaction with Philips people. This and his sense of humor contributed a lot to the serenity of this work.

Bo Zhang impressed me by his mastery of mathematical optimization, which he has been enthusiastic to share every time I asked. His help has been priceless.

Sherif Makram-Ebeid had it all figured from the beginning: he started investigating 4D reconstruction with regularization along time way before I even thought about it, and the idea of using a segmentation of the heart is also from him. Although he did not supervise me directly, he deserves his fair share of gratitude for the ideas I stole from him.

Michael Graß and Dirk Schäfer set me on tracks at the beginning of this work. They helped me discover the field of tomography, which was completely new for me. During these three years, they have been efficient, reliable and trustworthy. Convincing them that my crazy ideas were in fact reasonable was sometimes hard, but the challenge forced me to stay sharp.

Françoise Peyrin and Simon Rit provided relevant guidance during our frequent meetings in Grenoble. Simon has been very reactive whenever I have asked him for help, especially regarding RTK. I look forward to my post-doc, as I will be working with him and his team.

Philippe Douek and Loïc Boussel have made their best to supervise this work, encourage and challenge me. Their expertise in cardiac imaging and their support inside the hospital to make the experiments happen have been essential to the success of my PhD.

I'd also like to express my gratitude to Hélène Langet, who whispered me the initial idea of the mathematical considerations developed in section V.6.

I have to thank the radiologists and cardiologists of Lyon, and the whole staff of the cardiology hospital in Lyon, especially the X-ray technologists, for welcoming me in the interventional X-ray control room, enduring my numerous and pesky questions, and obviously for having acquired data for my experiments.

I am also grateful to Dr. Carroll and his team in Denver for providing the injected scans with which I have been working for three years.

Finally, I'd like to thank the patients who accepted to take part in the studies, both in Lyon and in Denver. Without them, I would have spent three years working with Shepp & Logan phantoms and scans of syringes filled with contrast medium, and I'm not sure I'd have made it to the end.

Bibliography

Afonso M V, Bioucas-Dias J M and Figueiredo M A T 2011 An Augmented Lagrangian Approach to the Constrained Optimization Formulation of Imaging Inverse Problems *IEEE Transactions on Image Processing* **20** 681 –695

Afonso M V, Bioucas-Dias J M and Figueiredo M A T 2010 Fast Image Recovery Using Variable Splitting and Constrained Optimization *IEEE Transactions on Image Processing* **19** 2345 –2356

Andersen A H and Kak A C 1984 Simultaneous algebraic reconstruction technique (SART): a superior implementation of the art algorithm *Ultrason Imaging* **6** 81–94

Babacan S D, Peng X, Wang X-P, Do M N and Liang Z-P 2011 Reference-guided sparsifying transform design for compressive sensing MRI *Conf Proc IEEE Eng Med Biol Soc* **2011** 5718–21

Badea C T, Johnston S M, Qi Y and Johnson G A 2011 4D micro-CT for cardiac and perfusion applications with view under sampling *Phys. Med. Biol.* **56** 3351–69

Bauschke H H and Combettes P L 2009 The Baillon-Haddad theorem revisited *arXiv preprint arXiv:0906.0807* Online: http://arxiv.org/abs/0906.0807

Bauschke H H, Martín-Márquez V, Moffat S M and Wang X 2012 Compositions and convex combinations of asymptotically regular firmly nonexpansive mappings are also asymptotically regular *Fixed Point Theory and Applications* **2012** 1–11

Bergner F, Berkus T, Oelhafen M, Kunz P, Pan T, Grimmer R, Ritschl L and Kachelrieß M 2010 An investigation of 4D cone-beam CT algorithms for slowly rotating scanners *Medical Physics* **37** 5044

Blondel C, Malandain G, Vaillant R and Ayache N 2006 Reconstruction of coronary arteries from a single rotational X-ray projection sequence *IEEE Trans Med Imaging* **25** 653–63

Blondel C, Vaillant R, Malandain G and Ayache N 2004 3D tomographic reconstruction of coronary arteries using a precomputed 4D motion field *Phys Med Biol* **49** 2197–208

Boussel L, Ribagnac M, Bonnefoy E, Staat P, Elicker B M, Revel D and Douek P 2008 Assessment of acute myocardial infarction using MDCT after percutaneous coronary intervention: comparison with MRI *AJR Am J Roentgenol* **191** 441–7

Bracewell R N 2003 *Fourier analysis and imaging*

Candes E J, Romberg J K and Tao T 2006a Stable signal recovery from incomplete and inaccurate measurements *Communications on pure and applied mathematics* **59** 1207–23

