

On the Use of Temporal Information for the Reconstruction of Magnetic Resonance Image Series

Dissertation
for the award of the degree
"Doctor rerum naturalium"
of the Georg-August-Universität Göttingen

within the doctoral program
Physics of Biological and Complex Systems
of the Georg-August University School of Science (GAUSS)

submitted by
Jakob Klosowski

Göttingen, 2019

Thesis Committee

First referee and supervisor:

Prof. Dr. Jens Frahm
Biomedizinische NMR
Max-Planck-Institut für biophysikalische Chemie, Göttingen

Second referee:

Prof. Dr. Jörg Enderlein
Drittes Physikalisches Institut
Georg-August-Universität Göttingen

Prof. Dr. Helmut Grubmüller
Abteilung Theoretische und Computergestützte Biophysik
Max-Planck-Institut für biophysikalische Chemie, Göttingen

Examination Board

Prof. Dr. Jens Frahm

Prof. Dr. Jörg Enderlein

Prof. Dr. Helmut Grubmüller

Prof. Dr. Marina Bennati
Forschungsgruppe Elektronenspinresonanz-Spektroskopie
Max-Planck-Institut für biophysikalische Chemie, Göttingen

Prof. Dr. Thorsten Hohage
Arbeitsgruppe Inverse Probleme
Institut für Numerische und Angewandte Mathematik,
Georg-August-Universität Göttingen

Prof. Dr. Ulrich Parlitz
Forschungsgruppe Biomedizinische Physik
Max-Planck-Institut für Dynamik und Selbstorganisation, Göttingen

Date of examination: 26 February 2020

Abstract

Standard clinical magnetic resonance imaging uses data sampled at the Nyquist rate which limits it to frame rates of 1-4 frames per second. Over the past decade different methods for the reconstruction of images from sub-Nyquist sampled data have been proposed, raising frame rates beyond 30 frames per second and opening the door to real-time cardiac imaging, the study of speech and many more. Practically all such methods exploit spatio-temporal structure in the data as prior knowledge to an inverse problem.

This work combines the method of non-linear inversion reconstruction and motion estimation by an optical flow, a model for in-plane motion that is frequently used in computer vision tasks. The model is quantitatively compared against simpler methods based on temporal finite difference schemes in simulations and real cardiac data.

Adding an optical flow constraint to the reconstruction's cost functional is shown to be advantageous in situations with purely in-plane motion while failure of the model is observed in the form of artifacts and blur during through-plane motion. A point that got little attention in related literature. The model's limits are examined in detail and pointers to possible extensions to overcome them are given. Further, a simple centered finite temporal difference constraint for a batch of frames is shown to provide better image quality than a previously established scheme that considers only the backward temporal difference at a single frame.

As in previous works, non-linear inversion reconstruction considers at most a small batch of frames for every time point. The method does not require availability of the entire image data series at any point during computation. It is thus conceptually capable of providing images in real-time simultaneously with the data acquisition, a prerequisite for interactive MRI.

Contents

1	Introduction	1
2	Parallel Magnetic Resonance Imaging	3
2.1	Basic Data Acquisition and -Description	3
2.2	Data Undersampling	6
2.3	Image Reconstruction	9
2.3.1	Linear Problem	9
2.3.2	Non-Linear Problem	11
3	Real-time MRI	15
3.1	Dynamic Sampling	16
3.2	Temporal Regularization	18
3.2.1	Time Difference Regularization	18
3.2.2	Optical Flow	19
4	Materials and Methods	25
4.1	Setup and Data Acquisition	25
4.2	Test Data	26
4.2.1	Numerical Phantoms	26
4.2.2	Surrogate Data	27
4.2.3	Image Quality Assessment	29
4.3	Reconstruction Algorithm	30
4.3.1	Framework and Parameters	30
4.3.2	Numerical Functional Minimization	32
4.3.3	Operator Discretization	34
5	Results	37
5.1	Image Reconstruction with Known Motion	37
5.1.1	Optical Flow Compatible Data	38
5.1.2	Optical Flow Model Violation	44
5.1.3	Applicability to Real Data	47
5.2	Motion Estimation in the Presence of Artifacts	52

6 Discussion	57
7 Summary	63
A Appendix	65
A.1 Prox Operators	65
A.2 Parameter Tables	66
Abbreviations	69
Bibliography	71
Acknowledgements	77

Introduction

Magnetic resonance imaging (MRI) offers non-invasive patient scanning with a large field of view (FOV), deep sample penetration and without damaging radiation. A rich set of contrasts allows for highlighting different soft tissue types depending on their fat and water content as well as visualizing metabolic activity. On the other hand, MRI suffers from low frame rates on the order of one frame per second (fps), depending on spatial resolution, which leads to long imaging times. Patients may be required to lie still for several minutes and sometimes hold their breath for an extended interval, lest motion artifacts render the images useless. Low patient throughput and associated higher costs per exam have spurred research in accelerating MRI. As MRI hardware reached electronics- and physiology related optimization limits data undersampling below the Nyquist rate in conjunction with advanced image reconstruction methods emerged as a route for further acceleration [1].

In recent years, MRI has been demonstrated at frame rates upward of 30 fps [2], allowing imaging of the heart in real-time. Established methods thus far relied on ECG synchronized data sorting after acquisition (gating) to form synthetic heart beats (cine imaging [3]), whereby aperiodic data is discarded which in turn precludes examinations of heart arrhythmia. Use of ECG or finger pulse triggers for data sorting also lengthens and complicates MRI exams and is not always accurate [4]. Methods for self-gating using only the acquired signal are still an active field of research.

The leading idea in real-time MRI (rt-MRI) is that data of anatomic motion exhibits strong redundancy along the time dimension. That is, neighboring frames differ little from each other. Consequently, updating information on an imaged object's motion state should require only a subset of the data that constitutes a whole frame. The situation bears resemblance to data compression and decoding in online video streaming. Unlike streaming, however, MRI is hardware constrained to encode 2D data as a set of lines in an incompletely filled 2D Fourier-space instead of using efficient, layered spatio-temporal sparsity transforms [5]. Another challenge is that artifacts from flawed decoding must not mislead diagnosis.

Image reconstruction methods that exploit temporal continuity can be divided into those processing the entire series in one large computation after acquisition was finished and those that start immediately and simultaneously with the acquisition and use only a limited amount of the most recent data to complement the current frame's incomplete data. The former method category is referred to as offline-reconstructions. It greatly benefits from the fact that long signal trains have a richer temporal structure than short ones. This work deals with the latter category of online-reconstructions which in principle enables interactive MRI, provided image reconstruction with little delay to the acquisition is feasible.

Another key point of this work is to combine the use of the temporal continuity property with parallel imaging, which refers to exploitation of data redundancy between multiple parallel receive coil channels. To that end, non-linear inversion reconstruction (NLINV) [6] is complemented with different motion priors. A focus here is to compare the quality of reconstructions from undersampled data with regard to image artifacts and temporal fidelity in a quantitative fashion against fully sampled reference images in simulated and human data. The applicability of an optical flow prior - a popular concept in computer vision [7] - to connect neighboring frames is of special interest. While a first effort to combine NLINV with optical flow in the context of cardiac imaging has previously been reported [8], an assessment of robustness of this motion prior to model violations is still missing and potential improvements remained unexplored.

Another focus of the present work is to compare the optical flow prior to simpler temporal finite difference schemes, an established and a proposed one. The effort to improve and confirm temporal fidelity is motivated by the fact that temporal and spatial resolution are inherently traded against each other in MRI.

The structure of this thesis is as follows: Chapter 2 revisits the signal generation mechanism in MRI - without detailing on contrast generation - and why it is slow. It then goes on to parallel imaging and how different undersampling schemes influence image quality before reviewing iterative image reconstruction from undersampled data. Chapter 3 covers the integration of motion priors, notably optical flow, in the cost functional of the reconstruction problem, while the remaining chapters cover methods, results and a discussion wrapped up at the end by a summary.

Parallel Magnetic Resonance Imaging

2.1 Basic Data Acquisition and -Description

Signal and Image Formation

MRI extends the technique of nuclear magnetic resonance spectroscopy (NMR) with a 3D spatial encoding mechanism. The principle of NMR and its extension to MRI is extensively covered in Brown et al. [9]. Following is a short summary of the points most relevant to this work. In brief, NMR probes the spin magnetic dipole moment - or simply, spin - of protons in atomic nuclei with a series of radio frequency pulses (rf-pulses). By virtue of the Pauli-exclusion-principle this quantity is non-zero in chemical elements with an odd proton count.

An ensemble of spins subjected to an external magnetic field will align with it in two energy states - parallelly (ground state) and anti-parallelly (excited state) to the field lines. Each spin is in a quantum superposition of both. The two Eigenstates are separated by an energy gap in the radio-frequency regime calculated as

$$\Delta E = \hbar\omega_L = \hbar\gamma B_0, \quad (2.1)$$

where γ is the element-dependent gyromagnetic ratio, B_0 the external magnetic field's strength and ω_L - the *Larmor* frequency - is the resonance frequency of the spins associated with the energy gap. An incident ω_L -pulse will excite spins from the ground state and create phase-coherence amongst all ensemble spins. The phase coherent ensemble then acts as a macroscopic magnetic dipole moment - referred to as magnetization - whose behavior can be understood in terms of classical physics. In clinical MRI ω_L is tuned to the resonance frequency of hydrogen (≈ 42.58 MHz/T in bulk water) - the most abundant element with odd proton count in human tissue as well as the one most sensitive to external fields, owed to a lack of shielding electrons.

Classical theory formalized in the Bloch equations now describes the magnetization as rotating with frequency ω_L in a plane transverse to the external magnetic

field while the rf-pulse is on. By virtue of Faraday's law this induces a voltage signal in nearby receive coils that is proportional to the proton spin density. After the rf-pulse is switched off the signal starts to decay. First, because phase coherence is lost amongst the spin ensemble on a time scale T_2 as individual spins drift in Brownian motion into slightly different chemical environments (fat, bulk water) with different local field strengths. Also, scanner field inhomogeneities further add to spin dephasing on a time scale $T_2^* < T_2$. Second, spins relax back into the equilibrium ground state through random collisions on a time scale $T_1 > T_2$. T_2^* can lie below ten milliseconds while T_2 and T_1 range from tens of milliseconds to seconds. The duration depends on tissue composition and thus opens a door to tissue specific image contrast.

Given that spins have no long-range interaction with each other the received time dependent signal $s(t)$ is a linear superposition of all rf-wave responses of spin ensembles from all positions \mathbf{r} in the sampled volume Ω

$$s(t) \propto \int_{\Omega} d\mathbf{r} c(\mathbf{r}) \rho(\mathbf{r}) e^{i(\omega_L t + \phi(\mathbf{r}, t))} + \eta(t). \quad (2.2)$$

Here c is the receive coil sensitivity field, ρ is the anatomic image and η is Gaussian white noise. Signal decay has been neglected in eq. 2.2 assuming signal readout time is short compared to T_2^* which is the case when using a fast low angle shot (FLASH) pulse sequence [10].

The rf-pulse played out over a finite interval τ will not have a δ frequency distribution and spin ensembles in different chemical environments will have slightly shifted Larmor frequencies. Spins may thus acquire an additional phase ϕ with respect to the reference phase

$$\phi(\mathbf{r}, t) = - \int_0^t dt' \omega(\mathbf{r}, t'). \quad (2.3)$$

Spatial variation of the phase is the key to spatial encoding. The simplest way is to introduce linear magnetic field gradients \mathbf{G} so the Larmor frequency is linearly varied

$$\omega(\mathbf{r}, t) = \omega_L + \delta\omega(\mathbf{r}, t) \stackrel{\text{eq. 2.1}}{=} \gamma(B_0 + \delta B(\mathbf{r})) \quad (2.4)$$

$$\delta B(\mathbf{r}, t) = xG_x(t) + yG_y(t) + zG_z(t). \quad (2.5)$$

Writing the gradient trajectory sampled by the scanner as

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_0^t dt' \mathbf{G}(t') \quad (2.6)$$

the phase term in eq. 2.2 can be re-expressed to yield the final MRI signal equation for a FLASH sequence in following convenient form

$$s(\mathbf{k}(t)) \propto \int_{\Omega} d\mathbf{r} c(\mathbf{r}) \rho(\mathbf{r}) e^{-2\pi i \mathbf{k}(t) \mathbf{r}} + \eta(t). \quad (2.7)$$

For slice-wise imaging signal generation can be reduced to a slice Ω_S by switching on a gradient along the slice selection direction during excitation and before signal acquisition. This way only spins from the slice where $\delta B = 0$ get excited.

In essence reading out the signal over period t with the gradients fixed along a certain direction corresponds to sampling a line of frequencies in the spatial frequency domain. Sampling an entire image entails repeated NMR experiments with different gradient settings each time. Once the entire \mathbf{k} -space (Fourier-space) of a frame has been sampled, the product $\rho \cdot c$ can be obtained by inverse Fourier transformation of the signal. Owing to the smoothness of the coil sensitivity field c this already reveals the anatomic structure.

It is important to note that even though the spin density is in principle real valued, image ρ and coil c are complex valued quantities. Aforementioned phase inhomogeneities due to different chemical environments could be factored out to separate maps but are here considered as part of ρ . Apart from physical processes, phase is also created as signal reception in the coil happens simultaneously on two orthogonal channels for better signal to noise ratio (SNR).

FLASH is a gradient echo sequence (GRE), and while not the only technique it is the fastest and therefore preferred where high frame rates are desired. GRE means spin coherence is destroyed following the excitation pulse and refocused to a signal echo at echo time TE which is determined by the gradient timing. The gradients remain on after TE until phase coherence is again destroyed.

Each \mathbf{k} -space line acquisition follows the same pattern and takes upward of 2 ms. For an image of size 128^2 - 256^2 px this adds up to ≈ 0.3 -1 seconds.

Further acceleration of line-sampling means even less sampling time will be spent per \mathbf{k} -space pixel and consequently, the SNR will deteriorate. Coarser binning of \mathbf{k} -space pixels while sampling time stays constant will increase the SNR but at the expense of spatial resolution. SNR, spatial and temporal resolution are intrinsically traded against each other.

Another limiting factor for the scanning speed is that too fast magnetic field gradient switching will cause patient discomfort through peripheral nerve stimulation and thus lower acceptance in clinical practice.

Parallel Imaging

One way to increase SNR while circumventing a trade-off in resolution is through receiver coil design. Early designs used a single receive coil that is easy to manufacture and collects signals from depth as far as its diameter. Deep reception of large coils also means a lot of background noise will enter the image.

A more sophisticated alternative is to use an array of smaller coils [11] that fit closer to the sample to combine better surface SNR with good volume coverage. Coil electronics are setup to decouple the individual channels as much as possible. Noise is approximately independent but the signal is overlapping and redundant between channels. Combining the individual channel images - for instance in a root sum square sense - then yields a better quality image.



Figure 2.1: Illustration of parallel imaging: Three simulated channels of a Shepp-Logan phantom (radially sampled) and root sum squares composition from a total of eight channels.

By introducing multiple channels the coil sensitivity map $c(\mathbf{r})$ in eq. 2.7 becomes a vector $\mathbf{c}(\mathbf{r})$. Modern receive coils in clinical use have up to 128 channels. SNR may vary from less than 1 to double digits depending on the exact position of the coil. To make data amounts manageable on a computer data compression by PCA into virtual channels is common.

2.2 Data Undersampling

As sequence and hardware optimization has reached electronics- and physiology related limits further acceleration for motion capture relies on \mathbf{k} -space undersampling. Measuring only every n th \mathbf{k} -space line with respect to the Nyquist limit ([9], ch. 12) yields a speed up of a factor n . However, any image sampled below the Nyquist rate and reconstructed via direct Fourier inversion will suffer from image artifacts that depend on the geometry of the sampling trajectory. For successful reconstruction of diagnostic quality images additional prior knowledge needs to be incorporated into the reconstruction to fill the \mathbf{k} -space gaps (section 2.3).

Cartesian Sampling

Cartesian sampling as of now is the clinical standard. The samples are collected in parallel lines by adding a phase shift in image space (eq. 2.7) corresponding to a

spatial shift in Fourier space between subsequent lines. Any sampling inaccuracy caused by electronics delays will be the same for all lines and therefore will not be noticeable in the image. The samples also map directly onto the Cartesian pixel grid without interpolation and the Nyquist limit is simply given by the number of lines. Leaving out parallel lines results in ghost copies of the object along the undersampling direction. Undersampling in this way by a factor 2 corresponds to a multiplication of the object \mathbf{k} -space with a 2 pixel periodic δ -comb along the undersampling direction or, equivalently, to a convolution in image space with a comb that has a period of half the image size.




Figure 2.2: Vertically undersampled Cartesian Shepp-Logan phantom with (left) two-fold, (mid-left) three-fold and (mid-right) three-fold undersampling plus 20 center lines as well as (right) the \mathbf{k} -space sampling pattern of the mid-right image with the \mathbf{k} -origin in the center.

Generally, low spatial frequencies correspond to low resolution details while high spatial frequencies correspond to fine details, including edges. Most signal energy of the image ρ is concentrated in the \mathbf{k} -space center. And this is even more true for the image-coil-product $\rho \cdot c$ since the smooth coil sensitivity field is entirely captured in very low frequencies. Image reconstruction procedures for Cartesian data therefore usually rely on densely sampled central lines [12],[13] as shown in Fig. 2.2.

Radial Sampling

Image reconstruction from radial data usually involves an additional conversion from polar coordinates onto a Cartesian grid via interpolation, called gridding [14]. The Nyquist limit is also $\pi/2$ times higher compared to Cartesian sampling. Here, not the line increment has to be one pixel in size but the azimuthal increment between the ends of neighboring radial lines (spokes) [15]. It is also more difficult to realize a trajectory in hardware where all radial lines meet exactly in the \mathbf{k} -space center point which mandates a gradient delay correction [16] as a further preprocessing step.

These disadvantages are offset by a high tolerance to oversampling. As in every tomographic imaging method all spokes have an equal share of high and low frequencies, thus it is irrelevant which ones are left out given a half-ways isotropic object and isotropic undersampling. The fact that all spokes go through the center also means that the central part of \mathbf{k} -space can be considered densely sampled up

to s/π samples away from the center, where s is the number of sampled spokes. Even though the Nyquist limit is higher, objects remain recognizable at low spoke counts where they would no longer be in Cartesian sampling.




Figure 2.3: Radially undersampled Shepp-Logan phantom reconstructed by Filtered Back-Projection from (left) 7 and (middle) 15 spokes as well as (right) the \mathbf{k} -space sampling pattern for the middle image. The artifacts extend beyond the shown field of view which is cropped to one third of the original matrix size (see section 4.1 Data Preprocessing).

By virtue of the Projection-slice theorem [9] each spoke corresponds to the 1D Fourier transform of a projection of the 2D object onto a line. The projection direction is perpendicular to the spoke angle. The most basic reconstruction algorithm for tomographic imaging is Filtered Back-Projection [9]. It first reweights the Fourier data spokes with a ramp filter, linearly decreasing the weight of low frequencies with respect to high frequencies, then transforms back to image space which corresponds to smearing the 1D projection over the 2D grid and finally sums over all projections.

