

Phase reconstruction from velocity-encoded MRI measurements – a survey of sparsity-promoting variational approaches

Martin Benning* Lynn Gladden Daniel Holland Carola-Bibiane Schönlieb
Tuomo Valkonen

October 17, 2013

Abstract

In recent years there has been significant developments in the reconstruction of magnetic resonance velocity images from sub-sampled k -space data. While showing a strong improvement in reconstruction quality compared to classical approaches, the vast number of different methods, and the challenges in setting them up, often leaves the user with the difficult task of choosing the correct approach, or more importantly, not selecting a poor approach. In this paper, we survey variational approaches for the reconstruction of phase-encoded magnetic resonance velocity images from sub-sampled k -space data. We are particularly interested in regularisers that correctly treat both smooth and geometric features of the image. These features are common to velocity imaging, where the flow field will be smooth but interfaces between the fluid and surrounding material will be sharp, but are challenging to represent sparsely. As an example we demonstrate the variational approaches on velocity imaging of water flowing through a packed bed of solid particles. We evaluate Wavelet regularisation against Total Variation and the relatively recent second order Total Generalised Variation regularisation. We combine these regularisation schemes with a contrast enhancement approach called Bregman iteration. We verify for a variety of sampling patterns that Morozov’s discrepancy principle provides a good criterion for stopping the iterations. Therefore, given only the noise level, we present a robust guideline for setting up a variational reconstruction scheme for MR velocity imaging.

1 Introduction

Phase-encoded Magnetic Resonance (MR) velocity imaging is used widely. In medical imaging, it is used to study the distribution and variation in flow in blood vessels and around the heart to both diagnose and understand congenital heart disease and the behaviour of heart valves [1]. In the physical sciences, MR velocity imaging has been used to study the rheology of complex fluids [2], liquids and gases flowing through packed beds [3, 4, 5], granular flows [6, 7] and multiphase turbulence [8]. The main advantage of MR for studying flow is that it is possible to non-invasively image systems without the use of a tracer. However, a major drawback of the technique, in both medical and non-medical applications, is the acquisition time of the measurement. In the medical field, long acquisition times (\sim min) necessitate breath hold or triggered acquisition techniques that, although impressive, are susceptible to artefacts and may be impractical in certain cases. In non-medical applications, long acquisition times require very stable systems and can prohibit the study of certain features, e.g. fine vortices in turbulent flow. To overcome these limitations many studies have explored methods to increase the temporal resolution of velocity-encoded imaging (e.g. [9, 10, 11]). These “ultra-fast” (< 1 s imaging time) techniques have all been demonstrated to provide a substantial improvement in the temporal resolution. However, despite these advances, each has limitations in terms of the systems that can be studied and the trade-off between spatial and temporal resolution. Therefore, based on the theory of compressed sensing (CS) [12, 13, 14] there has recently been increasing interest in introducing image reconstruction techniques that exploit prior knowledge of the signal [5, 8, 15, 16, 17].

The image reconstruction problem for phase-encoded MR velocity imaging is modelled as follows. Let $\Omega := \{1, \dots, n_1\} \times \{1, \dots, n_2\}$ be our discrete image domain. In velocity-encoded MRI as described in [5] the goal is to recover the phase of a complex signal $u \in \mathbb{C}^\Omega$ that is related to the k -space measurements via

$$S\mathcal{F}u + \nu = f. \quad (1)$$

Here, $f \in \mathbb{C}^m$ denotes the noisy subsampled k -space measurements with $m \ll n_1 n_2$. The noise $\nu \in \mathbb{C}^m$ is assumed to be normal-distributed with deviation σ and mean zero. The linear operator $\mathcal{F} : \mathbb{C}^\Omega \rightarrow \mathbb{C}^\Omega$

*Corresponding Author - Address: Magnetic Resonance Research Centre, Department of Chemical Engineering and Biotechnology, c/o Cavendish Stores, JJ Thomson Avenue, Cambridge, CB3 0HE, United Kingdom - E-Mail: mb941@cam.ac.uk - Telephone: +44 (0)1223 (7)61630 - Fax: +44 (0)1223 334796

represents the discrete two-dimensional Fourier transform and $S : \mathbb{C}^\Omega \rightarrow \mathbb{C}^m$ the sampling pattern. It maps the full two-dimensional k -space data $\mathcal{F}u$ to the m -dimensional vector f consisting of the samples present in the considered sampling pattern.