Candes E J, Romberg J and Tao T 2006b Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information *Information Theory, IEEE Transactions on* **52** 489 – 509

Candes E J and Wakin M B 2008 An Introduction To Compressive Sampling *IEEE Signal Processing Magazine* **25** 21–30

Chambolle A 2004 An Algorithm for Total Variation Minimization and Applications *J. Math. Imaging Vis.* **20** 89–97

Chan T F, Osher S and Shen J 2001 The digital TV filter and nonlinear denoising *Image Processing, IEEE Transactions on* **10** 231–41

Chen G-H, Tang J and Hsieh J 2009 Temporal resolution improvement using PICCS in MDCT cardiac imaging *Med Phys* **36** 2130–5

Chen G-H, Tang J and Leng S 2008 Prior Image Constrained Compressed Sensing (PICCS) *Proc Soc Photo Opt Instrum Eng* **6856** 685618

Chen G-H, Theriault-Lauzier P, Tang J, Nett B, Leng S, Zambelli J, Qi Z, Bevins N, Raval A, Reeder S and Rowley H 2012 Time-Resolved Interventional Cardiac C-arm Cone-Beam CT: An Application of the PICCS Algorithm *Medical Imaging, IEEE Transactions on* **31** 907 –923

Van Cittert P H 1931 Zum Einfluss der Spaltbreite auf die Intensitätsverteilung in Spektrallinien. II *Zeitschrift für Physik* **69** 298–308

Cormack A M 1963 Representation of a Function by Its Line Integrals, with Some Radiological Applications *Journal of Applied Physics* **34** 2722–7

Cormack A M 1964 Representation of a Function by Its Line Integrals, with Some Radiological Applications. II *Journal of Applied Physics* **35** 2908–13

Daubechies I, Defrise M and De Mol C 2004 An iterative thresholding algorithm for linear inverse problems with a sparsity constraint *Communications on Pure and Applied Mathematics* **57** 1413–57

Dhawan A P, Rangayyan R M and Gordon R 1985 Image restoration by Wiener deconvolution in limited-view computed tomography *Appl. Opt.* **24** 4013–20

Dong B, Li J and Shen Z 2013 X-Ray CT Image Reconstruction via Wavelet Frame Based Regularization and Radon Domain Inpainting *J Sci Comput* **54** 333–49

Feldkamp L A, Davis L C and Kress J W 1984 Practical cone-beam algorithm *J. Opt. Soc. Am. A* **1** 612–9

Figueiredo M, Bioucas-Dias J M and Afonso M V 2009 Fast frame-based image deconvolution using variable splitting and constrained optimization *Statistical Signal Processing, 2009. SSP'09. IEEE/SP 15th Workshop on* pp 109–12

Glöckler M, Halbfaβ J, Koch A, Achenbach S and Dittrich S 2013 Multimodality 3D-roadmap for cardiovascular interventions in congenital heart disease--a single-center, retrospective analysis of 78 cases *Catheter Cardiovasc Interv* **82** 436–42

Go A S, Mozaffarian D, Roger V L, Benjamin E J, Berry J D, Borden W B, Bravata D M, Dai S, Ford E S, Fox C S, Franco S, Fullerton H J, Gillespie C, Hailpern S M, Heit J A, Howard V J, Huffman M D, Kissela B M, Kittner S J, Lackland D T, Lichtman J H, Lisabeth L D, Magid D, Marcus G M, Marelli A, Matchar D B, McGuire D K, Mohler E R, Moy C S, Mussolino M E, Nichol G, Paynter N P, Schreiner P J, Sorlie P D, Stein J, Turan T N, Virani S S, Wong N D, Woo D and Turner M B 2013 Heart Disease and Stroke Statistics—2013 Update A Report From the American Heart Association *Circulation* **127** e6–e245

Gordon R, Bender R and Herman G T 1970 Algebraic Reconstruction Techniques (ART) for three-dimensional electron microscopy and X-ray photography *Journal of Theoretical Biology* **29** 471–81

Gordon R and Rangayyan R M 1983 Geometric Deconvolution: A Meta-Algorithm for Limited View Computed Tomography *IEEE Transactions on Biomedical Engineering* **BME-30** 806–10

Gottlieb D, Gustafsson B and Forssen P 2000 On the direct Fourier method for computer tomography *IEEE Trans. Med. Imaging* **19** 223–32

Grangeat P 2009 *Tomography* (London [u.a.: ISTE [u.a.])