The initial reweighting compensates the redundant sampling in the \mathbf{k} -space center and provided that as many back-projections as mandated by the Nyquist criterion have been sampled, individual projections and their directions can no longer be discerned. Filtered Back-Projection is accurate in that case. Otherwise, streak artifacts appear and are amplified as the weight of central \mathbf{k} -space lines is lowered too much with respect to the higher frequencies. Removal of these streak artifacts is the core task in this Thesis.

Spiral Sampling

Spiral sampling is not considered in this thesis, yet an interesting polar-coordinate alternative. Its main benefit is a longer readout time per excitation. Depending on the curvature a few spirals can cover the same number of \mathbf{k} -space samples as many radial spokes would. The ratio of low to high frequency sampling density can also be varied by varying the number and curvature of the spirals. In principle, since more time is spent in readout, better coverage and sharper images can be achieved in the same time as compared to radial sampling as long as the readout time is kept short enough to avoid off-resonance artifacts and signal decay. How-

ever, as in radial imaging, gradient delays occur and require more sophisticated corrections.

2.3 Image Reconstruction

2.3.1 Linear Problem

Let $\mathbf{y} \in \mathbb{C}^{Z \times N^2}$ be the measured multi-channel data with N^2 pixels and Z channels and $\rho \in \mathbb{C}^{N^2}$, $\mathbf{c} \in \mathbb{C}^{Z \times N^2}$ again the unknown anatomic image and the receive coils. Assuming an estimate for the coils exist and can be used as a parameter the reconstruction problem is to solve

$$\mathcal{A}\rho = \mathbf{y}, \quad \mathcal{A} := \mathcal{P}\mathcal{F}\mathcal{C} \quad (2.8)$$

for the unknown ρ . In this operator notation $\mathcal{C}\rho = (c_1 \cdot \rho, \dots, c_Z \cdot \rho)$, \mathcal{F} is the Fourier transform and \mathcal{P} is the sampling operator. When \mathcal{P} is invertible and \mathbf{c} has support in the entire spatial domain, then \mathcal{A} has an inverse and eq. 2.8 can be solved directly. If samples are omitted, however, \mathcal{P} is not invertible and neither is \mathcal{A} . In that case the reconstruction task can still be posed as an inverse problem where additional constraints \mathcal{R} are added to improve the problem condition

$$\rho^* = \underset{\rho}{\operatorname{argmin}} \quad \|\mathcal{A}\rho - \mathbf{y}\|^2 + \alpha\mathcal{R}(\rho). \quad (2.9)$$

The reference used for the remaining chapter on the theory of inverse problems and solution strategies is found in [17]. A penalty term \mathcal{R} weighted by $\alpha \in \mathbb{R}$ adds to the cost of undesired solutions in the functional minimization. Setting $\mathcal{R} = \|\rho\|^2$ (Tikhonov regularization) penalizes solutions with amplified noise and outliers, thus stabilizing variance at the cost of an introduced bias (blur in this case). With $\alpha = 0$ the analytic solution to eq. 2.9 boils down to applying the pseudo-inverse of \mathcal{A} to \mathbf{y} which again will not work numerically on undersampled data. A regularized iterative approach such as the conjugate gradient method [18] on the other hand will converge to a solution close to the true ρ .

From a signal processing point of view eq. 2.8 is a deconvolution problem in \mathbf{k} -space with known filters c_i as well as a deconvolution in image space with regard to the product $\mathbf{c}\rho$ and the sampling pattern \mathcal{P} . Deconvolving data \mathbf{y} and coils \mathbf{c} to yield ρ will recover non-measured \mathbf{k} -space frequencies located beside the sampled spokes to the extent of the coil support. Even though the coils are smooth they do have a spatial structure and thus finite extended support in \mathbf{k} -space. It increases with the number of independent coil channels in the coil array arrangement.




Figure 2.4: Radial Shepp-Logan phantom with 15 spokes with (left) one constant coil channel, (middle) 8 coil channels as shown right and (bottom) the corresponding \mathbf{k} -spaces (4-fold zoom on center). The phantom images are cropped to half of the original matrix size (see section 4.1 Data Preprocessing).

Carrying out the minimization in eq. 2.9 with a Tikhonov penalty leads to following normal equations

$$\left(\sum_{i=1}^Z c_i^* \mathcal{F}^{-1} \mathcal{P}^H \mathcal{P} \mathcal{F} c_i + \alpha \mathbf{1} \right) \rho = \sum_{i=1}^Z c_i^* \mathcal{F}^{-1} \mathcal{P}^H y_i, \quad (2.10)$$

where c_i^* are the complex conjugated coil channels. Iterative solution of eq. 2.10 is referred to as sensitivity encoding (SENSE) [19]. Contrary to Filtered Back-Projection image quality no longer suffers from over-amplified high frequencies. A technical issue to note is that oversampling of the object by a factor 2 is necessary for artifact-free deconvolution since the process involves repeated convolutions of the form $\mathcal{F}^{-1}[\mathcal{P}^H \mathcal{P}] * [C\rho]$ (carried out in Fourier space) the output of which extends to twice the image support. As a consequence the object support occupies only the inner half of the compute matrix which allows the introduction of a truncation in image space [20] (used in Fig. 2.4) as an additional constraint. In the absence of noise the minimization now will converge even if $\alpha = 0$.

SENSE type reconstructions that deconvolve estimated coils from the image are one of two main brands of reconstructions. An alternative \mathbf{k} -space based approach [12] iteratively estimates - in what is referred to as calibration - a linear correlation kernel mostly from densely sampled \mathbf{k} -space area that is assumed to

consistently connect samples throughout the entire \mathbf{k} -space. The estimated kernel is then applied to all other samples in a calibration consistency constraint introduced to the cost functional. The automatic occurrence of densely sampled parts of \mathbf{k} -space for coil estimation or kernel calibration required in linear methods adds another advantage to polar sampling schemes.

Recently, image space based coil deconvolution and the \mathbf{k} -space based calibration approach have been unified in a common framework [13].

2.3.2 Non-Linear Problem

Parallel imaging based reconstruction crucially relies on accurate estimates of \mathbf{c} . Especially image space based approaches tend to suffer from artifacts due to coil mismatches or failed deconvolution due to a too small FOV relative to the object size [13].

In real-time imaging patient breathing may alter coil positions during the scan. Further, internal motion of the heart or blood inflow will alter the dielectric properties of the sample which has an interaction with the coil sensitivity field [21],[6]. Coil sensitivity estimates thus need updating as soon as the object changes. Joint image and coil estimation in alternating linear reconstructions has been proposed [22] but a more accurate simultaneous minimization

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{arg\,min}} \quad \|\mathcal{A}\mathbf{x} - \mathbf{y}\|^2 + \alpha\mathcal{R}(\mathbf{x}), \quad \mathbf{x} = (\rho, \mathbf{c})^T \quad (2.11)$$

can be achieved via the iteratively regularized Gauss-Newton method [6] that will be used in this work. The reconstruction is also referred to as non-linear inversion (NLINV) and is a blind deconvolution generalizing SENSE type methods. An extension of NLINV to add the advantages of calibration methods such as robustness to insufficient matrix support has recently been proposed [23].

Instead of solving eq. 2.11 in one step, NLINV approximates the joint minimization by a series of linear problems each defined by a first order Taylor expansion of the measurement $\mathcal{A}\mathbf{x}$ given as

$$\begin{aligned} \mathbf{x}_{n+1} &= \mathbf{x}_n + d\mathbf{x} \\ \mathcal{A}\mathbf{x}_{n+1} &\approx \mathcal{A}\mathbf{x}_n + D\mathcal{A}(\mathbf{x}_n)d\mathbf{x}, \end{aligned} \quad (2.12)$$

where $D\mathcal{A}(\mathbf{x}_n)$ is the Jacobian of \mathcal{A} around \mathbf{x}_n . Starting from an initial guess $\mathbf{x}_0 = (\rho = 1, \mathbf{c} = 0)$ the joint variable \mathbf{x} is updated in every Newton step by an increment $d\mathbf{x}$. Inserting eq. 2.12 into eq. 2.11 yields the update rule with parameter \mathbf{x}_n

$$d\mathbf{x}^* = \underset{d\mathbf{x}}{\operatorname{arg\,min}} \quad \|D\mathcal{A}(\mathbf{x}_n)d\mathbf{x} - (\mathbf{y} - \mathcal{A}\mathbf{x}_n)\|^2 + \alpha_n \|\mathbf{x}_n + d\mathbf{x}\|^2. \quad (2.13)$$

The definition of $D\mathcal{A}$ and its adjoint DA^H is given in [6]. In eq. 2.13 ρ and \mathbf{c} are tied into a product by \mathcal{A} and $D\mathcal{A}$, thus additional regularization is required to separate the two variables. Subjecting \mathbf{c} to a Sobolev penalty $\hat{\mathbf{c}} = \mathcal{W}^{-1}\mathbf{c}$ with $\mathcal{W}^{-1} = (1+u\|\mathbf{k}\|^2)^{l/2}\mathcal{F}$ while leaving ρ unchanged will penalize all but the center \mathbf{k} -space frequencies of \mathbf{c} in accordance with the coil smoothness assumption. Here $u = 440$ and $l = 32$ are used. Solving eq. 2.13 for the transformed variables $\hat{\mathbf{x}}$ and $d\hat{\mathbf{x}}$ in conjunction with the modified operators $\mathcal{G} = \mathcal{A}\mathcal{W}$ and $D\mathcal{G} = D\mathcal{A}\mathcal{W}$ will avoid large numbers and is numerically more convenient.

An important ingredient to the method is that the penalty weight shrinks with $\alpha_n = (\alpha_0/2)^n$, where typically $\alpha_0 = 1$, but lower start values often also work. Decay rates slower than 0.5 will slow down convergence and necessitate more Newton steps while faster decay may lead to diverging solutions. The algorithm implementation is conceptually simple

```

1  inputs:
2      Y = data,    P = sampling_pattern
3  init:
4      rho = 1,    c = 0,    alpha = 1
5
6  for n = 1:newton_steps
7
8      //solve eq. 2.13 (e.g. with conjugate gradient method)
9      (d_rho, d_c) = calc_image&coil_update(Y,P,rho,c,alpha)
10
11     (rho, c) = (rho, c) + (d_rho, d_c)
12     alpha = alpha/2
13

```

Algorithm 2.1: Iteratively regularized Gauss-Newton method for Non-linear inversion reconstruction.

While optimization in the inner, linear problem will be stopped according to some fixed stopping rule, for instance when a fraction of the initial residual has been reached [18], the outer loop runs a predetermined number of iterations [6]. Too few iterations will result in blur and - in the case of undersampling - more remaining artifacts, whereas too many iterations result in increased noise and computing cost. Typically, good image quality is reached after 10 Newton steps. Availability of a better initial estimate, such as a body coil image or a previous frame in an image series may allow for earlier stopping.

When starting from a constant initial guess, images in early Newton steps will look blurry and the intensity distribution will likely be incorrect since the coil estimate is still inadequate. It has also been observed here that energy is exchanged between image and coils in a kind of swing phase until a stable distribution is reached from the 6th step onward. From there convergence proceeds more steadily.




Figure 2.5: Anatomy and coil sensitivity of third virtual channel at different Newton steps reconstructed with NLINV. Shown is a four chamber view of the heart with 75 spokes per frame obtained by resorting and rebinning undersampled data (see section 4.2.2). Intensity is windowed identically in an absolute sense for the image and relative sense for the coils.

Another point to stress is that NLINV only recovers the parts of the coil sensitivity field that overlaps with the object, not the physical coils. The part that extends beyond into the background cannot be recovered from the signal void.

Real-time MRI

Rt-MRI refers to the measurement and display of anatomic motion in frame series with high temporal fidelity at rates of 30 fps or even 50 fps [2]. Applications of interest include amongst others arrhythmic cardiac motion, human speech, joint movement, dynamic contrast enhanced MRI for tumor detection in abdominal scans and MRI of pediatric patients that otherwise require anesthesia to prevent motion artifacts.

Methods proposed in the literature can broadly be divided into either online or offline reconstructions. The former is in principle capable to output the reconstructed images simultaneously with the data acquisition or with a minimal - sub-second - delay. The latter requires the acquisition of the entire series to be finished before the start of reconstruction. Offline methods can exploit space-time structure in the signal as prior knowledge to the reconstruction. Proposed procedures usually follow the same pattern of finding some suitable transform along the time dimension under which the object becomes sparse and then add an L_1 -penalty term involving the transformed object to the cost functional to be minimized. The methods then differ mostly by the transform or set of transforms chosen which include discrete time differences (total variation regularization) [24], discrete time differences plus a data reordering to separate cardiac and breathing components [25], temporal Fourier-transform [26] (in combination with a low rank constraint on the data term), or one of several motion estimation based transforms [27], [28]. The great flexibility to realize temporal regularization, however, does not automatically safeguard against temporal blur [29] which requires careful parameter tuning. Online methods are far more limited as they cannot draw on information from future frames yet enable interventional MRI applications such as operator subject communication in speech studies.

The exact distinction between rt-MRI and non-rt-MRI is still subject to debate and sometimes the qualifying requirement is added that image reconstruction frame rates should also be real-time - or close to - with present day computing hardware [30]. This work targets an online approach. However, the definition of rt-MRI is relaxed to include reconstructions that conceptually allow real-time viewing with

a constant delay of a few frames even if this is not quite within the capacity of present day hardware yet.

3.1 Dynamic Sampling

Parallel imaging with non-Cartesian sampling is considered sufficient to restore images to acceptable quality at undersampling factors of 2-4 [31] in 2D. Higher factors of 7 [12],[32] to 12 [33] are only considered in conjunction with additional priors such as sparsity in dynamic contrast enhanced imaging, MR angiography or brain imaging. Compressed sensing methods that are used to exploit spatial sparsity, however, work best with randomly undersampled 3D data. Normal, regular breathing at a frequency of 0.2 Hz in abdominal scans may still be well captured with an undersampling factor of 3-5. Heart beats, however, reach rates of 1-1.8 Hz and heart contraction during the systolic interval lasts for less than 40% of the beat period, usually about 350 ms in total [34].

For an FOV of size 200^2 px the Nyquist rate is ≈ 314 spokes. Thus, realizing a frame rate of 30 fps with 2.2 ms per spoke or 15 spokes per frame requires undersampling by a factor of ≈ 21 in radial imaging, well beyond the capacity of parallel imaging alone. In a series of frames temporal continuity of anatomic motion can be added as a powerful prior. It assumes that most pixels in image space will not change much from one frame to the next given sufficiently small time steps.

Temporal continuity is exploited by splitting a “full” k -space sampling pattern into components with minimal overlap and spreading these out in time to form several undersampled frames.

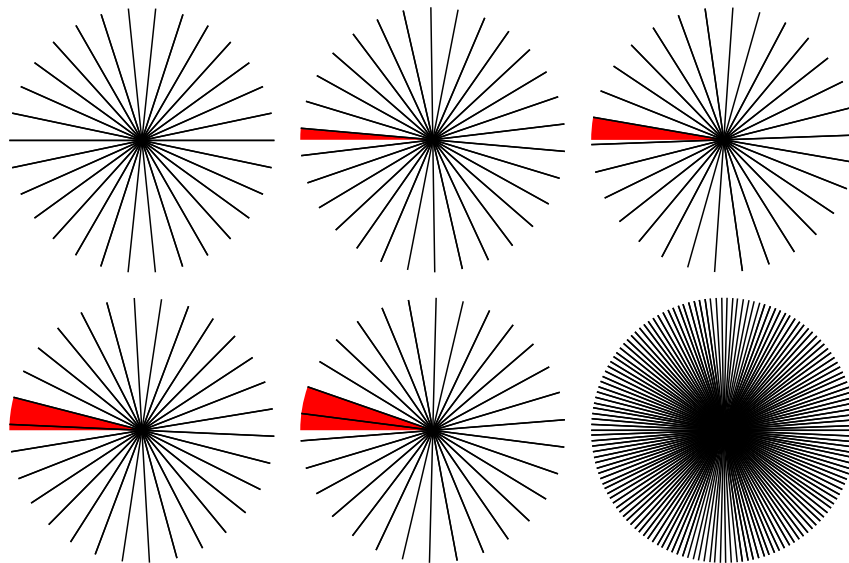


Figure 3.1: Schematic of a rotating sampling pattern with 15 spokes per frame and a finite repetition period of length 5. The angle increment with respect to the first pattern is indicated in red. The total k -space coverage after a full period is shown in the bottom right.

Most of the information held in one frame will still be up to date with that in the next frame and can be copied directly. The rest should require only minor adjustments to be incorporated into neighboring frames without causing temporal blur given sufficiently small time steps. Data sharing among frames is realized by adding a motion model to the cost function.

A natural way of splitting in radial imaging is to rotate the sampling pattern from frame to frame (Fig. 3.1). The pattern rotation period can be finite with a regular pattern and constant angle increment in between frames [35],[2] or infinite with an irregular pattern following a golden angle progression [36]. A larger repetition period results in better k -space coverage. Infinite periods are especially popular in offline reconstructions [24].

The accuracy of the motion model determines the amount of motion blur. Another consequence of an inaccurate model may be an uneven weighting of patterns of different frames, such that the samples of the current frame will have systematically higher energy than samples stemming from neighboring time points.




Figure 3.2: Four chamber view of the heart (magnitude image) reconstructed with NLINV with an affine time constraint (see section 3.2.1). Streak artifacts are less visible than without temporal regularization (see f. i. Fig. 4.3) but are still of similar intensity as the heart motion as shown in time difference images (real part of image). The arrows follow the clock-wise rotation of undersampling artifacts.

As a result, radial streak artifacts (Fig. 3.2) appear markedly less intense than when no temporal regularization is used but similar in shape and very visible due to their fast rotating motion driven by the pattern rotation. Their intensity is on a par with the intensity of the anatomic motion.

Returning to the example of a 200^2 px FOV with 15 spokes per frame, a motion model that correctly connects neighboring frames within a 5-frame time window can recover \mathbf{k} -space frequencies for any frame to 75 spokes. Adding the independent effect of parallel imaging will restore another factor 3-4 which rounded up brings frames close to the Nyquist rate of 314 spokes.

3.2 Temporal Regularization

3.2.1 Time Difference Regularization

The simplest temporal regularization scheme adds an affine penalty term to the cost function [2] which in the linear problem for the frame at time point t reads as

$$\rho_t^* = \arg \min_{\rho_t} \|\mathcal{A}\rho_t - \mathbf{y}_t\|^2 + \alpha \|\rho_t - \lambda \rho_{t-1}\|^2, \quad (3.1)$$

where $\lambda \in [0, 1]$ and the formulation for the non-linear problem is analogous. In the linear problem $\lambda = 1$ maximizes the connection between the current frame ρ_t and the preceding one, whereas in the non-linear problem for values $\lambda > 0.9$ an undesired steady - approximately linear - increase in image energy is observed from frame to frame. For online reconstructions this simple scheme is a common choice [37],[38],[39].

It borrows from “view sharing” and related ideas [1],[35]. On the pros’ side, it does not introduce additional complexity or computing cost as compared to simple non-temporal Tikhonov regularization. In fact, in the non-linear case fewer Newton steps are required to obtain sharp images (8 in [2] vs. 10 in [6]). Thus temporal regularization even speeds up computation, possibly because it effectively lowers the undersampling factor, thereby improving the problem condition. Also, the sole artifact introduced by temporal affine regularization is time blur. Spatial artifacts that may resemble pathologies are not introduced.