The problem (1) is ill-posed due to noise and sub-sampling. A naive approach to find a solution of (1) is to compute its least-squares solution, i.e.

$$u_z = \mathcal{F}^{-1} S^T f, \quad (2)$$

which simply aims at filling missing k -space entries with zero. Throughout this work we will refer to this solution as backprojection (BP). However, this solution is dominated by noise and sub-sampling artefacts, since no regularity is imposed on the solution. We therefore seek to find approximate solutions u to (1) by a variational regularisation approach. Such an approach consists of balancing between the model (1), and prior assumptions on the reconstruction u in terms of a regularisation functional J . Namely, we look for approximate solutions u_σ to

$$u_\sigma \in \arg \min_u \{J(u)\} \quad \text{subject to} \quad \|f - S\mathcal{F}u\|_2 \leq \sigma\sqrt{m}. \quad (3)$$

Here $\|v\|_2 := \sqrt{\sum_{i=1}^m |v(i)|^2}$ is the standard Euclidean 2-norm. The role of the inequality constraint $\|f - S\mathcal{F}u\|_2 \leq \sigma$ is to model (1), where the 2-norm corresponds to our assumption of normally distributed noise with standard deviation σ . The regularisation term J allows us to incorporate prior information on the reconstruction u_σ , that is essential to counteract the ill-posedness of the problem (1). Typically one either chooses a non-smooth regulariser, such as Total Variation [18] (TV), if the image to reconstruct contains non-smooth features such as edges, or a smooth regulariser, such as a Daubechies wavelet transform, if the intensities in the object change smoothly. However, it is challenging to find a regulariser that is suitable for images that contain both smooth and non-smooth variations in intensity. Therefore, in order to overcome deficiencies in first-order approaches like TV, we also consider higher-order approaches, in particular second-order Total Generalised Variation [19], $J = \text{TGV}^2$, for which an example reconstruction is provided in Figure 1.

It can be shown through a Lagrange multiplier approach [20, VII], see also [21, Chapter 5], that for σ larger than the realised noise level, there exists a parameter $\alpha > 0$ such that u_σ is equivalent to a signal u_α that is the minimiser of the Tikhonov regularisation approach

$$u_\alpha \in \arg \min_u \left\{ \frac{1}{2} \|f - S\mathcal{F}u\|_2^2 + \alpha J(u) \right\}. \quad (4)$$

The parameter $\alpha > 0$ is a positive regularisation parameter that weights the influence of the fidelity and the regularisation term.

Approaches within the scheme (4) have been shown to be powerful but there are challenges associated with the introduction of the adjustable "regularisation" parameter α and choice of prior knowledge. In this article we examine these challenges and present an approach to objectively determine the regularisation parameter. Further, solutions of (4) (and therefore also of (3)) suffer from a loss of contrast, see for instance [22, 23]. An approach to overcome this systematic bias is to replace the regulariser in (4) by its corresponding generalised Bregman distance [24, 25]. This modification, known as Bregman iteration, has been proposed in [26] and shown to significantly improve the signal contrast without accumulating the noise in the case of one-homogeneous regularisers. However, in order to ensure Bregman iteration to be a regularisation method, the iteration needs to be stopped with a suitable stopping criterion. In accordance with (3) a suitable stopping criterion is Morozov's discrepancy principle [27].

One of the main goals of this article is to demonstrate that the use of Bregman iteration in combination with the discrepancy principle will help to overcome contrast loss and facilitate robust choice of the unknown ideal regularisation parameter α . Moreover, we demonstrate that higher-order regularisation in the form of $J = \text{TGV}^2$ significantly outperforms first-order regularisation. We believe this is the first report of combining Bregman iteration and higher-order regularisation in the context of sub-sampled MRI reconstructions, though the combination has been investigated recently in a denoising application [28].

The paper is organised as follows. First in Section 2 we introduce our variational reconstruction approach in more detail. We provide a detailed discussion of different choices of the regularisation functional J , as well as the Bregman iteration. We also introduce the sampling schemes used in our numerical work that will be presented in Section 3, in which we do compare the different approaches on a computer-generated data set, as well as a real data set. The paper is concluded with an outlook and summary in Section 4.

2 Methods

As discussed in the introduction, our goal in this paper is to find approximate solutions to (1) by solving the variational scheme (4) for u_α , and to compare the solutions for different regularisation functionals J . First,




Figure 1: Example reconstruction of a sub-sampled signal to demonstrate the impact of regularisation on the reconstruction. The two images on the left show the ground truth magnitude and phase image of a computer-generated velocity encoded dataset. The two images in the middle show the backprojection of the corresponding noisy, sub-sampled k -space data, while the two images on the right show a regularised reconstruction (4) of the same data with $J = \text{TGV}^2$.

in Section 2.1, we are going to recall three popular choices for J . The choices that we consider for J are, more specifically, the Daubechies-3 wavelet, total variation, and the relatively new total generalised variation. Afterwards, in Section 2.2, we are going to discuss different sampling strategies and will discuss reasons to favour some schemes over others. We then introduce in Section 2.3 a strategy called Bregman iteration which is used to enhance the contrast in solutions of (4). In Section 2.4, we introduce the discrepancy principle for stopping the Bregman iterations and the measures of image quality used. Finally, we conclude the description of methodology in Section 2.5 with a discussion of the advantages and disadvantages of the different regularisers.