Grass M, Koppe R, Klotz E, Proksa R, Kuhn M H, Aerts H, Op de Beek J and Kemkers R 1999 Three-dimensional reconstruction of high contrast objects using C-arm image intensifier projection data *Comput Med Imaging Graph* **23** 311–21

Gullberg G T 1979 The reconstruction of fan-beam data by filtering the back-projection *Computer Graphics and Image Processing* **10** 30–47

Hansis E, Carroll J D, Schäfer D, Dössel O and Grass M 2010 High-quality 3-D coronary artery imaging on an interventional C-arm x-ray system *Med Phys* **37** 1601–9

Hansis E, Schafer D, Dossel O and Grass M 2008 Evaluation of Iterative Sparse Object Reconstruction From Few Projections for 3-D Rotational Coronary Angiography *IEEE Transactions on Medical Imaging* **27** 1548–55

Hao J, Zhang L, Li L and Kang K 2012 An improved non-local means regularized iterative reconstruction method for low-dose dental CBCT *2012 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)* 2012 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC) pp 3422–5

Henry Stark J W W 1981 An Investigation of Computerized Tomography by Direct Fourier Inversion and Optimum Interpolation *Biomedical Engineering, IEEE Transactions on* 496 – 505

Hsieh J, Londt J, Vass M, Li J, Tang X and Okerlund D 2006 Step-and-shoot data acquisition and reconstruction for cardiac x-ray computed tomography *Med Phys* **33** 4236–48

Jia X, Dong B, Lou Y and Jiang S B 2011 GPU-based iterative cone-beam CT reconstruction using tight frame regularization *Physics in medicine and biology* **56** 3787

Jia X, Lou Y, Dong B, Tian Z and Jiang S 2010 4D Computed Tomography Reconstruction from Few-Projection Data via Temporal Non-local Regularization *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2010* Lecture Notes in Computer Science ed T Jiang, N Navab, J P W Pluim and M A Viergever (Springer Berlin Heidelberg) pp 143–50 Online: http://link.springer.com/chapter/10.1007/978-3-642-15705-9_18

Kachelriess M and Kalender W A 1998 Electrocardiogram-correlated image reconstruction from subsecond spiral computed tomography scans of the heart *Medical physics* **25** 2417

Kaczmarz S 1937 Angenäherte Auflösung von Systemen linearer Gleichungen *Bulletin International de l'Académie Polonaise des Sciences et des Lettres* **35** 355–7

Kak A and Slaney M 1988 *Principles of Computerized Tomographic Imaging* (IEEE Press)

Lalush D S and Tsui B M W 1994 Improving the convergence of iterative filtered backprojection algorithms *Medical Physics* **21** 1283–6

Langet H, Riddell C, Trousset Y, Tenenhaus A, Lahalle E, Fleury G and Paragios N 2011 Compressed sensing based 3D tomographic reconstruction for rotational angiography *Med Image Comput Comput Assist Interv* **14** 97–104

Langet H, Riddell C, Trousset Y, Tenenhaus A, Lahalle E, Fleury G and Paragios N 2012 Compressed sensing dynamic reconstruction in rotational angiography *Med Image Comput Comput Assist Interv* **15** 223–30

Lauritsch, Boese J, Wigström L, Kemeth H and Fahrig R 2006 Towards cardiac C-arm computed tomography *IEEE Trans Med Imaging* **25** 922–34

Lauzier P T, Tang J and Chen G-H 2012a Prior image constrained compressed sensing: Implementation and performance evaluation *Medical Physics* **39** 66

Lauzier P T, Tang J and Chen G-H 2012b Time-resolved cardiac interventional cone-beam CT reconstruction from fully truncated projections using the prior image constrained compressed sensing (PICCS) algorithm *Phys Med Biol* **57** 2461–76

Li M, Yang H and Kudo H 2002 An accurate iterative reconstruction algorithm for sparse objects: application to 3D blood vessel reconstruction from a limited number of projections *Phys Med Biol* **47** 2599–609

Mc Kinnon G C and Bates R H 1981 Towards imaging the beating heart usefully with a conventional CT scanner *IEEE Trans Biomed Eng* **28** 123–7

Medoff B P, Brody W R, Nassi M and Macovski A 1983 Iterative convolution backprojection algorithms for image reconstruction from limited data *J. Opt. Soc. Am.* **73** 1493–500

Moret J, Kemkers R, Op de Beek J, Koppe R, Klotz E and Grass M 1998 3D rotational angiography: Clinical value in endovascular treatment *medicamundi* **42** 8–14