The downside, as in any view sharing method, is degraded temporal resolution. With a pattern repetition period of 5 frames and a temporal binning of 33 ms the frame update rate is still every 33 ms but each frame now has a footprint several times longer. The nominal temporal resolution is only an apparent one while the true, lower resolution depends on α and λ .

With $\lambda = 1$, the effect of affine regularization is similar in character to a sliding window box average applied in post-processing. The value of α , however, will usually be too low to produce a time average with equal weights on every time

point since the deconvolution in the data term will not be properly resolved with a large α , which leads to a blurry image. In the affine scheme the most current frame gets the largest weight while past frames enter the sliding window average with decreasingly lower weights.

As a consequence undersampling artifacts will never vanish entirely (see Fig. 3.2). Therefore, an additional median filter is used along the time dimension in a post-processing step. Over a slowly moving background artifacts will appear as outliers and be filtered away. But this also holds for any movement faster than half the filters width, which is matched to the pattern repetition period. Another drawback of this scheme is that it reaches only into the past, thus the center of the average is not the current frame.

A straight forward way to extend the affine regularization that did not yet get much attention in MRI is to consider a batch of frames $\rho := (\rho_{t-\tau}, \dots, \rho_{t+\tau})$ that is estimated simultaneously instead of a single frame ρ_t

$$\rho^* = \underset{\rho}{\operatorname{argmin}} \sum_{t'=t-\tau}^{t'+\tau} \|\mathcal{A}_{t'}\rho_{t'} - \mathbf{y}_{t'}\|^2 + \beta \|\partial_t \rho_{t'}\|^2. \quad (3.2)$$

Temporal regularization of the coils can be added, but was not observed to have much effect. Obviously, the problem to compute is now $2\tau+1$ times as large and costly. More importantly though it is now centered with respect to the current frame. Therefore delays or temporal blur will not range beyond τ frames. Another benefit is that the time differences summarized in the ∂_t -operator do not require a damping factor λ in the non-linear problem. Reconstruction of the entire series can either proceed in time increments of single frames or whole batches to avoid 2τ -fold recomputation of every frame. In the latter case, however, all but the batch's center frame will be of poorer quality due to their asymmetric neighborhood. Therefore, only the center frame is stored in the output buffer. The centrality also promises a more even weighting of the different patterns combined in the batch's center frame such that unlike in the affine regularization no further post-processing is required.

As a variant, the temporal penalty in eq. 3.2 can also be put in an L_1 -penalty, corresponding to a temporal total variation constraint as in [24] but on a shorter time interval. The outcome, however, is unlikely to match the results of corresponding whole series offline reconstructions. Usually a sufficiently long series is regarded as a necessity to make the undersampling artifacts appear incoherent and noise-like in the context of the whole series [40].

3.2.2 Optical Flow

The name "optical flow" implies the assumption that intensity or total brightness stays constant in between two frames of a series and that only its spatial distri-

bution changes. Like a fluid of constant volume that flows continuously and does not jump between places. As derived in the original paper [7], the assumption of a constant energy and application of the chain rule straight forwardly lead to the model

$$\begin{aligned} \frac{d}{dt}\rho &= 0 \\ &= \partial_t\rho + v_x\partial_x\rho + v_y\partial_y\rho, \end{aligned} \quad (3.3)$$

$$= \partial_t\rho + \mathbf{v}\nabla\rho \quad (3.4)$$

$$v_{x,y} = \frac{dx,y}{dt}, \quad (3.5)$$

which is a divergence-free continuity equation. Translated into image processing terms in a discrete setting it means that the follow up to a frame is obtained by adding to the previous frame some more intensity at the object edges

$$\rho_{t+1} = \rho_t + \mathbf{v} \underbrace{\nabla\rho}_{\text{edges}}. \quad (3.6)$$

Since eq. 3.3 is one linear equation for two variables $v_{x,y}$ a further constraint is required. The original paper proposes a smoothness constraint $\|\underline{\nabla}\mathbf{v}\|^2$ on the flow field which amounts to the assumption that neighboring pixels have similar velocities enforced by blurring the flow field gradient. The operator $\underline{\nabla}$ here is meant as a 2D-gradient applied to v_x and v_y respectively not a divergence operator on \mathbf{v} . The paper goes on to note that a smoothness penalty complicates the handling of occlusions, which is a major concern in camera imagery but not in MRI. Putting together the optical flow- (intensity conservation) and the smoothness constraint yields the minimization problem

$$\mathbf{v}^* = \underset{\mathbf{v}}{\operatorname{arg\,min}} \quad \|\partial_t\rho + \mathbf{v}\nabla\rho\|^2 + \|\underline{\nabla}\mathbf{v}\|^2. \quad (3.7)$$

Why this is expected to help with the streak artifacts can already be seen in eq. 3.6: only values of \mathbf{v} that overlap with edges translate between frames. Clearly, this is restricted to continuous deformations which is typical of anatomic motion. All other values of \mathbf{v} will be suppressed to zero, especially where the abrupt, teleport-like motion of artifacts occurs. Another way to see this is in the normal equations to eq. 3.7 after carrying out the minimization component-wise

$$(\partial_x^H\partial_x + |\partial_x\rho|^2)v_x = -(\partial_x\rho)^H\partial_t\rho - v_y(\partial_x\rho)^H\partial_y\rho \quad (3.8)$$

and similarly for v_y . In an iterative procedure where v_x is computed before v_y and that is initialized with 0 the solution is build successively on the correlation between the spatial and temporal derivative of the object ρ . For jumping arti-

facts this correlation should be considerably smaller than for an expanding heart muscle- or chest wall.

The combination of flow estimation with image reconstruction could be realized in one joint minimization over three variables $\rho, \mathbf{c}, \mathbf{v}$. However, alternating minimization between image reconstruction and flow estimation facilitates debugging and output analysis. It is therefore common to implement a two step minimization.

In this work an approach used by Dirks [41] in camera footage and Burger [42] - linear CT imaging - is adapted to NLINV (compare eq. 2.13)

$$\begin{aligned} \hat{\mathbf{x}}_{n+1} = \hat{\mathbf{x}}_n + \arg \min_{d\hat{\mathbf{x}}=(d\rho, d\hat{\mathbf{c}})} & \sum_{t'=t-\tau}^{t'=t+\tau} \|D\mathcal{A}_{t'}(\hat{\mathbf{x}}_{t',n})d\hat{\mathbf{x}}_{t'} - (\mathbf{y}_{t'} - \mathcal{A}_{t'}\hat{\mathbf{x}}_{t',n})\|^2 \\ & + \alpha_n \|\hat{\mathbf{x}}_{t',n} + d\hat{\mathbf{x}}_{t'}\|^2 + \gamma_n \|(\partial_t + \mathbf{v}_{t',n}^{t'+1}\nabla)(\rho_{t',n} + d\rho_{t'})\|^2 \end{aligned} \quad (3.9)$$

$$\begin{aligned} \mathbf{v}_{n+1} = \arg \min_{\mathbf{v}} & \sum_{t'=t-\tau}^{t'=t+\tau} \|\partial_t \rho_{t',n+1} + \mathbf{v}_{t',n+1}^{t'+1} \nabla \rho_{t',n+1}\|^2 + \frac{\beta_n}{\gamma_n} \|\nabla \mathbf{v}_{t',n+1}^{t'+1}\|^2 \\ & + \frac{1}{\gamma_n} \mathcal{R}(\mathbf{v}), \end{aligned} \quad (3.10)$$

where \mathbf{v}_t^{t+1} designates flow from ρ_t to ρ_{t+1} with $\mathbf{v}_{t+\tau}^{t+\tau+1} = 0$ and $\hat{\mathbf{x}} = (\rho, \mathcal{W}^{-1}\mathbf{c})$ is again the joint variable with Sobolev-weighted coils. The coils are not subjected to the optical flow constraint since their edges are already smoothed by the Sobolev regularization. Omitting the flow estimation and keeping $\mathbf{v} = 0$ at its initial value leads to the batch time difference regularization of the previous section, again avoiding extra coil regularization. All penalty norms except the data term in the \mathbf{x} -problem may also be changed into L_1 . As will be explained in section 4.3.1 the computation starts as regular NLINV without temporal regularization and adds optical flow only after a few Newton steps.

By splitting the problem into two alternating parts it is easier to separate the shaping of the flow field as a variable in the \mathbf{v} -problem from the application of the flow field as a parameter in the \mathbf{x} -problem. \mathcal{R} summarizes further regularization terms. Here, an additional L_1 smoothness penalty is considered and/or masking of the flow field where either $|\rho|$ or $|\partial_t \rho|$ lies below noise level. Another idea is to penalize the field where $|\partial_t \rho|$ was large at step n and considerably lower at $n+1$. This assumes that most of the intensity change before frames are connected stems from artifacts which are reduced after applying the motion model

$$\mathcal{R}(\mathbf{v}) = \beta_{2,n} \|\nabla \mathbf{v}\|_1 \quad (3.11)$$

$$\wedge \quad \mathcal{R}(\mathbf{v}) = \delta_n \mathbf{1}_{(|\partial_t \rho_{t,n}| > 3\sigma)} \cdot \mathbf{1}_{(|\rho_{t,n}| > 3\sigma)} \cdot \|\mathbf{v}\|^2 \quad (3.12)$$

$$\wedge \quad \mathcal{R}(\mathbf{v}) = \delta_n \|\partial_t \rho_{t,n} - \partial_t \rho_{t,n-1}\|^2 \cdot \|\mathbf{v}\|^2. \quad (3.13)$$

In eq. 3.12 $\mathbb{1}$ denotes the support of values above noise level, whereby the standard deviation σ is estimated using the median absolute deviation value of a wavelet-high-pass filtered ρ [43]. For large values of δ this constraint acts like a hard thresholding operation and will be referred to as such in the later part.

The approach in eq. 3.9-3.10 estimates an entire batch at once $\mathbf{x} \in \mathbb{C}^{N^2 \times Z \times (2\tau+1)}$, all of which could be saved as output but, again, as in the batch time difference regularization only the center frame will be used.

An important point to note is that only differences between nearest neighbors matter in this approach. As long as these are small for all frame pairs the motion model should hold irrespective of the total batch size and the total time difference between the first and last frame in the batch.

Li et al. [8] presented an alternative combination of optical flow with NLINV dubbed aggregated motion estimation. First, they reconstructed initial images in a batch by NLINV without temporal regularization. The batch size again matched the pattern repetition period. They then used the obtained estimates of $\rho_{t+\tau}$, $\tau \in [-2, 2]$ to compute the flow field \mathbf{v} . In their case all penalty terms were in L_1 and they added an auxiliary regularized variable to eq. 3.7 to model undersampling artifacts. Importantly, and contrary to the approach chosen in the present work, their flow modeled displacements between frames more than a frame apart. This somewhat strains the assumption that temporal connections between frames are restricted to the support of the image edges (eq. 3.6). The deformation of ρ_t to $\rho_{t+\tau}$ is denoted as

$$\Phi_t^{t+\tau}(\rho_t) := \rho_t + \mathbf{v}_t^{t+\tau} \nabla \rho_t \approx \rho_{t+\tau}. \quad (3.14)$$

It is included in the minimization problem over a *single* frame at time t to map to other frames in the batch

$$\begin{aligned} \hat{\mathbf{x}}_{t,n+1} = \hat{\mathbf{x}}_{t,n} + \operatorname{argmin}_{d\hat{\mathbf{x}}} & \left(\sum_{t'=t-\tau}^{t'=t+\tau} \|D\mathcal{A}_{t'}(\Phi_t^{t'}(\hat{\mathbf{x}}_{t,n}))d\hat{\mathbf{x}} - (\mathbf{y}_{t'} - \mathcal{A}_{t'}\Phi_t^{t'}(\hat{\mathbf{x}}_{t,n}))\|^2 \right. \\ & + \alpha_n \|\hat{\mathbf{c}}_{t',n} + d\hat{\mathbf{c}}_{t',n} - \hat{\mathbf{c}}_{t',0}\|^2 \\ & \left. + \alpha_n \|\rho_{t,n} + d\rho_{t,n} - \rho_{t,0}\|^2, \right) \end{aligned} \quad (3.15)$$

whereby $\hat{\mathbf{x}}_{t,n} = (\rho_{t,n}, \hat{\mathbf{c}}_{t-\tau,n}, \dots, \hat{\mathbf{c}}_{t+\tau,n})$ is now one anatomic image jointly with a time batch of (Sobolev-weighted) coils.

Yet another difference between the two approaches is that flow estimation here happens only once based on images reconstructed without temporal regularization and with strong artifacts. The first approach on the other hand re-estimates the flow field repeatedly in every newton step expecting an improved flow field at each further step.

Lastly, in this work, similar to Li et al. the most basic form of optical flow is considered since deformations between frames are expected to range on the order of a few pixels in cardiac imaging at 30 fps. More sophisticated forms can be employed such as affine or locally affine optical flow [44] for higher accuracy as well as multiscale- and patch-based versions to cover larger displacements [45],[46]. However, it is not expected that these would better deal with the challenge of inflowing intensity which would be avoided with a matching source term in the optical flow continuity equation or if the measurement was 3D or multislice.

Materials and Methods

4.1 Setup and Data Acquisition

Scanning Hardware

All human heart data in this work was acquired with a 3T scanner (Magnetom Prisma fit, Siemens Healthcare, Erlangen) using a standard 32-channel spine coil and 18-channel thorax coil. Written informed consent by the subject (healthy volunteer) according to the recommendations of the local ethics committee was obtained prior to MRI.

Pulse Sequence

The pulse sequence used was a single-shot radial FLASH [10] with 1.44 ms echo time, 2.22 ms repetition time, 1.6 mm^2 in-plane resolution, 6 mm slice thickness, 8° flip angle and an FOV of 200 px. Each image frame comprised 15 radial spokes, with the pattern rotated between frames as illustrated in Fig. 3.1. This results in a frame rate of $(2.22 \times 15 \text{ ms})^{-1} \cong 30 \text{ fps}$ and a total Fourier-space coverage of 75 spokes. To counter aliasing an oversampling ratio of a factor 2 was applied, doubling the number of pixels per spoke. Imaging parameters were identical across all measurements.

Data Simulation

Phantom multi-channel data was simulated as the analytical Fourier-transform of a grid of ellipses, again with two-fold data oversampling. Its corresponding 8-channel coil sensitivities were simulated with a Biot-Savart model [47]. Data was simulated both noiseless and with additive white Gaussian noise with a standard deviation of five percent of the mean intensity.

Data Preprocessing

The raw radial NMR data was interpolated onto a Cartesian grid using a Kaiser-Bessel kernel appropriate for an oversampling factor of 2 [14]. To reduce interpo-

lation errors the Cartesian pixel grid was stretched by a factor 1.5, yielding an overall image matrix size three times that of the FOV. Human data was compressed by PCA into 10 virtual coil channels and a magnetic field gradient delay correction [16] was performed prior to gridding to account for deviations between the desired and the traversed sampling trajectory, induced by eddy currents and scanner electronics delays. The data was then multiplied by a factor of 100 divided by the L_2 -norm of the first frame. Hence, all parameter settings here are with respect to a fixed scale.

Reconstruction Soft- and Hardware

Data Preprocessing scripts and simulation software were written in MatLab and the reconstruction framework in C++ using an in-house library of the BiomedNMR workgroup for multi-GPU support. Reconstructions were run on 4 NVIDIA GTX Titan Black with Cuda 9.2, whereby inter-GPU parallelization was along the channel dimension and inner-GPU parallelization over pixels. Runtime depended on image size, number of run iterations, and reconstruction type (known/unknown coils, with/without flow constraint, frame batch size) and ranged between ≈ 100 ms and 30 seconds per frame.

4.2 Test Data

4.2.1 Numerical Phantoms

Numerical phantom data has been created to simulate the effect of pure in-plane and pure through-plane motion and each motion model's response to it. It is less suited to assess structural blur. The phantoms show a grid of circles, either blinking or pulsating, whereby the blinking/pulsation frequency ω increases linearly from the bottom right ($\omega = 0$, static) to the top left circle ($\omega = 2\pi/60$ frames $^{-1}$).




Figure 4.1: Noiseless numerical phantom simulated with 75 spokes for (left) through-plane and (mid) in-plane motion as well as an undersampled in-plane phantom with 15 spokes per frame. Images were reconstructed with known coils calculated by a prior NLINV reconstruction. Parameters see Tab. A.2 first row. Left is a real-valued image with negative values in black and positive in white while right are magnitude images.

A phantom simulated with 75 spokes per frame is referred to as a “fully” sampled reference even though it is still below the Nyquist sampling rate of $256 \times \pi/2 \approx 403$ spokes for a field of view of 256^2 px [15]. This is to be consistent with the human data which has 15 spokes per frame and visits 75 different spoke angles in total (Fig. 3.1). For the same reason, the simpler, linear reconstruction with known coils is always computed with coils estimated from a previous NLINV reconstruction instead of the original simulated coils (compare Fig. 2.4 top right) that have no analogue in human data. Since the coils are known to be static, only the first frame where the circles are all of equal size and the grid is most regular is considered and these coils used for the entire image series.




Figure 4.2: The 8 complex valued coil sensitivities estimated by NLINV for the circle grid phantom in Fig. 4.1. Complex phase is color coded.

4.2.2 Surrogate Data

Heart data was synthesized from measured data to assess structural blur and to rate its compliance with optical flow within the scale set by above phantoms. In cardiac imaging, unlike other applications such as joint movement or speech, creating ground truth reference images is greatly simplified by the periodicity of the motion. As in standard clinical cardiac cine imaging [3] the volunteer was asked to hold her breath during the MRI exam for 15 seconds to reduce all motion but that of the heart.

In cine an ECG trigger or a biomarker derived from the raw data is used to eliminate irregular heart beats - those too long or short or with abnormal heart motion - from the series. The rest is averaged into one synthetic heart beat. Provided that enough data remains after discarding, and the remaining frames form a periodic series no artifacts will appear.

Data time stamps from a finger pulse trigger have been recorded alongside MRI,

but reconstructions from cine data always turned up with at least slight artifacts that resembled ghosting edges on both sides of the heart wall, probably due to slight movements, such as chest expansion, by the volunteer.



Figure 4.3: (Left) magnitude images reconstructed from surrogate heart data (75 spokes per frame) obtained by resorting and combining MR data of an undersampled acquisition with 15 spokes per frame (corresponding images on the right).

Therefore, instead data was reconstructed by NLINV with affine temporal regularization + median filter first, assuming that any introduced temporal distortions would be the same among all frames showing the same heart phase. Subsequently, every frame of 15 spokes was combined with those frames in the series that were the most similar in a mean square error sense (MSE) - indicating the same motion state - yet had a different sampling pattern orientation. The result was a full frame for every time point. The sorting order obtained that way was then applied to combine the raw MR data before gridding interpolation. Subsequent reconstruction from that data yielded largely artifact free heart movies with naturally appearing motion.