2.1 Regularisation Methods

We now introduce our choices for J , starting with wavelets and then moving on to first- and higher-order total variation schemes. The drawbacks and benefits of each of the schemes will be discussed in detail in Section 2.5.

Wavelets One standard regularisation approach for CS-MRI reconstructions is to penalise the 1-norm of the wavelet-transformed signals (cf. [29]). That is, we take

$$J(u) = \text{WAV}(u) = \|Wu\|_1.$$

We recall that the 1-norm $\|v\|_1$ is defined by $\|v\|_1 = \sum_{(i,j) \in \Omega} |v(i,j)|$. The operator $W : \mathbb{C}^\Omega \rightarrow \mathbb{C}^\Omega$ represents a discrete wavelet transform. With this choice of regulariser one forces the reconstruction u_α to be sparse in the wavelet domain. The rationale for this approach is that many signals are “compressible” in the sense that they are (almost) sparse in a suitable wavelet domain. In the present work we concentrate for W on the two-dimensional six-level Daubechies-3 wavelet transform (cf. [30]). The variational scheme (4) with $J = \text{WAV}$ may then be written as

$$u_\alpha \in \arg \min_u \left\{ \frac{1}{2} \|f - S\mathcal{F}u\|_2^2 + \alpha \text{WAV}(u) \right\}. \quad (5)$$

Total Variation Total variation regularisation, i.e., the 1-norm penalty on a discrete finite difference approximation of the two-dimensional gradient $\nabla : \mathbb{C}^\Omega \rightarrow (\mathbb{C}^2)^\Omega$ with zero Neumann boundary conditions, that is $\nabla u(i,j) = (\nabla_1 u(i,j), \nabla_2 u(i,j))^T$, with

$$\nabla_1 u(i,j) = \begin{cases} u(i+1,j) - u(i,j) & \text{if } i < n_1 \\ 0 & \text{if } i = n_1, \end{cases}$$

$$\nabla_2 u(i,j) = \begin{cases} u(i,j+1) - u(i,j) & \text{if } j < n_2 \\ 0 & \text{if } j = n_2, \end{cases}$$

for $i = 1, \dots, n_1$ and $j = 1, \dots, n_2$. Then the discrete total variation functional is given by

$$J(u) = \text{TV}(u) = \|\nabla u\|_{2,1}.$$

This constitutes a popular regularisation functional in image processing due to its edge-preserving properties (note that we fill up ∇ with zero rows so that the number of pixels is not decreased). Here, $\|\nabla u\|_{2,1}$ is defined as

$$\|\nabla u\|_{2,1} = \sum_{(i,j) \in \Omega} \sqrt{|\nabla_1 u(i,j)|^2 + |\nabla_2 u(i,j)|^2},$$

to obtain isotropic (direction-invariant) total variation regularisation, which is more suitable compared to anisotropic total variation (as it is used for instance in [14]) when dealing with circular structures as they appear in the data we are going to analyse in Section 3. The variational scheme (4) with $J = \text{TV}$ may then be written as

$$u_\alpha \in \arg \min_u \left\{ \frac{1}{2} \|f - S\mathcal{F}u\|_2^2 + \alpha \text{TV}(u) \right\}. \quad (6)$$

Total Generalised Variation Simplifying the mathematical intricacies, second-order Total Generalised Variation [19] may for parameters $(\beta, \alpha) > 0$ be written [31, 32] as the “differentiation cascade”

$$\text{TGV}_{(\beta, \alpha)}^2(u) = \min_w \alpha \|\nabla u - w\|_{2,1} + \beta \|\mathcal{E}w\|_{F,1}. \quad (7)$$