Mory B, Somphone O, Prevost R and Ardon R 2012 Real-Time 3d image segmentation by user-constrained template deformation *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2012* MICCAI 2012 (Springer) pp 561–8 Online: http://link.springer.com/chapter/10.1007/978-3-642-33415-3_69

Müller K, Rohkohl C, Lauritsch G, Schwemmer C, Heidbüchel H, De Buck S, Nuyens D, Kiriakou Y, Köhler C and Hornegger J 2012 4-D Motion Field Estimation by Combined Multiple Heart Phase Registration (CMHPR) for Cardiac C-arm Data IEEE MIC

Müller K, Schwemmer C, Hornegger J, Zheng Y, Wang Y, Lauritsch G, Rohkohl C, Maier A K, Schultz C and Fahrig R 2013 Evaluation of interpolation methods for surface-based motion compensated tomographic reconstruction for cardiac angiographic C-arm data *Med Phys* **40** 031107

Nassi M, Brody W R, Medoff B P and Macovski A 1982 Iterative reconstruction-reprojection: an algorithm for limited data cardiac-computed tomography *IEEE Trans Biomed Eng* **29** 333–41

Nett B, Tang J, Leng S and Chen G-H 2008 Tomosynthesis via Total Variation Minimization Reconstruction and Prior Image Constrained Compressed Sensing (PICCS) on a C-arm System *Proc Soc Photo Opt Instrum Eng* **6913** nihpa92672

Pan X, Sidky E Y and Vannier M 2009 Why do commercial CT scanners still employ traditional, filtered back-projection for image reconstruction? *Inverse Problems* **25** 123009

Prümmer M 2009 *Cardiac C-Arm Computed Tomography: Motion Estimation and Dynamic Reconstruction* (Erlangen) Online: http://www.opus.ub.uni-erlangen.de/opus/volltexte/2010/1931/pdf/MarcusPruemmerDissertation.pdf

Prümmer M, Hornegger J, Lauritsch G, Wigström L, Girard-Hughes E and Fahrig R 2009 Cardiac C-arm CT: a unified framework for motion estimation and dynamic CT *IEEE Trans Med Imaging* **28** 1836–49

Radon J 1917 Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten *Akad. Wiss.* **69** 262–77

Rohkohl C, Lauritsch G, Biller L, Prümmer M, Boese J and Hornegger J 2010 Interventional 4D motion estimation and reconstruction of cardiac vasculature without motion periodicity assumption *Medical Image Analysis* **14** 687–94

Saint-Félix D, Trousset Y, Picard C, Ponchut C, Roméas R and Rougée A 1994 In vivo evaluation of a new system for 3D computerized angiography *Phys Med Biol* **39** 583–95

Schwartz J G, Neubauer A M, Fagan T E, Noordhoek N J, Grass M and Carroll J D 2011 Potential role of three-dimensional rotational angiography and C-arm CT for valvular repair and implantation *Int J Cardiovasc Imaging* **27** 1205–22

Shepp L A and Logan B F 1974 The Fourier reconstruction of a head section *IEEE Transactions on Nuclear Science* **21** 21–43

Sidky E Y and Pan X 2008 Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization *Physics in Medicine and Biology* **53** 4777–807

Starck J-L, Donoho D L and Candes E J 2001 Very high quality image restoration by combining wavelets and curvelets vol 4478 pp 9–19 Online: http://dx.doi.org/10.1117/12.449693

Sunnegaardh J 2009 *Iterative Filtered Backprojection Methods for Helical Cone-Beam CT* (Linköping University)

Sunnegardh J and Danielsson P-E 2008 Regularized iterative weighted filtered backprojection for helical cone-beam CT *Medical Physics* **35** 4173–85

Tibshirani R 1994 Regression Shrinkage and Selection Via the Lasso *Journal of the Royal Statistical Society, Series B* **58** 267–88

Turbell, H. 2001 *Cone-Beam Reconstruction Using Filtered Backprojection* (Linköping University) Online: http://people.csail.mit.edu/bkph/courses/papers/Exact_Conebeam/Turbell_Thesis_FBP_2001.pdf

Tuy H K 1983 An Inversion Formula for Cone-Beam Reconstruction *SIAM Journal on Applied Mathematics* **43** 546–52

Valton S 2007 *Reconstruction tomographique 3D en géométrie conique à trajectoire circulaire pour des prototypes d'imageur bimodal pour le petit animal* Online: http://www.theses.fr/127669795

Wallace M J, Kuo M D, Glaiberman C, Binkert C A, Orth R C and Soulez G 2008 Three-Dimensional C-arm Cone-beam CT: Applications in the Interventional Suite *Journal of Vascular and Interventional Radiology* **19** 799–813

Zhang B and Zeng G L 2007 Two-dimensional iterative region-of-interest (ROI) reconstruction from truncated projection data *Medical Physics* **34** 935–44

Appendix

The following is the proof that averaging along time is a firmly non-expansive mapping.