4.2.3 Image Quality Assessment

Quality assessment of reconstructed images with respect to a reference by a scalar score conveniently allows for comparing different methods at a glance. A fair comparison of different temporal regularization types, however, requires the coil deconvolution part in the inverse problem to be identical amongst the compared methods. Therefore, only the simpler, linear problem with known coils as a parameter is open to such a comparison. In the non-linear problem temporal regularization influences the coil sensitivity estimate and thus indirectly the intensity distribution in the low spatial frequency regime. Consequently, temporal fidelity will no longer be the determining factor in the calculation of any quality score. Thus, quality assessment of non-linear reconstructions in this work will be limited to a qualitative visual comparison.

It must be noted that quantitative assessment in the simpler linear reconstruction is always with respect to a certain set of reconstruction parameters, as the reference must first be reconstructed itself. In the case of a noiseless phantom no bias-variance trade-off has to be made and an ideal reference - the one with least bias (small α in eq. 2.10) - is obtained. Yet in the presence of noise too many iterations with a small α will yield a result with crisp edges but excessively amplified noise. Parameter selection thus has to be done upon visual inspection for the phantoms and heart images and the optimal choice will vary with image noise level. The corresponding parameter sets are listed in the appendix (Tab. A.2-A.12). Another issue with noise is that while the simple mean square error (MSE) between a ground truth and different reconstructed results may adequately rank the performance of the tested methods in a noiseless case, it may put methods with better noise suppression above those with better feature preservation when that is not desired.

Quality metrics with an emphasis on feature similarity include the well established structural similarity index (SSIM) [48] as well as the more elaborate feature similarity index (FSIM) [49]. For simplicity the former will be used in quantitative com-

parisons. After all, both indicate only to what extent feature fidelity with respect to the ground truth was missed and not how. Therefore, comparisons between different image reconstructions must still include a qualitative visual comparison on typical examples. SSIM has been designed for camera imagery with values between 0 and 255 and is therefore calculated on rescaled magnitude images.

4.3 Reconstruction Algorithm

4.3.1 Framework and Parameters

The reconstruction algorithm developed here extends the non-linear inversion reconstruction algorithm of section 2.3.2 by adding the optical flow estimation as a separate subproblem as described in section 3.2.2.

```

1  for t = 1:frame_count
2
3  //repeat frames on boundary
4  tw := current_time_window = [t-2:t+2]
5
6  inputs:
7  Y = data(tw), P = sampling_pattern(tw)
8  init:
9   $\rho = 1$ ,  $\mathbf{c} = 0$ ,  $\mathbf{v} = 0$ ,
10  $\eta = (\alpha_0, \beta, \beta_2, \gamma, \delta)$ ,
11  $n_v = 7$ , if opt. flow constraint
12     = 0, else
13
14  for n = 1:newton_steps
15
16  if n <  $n_v$ :
17  (d $\rho$ , d $\mathbf{c}$ ) = calc_image&coil_update( $\rho, \mathbf{c}, \eta$ )
18
19  else: //only when optical flow is used
20   $\mathbf{v} = \text{calc\_flow}(\rho)$ 
21  (d $\rho$ , d $\mathbf{c}$ ) = calc_image&coil_update(Y, P,  $\rho, \mathbf{c}, \mathbf{v}, \eta$ )
22
23  ( $\rho, \mathbf{c}$ ) = ( $\rho, \mathbf{c}$ ) + (d $\rho, d\mathbf{c}$ )
24   $\eta = \text{reduce\_parameters}(\eta, n)$  //  $\alpha \rightarrow \alpha/2$ 
25

```

Algorithm 4.1: Non-linear inversion reconstruction with optical flow constraint. The flow estimation subproblem alternates with the image reconstruction subproblem. At each time point t a whole batch of frames is considered with a repeating frame boundary condition (constant time derivative). Solvers for the subproblems in line 17, 20 and 21 are explained in the next section.

Using an optical flow constraint is optional in this framework and can be replaced by a ∂_t -operator (batch time difference), affine temporal regularization or no tem-

poral regularization. It is further possible to change the regularization type and the penalty norm in the subproblem solvers (section 4.3.2) between Newton steps.

When the optical flow \mathbf{v} is used, no temporal regularization is applied during the first newton steps ($n < n_v$) and the focus is entirely on deconvolving ρ and \mathbf{c} , corresponding to standard NLINV. In principle optical flow estimation and application from the first newton step on should work. For the first step ρ is the constant initial guess, flow estimation results in zeros and the flow constraint amounts to a batch time difference operation. Adding the optical flow constraint early, however, was observed to cause instability in later steps and it is therefore only enabled after 7 newton steps, whereas the affine or the batch time difference constraint start immediately.

One possible explanation for the instability is that during the first newton steps image and coil L_2 -norms vary strongly between steps, and so does the intensity distribution in both variables. Therefore a flow field estimated from ρ_n is likely not to fit in the computation of ρ_{n+1} in early newton steps.

Given a perfect optical flow estimate, only one further Newton step with the flow constraint enabled is necessary (Newton steps = n_v) to combine information of neighboring frames and remove undersampling artifacts. Without such an estimate the first optical flow computed will stem from images with strong artifacts and require several iterative refinements of the flow field with every Newton step beyond n_v .

In order to simplify the above algorithm to the case where the flow is known n_v is set to 0 and the flow estimation omitted. Likewise, if the coils are known the number of Newton steps is set to 1. As another minor variant, it is possible to initialize ρ and \mathbf{c} with the result of the previous frame (batch) starting from the second frame. Especially when the affine regularization is used, this may stabilize and speed up convergence.

Which parameter settings are best depends on the chosen reconstruction type. For Figs. 5.1-5.22 parameters are given in Tab. A.2-A.12 (appendix). For reference, parameters for joint image and coil reconstruction with different time constraints are given in Tab. 4.1.

Their exact choice seemed not to be critical and best values will vary with noise level, which depends on the data acquisition sequence. Usually, the regularization parameters in the iteratively regularized Gauss-Newton method are supposed to decrease with the newton step count. Yet, in the case of γ , iterative reduction merely led to gradual waning of the temporal regularization effect. Therefore it is kept constant. The flow regularization parameters β, β_2 and δ are also not reduced as they act on the image only in later Newton steps, where the image norm has leveled.

joint $\rho + \mathbf{c}$ with flow or ∂_t -op constraint:	Newton steps = 13, $\alpha_n = (1/2)^n, \gamma_n = 0.3 \quad \forall n$
joint $\rho + \mathbf{c}$ with affine constraint:	Newton steps = 7, $\alpha_n = (1/2)^n, \lambda = 0.9$
optical flow estimation:	
$\ \underline{\nabla \mathbf{v}}\ ^2$:	$\beta = 0.001,$ $\delta = 100$ (hard thresholding, eq. 3.12)
$\ \underline{\nabla \mathbf{v}}\ ^2 + \ \underline{\nabla \mathbf{v}}\ _1$:	$\beta = 0.0001, \beta_2 = 0.0001$
L_1 -transport & $\ \underline{\nabla \mathbf{v}}\ _1$:	no stable convergence

Table 4.1: Recommended parameters for joint image and coil estimation with optional optical flow constraint.

4.3.2 Numerical Functional Minimization

Along the lines of Dirks [41] and Burger et al. [42] this work used the Chambolle-Pock algorithm [50] for minimizing functionals (lines 17, 20 and 21 in Alg. 4.1). It solves problems of the form

$$\text{Find } \hat{x} \in \arg \min_x G(x) + \sum_i^m H_i(K_i x), \quad (4.1)$$

which include the image reconstruction and optical flow estimation problem. Here G and the H_i may be non-smooth penalty terms, such as an L_1 -norm and the K_i are continuous, bounded linear operators with Hermitian adjoint K_i^H . Originally, the algorithm did not focus on multiple constraints ($i = \{1\}$), yet by now suitable extensions exist ([51],[52],[53]) as recently reviewed by Komodakis and Pesquet [54]. The algorithm follows a primal dual approach - see Boyd [55] for a reference on duality. It is related to proximal point algorithms such as Douglas Rachford splitting (Eckstein and Bertsekas [56]) and sub-variants including the popular alternating direction method of multipliers (ADMM), a comprehensive introduction to which is given by Parikh and Boyd [57]. Instead of solving eq. 4.1 as one large problem, minimization is decomposed into a set of subproblems solved iteratively starting from an initial guess by following procedure:

for $k = 1$ until converged do

$$\forall i = 1 \dots m, \quad \tilde{y}_i = \text{prox}_{\sigma H_i^*}(y_i^k + \sigma K_i x^k) \quad (4.2)$$

$$y_i^{k+1} = 2\tilde{y}_i - y_i^k \quad (4.3)$$

$$\tilde{x} = \text{prox}_{\tau G_i}(x_k - \tau \sum_i^m K_i^H y_i^{k+1}) \quad (4.4)$$

$$x^{k+1} = 2\tilde{x} - x^k. \quad (4.5)$$

Eq. 4.2 holds the m solutions to aforementioned subproblems, each referring to a different penalty term H_i . These intermediate solutions are stored in the *dual* variables \tilde{y}_i . And the *primal* objective x is the variable of interest. Evaluation of the \tilde{y}_i as the solution to

$$\text{prox}_{\sigma H_i^*}(y_i^k + \sigma K_i x^k) := \arg \min_{\tilde{y}} \sigma H_i^*(\tilde{y}) + \|\tilde{y} - (y_i^k + \sigma K_i x^k)\|^2 \quad (4.6)$$

is called the dual problem, where the “prox”-operator is a quadratic relaxation of the minimization problem of H_i^* . H_i^* itself is the (Fenchel) conjugate of H_i defined by [57]

$$H_i^*(y) := \sup_{x \in \text{dom} H_i} \langle y, K_i x \rangle - H_i(K_i x). \quad (4.7)$$

An important advantage of the dual approach beside problem splitting is that the operators K_i appear only with the parameters not the variables. The prox evaluation (eq. 4.6) therefore will not entail further matrix applications or the need to perform matrix inversions, which sometimes might be hard or impossible. Of course, this operator splitting is of use only if evaluation of eq. 4.7 is simple. For some of the most frequent cases look-up references exist [57],[41]. A list of the prox functions required for this work is given in the appendix.

As in gradient based methods, convergence rates depend on the step width denoted here by σ for the dual ascent step in eq. 4.2 and τ for the primal descend step in eq. 4.4 respectively. Convergence of the Chambolle-Pock algorithm is guaranteed [52] when

$$\sigma \tau \left\| \sum_i^m K_i^H K_i \right\| < 1, \quad (4.8)$$

thus $\sigma = \tau = 1/(1.1 \cdot \left\| \sum_i^m K_i^H K_i \right\|^{1/2})$ is used here, and the summed operator norm is computed using power iterations.

Lastly, eq. 4.3 and 4.5 are an optional inexpensive addition to speed up convergence in many practical cases [50].

Similar to ADMM, the algorithm converges slower in most problems than the classical conjugate gradient method. Its main benefit lies in the plug-and-play fashion in which regularization terms can be exchanged by simply replacing the prox function, which is typically one line of code. Proposed ways to speed up convergence include using a vector of step widths with an entry σ_i for each operator K_i instead of an average step width [58], computing an optimal step size with a line search [59] or adapting the step size automatically and iteratively based on the evolution

of the residual [60]. Most methods require setting additional parameters and while noticeable speed-ups were observed here in some examples, divergence occasionally occurred in others. Therefore the basic Chambolle-Pock algorithm was used.

Apart from the step size, reconstruction runtime is also determined by the iteration count. That in turn depends on the stopping criterion. A common choice is to stop optimization once the functional residual norm has fallen below a fraction of the norm of the initial residue or once the residual norm is no longer changing by more than a preset fraction with respect to the previous iteration.

Here, however, a fixed number of iterations was used eventually to terminate optimization, which may have allowed more iterations than were necessary. In the case of image reconstruction with an optical flow constraint for instance, minor model violations (through-plane motion) occasionally caused a slight increase of the norm of the total residual (root summed squares of residual norm of all penalty terms). This would lead to termination with a relative norm reduction criterion while the images were still gaining sharpness, bringing them closer to the reference. A fixed number of iterations chosen through a parameter grid search and subsequent comparison between reference and reconstruction prevented premature termination and simplified comparisons between reconstructions with different regularizations.

Concerning the other temporal regularizations, termination as soon as the residual norm has changed by less than 0.1% with respect to the previous iteration seemed a sound criterion. Nevertheless, for sake of consistency again a fixed number of iterations chosen after a parameter grid search was used.

4.3.3 Operator Discretization

The discretization scheme chosen here, largely coincides with that adopted by Dirks [41] and Burger et al. [42]. For the time derivative operator ∂_t a forward finite difference is used and correspondingly a backward difference for the adjoint operator ∂_t^H (three frame time window example)

$$\partial_t = \begin{bmatrix} -1 & 1 & \\ & -1 & 1 \\ & & 0 \end{bmatrix}, \quad \partial_t^H = \begin{bmatrix} -1 & & \\ 1 & -1 & \\ & 1 & 0 \end{bmatrix}.$$

And the same discretization applies for the spatial gradient operators in the flow smoothness penalty $\|\underline{\nabla}\mathbf{v}\|^p$ during optical flow estimation. A different choice is made, however, with regard to the gradient discretization inside the transport term $\|\partial_t\rho + \mathbf{v}\nabla\rho\|^p$, where they propose to keep the forward difference operator for time but use a centered gradient in space. Without any objections regarding

the ∂_t -discretization an upwind scheme [61] is used here instead of the centered gradient.

In its basic form it is an average of a forward difference operator in the current frame and a backward difference in the next frame or vice versa, depending on the direction of the optical flow. Same as for the centered gradient, the advantage over a conventional forward or backward difference operator is the increased spatial support. Coverage of a larger area per time step means larger time steps, or equivalently, faster motion is better captured. The centered difference, however, still has a smaller support than the upwind gradient, which considers the starting position of an edge in one frame and its end position in the next. This also adds to accuracy on sub-pixel sized shifts.




Figure 4.4: Noiseless in-plane motion phantom during circle expansion subjected to different gradient functions. The upwind gradient has the largest support in fast motion increasing the support area for the optical flow.

Dirks and Burger pointed to the advantages of the upwind scheme yet observed artifacts when using it. Artifacts in the form of a blow up of optical flow values were also observed here, but could be completely eliminated by adding a Gauss filter with width 5 and standard deviation 1 to the transport term evaluation (and only there) in both the image- and flow estimation problem (eq. 3.9-3.10).

The classical approach proposed by Horn and Schunck [7] - not considered here - used centered cell gradients, which is robust but according to Dirks [41] yields less accurate flow fields.

Results

This chapter is divided into two parts:

- The first part compares different temporal regularizations both quantitatively and qualitatively and therefore assumes that coils are known. Optical flow fields are always estimated from images with the full 75 spoke data set. It proceeds by **(i)** comparing the performance of all considered temporal regularization types in a phantom compliant with the optical flow model and without noise, **(ii)** considering the influence of noise in the same phantom, **(iii)** addressing the two most likely types of optical flow model violation in phantom data - long-range in-plane motion and through plane motion - and **(iv)** comparing reconstruction types in human heart data as well as assessing to what extent the data corresponds to an optical flow compliant phantom.
- The second part deals with the case where the coils need to be reconstructed jointly with the image and the flow field needs to be estimated from under-sampled instead of fully sampled data.
- A brief summary is given at the end

5.1 Image Reconstruction with Known Motion

In this part the receive coil sensitivities are assumed known and optical flow fields are estimated from fully sampled data. It will be shown that an optical flow field can be estimated perfectly in the absence of noise and still reasonably well when noise is present. If the measured object is compatible with the motion model application of the flow field can markedly improve image reconstructions from undersampled data.

Object motion that violates the optical flow model will cause the flow field estimation to converge at a slower rate and to a shallower minimum, possibly causing blur upon application in image reconstruction.

5.1.1 Optical Flow Compatible Data

Noiseless Data

In order to compare the effect of different temporal penalties in the cost functional, the in-plane phantom with 75 spokes in each frame has been reconstructed as a reference without added noise. Reconstructions of a five-fold undersampled (15 spokes) noiseless version of the phantom (Fig. 5.1) that were compared against the reference used (i) no temporal penalty, (ii) an affine temporal con-




Figure 5.1: Upper left corner of the in-plane phantom (fastest area). a) Reference with 75 spokes and b-h) different reconstructions. Reconstruction parameters see Tab. A.2.

straint (eq. 3.1), (iii) a batch time difference constraint (eq. 3.2) and (iv) the optical flow constraint (linear version of eq. 3.9). In addition, the effect of a temporal median filter as well as a sliding window average in post processing was also tested. Adding a temporal regularization or a filter improved over no regularization in every case, with the optical flow model unsurprisingly performing best in this phantom. Apart from slight ring-like artifacts (Fig. 5.1b) it could well reproduce the reference.

The temporal median filter preserved the shape of the circles surprisingly well (Fig. 5.1d) while the affine regularization, sliding window averaging and batch- ∂_t regularization introduced visible motion blur. The difference between the affine regularization and the rest is mostly that regularization is with respect to a history of up to five frames - the pattern repetition period - versus a temporally centered combination of frames in the other methods.

For a more quantitative analysis SSIM values were computed from the magnitude part of each complex image for the image itself (Fig. 5.2a) and the temporal derivative to the next image (Fig. 5.2b). The derivative is to show the dynamic effect of the regularization.

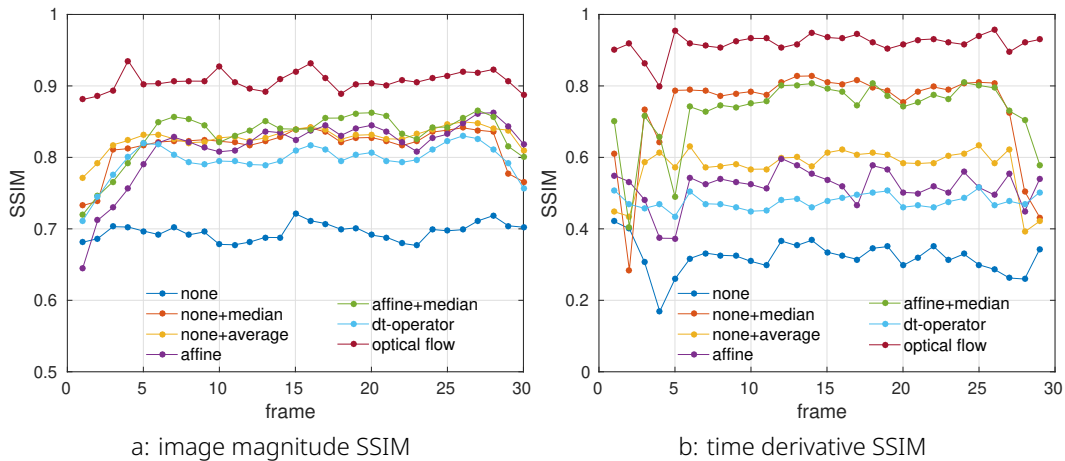


Figure 5.2: SSIM values from a 30 frame sequence of the noiseless in-plane phantom for (a) image magnitude and (b) image magnitude time derivative. Reconstruction parameters see Tab. A.2.

As each regularization type converges slightly differently parameter settings are not identical between different reconstructions. For each regularization type a parameter grid search was performed and parameters selected according to the best outcome by SSIM value. In case similar results were reached with different parameters the smaller regularization weight and iteration count were chosen as best parameters.