Here

$$\mathcal{E}w(i, j) = \frac{1}{2} (\nabla w(i, j) + (\nabla w(i, j))^T)$$

denotes a finite difference approximation of the symmetrised gradient of $w = (w_1, w_2) \in (\mathbb{C}^2)^\Omega$, where $\nabla w(i, j) = (\nabla w_1(i, j), \nabla w_2(i, j))$. Note that $\|\nabla u - w\|_{2,1}$ is defined above, while $\|\mathcal{E}w\|_{F,1}$ is defined analogously as $\sum_{(i,j) \in \Omega} \|\mathcal{E}w(i, j)\|_{F,1}$, with $\|\mathcal{E}w(i, j)\|_{F,1} = \sqrt{\sum_{k=1}^2 \sum_{\ell=1}^2 |(\mathcal{E}w(i, j))_{k,\ell}|}$ being the Frobenius norm of a matrix. If we fix the ratio β/α , the variational scheme (4) with $J = \text{TGV}_{(\beta/\alpha, 1)}^2$ may then be written as

$$u_\alpha \in \arg \min_u \left\{ \frac{1}{2} \|f - S\mathcal{F}u\|_2^2 + \alpha \text{TGV}_{(\beta/\alpha, 1)}^2(u) \right\}. \quad (8)$$

From the expression (7) we may observe how TGV^2 balances between first and second order features, by splitting low-cost second order features, as defined by the ratio β/α , into the variable w . This is the crucial advantage of TGV^2 over first-order TV, as we will discuss in further detail in Section 2.5.

2.2 Sampling Schemes

Various different sampling strategies have been proposed in the compressed sensing literature. Compressed sensing theory suggests that data should be sampled randomly throughout k -space [12, 13]. However, conventional MR images are obtained using frequency encoding in one direction, which means that an entire line in k -space is obtained from a single excitation. In this type of acquisition, an N -dimensional k -space is only undersampled in $N - 1$ dimensions. Ultra-fast imaging techniques utilising more complex trajectories are also possible, such as variable density spiral acquisitions [33]. These potentially enable undersampling in all N dimensions. However, variable density spiral acquisitions require all locations sampled to be connected smoothly and result in long read-out times, which can be problematic in some heterogeneous systems. We make a choice of eight schemes for sub-sampling velocity-encoded image data, including two purely random sampling schemes that guarantee large incoherence, as well as one line and one spiral sampling scheme. Random sampling would correspond to either pulse sequence implementations such as a purely phase-encoded scheme as might be used on a sample that would yield a low signal-to-noise ratio, or to the acquisition in the 2nd and 3rd dimensions of a three-dimensional image. Line sampling would correspond to a conventional two-dimensional imaging technique. The spiral scheme would be most suitable for very fast imaging (e.g. [8]). We have designed two variants of each of these schemes; the first variant samples 33% of the entire k -space, while the second variant samples 15%.

We have visualised the variants with 15% coverage of the k -space in Figure 2. For brevity, in the figure and the rest of this paper, we use the notation $\mathcal{N}_\sigma^{q\%}$ for two-dimensional Gaussian sampling with variance 128σ and $q\%$ coverage of the 256×256 total k -space. The notation $\mathcal{L}_\sigma^{q\%}$ stands for one-dimensional Gaussian sampling of lines with standard deviation 128σ , while $\mathcal{S}^{q\%}$ denotes a spiral pattern with $q\%$ of k -space frequencies covered.




Figure 2: Sub-sampling patterns employed. $\mathcal{N}_{\sigma}^{q\%}$ denotes two-dimensional Gaussian sampling with variance 128σ and $q\%$ coverage of the 256×256 total space. $\mathcal{L}_{\sigma}^{q\%}$ denotes one-dimensional Gaussian sampling of lines with variance 128σ , while $\mathcal{S}^{q\%}$ denotes a spiral pattern with $q\%$ coverage.

2.3 Contrast-Enhancement via Bregman Iteration

It is well known that for one-norm based regularisations as the ones described in Section 2.1 the solutions of variational schemes like (4) suffer from a loss of contrast (cf. [22, 23]). In order to compensate for this systematic bias, in [26] the use of the so-called Bregman iteration has been proposed. This procedure is a modification of (4) where the regulariser J has been replaced by what is called the generalised Bregman distance with respect to J (cf. [24, 25]).

It has been pointed out in [34] that Bregman iteration is equivalent to solving

$$u \in \arg \min_u \{J(u)\} \quad \text{subject to} \quad f = \mathcal{S}\mathcal{F}u \quad (9)$$

numerically via a scaled alternating direction method of multipliers (ADMM) of the augmented Lagrangian formulation of (9). Thus, after appropriate rescaling and substitution of variables, Bregman iteration equals

$$u_{\alpha}^k \in \arg \min_u \left\{ \frac{1}{2} \|f^{k-1} - \mathcal{S}\mathcal{F}u\|_2^2 + \alpha J(u) \right\}, \quad (10a)$$

$$f^k = f^{k-1} + f - \mathcal{S}\mathcal{F}u_{\alpha}^k, \quad (10b)$$

for $f^0 = f$. Hence, instead of solving (4) once for a fixed α we solve a sequence of similar problems but add back the residual to the data in each iterate. This intuitively explains why the contrast is enhanced. However, one might argue that by adding back the residual, errors will also accumulate. Remarkably, this is not the case for one-norm based regularisers, as it has been pointed out in [23, 35] (for appropriately chosen α and k).