Let us consider a mapping

$$A : \mathbb{R}^{MN} \longrightarrow \mathbb{R}^{MN}$$

$$x = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_M \\ x_{M+1} \\ \vdots \\ x_{NM} \end{pmatrix} \longrightarrow \frac{1}{N} \begin{pmatrix} x_1 + x_{1+M} + \cdots + x_{1+(N-1)M} \\ x_2 + x_{2+M} + \cdots + x_{2+(N-1)M} \\ \vdots \\ x_M + x_{2M} + \cdots + x_{NM} \\ x_1 + x_{1+M} + \cdots + x_{1+(N-1)M} \\ x_2 + x_{2+M} + \cdots + x_{2+(N-1)M} \\ \vdots \\ x_M + x_{2M} + \cdots + x_{NM} \\ \vdots \\ \vdots \\ x_1 + x_{1+M} + \cdots + x_{1+(N-1)M} \\ x_2 + x_{2+M} + \cdots + x_{2+(N-1)M} \\ \vdots \\ x_M + x_{2M} + \cdots + x_{NM} \end{pmatrix} = \frac{1}{N} \begin{pmatrix} \sum_{i=0}^{N-1} x_{1+iM} \\ \sum_{i=0}^{N-1} x_{2+iM} \\ \vdots \\ \sum_{i=0}^{N-1} x_{M+iM} \\ \sum_{i=0}^{N-1} x_{1+iM} \\ \vdots \\ \sum_{i=0}^{N-1} x_{M+iM} \\ \vdots \\ \vdots \\ \sum_{i=0}^{N-1} x_{1+iM} \\ \vdots \\ \sum_{i=0}^{N-1} x_{M+iM} \end{pmatrix}$$

If x is a 3D + time series of N volumes, then the mapping A performs averaging along time.

A is firmly non-expansive if $\forall x, y \in \mathbb{R}^{MN}, \|Ax - Ay\|_2^2 \leq \langle Ax - Ay, x - y \rangle$. Let us compute both quantities:

$$\|Ax - Ay\|_2^2 = \frac{1}{N} \sum_{j=1}^{M} \left[\sum_{i=0}^{N-1} \left(x_{j+iM} - y_{j+iM} \right) \right]^2$$

$$\langle Ax - Ay, x - y \rangle = \frac{1}{N} \begin{pmatrix} \sum_{i=0}^{N-1} x_{1+iM} \\ \sum_{i=0}^{N-1} x_{2+iM} \\ \vdots \\ \sum_{i=0}^{N-1} x_{M+iM} \\ \sum_{i=0}^{N-1} x_{1+iM} \\ \vdots \\ \sum_{i=0}^{N-1} x_{M+iM} \\ \vdots \\ \vdots \\ \sum_{i=0}^{N-1} x_{1+iM} \\ \vdots \\ \sum_{i=0}^{N-1} x_{M+iM} \end{pmatrix} \cdot \begin{pmatrix} x_1 - y_1 \\ x_2 - y_2 \\ \vdots \\ x_M - y_M \\ x_{M+1} - y_{M+1} \\ \vdots \\ x_{NM} - y_{NM} \end{pmatrix}$$

$$= \frac{1}{N} \sum_{j=1}^{M} \left[\left[\sum_{i=0}^{N-1} (x_{j+iM} - y_{j+iM}) \right] (x_j - y_j + x_{j+M} - y_{j+M} + \cdots + x_{j+(N-1)M} - y_{j+(N-1)M}) \right]$$

$$= \frac{1}{N} \sum_{j=1}^{M} \left[\left[\sum_{i=0}^{N-1} (x_{j+iM} - y_{j+iM}) \right] \left[\sum_{i=0}^{N-1} (x_{j+iM} - y_{j+iM}) \right] \right]$$

$$= \frac{1}{N} \sum_{j=1}^{M} \left[\sum_{i=0}^{N-1} (x_{j+iM} - y_{j+iM}) \right]^2 = \|Ax - Ay\|_2^2$$

This proves that A is firmly non-expansive.