Interestingly, even though the simpler and cheaper sliding average scores a higher SSIM than the batch- ∂_t regularization, its corresponding image is blurrier. This may be attributed to the fact that all frames in the time window are combined

with equal weights whereas more emphasis is put on the central frame especially at higher spatial frequencies in the regularized approach, resulting in less motion blur.

Looking to the dynamic behavior of each reconstruction (Fig. 5.3) the blur becomes more apparent as well as the fact that motion may be obliterated by the median filter.




Figure 5.3: Upper left corner of the time derivative of the in-plane phantom image magnitude, absolute intensities windowed identically among subfigures. Reconstruction parameters see Tab. A.2.

Parameter Influence on Optical Flow Estimation

Adding an optical flow constraint enables deeper convergence in the image reconstruction problem of the numerical phantom albeit at a slow rate. In the absence of noise one could go to thousands of iterations and still observe a decline of the residual (Fig. 5.4a, blue curve). Apart from the number of iterations in the image reconstruction problem, the number of iterations spent in the flow estimation problem also influences the resulting image quality and again, convergence is rather slow (Fig. 5.4b). The optical flow is initialized with zeros and in the ex-

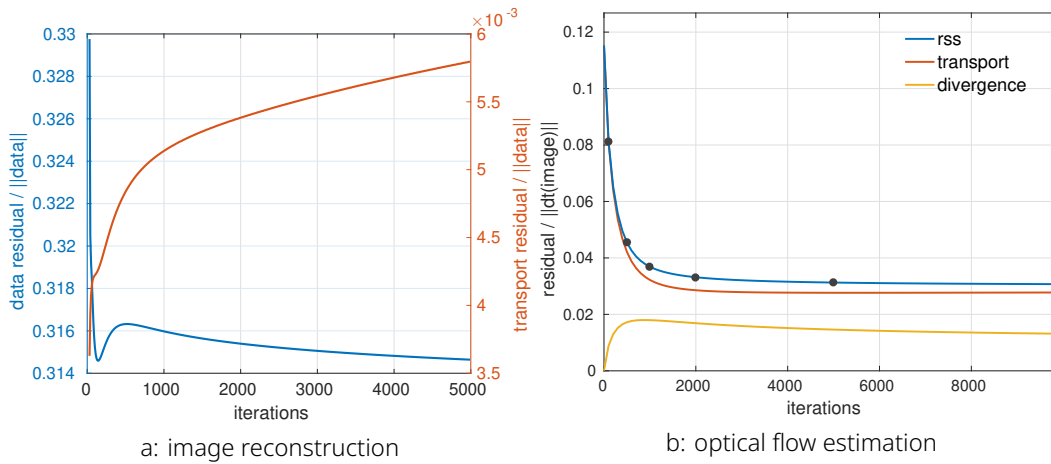


Figure 5.4: Residual evolution over iteration count. The residual is normalized by the norm of a) the data norm and b) the time derivative so the main residual obtained with the initial guess for the image and optical flow (ones and zeros, respectively) evaluates to one. The iteration axis starts at four for better visibility, so the y-axis does not go up to one. Black dots in b) mark the flow fields displayed in Fig. 5.5. Parameters see Tab. A.3.

treme case of terminating after zero iterations subsequent image reconstruction amounts to the same as batch- ∂_t regularization. Ring artifacts such as in Figs. 5.1b and 5.3b appear as a result of too low optical flow values and vanish with more iterations of the flow estimation. These rings are no Gibbs-ringing. Rather, the spatial gradients (edges) of the circles which are used to model the time derivative are amplified with the wrong optical flow amplitude. In the presence of noise early termination after 1000 iterations - ≈ 1 sec computing time for 384^2 px - usually sufficed to keep artifacts below noise level.




Figure 5.5: Horizontal component (v_x) of optical flow. Flow amplitude is build up successively in the first few hundred iterations while the effect of the smoothness penalty materializes only in later iterations, extending the spatial reach of the flow field. Parameters see Tab. A.3.

Plugging the flow field (Fig. 5.5) into $\partial_t \rho = -\mathbf{v} \nabla \rho$ shows how well the time derivative can be approximated by the spatial derivative (Fig. 5.6).




Figure 5.6: Transport term residue computed with flow fields from Fig. 5.5 and the ground truth reference image during circle contraction. At too few iterations the residual resembles the time derivative. Parameters see Tab. A.3.

In the shown frame circles are contracting therefore the temporal derivative is a subtraction of intensity from the past to the present frame (black signal = negative values). With too few iterations the transport residue qualitatively still looks like the temporal derivative. Only with sufficient iterations does it approach zero and the ring artifact vanishes (Fig. 5.7).




Figure 5.7: Images reconstructed from undersampled data with the optical flow constraint and the flow fields from Fig. 5.5. Insufficiently converged flow fields cause artifacts and slow convergence in the image reconstruction problem. Parameters see Tab. A.3.

Convergence rates are determined by the divergence weight β in the flow estimation problem and the transport weight γ in the image reconstruction problem. $\beta \approx 0.001$ led to stable convergence in all tested examples, while larger values significantly slowed down convergence and values < 0.0001 could result in blow ups. Without noise, image reconstruction was stable on the interval $\gamma \in [0.03, 0.5]$, where values on the upper end enforce stronger inflow of image information from neighboring frames into the current frame. In this example, however, results with different γ values differed little.

Noisy Data

The ground truth optical flow field from noiseless images also performs best when applied to noisy data (Fig. 5.8). Among all other reconstructions the rank order of SSIM values changed with respect to Fig. 5.2 with the introduction of noise. Notably, batch- ∂_t regularization on average now ranks second.

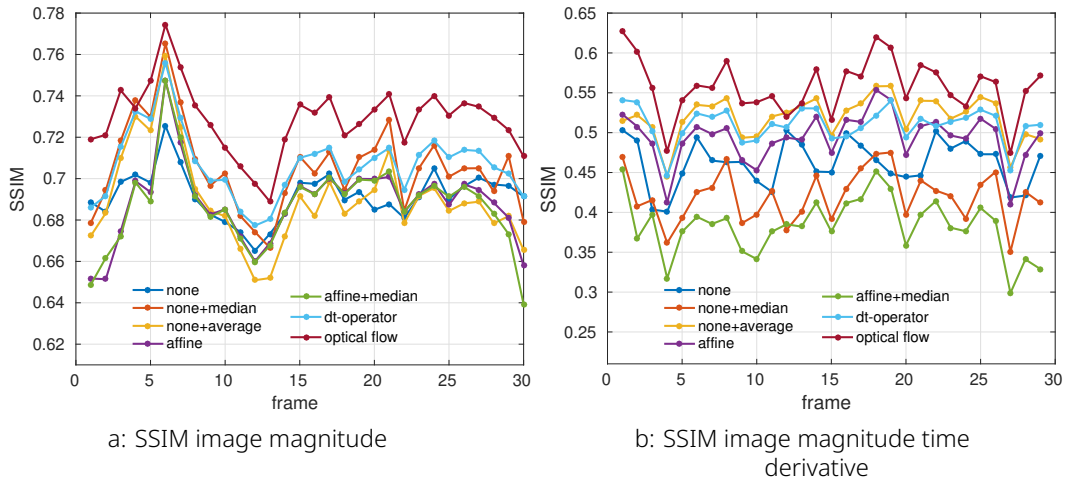


Figure 5.8: SSIM values representing the similarity score with respect to the noisy reference contrary to Fig. 5.2, which results in overall lower similarity since reconstructions do not recreate noise accurately. Parameters see Tab. A.4.

In practice the flow field obtained from noisy images will carry that noise itself even in areas without motion. This might raise the concern that artifacts in those areas multiplied by the noisy flow field may not get averaged out as a consequence. Noise suppression in the flow field by additional regularization can be realized by regularizing the smoothness term $||\nabla \mathbf{v}||$, transport term or both terms of the cost functional in the L_1 -norm. This is a popular method for sparse signal recovery that is well suited here as the flow field, temporal and spatial gradients are all sparse. However, pure L_1 -regularization may yield patchy textures, thus a mixture of an L_1 and L_2 penalty can be a good compromise between patchiness and noisiness. Alternatively, an additional term can be introduced to the cost functional masking all values where the image magnitude and its temporal gradient lie below the estimated noise level corresponding to hard thresholding (eq. 3.12).




Figure 5.9: (Top) Optical Flow fields estimated from noisy reference in-plane phantom (upper left corner shown) and (bottom) phantom with added noise, reference and various reconstructions. Parameters see Tab. A.5. Images reconstructed with L_1 -optical flow were visually indistinguishable from those reconstructed with other flow fields.

Yet, despite the change of appearance of the flow field under different regularizations the resulting reconstructed images and their dynamics change little.

With regard to noise, the choice of flow regularization has little effect on the resulting SSIM values (Fig. 5.10). Hard thresholding performs a little worse. The noise present in the simply L_2 -regularized optical flow field seems not to interfere with the averaging of undersampling artifacts across the reconstruction time window. Also, as in the noiseless case, image reconstruction was stable with the transport term weight $\gamma \in [0.03, 0.5]$ and results differed little over this interval.

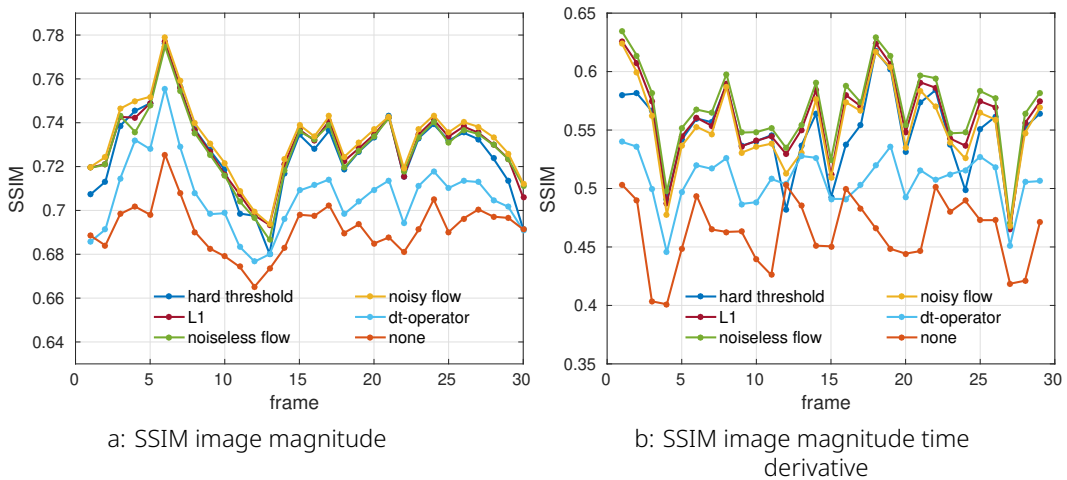


Figure 5.10: SSIM values for image reconstructions using optical flow fields obtained with different regularization types (see Fig. 5.9), batch- ∂_t - and no temporal regularization. Parameters in Tab. A.5

As soon as noise is included in the data, there is no truly objective reference anymore. The reference compared to depends on the choice of the Tikhonov-penalty weight α which trades noisiness against blur and has to be decided once upon visual inspection (Tab. A.4 for the noisy phantom).

5.1.2 Optical Flow Model Violation

So far the tested data has been compliant with the optical flow model. Real data on the other hand is expected to violate the model during through-plane motion or sudden long-range motion reaching beyond the overlap region of the spatial gradients between subsequent time points. The former is necessarily seen in heart imaging and interactive applications were the imaging plane is moved during measurement. The latter may occur in speech studies with rapid tongue motion.

Fast, Long-range Motion

The long range motion phantom (Fig. 5.11) is identical with the noiseless in-plane phantom except that the pulse rate has been doubled and the pulse amplitude

increased by 60%. As a result, the time difference between neighboring frames covers a larger area than the spatial gradients of those frames which leaves an area where - conflicting with the continuity equation - no linear relationship between the gradients can be established and the optical flow is reduced to zero by means of the smoothness penalty. Those areas experience averaging as if a simple finite difference operator were used.




Figure 5.11: Noiseless in-plane phantom with doubled pulse rate and 60% increased pulse amplitude. (top-left) reference image and (top-mid and -right) reconstructions with optical flow constraint (identical iteration count in image reconstruction problem and different iteration count in flow estimation problem). (bottom-left) time difference, and (bottom-mid and -right) horizontal components of the flow field (v_x), estimated with a simple L_2 -penalty. The top middle reconstruction has more rings than the shown flow field because of the flow fields in neighboring time points within the same batch. Parameters see Tab. A.6

Thus, in principle, the optical flow should not do worse than the batch- ∂_t regularization on fast objects. In practice, this is also true but convergence is slowed down further such that more iterations become necessary or severe artifacts appear even after 1000 iterations (Fig. 5.11 top-middle).

Through-plane Motion

Through plane-motion has an interaction with noise. Wherever an inflow of intensity can not be explained by anatomic edges the model tries to amplify noise to explain the inflow. This is also true in the absence of actual noise, when only minimal numeric noise is present.

As in the model violation by fast motion, convergence of the flow estimation is slower than when the model is satisfied. The attempt to explain through-plane motion by noise leads to amplified noise artifacts in the image reconstructed from undersampled data. The more iterations are run, the better the noise fits the motion. However, the blown up flow values after 5000 iterations (Fig. 5.12 middle)

cause problems during application of the flow field in image reconstruction, resulting in blurrier images.


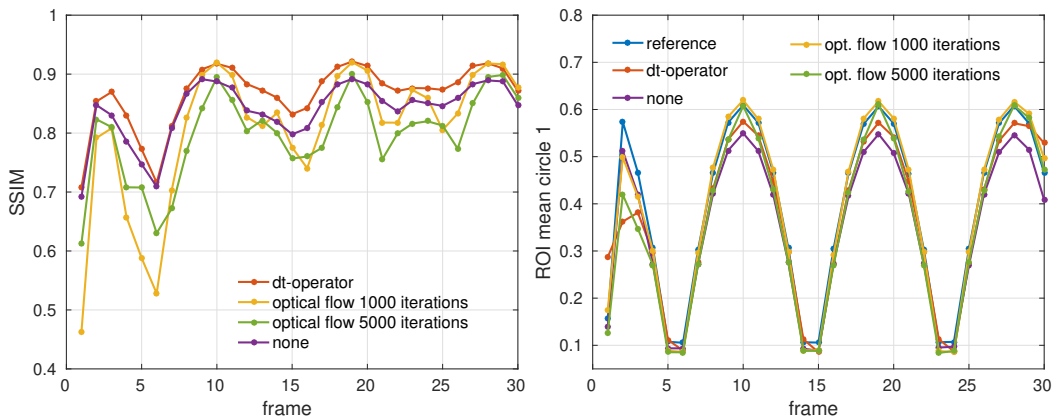


Figure 5.12: Noisy Through-plane motion phantom (real part of upper left corner shown, fastest area). Reference, reconstructions and flow fields with negative values in black. For better flow field contrast to noise ratio, flow fields were reconstructed with a hard thresholding constraint (eq. 3.12). Parameters see Tab. A.7.

SSIM values of those reconstructions (Fig. 5.13a) are lower than where batch- ∂_t regularization is used. In the latter case the images look exactly as the reference yet suffer from motion blur. The temporal fidelity of each reconstruction method is assessed in Figs. 5.13b-5.13d, by comparing region of interest average intensities in three different circles (upper left, middle left and lower right circle) each with different amounts of intensity inflow over time. Differences between reconstructions become apparent at points where the through-plane motion changes direction, rather than at zero-crossing points, which is why the magnitude is shown. The noise artifact in the optical flow based solution with 1000 iterations actually has ROI-mean values closer to the reference than the batch- ∂_t regularization based solution which suffers larger motion blur.

In summary, optical flow will - contrary to the case of too fast motion - not perform as well as batch- ∂_t regularization if too much through-plane motion takes place. This example, however, is to be understood as a limiting case.



a: SSIM values through-plane motion phantom (real-part)

b: Time course of magnitude signal ROI average upper left circle

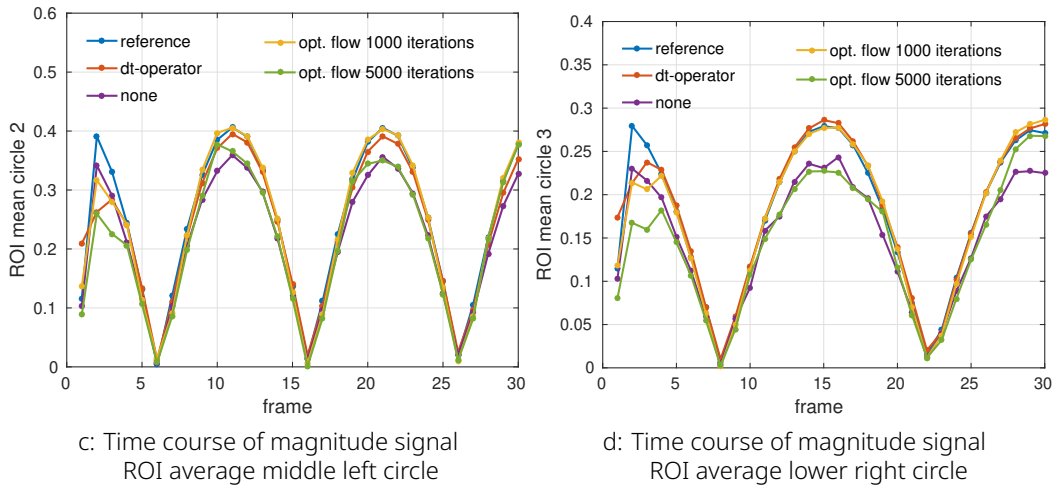



Figure 5.13: SSIM values and mean intensities of three different circles in Fig. 5.12. The real part of the phantom was used. Magnitude curves are shown in 5.13b-5.13d to highlight the motion turning points. Parameters see Tab. A.7.

5.1.3 Applicability to Real Data

While through-plane intensity inflow will be absent or negligible in most head imaging applications, which are less open to emulation by surrogate data, it is always seen in heart imaging. The extent varies depending on the imaging slice orientation. To gain insight into how well this is dealt with by the optical flow model the transport term residual is shown for a short axis-, four-chamber- and two-chamber-view (Fig. 5.14).

In the upper row a time point is shown where contracting motion of the myocardium is exclusively in-plane, while both lower rows also feature through-plane motion from (mid) in- and (bottom) out-flowing blood. Yet, whereas in the mid row inflow






Figure 5.14: Reference magnitude images, time differences of the image real part and transport residuals for (top) a short axis (middle) four-chamber and (bottom) two-chamber view of the heart. Shown image rows from top to bottom correspond to frames 13, 7, 10 in Fig. 5.15. Time difference images and transport residuals are shown after application of a Gauss-filter (width = 5, SD = 1) used by default for stability before estimation and application of the optical flow. Transport residuals have been computed using an optical flow with 5000 iterations and hard thresholding (not shown). Ideal model compliance is achieved were the residual stays below the noise level. Parameters see Tab. A.8

cannot be related to edge displacements by the flow field, the bottom row residual shows less extreme values despite considerable outflow in the aorta. The spatial distribution of inflowing intensity appears to be at least as important as the absolute amount.