2.4 Discrepancy Principle and Image Quality Measures

Discrepancy Principle A significant issue in the application of variational regularisation schemes is the correct choice of the regularisation parameter α . It depends on both the noise level σ and the sampling scheme in a non-linear fashion. One standard strategy is Morozov’s discrepancy principle [27]. The idea is to choose as large α as possible that $u = u_{\alpha}$ satisfies

$$\|f - \mathcal{S}\mathcal{F}u\|_2 \leq \sigma\sqrt{m}, \quad (11)$$

where m is the number of samples. It has been used widely in the inverse problems literature (e.g. [36]), although often visually satisfactory results can be obtained using alternative stopping criteria; see, e.g., [37]. We will, however, find that the discrepancy principle works very well in our application in combination with the Bregman iteration. Here the idea, which also goes back to [26], is to choose a significantly large fixed α and stop iterations (10) as soon as $u = u_{\alpha}^k$ satisfies (11). Observe that this involves solving the variational problem (4) multiple times. However, this additional computational effort allows us to find solutions that can be superior to solutions of (3). Assume for instance that we want to solve (3) for $J(u) = \|u\|_1$, $f = (1, 0)^T$, $\sigma = 0$ and $\mathcal{S}\mathcal{F}$ given in terms of

$$\mathcal{S}\mathcal{F} = \begin{pmatrix} \frac{1}{\sqrt{2}} & 1 & 0 \\ \frac{1}{\sqrt{2}} & 0 & 1 \end{pmatrix}.$$

Following [23] it can easily be seen that the one-norm minimizing solution u of $\mathcal{S}\mathcal{F}u = f$ is given by $u = (0, 1, 0)^T$. On the other hand, u does not fulfil the optimality condition of (4) for any $\alpha > 0$. It would satisfy it in the limit $\alpha = 0$, but due to the ill-posedness of the problem we cannot in practice solve (3) via (4) in that case. In contrast to that, for the Bregman iteration (10) with $1 > \alpha > 0$, we would obtain $u_{\alpha}^1 = (1 - \alpha)u$ and $u_{\alpha}^l = u$ for $l \geq 2$, thus solving (3). Thus, the Bregman iteration-based solution does not suffer loss of contrast. In Section 3.3 we will further support this result using computational simulation of noisy data.

Image Quality Measures In order to compare the different regularisation functionals J , and to justify the use of the discrepancy principle for stopping the Bregman iterations, we have to be able to compare the quality of solutions u_α^k to the scheme (10) with respect to a ground-truth u . The standard quality measure is the so-called peak signal-to-noise ratio (PSNR), which is defined with the dynamic range $L = \max_i |u_i|$ as

$$\text{PSNR}(u, u_\alpha) = 10 \log_{10} \left(\frac{L}{\|u - u_\alpha\|_2^2 / (n_1 n_2)} \right).$$

The perfect reconstruction would have $\text{PSNR} = \infty$.

Another recently proposed measure for assessing the quality of the reconstruction u_α is the structural similarity index (SSIM). It is based on computing in a window around every pixel of the image the local structural similarity measure

$$\text{LSSIM}(u, u_\alpha) = \frac{(2\mu_u \mu_{u_\alpha} + C_1)(2\sigma_{uu_\alpha} + C_2)}{(\mu_u^2 + \mu_{u_\alpha}^2 + C_1)(\sigma_u^2 + \sigma_{u_\alpha}^2 + C_2)}.$$

Here μ_u, σ_u^2 are the mean and variance of u in the local Gaussian window, and σ_{uu_α} is, likewise, the covariance between u and u_α in the local window. The global SSIM is then computed by averaging the LSSIM estimates over the whole image; for the exact expression we refer to [38]. In contrast to traditional quality measures like the PSNR, the SSIM index also assesses the conservation of the structural information of the reconstructed image, compare [38] for details. Note that the perfect reconstruction would have SSIM value equal to 1.

2.5 Comparison of TV, TGV and WAV

The regularisers presented in Section 2.1 all have different benefits and drawbacks. As a prototype wavelet basis we choose the six-level Daubechies wavelets. This constitutes an orthonormal transform that allows to sparsify relatively smooth signals. Moreover, due to the orthonormality, the effort required to implement the corresponding regulariser is comparably low. However, this wavelet basis might introduce scaling artefacts for large regularisation parameters, which is a common problem within the class of complete wavelet bases (cf. [39, page 229]). This may be due to the commitment of the so-called wavelet crime [39, page 232]; however, discussing effective pre-filtering of our signal to avoid the wavelet-crime is beyond the scope of this paper. Moreover, the Daubechies wavelet basis is not capable of recovering sharp transitions like edges. This is often desirable in imaging applications, including the present application of a packed bed of solid particles.

Some of the drawbacks of the Daubechies wavelet basis could be overcome by the use of another basis, such as the Haar basis, which is capable of recovering edges, to an extent. A drawback is that it does not preserve smooth parts of the signal. In this way the Haar basis is similar to Total Variation (TV). In theory and in practice TV however has advantages over the Haar wavelet. This is why we have chosen it as the representative first-order non-smooth regulariser.

Total Variation however suffers from the *stair-casing effect*. It promotes sparsity of the gradient, which results in flat areas with sharp transitions in the reconstruction. In order to overcome this problem, higher-order approaches have been recently proposed. These move the sparsity to higher-order gradients, while still preserving edges and other non-smooth features. The earliest work in this category is to our knowledge [40]. Total Generalised Variation, as defined above, is a more recent and sophisticated approach that has experimental and theoretical advantages over approaches proposed earlier; compare, e.g., [41, 42, 43]. It does, however, have the drawback of causing flat regions to slant. This will also be observed around the solid particles in our magnitude reconstructions below.

3 Results

To test the regularisation functions and optimisation procedures with a noise-free data set, a velocity map was obtained from a lattice-Boltzmann (LB) simulation of water flowing through a 39 mm diameter column that was randomly packed with 3 mm diameter spheres. The geometry for the simulation was derived from a three-dimensional spin-echo MR image of the bed that was obtained at a resolution of $164 \mu\text{m} \times 164 \mu\text{m} \times 164 \mu\text{m}$. This data set was converted to a binary image and used as input to the LB simulation. The LB method models the hydrodynamics of a fluid on a mesoscopic scale, where fluid kinetics are described by probability distribution functions of the location of particles. The LB code used has been described and validated previously [44, 45]. The time step for the simulations was 4.5×10^{-3} s, such that the dimensionless relaxation was 1, which optimises the simulation accuracy and stability [45]. The velocity map resulting from the LB simulation was converted to a phase-encoded image suitable for testing the CS algorithm by selecting a single axial (z) slice from the centre of the image. This velocity image was converted to a phase map by assuming that the maximum velocity in the image corresponded to a phase shift of 1.8π radians, thus ensuring that the total range of velocities was less than 2π radians. The phase map and the binary signal intensity map were then converted




Figure 3: The computer-generated test data, magnitude on the left and phase on the right. The rectangles indicate the detail on which we concentrate our attention on.

to a complex image, thus giving real and imaginary intensity data. The data are discretised on a grid of size $256 \times 256 \times 256$.

Out of this 3D data set, we extract slice 100 in the z direction. This is displayed in Figure 3, where we also indicate a detail on which we concentrate our attention in the later figures containing computational results. We add normal-distributed noise to the computer-generated 2D image constructed this way. This is done for two different noise levels, the “low-noise case” with pointwise standard deviation $\sigma = 0.2$ and the “high-noise case” with $\sigma = 0.5$. We then sub-sample the noisy image using all the sampling patterns described in Section 2.2. For this data set, $\Omega = \{1, \dots, n\}^2$ with $n = 256$.

We begin by describing how we have obtained our computational results in Section 3.1. We then move on to analysing these results. First, in Section 3.2, we study the effect of the different sampling schemes on the reconstruction quality. Then, in Section 3.3 we verify the contrast-enhancing effect of the Bregman iteration for our application. We are then going to focus on justifying the use of the discrepancy principle as a stopping criterion in Section 3.4. To conclude our analysis on the computer-generated data set, we compare the reconstruction quality of the different regularisers in Section 3.5. We want to emphasise that we claim validity of all our conclusions only for the particular data set described above. However, we believe that our results are a good indicator of the overall performance of the considered approaches, working as a guideline for the selection of regularisation methods and sampling patterns for velocity-encoded MRI imaging. To support this claim, we finish our comparative evaluation in Section 3.6 by applying the methods on a real data set.

3.1 Regularisation Methods

We are now going to describe the details of the computational realisation of each of the variational regularisation approaches, whose application to MRI phase recovery we evaluate. As we recall, these approaches are Bregmanised regularisation by Daubechies wavelets, total variation, and second-order total generalised variation. For each of these approaches, we consider a range of regularisation parameters and fix the maximum number of Bregman iterations. Moreover, for each approach we pick a numerical method for the approximate solution of the corresponding minimisation problem in each Bregman iteration.