In order to place the surrogate data examples between the two extremes of perfectly compliant motion and motion incompatible with the optical flow model the L_2 -norm of the transport residual for each data set is plotted with the residual norm of the through-plane and in-plane phantom normalized by the norm of the respective time derivative (Fig. 5.15). A value close to one indicates pure through-plane motion as the time derivative could not at all be related to the object edges. Likewise a value close to zero means all motion was in-plane and short-range. Due to noise neither extreme will be reached in practice, but the in-plane and the through-plane phantom are clearly separated from each other.

The error-bars in Fig. 5.15 quantify the motion SNR of each frame. That is, the signal to noise power ratio of the time derivative restricted to the support of meaningful (above noise level) values of the time derivative. A support mask is conveniently provided by the hard-threshold regularized flow field. In essence, time points with large error-bars belong to frames where less motion above noise level appears in the time derivative or transport residual and uncertainty about whether it is actually through-plane motion or just noise is higher. The reciprocal motion SNR values serve as weights in the weighted averages denoted by solid lines.

Time points tend to have higher values when the transport residual is mostly noise. The curves in Fig. 5.15 are, however, not random but follow the heart beat cap-

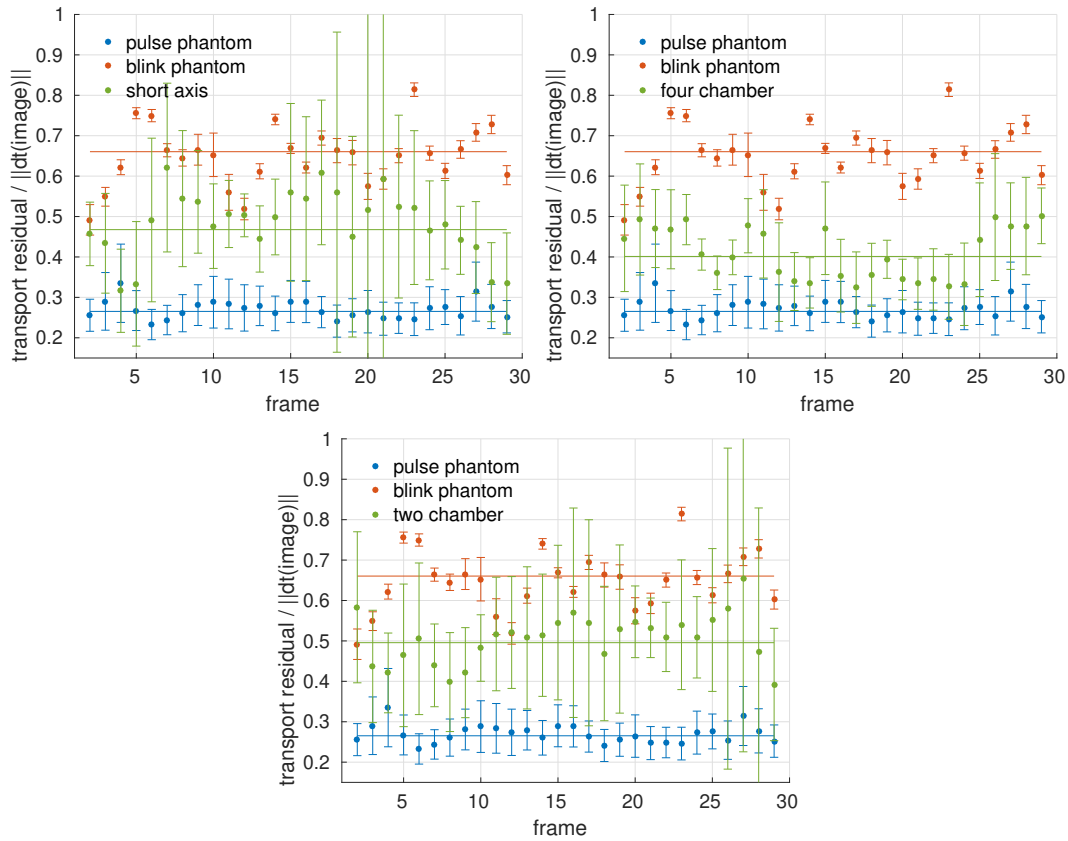


Figure 5.15: L_2 -norm of the transport residual divided by L_2 -norm of the time derivative for (left) short axis (middle) four-chamber and (right) two chamber heart series, shown with the in-plane and through-plane phantom. Error-bars are computed as the signal to noise power ratio (“motion SNR”) over the time derivatives support (above noise level). Solid lines denote the weighted average over time, with the reciprocal motion SNR as weights. Flow fields have been computed with the hard threshold regularization (eq. 3.12) and 5000 iterations. Parameters see Tab. A.8.

tured in the series of 30 frames. In the left plot of Fig. 5.15 (short-axis-view) a beat starts at the 4th frame with in-plane motion followed by heart expansion - inflow and model violation - and ends at the 28th frame. A similar pattern exists in the other examples. Apart from the four-chamber-view, the surrogate data behaves like a half-half mix of both phantoms.

Again, as in the phantom data (Fig. 5.10), the particular choice of the flow-field regularization - L_2 (+ hard thresholding) or L_1 - seemed not to affect image quality or temporal fidelity much, judging from SSIM values (not shown).

Upon inflow significant signal remains in the transport residual, irrespective of the flow regularization. As in the phantom data this may lead to through-plane artifacts in real human data (Fig. 5.16, compare mid- and left-bottom). The artifact - which might be misinterpreted as an actual myocardial dysfunction - comes along with blur due to a slower convergence rate.




Figure 5.16: (Top) horizontal component of flow fields with different regularization, 7th frame of the four-chamber-view as in Fig. 5.14. (Bottom) Reference and reconstructed image using the L_2 -optical flow constraint. The reconstructed image has a through-plane artifact where the flow field is abnormally bright. The frame is taken at a time point of blood inflow into the (anatomically) left ventricle shown on the right. Parameters see Tab. A.9.

It is also possible to track the time interval where inflow occurs in a plot of SSIM values. The through-plane artifact seen in Fig. 5.16 occurs in the 7th frame of the series in the middle of an interval where the blue curve of SSIM values in Fig. 5.17a takes a dip when the data no longer fits the model.

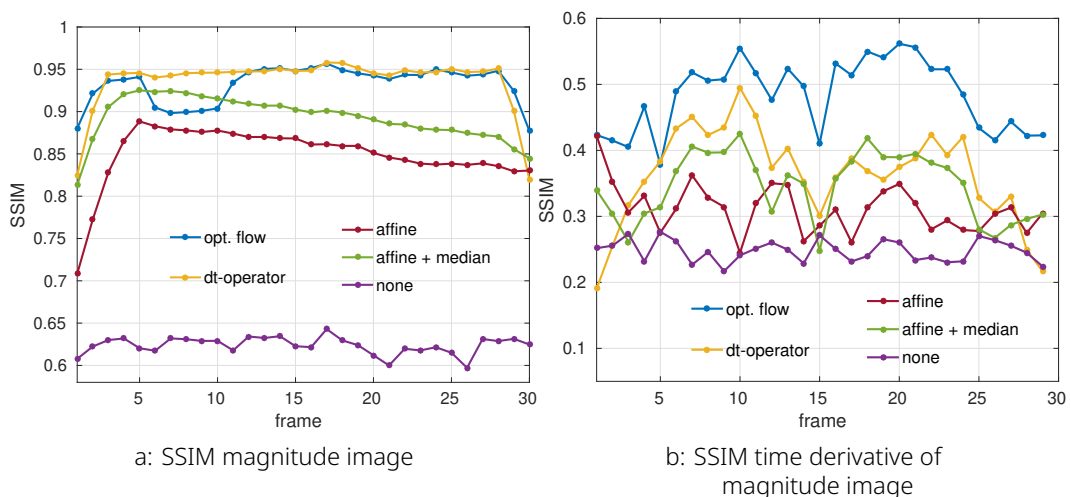


Figure 5.17: SSIM values of four-chamber-view data for several temporal regularization types. The optical flow used again has 5000 iterations. A dip in the blue curve of Fig. 5.17a marks an interval with blood inflow into the left (large) ventricle in violation of the optical flow model. Parameters see Tab. A.9.

The drop in SSIM values is not due to the artifact itself but due to the blur and incomplete deconvolution of the coils from the anatomic object. In the time derivative of the image magnitude which is unaffected by the global intensity distribution the optical flow reconstruction still scores highest throughout all frames. Partly, because a blurry image is less noisy. Blur, however, only occurs in the dip interval. The main reason explaining the score is better handling of in-plane motion. This is further exemplified in Fig. 5.18 in magnitude difference images between reference and reconstructions. During inflow, the batch- ∂_t regularization keeps spatial sharpness and suffers motion blur, where optical flow based regularization degrades spatial sharpness yet better keeps temporal fidelity. The affine temporal regularization + median filter has the lowest temporal fidelity, strongest background artifacts and a different intensity distribution above the heart (dark area). During in-plane motion little difference is seen between reference and the reconstruction with optical flow constraint, while batch- ∂_t regularization ranks second and affine regularization clearly third.




Figure 5.18: Difference between reference magnitude image and different reconstructions unrelated to Fig. 5.17. Through-plane motion in the top row causes a temporal delay in the ∂_t -operator based reconstruction, seen as a bright residue on the heart in the difference image (top left). The residue on the heart is smaller in the optical flow constraint reconstruction. However, violation of the optical flow model impedes convergence reducing image sharpness in the reconstructed image. Therefore, all edges in the difference image (top middle) appear bright. Conversely, the optical flow model is fulfilled in the bottom row which shows the heart during in-plane motion and the residue in the difference image (bottom middle) is minimal. Parameters see Tab. A.9.

5.2 Motion Estimation in the Presence of Artifacts

Up until now coil sensitivities have been assumed known and optical flow fields were computed from images with fully sampled data sets. In the presence of undersampling artifacts, object motion may no longer be the dominating feature in the time derivative of the image. Artifacts, despite their rapid, discontinuous dynamic, will still enter the optical flow field and visibly distort object movement when applied during reconstruction. The artifacts are smoothed out by the smoothness penalty in the flow estimation problem but retain their rotating dynamic. After applying the flawed flow field, the artifact no longer looks like streaks moving on top of the object. Rather the object itself now looks as if its walls were “wobbling” because the streaks are frozen in areas with flat intensity gradients and only move near object edges. Repeated flow estimation and application on the same image gradually improves both flow field and image (Fig. 5.19).




Figure 5.19: In-plane phantom (top row) fully and undersampled (2nd row) corresponding horizontal L_2 -regularized flow fields (3rd row) NLINV reconstruction with flow constraint (bottom row) flow re-estimation from reconstructed image (3rd update) and resulting reconstruction. Parameters see Tab. A.10.

The gradual improvement seen in the phantom's flow field eventually even eliminates noise and leaves no trace of artifacts. That may be owed to the phantom's geometric regularity and homogeneous intensity. From the very beginning the motion of the phantom stands out from the artifacts. In anatomic images with bright fatty tissue and less regular intensity distribution that rarely ever holds. In fact as shown in Fig. 5.20, even though several flow estimation and application cycles result in significant artifact reduction some will persist - especially in bright, fatty tissue - irrespective of the cycle count. And likewise does the "wobbling" remain prominent (time difference Fig. 5.20).




Figure 5.20: (Top row) NLINV reconstruction with flow constraint (again 7th frame) and time derivative after 15th Newton step, (2nd row) 0th and 7th optical flow update (horizontal component) started after 7th Newton step, (3rd row) flow with $||\nabla \mathbf{v}||$ regularized by L_2 , L_2 + hard threshold and $L_1 + L_2$, (bottom) flow with all penalties in L_1 -norm. Flow fields are windowed identically. Parameters see Tab. A.11.

Not all regularization types performed equally but none of them did satisfyingly. The L_1-L_1 total variation constraint as proposed by Li et al. and Dirks has also been tested here. Indeed, the support of non-zero values it leaves is the smallest of all but still not small enough. More importantly, with L_1-L_1 , image reconstruction diverged in a few frames of the series. No set of parameters could be found in the grid search to yield reliable convergence for the entire series.

Experiments with further added regularization terms in the flow estimation problem, tailored to reduce the flow values where the time difference values are large before and small after application of the flow field in several variations (eq. 3.13), also turned out unsuccessful.

In summary, reliable discrimination between artifacts and motion near edges or areas of extreme contrast using information from the time window of a single repetition period so far has failed no matter the regularization. A possible workaround, still under investigation at the time of writing, is to add another frame to the time window such that two interleaving sets of full patterns can each be combined into time blurred full frames and their magnitudes be subtracted (Fig. 5.21).




Figure 5.21: (Top) Data acquired with a pattern repetition period of five combined in a sliding window average before first estimation of optical flow (e.g. 7th newton step, complex images shown) forms full frames (middle row) without artifacts but with motion blur. The difference of two consecutive full magnitude images has much less artifacts in static areas and still contains a temporally coarse-grained picture of motion within the batch (bottom right, non-averaged reference time difference, spatially blurred for feature visibility).

The time window considered would span one frame beyond the pattern repetition period, thus artifacts in areas which are static throughout that interval could be better recognized. On the other hand, artifacts do not vanish completely as seen in the block average time difference in Fig. 5.21. Likewise, motion which lasts for a mere two or three frames might be averaged to noise level and mistakenly removed. Whether or not such issues arise in practice and impede the success of such an approach needs further investigation.

Until the open issues of inflow and artifact propagation through the flow field are resolved, the batch time difference regularization stands as the best reconstruction mode. It is easily integrated into the joint non-linear inversion reconstruction of image and coils and offers higher temporal fidelity than standard affine regularization + median filter (Figs. 5.17, 5.18).




Figure 5.22: Four-chamber-view reconstructed by NLINV with (left) batch- ∂_t constraint and (right) affine temporal regularization. Images have same relative intensity windowing. Absolute intensities differ due to different regularization and resulting coil estimation. Therefore a comparison to a common reference is not possible. The affine regularization has a background artifact. Parameters see Tab. A.12

Apart from temporal fidelity loss, affine regularization results in background artifacts (bright haze, Fig. 5.22) that build up more prominently during joint reconstruction than when the coils are a known parameter. The artifact grows over several frames as a result of inconsistent superpositions of image projections between frames while the object moves: Constraining the current frame to look like the previous one by adding $\alpha \|\rho_t - \lambda \rho_{t-1}\|$, $\lambda \in [0, 1]$ to the cost function does not result in an equal weighting of projections. The data of the current time point will always have a higher weight than that of the previous time point, even if $\lambda = 1$ yet especially when $\lambda < 1$. As a consequence, the projections will not cancel out to form a black background. In NLINV $\lambda \leq 0.9$ (here 0.9) must be set to prevent the haze from growing indefinitely.

In brief:

In a comparison between the established simple affine temporal regularization, optical flow- and batch time difference regularization the latter performs best in general on human data.

Adding an optical flow constraint to the cost function demonstrably improves the preservation of in-plane motion. In its traditional sourceless form, however, the model may create artifacts or hamper overall convergence in heart imaging - blurring the output images - when through-plane motion occurs.

Optical flow fields estimated from undersampled data may also be heavily distorted by undersampling artifacts. Subsequent application of flawed flow fields then results in visible distortions of object movement. As long as intensity change in time is dominated by anatomic motion rather than artifacts an iterative cycle of flow field estimation and image quality improvement through flow field application may succeed in eliminating artifacts. However, this is often not the case in real human data with high intensity signal (fat tissue). Additional regularization did not help to purge artifacts from the flow field.

The SSIM scores of the batch time difference constraint consistently exceed those of the affine regularization in human data. While the affine term incorporates a long history of T frames (the pattern repetition period) strictly from the past into the current time point, the ∂_t -operator combines data of a batch centered around the current time point, yielding higher temporal fidelity. In addition, affine regularization itself without the median filter suppresses artifacts inefficiently, allowing static undersampling artifacts - which cannot be removed by the filter - to build up in the image background.

Discussion

Comparison of Temporal Regularizations

This work compared and rated two reconstruction schemes for real-time MRI: Minimization with respect to a single frame under an affine temporal regularization constraint - the established approach - and simultaneous minimization of a batch of frames centered around the current time point. The latter was found to outperform the former in human data in terms of temporal fidelity (Figs. 5.17, 5.18) and, notably, artifact suppression in the image background (Fig. 5.22). It also does not rely on an additional post-processing temporal median filter that may distort or destroy rapid motion that lasts for half the median filter width (Fig. 5.3).

Amongst the two batch regularized reconstruction types the one constrained by optical flow delivered higher temporal fidelity with regard to in-plane motion and even worked on far-reaching motion. An advantage that is confounded, however, by the method's limited ability to process through-plane motion.

Human Data Compliance with Optical Flow

Even though optical flow constraint cardiac MRI reconstruction has been studied before this issue has never been raised. Li et al. [8] (see section 3.2.2) made a qualitative visual comparison to affine regularized NLINV without showing flow fields. Zhao et al. [44] used an offline approach with a golden angle acquisition and compressed sensing reconstruction of the entire data stack at once. Their reconstruction was (by necessity) done post-acquisition with known coils and lower undersampling factors of 9-12. Comparisons to other motion models or the estimated flow fields were not shown. Burger et al. [62] who studied CT image reconstruction used an experimental phantom with purely in-plane motion.

Heart data, however, has been shown here to lie midway between pure in- and pure through-plane motion (Fig. 5.15). Through-plane artifacts such as in Fig. 5.16 thus are a likely occurrence that requires a solution if radiologists are to trust the images.

A possible route is to add a source to the continuity equation as an additional variable with separate regularization. Inflowing blood as seen in the heart mostly enters locally in bulk, which should be discernible in the lower spatial frequencies of the time derivative. Through-plane sources might thus be regularized to resemble the time derivative under a smoothness constraint similar to the Sobolev penalty applied to the coils. This would shift the problem from the current frame to the boundaries of the frame batch around it beyond which there is no information. Sudden inflow events could be captured in this way.

Regarding the second problem, how to estimate flow in the presence of artifacts, successful utilization of flow fields estimated with an L_1 - L_1 constraint as reported by Li et al. could not be reproduced, even though repeated flow estimations over the course of several Newton steps were used vs only one estimation by Li. In Zhao et al.'s work the undersampling factor of around 10 vs 27 here plus the fact that their golden angle acquisition visited far more positions in Fourier-space and frames were combined all at once probably improved the condition of the problem. This is, however, no option in an online reconstruction.

Since most artifacts that enter the optical flow field occur near strong contrast lines, application of a fat saturation pulse has been tried to remove the bright fat signal and make intensities more uniform across the image similar to the in-plane phantom. Fat suppression, however, proved insufficient to remove the problem entirely and may also not always be desired.

One intended future direction toward a solution - illustrated before in Fig. 5.21 - is temporal coarse graining within the frame batch to localize areas where intensity changes in time stem purely from artifacts. Complementary to that, motion vector field estimation in a pyramidal multi-resolution approach corresponding to spatial coarse graining has been used by Rank et al. [63]. Both approaches could in principle be combined.

Applications free of through-plane motion that could benefit from the presented optical flow constrained reconstruction without further modifications include studies of temporomandibular joint disorder [38] or of tongue movement [37] and swallowing [39].