Computational Realisation for WAV We compute (10) with $J = \text{WAV}$ for 100 values of α from the interval $[0.1, 10]$ with equidistant spacing 0.1. The Daubechies wavelet transformation is calculated with WaveLab [46]. We perform 20 Bregman iterations for each α . The computations are done by solving each iterate of (10) via ADMM applied to the augmented Lagrangian formulation of (4), see for instance [47]. With a slightly different scaling this is also known as the split Bregman method [48]. We choose the additional stepsize parameter of the ADMM method according to a heuristic rule presented in [49]. For each α , the computation for the first Bregman iterate is initialised with the least squares solution $\mathcal{F}^{-1}S^T f$, also known as the backprojection or zero-filling solution. The computations for all subsequent Bregman iterates are initialised with the previous Bregman iterate. Computations are stopped as soon as two subsequent ADMM iterates satisfy $\|u_j - u_{j-1}\|_2 / \|u_j\|_2 \leq 10^{-4}$.

Computational Realisation for TV and TGV² For the numerical solution of (10) with $J = \text{TV}$ and $J = \text{TGV}^2$, we perform the computations for

$$\alpha \in \{0.05, 0.075, 0.1, 0.2, 0.25, 0.5, 1.0, 1.5, 2.0, 2.5, 3, 4, 5\}$$

and, for TGV^2 , the second-order parameter

$$\beta \in \{1, 1.1, 2, 10\}\alpha.$$

We use the formulation (10a)–(10b) of the Bregman iterative scheme. For the solution of (10a) we write it in saddle-point form and use the Chambolle-Pock primal-dual hybrid gradient method (PDHGM) [50]. Details of our implementation of the method and parameter choice are described in [43, 51, 52].

In each Bregman iteration, we perform a maximum of 10000 iterations of the Chambolle-Pock algorithm, stopping earlier if the fractional pseudo-duality gap, described in detail in [43], is less than 0.1% of the fractional pseudo-duality gap for a zero initialisation. We use this zero initialisation of the Chambolle-Pock algorithm for the first Bregman step, $k = 0$. As in the wavelet case, for the subsequent steps, $k \geq 1$, we initialise the method with the results of the previous Bregman step. Experiments suggest that this provides faster convergence than zero initialisation.

We want to point out that the use of different computational methods (ADMM versus PDHGM) for the different regularisers does not affect the results. Moreover, the Chambolle-Pock method can be seen as a preconditioned ADMM [53]. Both methods applied to (10a) yield the same results up to computational precision for the same choice of J and α . The use of differing methods is based on the availability of existing implementations for the different regularisation functionals.

3.2 Sampling Schemes

For evaluating the quality of reconstruction for the different sampling schemes from Section 2.2, we employ the two error measures, PSNR and SSIM, introduced in Section 2.4. We report the value of the chosen error measure (SSIM or PSNR of phase or magnitude) at both the optimal Bregman iterate and at the iterate where the discrepancy principle is violated. The numerical values are presented in Tables 1–4 for all of the variational approaches and the backprojection, explained in Section 3.1. For $J = \text{TV}, \text{TGV}^2$ we use $\alpha = 1.6$ and for $J = \text{WAV}$ we use $\alpha = 6$. These choices will be justified in Section 3.4, as we study the performance of the discrepancy principle. Note that we have decided not to include values for the 15% line sampling strategy, as it did not perform well in practice.

Studying Tables 1–4, we observe that the 33% line sampling scheme $\mathcal{L}_{0.35}^{33\%}$ and the widely spread 15% normal-distributed scheme $\mathcal{N}_{0.35}^{15\%}$ perform worse than the other sampling strategies. For the latter, this might be expected as we miss too much of the centre of k -space. For the former, the samples are too coherent in one direction of k -space, while missing the other. In fact, several 15% schemes perform better than $\mathcal{L}_{0.35}^{33\%}$, although less of k -space is sampled. As a consequence, we are going to eliminate these two poorly-performing schemes from further consideration.

Our second observation is that the sample count to quality ratio of $\mathcal{N}_{0.15}^{15\%}$ is clearly the best. Overall $\mathcal{N}_{0.15}^{33\%}$ provides slightly better performance, but at the cost of significantly more samples.

Based on the above observations, and for the sake of conciseness, we mainly focus on just two sampling patterns in the remainder of this comparative analysis. As the first sampling pattern we pick $\mathcal{N}_{0.15}^{15\%}$ that we have observed to provide the best sample count to quality ratio. We were unable to detect significant differences in performance within the 33% sampling patterns, with the exception of line sampling. As spiral patterns are in practise widely used for velocity-encoded MRI, we pick $\mathcal{S}^{33\%}$ as our second pattern of choice.