Reconstruction Process, Regularization and Convergence

Combined optical flow estimation and image reconstruction has been implemented here as an alternating process. It is thus not a truly "joint" reconstruction though sometimes referred to as such [42]. Merging of the alternating minimizations into one minimization over ρ , \mathbf{c} and \mathbf{v} might enable use of the optical flow field from the first Newton step on. In the current scheme this does not work because by virtue of alternation there is always one Newton update step $x_{n+1} = x_n + dx$ between

estimation and application of the flow. In early Newton steps this means flow is computed on images with a different intensity distribution than that of those it is applied to.

On the other hand, merging may be complicated by the fact that the optical flow estimate initially has very low amplitudes making it effectively a time difference operator which leads to averaging amongst images (Fig. 5.6). Further, given the presently slow convergence of the optical flow estimation (Figs. 5.4b-5.5) it is unclear whether merging will improve reconstruction. One potential benefit of the alternating approach with respect to computing time is that such strategies as coarse-graining plus image warping can be used more easily to lay off iterations to a smaller matrix size [45].

Another interesting finding is that L_2 -penalties did not perform worse than L_1 penalties (Fig. 5.10a). Superior performance of L_1 -penalties has been observed in camera footage, especially regarding robustness to outliers and preservation of intensity discontinuities [64]. Yet this seems to be no issue here, possibly, because images with only 75 spokes underlying have less sharply defined edges and motions in general. Another reason for L_1 -penalties performing below expectation might be due to the small batch size required for online computation. Compressed sensing approaches such as temporal TV-reconstruction require undersampling artifacts to have a high degree of incoherence. That is, in contrast to the signal they should appear noise-like along the time dimension [40]. Consequently Zhao et al. used the entire time series in time-Fourier domain where the breathing and cardiac frequencies stand out from noise, while artifacts of a sampling pattern in a golden angle progression have no periodicity.

From a numeric point of view using the L_2 -norm is advantageous as the problem remains convex, smooth and thus open to efficient methods such as conjugate gradients [18]. This, along with the above mentioned coarse-graining may also bring down the computing cost of optical flow estimation from about 1 second per Newton step in the used implementation. Currently the cost is slightly more than the cost of the image estimation and it should be way lower given that only finite difference operators are involved.

Lowering the number of excess iterations is another option to speed up computing. A suitable stopping criterion for optical flow constrained reconstruction where only in-plane motion is involved is to stop as soon as the total residual at any iteration has fallen by less than 0.1% with respect to the previous iteration. For general use of such a criterion the model should, however, first be adjusted to properly handle through-plane motion to prevent early stopping and blurred results.

The choice of parameters was uncomplicated as, once set, parameters did not have to be changed between different measurements and slice orientations (see

section 4.3.1). Surprisingly, the γ parameter that controls temporal regularization strength in batch reconstruction was best kept constant throughout all Newton steps. Only the Tikhonov penalty controlled by α had to be decreased successively, whereas reduction of γ merely led to a loss of temporal regularization effect.

Better Data Exploitation: Aggregation

Batch reconstruction entails a higher computing cost that is linear in the batch size. It also yields as many output frames per time point as fit in the batch. The central frame is expected to be of highest quality since the fringe frames in the batch can draw on information only from one neighbor and that neighbor in turn will have a higher quality neighbor toward the center and a poor quality neighbor toward the fringe. Therefore, only the center frame has been considered in this work.

Nonetheless, every time point gets reconstructed several times with partly different member frames in the batch each time. As each such instance of a time point should always show the object in the same motion state, those instances could be aggregated for instance by averaging to yield a better final image.

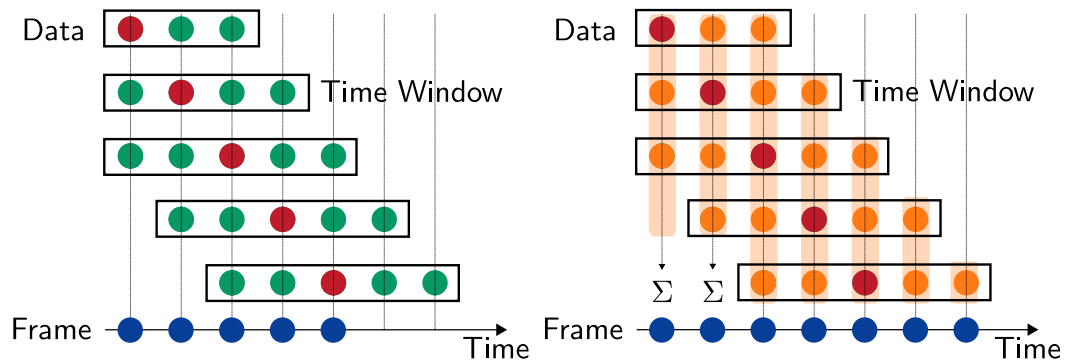


Figure 6.1: Data processing scheme in the used reconstruction: Data batches over a time window of five frames are processed together resulting in five images. However, (left) only the batch's center frame is written to the output buffer. Since reconstructed frames referring to the same time point (orange columns right) should show the same motion state their average should as well. Batch overlap means the maximum number of different frames in the average exceeds the number of different frames in a batch, potentially increasing k -space coverage.

This will be the more effective the larger the pattern repetition period and the more spokes in Fourier-space are visited. Given an acquisition with 9 different pattern orientations a reconstruction with a batch size of $T=5$ frames should still be able to combine the information of 9 frames by virtue of aggregation as the overlap of same motion state frames is $2T-1$ (Fig. 6.1). Potential benefits include increased image sharpness and SNR without increased reconstruction time. Whether the necessary requirement that different motion state instances of each time point indeed do not vary visibly amongst each other is fulfilled remains to be investigated. An optical flow model is highly likely to perform better in this

than batch- ∂_t regularization. The assessment framework based on surrogates presented in this work provides a good bases to continue along that direction.

Summary

In this thesis non-linear inversion reconstruction for joint image and coil sensitivity estimation was combined with different temporal regularization schemes. The method was tested for 2D cardiac magnetic resonance imaging using undersampled data at 30 frames per second. Comparisons to ground truth references of numerical phantoms and human heart data showed that adding an optical flow constraint leads to superior performance if motion is in-plane. Through-plane motion which violates the optical flow model markedly slowed down reconstruction convergence, caused blur and in some instances image artifacts at the locus of intensity inflow. A result that contrasts with recent literature and stimulates future model extensions for optical flow to capture local intensity inflow. Further, flow fields estimated from undersampled data still contained artifacts, even after iterative refinement, prompting the need for a spatio-temporal coarse graining approach that is left for future work. Meanwhile, a proposed plain batch time difference regularization in the L_2 -norm yielded visibly improved image quality as well as better temporal fidelity compared to the established temporal affine constraint. As another surprising finding, L_2 -norm penalties performed at least as well in batch regularization as non-smooth L_1 -penalties which greatly simplifies functional minimization.

Appendix

A.1 Prox Operators

The penalty terms used in this work are summarized in Tab. A.1. They either fall in the linear or affine category. The prox-functions to the penalties $F(x)$ with non-transformed variables are computed using eq. 4.6. The prox-functions involving transformed variables $F(\mathcal{A}x)$ do not need to be computed for Chambolle-Pock's primal-dual scheme. Instead the prox-functions to the dualized penalties $F^*(y)$ are used and computed by applying Moreau's decomposition

$$y = \text{prox}_{\gamma F}(y) + \gamma \text{prox}_{\frac{F^*}{\gamma}}\left(\frac{y}{\gamma}\right)$$

$F(x)$	$\text{prox}_{\sigma F}(y)$	$\text{prox}_{\tau F^*}(y)$
$\frac{\lambda}{2} \ x\ _2^2$	$\frac{1}{1 + \sigma\lambda} y$	$\frac{1}{1 + \tau/\lambda} y$
$\lambda \ x\ _1$	$\text{sign}(y) \max(y - \sigma\lambda, 0)$	$\text{sign}(y) \min(y , \lambda)$
$\frac{\lambda}{2} \ ax - b\ _2^2$	$\frac{y + \sigma\lambda a^* b}{1 + \sigma\lambda a ^2}$	$(a^* \lambda) \frac{ay - \tau b}{\tau + a ^2 \lambda}$
$\lambda \ ax - b\ _1$	$\text{sign}(y - \frac{b}{a}) \max(y - \sigma\lambda a , 0)$	$\text{sign}(y - \frac{b}{a}) \min(y , \lambda a)$

Table A.1: Prox functions for the Chambolle-Pock algorithm. $\text{sign}(x) = x/|x|$, $|x| > 0$; 0 else is used in the complex sense of the function. Stabilization of $\frac{b}{a}$ around small values of a is advisable. An in-depth derivation is given in [41].

A.2 Parameter Tables

Figs. 5.1-5.3

Fig. 5.1a	reference image:	image iter. = 50, $\alpha = 0.0001$
Fig. 5.1b	opt. flow constraint:	image iter. = 1000, $\gamma = 0.03$
Fig. 5.1c	no temp. regularization:	image iter. = 50, $\alpha = 0.0001$
Fig. 5.1d	no temp. reg. + median:	image iter. = 50, $\alpha = 0.0001$
Fig. 5.1e	affine temp. reg.:	image iter. = 50, $\alpha = 0.02$, $\lambda = 1$
Fig. 5.1f	aff. temp. reg. + median:	image iter. = 50, $\alpha = 0.02$, $\lambda = 1$
Fig. 5.1g	no temp. reg. + sliding average:	image iter. = 50, $\alpha = 0.0001$
Fig. 5.1h	finite diff. operator:	image iter. = 200, $\gamma = 0.1$

Table A.2: Reconstruction parameters in screen-shots of noiseless pulsating phantom. Parameters in the same order identical in Figs. 5.2 and 5.3.

Figs. 5.4-5.7

Fig. 5.4a	image reconstruction convergence:	image iter. = 5000, $\gamma = 0.05$
Fig. 5.4b	opt. flow estimation convergence:	flow iter. = 10000, $\beta = 0.001$
	reference image:	image iter. = 50, $\alpha = 0.0001$

Table A.3: Reconstruction parameters in plot of convergence curve: Cost function residual of image reconstruction and optical flow estimation in noiseless pulsating phantom. Parameters identical in Figs. 5.5-5.7.

Figs. 5.8

Fig. 5.8	reference:	image iter. = 50, $\alpha = 0.01$
Fig. 5.8	no temp. reg.:	image iter. = 50, $\alpha = 0.01$
Fig. 5.8	no temp. reg. + median:	image iter. = 50, $\alpha = 0.01$
Fig. 5.8	no temp. reg. + sliding average:	image iter. = 50, $\alpha = 0.01$
Fig. 5.8	affine temp. reg.	image iter. = 50, $\alpha = 0.07$, $\lambda = 1$
Fig. 5.8	affine temp. reg. + median:	image iter. = 50, $\alpha = 0.07$, $\lambda = 1$
Fig. 5.8	finite diff. operator:	image iter. = 200, $\gamma = 0.1$
Fig. 5.8	opt. flow constraint:	image iter. = 500, $\gamma = 0.03$

Table A.4: Reconstruction parameters in SSIM plot of noisy pulsating phantom.

Figs. 5.9-5.10

Fig. 5.9 (top)	reference flow:	flow iter. = 1000, $\beta = 0.001$
Fig. 5.9 (top)	hard thresh flow:	flow iter. = 1000, $\beta = 0.001$
Fig. 5.9 (top)	L_1 -reg. flow:	flow iter. = 1000, $\beta_2 = 0.001$
Fig. 5.9 (top)	L_2 -reg. (noisy) flow	flow iter. = 1000, $\beta = 0.001$
Fig. 5.9 (bottom)	reference image:	image iter. = 50, $\alpha = 0.01$
Fig. 5.9 (bottom)	L_1 -reg. flow constraint:	image iter. = 500, $\gamma = 0.05$
Fig. 5.10 (bottom)	other flow constraints:	image iter. = 500, $\gamma = 0.05$
Fig. 5.9 (bottom)	finite diff. operator:	image iter. = 200, $\gamma = 0.1$
Fig. 5.9 (bottom)	no temp. reg.:	image iter. = 50, $\alpha = 0.01$

Table A.5: Reconstruction parameters in screen-shots of differently regularized flow fields and reconstructed images of the noisy pulsating phantom.

Figs. 5.11

Fig. 5.9 (top)	reference image:	image iter. = 50, $\alpha = 0.0001$
Fig. 5.9 (top)	opt. flow constraint 1000 iter.:	image iter = 500, $\gamma = 0.05$
Fig. 5.9 (top)	opt. flow constraint 5000 iter.	image iter = 500, $\gamma = 0.05$
Fig. 5.9 (bottom)	time difference	image iter. = 50, $\alpha = 0.01$
Fig. 5.9 (bottom)	opt. flow:	flow iter. = 1000, $\beta = 0.001$
Fig. 5.9 (bottom)	opt. flow:	flow iter. = 5000, $\beta = 0.001$

Table A.6: Reconstruction parameters in screen-shots of noiseless fast long-range phantom.

Figs. 5.12-5.13

Fig. 5.12	reference image:	image iter. = 50, $\alpha = 0.01$
Fig. 5.12	opt. flow constraint 1000 iter.:	image iter = 500, $\gamma = 0.05$
Fig. 5.12	opt. flow constraint 5000 iter.	image iter = 500, $\gamma = 0.05$
Fig. 5.12	opt. flow 1000 iter:	flow iter. = 1000, $\beta = 0.001$
Fig. 5.12	opt. flow 5000 iter:	flow iter. = 5000, $\beta = 0.001$

Table A.7: Reconstruction parameters in screen-shots and SSIM plots of the noisy through-plane phantom.

Figs. 5.14-5.15

Fig. 5.14	reference images (heart):	image iter. = 200, $\alpha = 0.002$
Fig. 5.14	reference images (heart):	image iter. = 200, $\alpha = 0.002$
Fig. 5.15	reference images (phantoms):	image iter. = 50, $\alpha = 0.01$
Figs. 5.14 - 5.15	opt. flow fields:	flow iter = 5000, $\beta = 0.001$

Table A.8: Reconstruction parameters in screen-shots and transport residual norm plots of the heart surrogate data and the noisy through- and in-plane phantoms.

Figs. 5.16-5.18

Fig. 5.16	L_2 opt. flow:	flow iter. = 5000, $\beta = 0.001$
Fig. 5.16	$L_2 + L_1$ opt. flow:	flow iter. = 5000, $\beta = 0.0001$, $\beta_2 = 0.0001$
Fig. 5.16	$L_2 +$ hard thresholding opt. flow:	flow iter. = 5000, $\beta = 0.001$
Fig. 5.16	reference:	image iter. = 200, $\alpha = 0.01$
Fig. 5.16-5.18	opt flow. constraint:	image iter. = 500, $\gamma = 0.3$
Fig. 5.17	no temp. reg.:	image iter. = 200, $\alpha = 0.002$
Fig. 5.17-5.18	dt-operator:	image iter. = 500, $\gamma = 0.3$
Fig. 5.17-5.18	affine reg. (+ median):	image iter. = 200, $\alpha = 0.01$, $\lambda = 1$

Table A.9: Reconstruction parameters for Four-chamber-view images: (5.16) Differently regularized flow fields, through-plane artifact demonstration, (5.17) SSIM comparison with model violation; short-axis-views: (5.18) comparison to reference for different reconstruction types.

Figs. 5.19

Fig. 5.19	reference:	known coils, img iter. = 200, $\alpha = 0.002$
Fig. 5.19	undersampled:	unknown coils, $n = 7$, $\alpha = 0.5^n$, image iter. = 40 per n
Fig. 5.19	(L_2) flow from reference:	flow iter. = 1000, $\beta = 0.001$
Fig. 5.19	(L_2) flow from undersampled:	flow iter. = 1000, $\beta = 0.001$
Fig. 5.19	reconstr. with ref. flow:	unknown coils, $n = 11$, $\alpha = 0.5^n$, $\gamma = 0.3$, image iter. = 40 per n
Fig. 5.19	reconstr. with usmp. flow:	unknown coils, $n = 7$, $\alpha = 0.5^n$, $\gamma = 0.3$, image iter. = 40 per n
Fig. 5.19	reconstr. with 3rd flow upd.:	unknown coils, $n = 11$, $\alpha = 0.5^n$, $\gamma = 0.3$, image iter. = 40 per n
Fig. 5.19	3rd flow update:	flow iter. = 1000, $\beta = 0.001$

Table A.10: Reconstruction parameters (non-linear inversion) of in-plane phantom. Iterative improvement of flow fields. n denotes newton step index.

Figs. 5.20

Fig. 5.20	reconstr. & time diff.:	unknown coils, $n = 15$, img iter. = 200, $\alpha = 0.002$
Fig. 5.20	L_2 -flow w/o hard thresh.:	flow iter. = 1000, $\beta = 0.001$
Fig. 5.20	$L_1 + L_2$ -flow:	flow iter. = 1000, $\beta = 0.0001$, $\beta_2 = 0.0001$
Fig. 5.20	$L_1 - L_1$ -flow:	flow iter. = 1000, $\beta_2 = 0.001 \cdot 0.5^m$, $m = \max(n - 7, 0)$

Table A.11: Reconstruction parameters (non-linear inversion) of four-chamber-view. Comparison of flows, flow field artifacts. n denotes newton step index.

Figs. 5.22

Fig. 5.22	dt-operator regularization:	unknown coils, $n = 13$, $\gamma = 0.3$ image iter. = 50 per n
Fig. 5.22	affine reg. + median filter:	unknown coils, $n = 13$, $\alpha = 0.5^n$, image iter. = 20 per n , $\lambda = 0.9$

Table A.12: Reconstruction parameters (non-linear inversion) of four-chamber-view. Comparison of finite time difference operator and affine regularization. n denotes newton step index

Abbreviations

ECG: Electrocardiography

FLASH: Fast Low Angle Shot

FOV: Field of View

fps: frames per second

GRE: Gradient Echo

MSE: Mean Square Error

NLINV: Non-linear Inversion Reconstruction (joint estimation of anatomic images and coil sensitivity function)

NMR: Nuclear Magnetic Resonance

PCA: Principle Component Analysis

(rt-)MRI: real-time Magnetic Resonance Imaging

rf-pulse: radio frequency pulse

SENSE: Sensitivity Encoding (image space based deconvolution of multi-channel MRI data with a known coil sensitivity function)

SNR: Signal to Noise Ratio

SSIM: Structural Similarity (Index)

TV: Total Variation

Bibliography

- [1] Jeffrey Tsao and Sebastian Kozerke. Mri temporal acceleration techniques. *Journal of Magnetic Resonance Imaging*, 36(3):543–560, 2012. doi: 10.1002/jmri.23640. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.23640>.
- [2] Martin Uecker, Shuo Zhang, Dirk Voit, Alexander Karaus, Klaus-Dietmar Merboldt, and Jens Frahm. Real-time mri at a resolution of 20 ms. *NMR in Biomedicine*, 23(8):986–994, 2010. doi: 10.1002/nbm.1585. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/nbm.1585>.
- [3] Andrew C. Larson, Richard D. White, Gerhard Laub, Elliot R. McVeigh, Debiao Li, and Orlando P. Simonetti. Self-gated cardiac cine mri. *Magnetic Resonance in Medicine*, 51(1):93–102, 2004. doi: 10.1002/mrm.10664. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.10664>.
- [4] Grace M. Nijm, Alan V. Sahakian, Steven Swiryn, James C. Carr, John J. Sheehan, and Andrew C. Larson. Comparison of self-gated cine mri retrospective cardiac synchronization algorithms. *Journal of Magnetic Resonance Imaging*, 28(3):767–772, 2008. doi: 10.1002/jmri.21514. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.21514>.
- [5] Y. Huo, C. Hellge, T. Wiegand, and L. Hanzo. A tutorial and review on inter-layer fec coded layered video streaming. *IEEE Communications Surveys Tutorials*, 17(2):1166–1207, Secondquarter 2015. ISSN 2373-745X. doi: 10.1109/COMST.2015.2392378.
- [6] Martin Uecker, Thorsten Hohage, Kai Tobias Block, and Jens Frahm. Image reconstruction by regularized nonlinear inversion—joint estimation of coil sensitivities and image content. *Magnetic Resonance in Medicine*, 60(3):674–682, 2008. doi: 10.1002/mrm.21691. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.21691>.
- [7] Berthold K. P. Horn and Brian G. Schunck. Determining optical flow. *ARTIFICIAL INTELLIGENCE*, 17:185–203, 1981.
- [8] Housen Li, Markus Haltmeier, Shuo Zhang, Jens Frahm, and Axel Munk. Aggregated motion estimation for real-time mri reconstruction. *Magnetic Resonance in Medicine*, 72(4):1039–1048, 2014. doi: 10.1002/mrm.25020. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.25020>.
- [9] Robert W. Brown, Yu-Chung N. Cheng, E. Mark Haacke, Michael R. Thompson, and Ramesh Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. John Wiley & Sons, second edition, April 2014. ISBN 9781118633953. doi: 10.1002/9781118633953. URL <https://onlinelibrary.wiley.com/doi/book/10.1002/9781118633953>.
- [10] A Haase, J Frahm, D Matthaei, W Hanicke, and K.-D Merboldt. Flash imaging. rapid nmr imaging using low flip-angle pulses. *Journal of Magnetic Resonance (1969)*, 67(2):258 – 266, 1986. ISSN 0022-2364. doi: [https://doi.org/10.1016/0022-2364\(86\)90433-6](https://doi.org/10.1016/0022-2364(86)90433-6). URL <http://www.sciencedirect.com/science/article/pii/0022236486904336>.

- [11] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller. The nmr phased array. *Magnetic Resonance in Medicine*, 16(2):192–225, 1990. doi: 10.1002/mrm.1910160203. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1910160203>.
- [12] Michael Lustig and John M. Pauly. Spirit: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magnetic Resonance in Medicine*, 64(2):457–471, 2010. doi: 10.1002/mrm.22428. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.22428>.
- [13] Martin Uecker, Peng Lai, Mark J. Murphy, Patrick Virtue, Michael Elad, John M. Pauly, Shreyas S. Vasanawala, and Michael Lustig. Espirit—an eigenvalue approach to autocalibrating parallel mri: Where sense meets grappa. *Magnetic Resonance in Medicine*, 71(3):990–1001, 2014. doi: 10.1002/mrm.24751. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24751>.
- [14] P. J. Beatty, D. G. Nishimura, and J. M. Pauly. Rapid gridding reconstruction with a minimal oversampling ratio. *IEEE Transactions on Medical Imaging*, 24(6):799–808, June 2005. ISSN 1558-254X. doi: 10.1109/TMI.2005.848376.
- [15] Matt A. Bernstein, Kevin F. King, and Xiaohong Joe Zhou. Chapter 17 - advanced pulse sequence techniques. In Matt A. Bernstein, Kevin F. King, and Xiaohong Joe Zhou, editors, *Handbook of MRI Pulse Sequences*, pages 802 – 954. Academic Press, Burlington, 2004. ISBN 978-0-12-092861-3. doi: <https://doi.org/10.1016/B978-012092861-3/50024-8>. URL <http://www.sciencedirect.com/science/article/pii/B9780120928613500248>.
- [16] M Uecker KT Block. Simple method for adaptive gradient-delay compensation in radial mri. volume 19, pages 2816–2816, May 2011.
- [17] A. B. Bakushinsky and M.Yu. Kokurin. *Iterative Methods for Approximate Solution of Inverse Problems*, volume 577. Springer Netherlands, 1 edition, 2004. ISBN 978-90-481-6798-2. doi: 10.1007/978-1-4020-3122-9. URL <https://www.springer.com/gp/book/9781402031212>.
- [18] Jonathan Richard Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. Technical report, School of Computer Science Carnegie Mellon University Pittsburgh, PA 15213, 1994.
- [19] Klaas P. Pruessmann, Markus Weiger, Markus B. Scheidegger, and Peter Boesiger. Sense: Sensitivity encoding for fast mri. *Magnetic Resonance in Medicine*, 42(5):952–962, 1999. doi: 10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S. URL [https://doi.org/10.1002/\(SICI\)1522-2594\(199911\)42:5<952::AID-MRM16>3.0.CO;2-S](https://doi.org/10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S).
- [20] Martin Uecker. *Nonlinear Reconstruction Methods for Parallel Magnetic Resonance Imaging*. PhD thesis, Institut für Numerische und Angewandte Mathematik, Universität Göttingen, 2009. URL <https://ediss.uni-goettingen.de/handle/11858/00-1735-0000-0006-B3C6-3>.
- [21] Olaf Dietrich, Maximilian F. Reiser, and Stefan O. Schoenberg. Artifacts in 3-t mri: Physical background and reduction strategies. *European Journal of Radiology*, 65(1):29 – 35, 2008. ISSN 0720-048X. doi: <https://doi.org/10.1016/j.ejrad.2007.11.005>. URL <http://www.sciencedirect.com/science/article/pii/S0720048X07005633>. 3 Tesla MR.
- [22] Leslie Ying and Jinhua Sheng. Joint image reconstruction and sensitivity estimation in sense (jsense). *Magnetic Resonance in Medicine*, 57(6):1196–1202, 2007. doi: 10.1002/mrm.21245. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.21245>.

- [23] H. Christian M. Holme, Sebastian Rosenzweig, Frank Ong, Robin N. Wilke, Michael Lustig, and Martin Uecker. Enlive: An efficient nonlinear method for calibrationless and robust parallel imaging. *Scientific Reports*, 9(1):3034, 2019. ISSN 2045-2322. doi: 10.1038/s41598-019-39888-7. URL <https://doi.org/10.1038/s41598-019-39888-7>.
- [24] Li Feng, Robert Grimm, Kai Tobias Block, Hersh Chandarana, Sungheon Kim, Jian Xu, Leon Axel, Daniel K. Sodickson, and Ricardo Otazo. Golden-angle radial sparse parallel mri: Combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric mri. *Magnetic Resonance in Medicine*, 72(3):707–717, 2014. doi: 10.1002/mrm.24980. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24980>.
- [25] Li Feng, Leon Axel, Hersh Chandarana, Kai Tobias Block, Daniel K. Sodickson, and Ricardo Otazo. Xd-grasp: Golden-angle radial mri with reconstruction of extra motion-state dimensions using compressed sensing. *Magnetic Resonance in Medicine*, 75(2):775–788, 2016. doi: 10.1002/mrm.25665. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.25665>.
- [26] Ricardo Otazo, Emmanuel Candès, and Daniel K. Sodickson. Low-rank plus sparse matrix decomposition for accelerated dynamic mri with separation of background and dynamic components. *Magnetic Resonance in Medicine*, 73(3):1125–1136, 2015. doi: 10.1002/mrm.25240. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.25240>.
- [27] S. G. Lingala, E. DiBella, and M. Jacob. Deformation corrected compressed sensing (dc-cs): A novel framework for accelerated dynamic mri. *IEEE Transactions on Medical Imaging*, 34(1): 72–85, Jan 2015. ISSN 1558-254X. doi: 10.1109/TMI.2014.2343953.
- [28] M. Salman Asif, Lei Hamilton, Marijn Brummer, and Justin Romberg. Motion-adaptive spatio-temporal regularization for accelerated dynamic mri. *Magnetic Resonance in Medicine*, 70(3): 800–812, 2013. doi: 10.1002/mrm.24524. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24524>.
- [29] Sungheon G. Kim, Li Feng, Robert Grimm, Melanie Freed, Kai Tobias Block, Daniel K. Sodickson, Linda Moy, and Ricardo Otazo. Influence of temporal regularization and radial undersampling factor on compressed sensing reconstruction in dynamic contrast enhanced mri of the breast. *Journal of Magnetic Resonance Imaging*, 43(1):261–269, 2016. doi: 10.1002/jmri.24961. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.24961>.
- [30] Bryson Dietz, B. Gino Fallone, and Keith Wachowicz. Nomenclature for real-time magnetic resonance imaging. *Magnetic Resonance in Medicine*, 81(3):1483–1484, 2019. doi: 10.1002/mrm.27487. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27487>.
- [31] Katherine L. Wright, Jesse I. Hamilton, Mark A. Griswold, Vikas Gulani, and Nicole Seiberlich. Non-cartesian parallel imaging reconstruction. *Journal of Magnetic Resonance Imaging*, 40(5): 1022–1040, 2014. doi: 10.1002/jmri.24521. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.24521>.
- [32] Tao Zhang, Shilpy Chowdhury, Michael Lustig, Richard A. Barth, Marcus T. Alley, Thomas Grafendorfer, Paul D. Calderon, Fraser J.L. Robb, John M. Pauly, and Shreyas S. Vasanawala. Clinical performance of contrast enhanced abdominal pediatric mri with fast combined parallel imaging compressed sensing reconstruction. *Journal of Magnetic Resonance Imaging*, 40(1):13–25, 2014. doi: 10.1002/jmri.24333. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.24333>.

- [33] Oren N Jaspan, Roman Fleysner, and Michael L Lipton. Compressed sensing mri: a review of the clinical literature. *The British Journal of Radiology*, 88(1056):20150487, 2015. doi: 10.1259/bjr.20150487. URL <https://doi.org/10.1259/bjr.20150487>. PMID: 26402216.
- [34] J Fabian, E J Epstein, and N Coulshed. Duration of phases of left ventricular systole using indirect methods. i. normal subjects. *Heart*, 34(9):874–881, 1972. ISSN 0007-0769. doi: 10.1136/hrt.34.9.874. URL <https://heart.bmj.com/content/34/9/874>.
- [35] Shuo Zhang, Kai Tobias Block, and Jens Frahm. Magnetic resonance imaging in real time: Advances using radial flash. *Journal of Magnetic Resonance Imaging*, 31(1):101–109, 2010. doi: 10.1002/jmri.21987. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.21987>.
- [36] S. Winkelmann, T. Schaeffter, T. Koehler, H. Eggers, and O. Doessel. An optimal radial profile order based on the golden ratio for time-resolved mri. *IEEE Transactions on Medical Imaging*, 26(1):68–76, Jan 2007. ISSN 1558-254X. doi: 10.1109/TMI.2006.885337.
- [37] Aaron Niebergall, Shuo Zhang, Esther Kunay, Götz Keydana, Michael Job, Martin Uecker, and Jens Frahm. Real-time mri of speaking at a resolution of 33 ms: Undersampled radial flash with nonlinear inverse reconstruction. *Magnetic Resonance in Medicine*, 69(2):477–485, 2013. doi: 10.1002/mrm.24276. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24276>.
- [38] Sebastian Krohn, Nikolaus Gersdorff, Torsten Wassmann, Klaus-Dietmar Merboldt, Arun A. Joseph, Ralf Buegers, and Jens Frahm. Real-time mri of the temporomandibular joint at 15 frames per second—a feasibility study. *European Journal of Radiology*, 85(12):2225 – 2230, 2016. ISSN 0720-048X. doi: <https://doi.org/10.1016/j.ejrad.2016.10.020>. URL <http://www.sciencedirect.com/science/article/pii/S0720048X16303254>.
- [39] Shuo Zhang, Arno Olthoff, and Jens Frahm. Real-time magnetic resonance imaging of normal swallowing. *Journal of Magnetic Resonance Imaging*, 35(6):1372–1379, 2012. doi: 10.1002/jmri.23591. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.23591>.
- [40] Michael Lustig, David Donoho, and John M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007. doi: 10.1002/mrm.21391. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.21391>.
- [41] Hendrik Meinert Dirks. *Variational Methods for Joint Motion Estimation and Image Reconstruction*. PhD thesis, Westfaelische Wilhelms-Universitaet Muenster, 2015.
- [42] Martin Burger, Hendrik Dirks, and Carola-Bibiane Schönlieb. A variational model for joint motion estimation and image reconstruction, 2016.
- [43] David L Donoho and Iain M Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455, 09 1994. ISSN 0006-3444. doi: 10.1093/biomet/81.3.425. URL <https://doi.org/10.1093/biomet/81.3.425>.
- [44] N. Zhao, D. O’Connor, A. Basarab, D. Ruan, and K. Sheng. Motion compensated dynamic mri reconstruction with local affine optical flow estimation. *IEEE Transactions on Biomedical Engineering*, 66(11):3050–3059, Nov 2019. ISSN 1558-2531. doi: 10.1109/TBME.2019.2900037.
- [45] D. Sun, S. Roth, and M. J. Black. Secrets of optical flow estimation and their principles. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2432–2439, June 2010. doi: 10.1109/CVPR.2010.5539939.

- [46] Denis Fortun, Patrick Boutheymy, and Charles Kervrann. Optical flow modeling and computation: A survey. *Computer Vision and Image Understanding*, 134:1 – 21, 2015. ISSN 1077-3142. doi: <https://doi.org/10.1016/j.cviu.2015.02.008>. URL <http://www.sciencedirect.com/science/article/pii/S1077314215000429>. Image Understanding for Real-world Distributed Video Networks.
- [47] M. Guerquin-Kern, L. Lejeune, K. P. Pruessmann, and M. Unser. Realistic analytical phantoms for parallel magnetic resonance imaging. *IEEE Transactions on Medical Imaging*, 31(3):626–636, March 2012. ISSN 1558-254X. doi: 10.1109/TMI.2011.2174158.
- [48] Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, April 2004. ISSN 1941-0042. doi: 10.1109/TIP.2003.819861.
- [49] L. Zhang, L. Zhang, X. Mou, and D. Zhang. Fsim: A feature similarity index for image quality assessment. *IEEE Transactions on Image Processing*, 20(8):2378–2386, Aug 2011. ISSN 1941-0042. doi: 10.1109/TIP.2011.2109730.
- [50] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, May 2011. ISSN 1573-7683. doi: 10.1007/s10851-010-0251-1. URL <https://doi.org/10.1007/s10851-010-0251-1>.
- [51] Patrick L. Combettes and Jean-Christophe Pesquet. Primal-dual splitting algorithm for solving inclusions with mixtures of composite, lipschitzian, and parallel-sum type monotone operators. *Set-Valued and Variational Analysis*, 20(2):307–330, Jun 2012. ISSN 1877-0541. doi: 10.1007/s11228-011-0191-y. URL <https://doi.org/10.1007/s11228-011-0191-y>.
- [52] Laurent Condat. A primal-dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms. *Journal of Optimization Theory and Applications*, 158(2):460–479, Aug 2013. ISSN 1573-2878. doi: 10.1007/s10957-012-0245-9. URL <https://doi.org/10.1007/s10957-012-0245-9>.
- [53] Peijun Chen, Jianguo Huang, and Xiaoqun Zhang. A primal-dual fixed-point algorithm for minimization of the sum of three convex separable functions. *Fixed Point Theory and Applications*, 2016, 12 2015. doi: 10.1186/s13663-016-0543-2.
- [54] N. Komodakis and J. Pesquet. Playing with duality: An overview of recent primal-dual approaches for solving large-scale optimization problems. *IEEE Signal Processing Magazine*, 32(6):31–54, Nov 2015. ISSN 1558-0792. doi: 10.1109/MSP.2014.2377273.
- [55] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004. ISBN 0521833787.
- [56] Jonathan Eckstein and Dimitri P. Bertsekas. On the douglas—rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1):293–318, Apr 1992. ISSN 1436-4646. doi: 10.1007/BF01581204. URL <https://doi.org/10.1007/BF01581204>.
- [57] Neal Parikh and Stephen Boyd. Proximal algorithms. *Found. Trends Optim.*, 1(3):127–239, 2014. ISSN 2167-3888. doi: 10.1561/2400000003. URL <https://doi.org/10.1561/2400000003>.
- [58] Hendrik Dirks. A flexible primal-dual toolbox. arXiv preprint, 2016. URL <http://www.flexbox.im>.

- [59] Yura Malitsky and Thomas Pock. A first-order primal-dual algorithm with linesearch. *SIAM Journal on Optimization*, 28, 08 2016. doi: 10.1137/16M1092015.
- [60] Tom Goldstein, Min Li, Xiaoming Yuan, Ernie Esser, and Richard Baraniuk. Adaptive primal-dual hybrid gradient methods for saddle-point problems, 2013.
- [61] Mohammed Hussaini, Bram van Leer, and John Van Rosendale. *Upwind and High-Resolution Schemes*. Springer, Berlin, Heidelberg, 03 2015. ISBN ISBN-10: 364264452X. doi: 10.1007/978-3-642-60543-7.
- [62] Martin Burger, Hendrik Dirks, Lena Frerking, Andreas Hauptmann, Tapio Helin, and Samuli Siltanen. A variational reconstruction method for undersampled dynamic x-ray tomography based on physical motion models. *Inverse Problems*, 33(12):124008, nov 2017. doi: 10.1088/1361-6420/aa99cf. URL <https://doi.org/10.1088/1361-6420/aa99cf>.
- [63] Christopher M. Rank, Thorsten Heußer, Maria T. A. Buzan, Andreas Wetscherek, Martin T. Freitag, Julien Dinkel, and Marc Kachelrieß. 4d respiratory motion-compensated image reconstruction of free-breathing radial mr data with very high undersampling. *Magnetic Resonance in Medicine*, 77(3):1170–1183, 2017. doi: 10.1002/mrm.26206. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26206>.
- [64] Andreas Wedel, Thomas Pock, Christopher Zach, Horst Bischof, and Daniel Cremers. An improved algorithm for tv-l1 optical flow. In Daniel Cremers, Bodo Rosenhahn, Alan L. Yuille, and Frank R. Schmidt, editors, *Statistical and Geometrical Approaches to Visual Motion Analysis*, pages 23–45, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. ISBN 978-3-642-03061-1.

all time stamps @ 31.12.2019, 15:00:00

Acknowledgements

I would first like to thank my supervisor Jens Frahm for offering this project and absolute freedom to pursue its completion. Further thanks go to Martin Uecker for important pointers and advice when I struggled to find which step best to take next. And of course many thanks go to Kurt Böhm for making the servers run impeccably and fixing crashes even when he had a day off. I am further indebted to Sasha Kalentev for an always friendly ear to questions on programming and for providing and extending his CUTEL library for easy multi-GPU programming. Adding to the long list I further have to mention Dietmar Merboldt, Dirk Void and Arun Joseph for helping me with the data acquisition. Last but not least my thanks also go to all others at the Frahm lab and the Uecker lab, especially Jost Kollmeier, Volkert Roeloffs, Andreas Merrem, Zhengguo Tan and Xiaoqing Wang.