3.3 Contrast-Enhancement via Bregman Iteration

As described in Section 2.3, we employ Bregman iterations in order to correct for the inevitable loss of contrast, characteristic of quadratic fidelity terms. That is, we compute solutions to (10) rather than (4). To visualise the improvement of contrast, in Figure 4 we compare a solution u_2 of (10), which has been stopped via (11), to a solution u_1 of (4). We want to emphasize that in order to guarantee an absolute fair comparison, we have computed u_1 such that $\|f - \mathcal{S}\mathcal{F}u_1\|_2 = \|f - \mathcal{S}\mathcal{F}u_2\|_2 = \bar{\sigma} \leq \sigma$, instead of $\|f - \mathcal{S}\mathcal{F}u_1\|_2 = \sigma$. Note that in order to do so, we need to pick an optimal regularisation α parameter to guarantee this discrepancy, which can be computationally as complex as performing several Bregman iteration steps.

It can be observed that the contrast of the Bregman iteration-based reconstruction is better than the Tikhonov-based reconstruction, in particular by looking at the line profile in Figure 4b. This is due to the fact that the TV value of the Bregman iteration-based reconstruction is $\text{TV}(u_2) = 27.96$, while the TV value of the Tikhonov-based reconstruction is $\text{TV}(u_1) = 25.03$, though both reconstructions have exactly the same discrepancy with respect to the noisy measurement data. By allowing for larger values of the regularisation functional, Bregman iteration realises contrast-enhanced reconstructions compared to solutions of (3) for identical σ . Further, according to [23] this holds true for all the regularisers considered in this work.

Table 1: PSNR values of the complex signal magnitude for various sampling patterns: optimal value over Bregman iterations (OV), and value at violation of discrepancy principle (DV). For $J = \text{TV}, \text{TGV}^2$ we use parameter $\alpha = 1.6$ and for $J = \text{WAV}$ we use $\alpha = 6$.

Method	Crite- rion	$\mathcal{N}_{0.35}^{33\%}$	$\mathcal{N}_{0.35}^{15\%}$	$\mathcal{N}_{0.15}^{33\%}$	$\mathcal{N}_{0.15}^{15\%}$	$\mathcal{S}^{33\%}$	$\mathcal{S}^{15\%}$	$\mathcal{L}_{0.35}^{33\%}$
BP		18.9	15.7	20.4	20.3	19.4	17.8	16.0
WAV	OV	19.8	15.9	20.4	19.6	19.5	17.3	15.9
WAV	DV	19.6	15.9	20.3	19.6	19.5	17.3	15.9
TV	OV	23.7	19.1	24.1	22.2	23.3	20.6	20.2
TV	DV	23.7	19.0	24.0	22.1	23.1	20.4	20.0
TGV ²	OV	23.4	15.8	24.0	21.5	22.8	19.4	18.6
TGV ²	DV	23.2	14.4	23.7	21.1	22.2	19.1	18.0

Table 2: PSNR values of the complex signal phase for various sampling patterns: optimal value over Bregman iterations (OV), and value at violation of discrepancy principle (DV). For $J = \text{TV}, \text{TGV}^2$ we use parameter $\alpha = 1.6$ and for $J = \text{WAV}$ we use $\alpha = 6$.

Method	Crite- rion	$\mathcal{N}_{0.35}^{33\%}$	$\mathcal{N}_{0.35}^{15\%}$	$\mathcal{N}_{0.15}^{33\%}$	$\mathcal{N}_{0.15}^{15\%}$	$\mathcal{S}^{33\%}$	$\mathcal{S}^{15\%}$	$\mathcal{L}_{0.35}^{33\%}$
BP		29.3	24.7	31.8	31.5	29.7	26.8	24.6
WAV	OV	29.2	23.5	30.4	29.3	28.7	25.4	23.2
WAV	DV	29.1	23.5	30.4	29.2	28.7	25.3	23.2
TV	OV	31.4	26.7	32.2	31.1	31.0	28.0	27.1
TV	DV	30.3	26.2	30.4	30.1	30.0	27.5	26.5
TGV ²	OV	31.4	22.6	32.4	31.3	30.9	27.5	26.0
TGV ²	DV	31.1	20.4	31.6	30.8	30.5	27.5	24.7

Table 3: SSIM values of the complex signal magnitude for various sampling patterns: optimal value over Bregman iterations (OV), and value at violation of discrepancy principle (DV). For $J = \text{TV}, \text{TGV}^2$ we use parameter $\alpha = 1.6$ and for $J = \text{WAV}$ we use $\alpha = 6$.

Method	Crite- rion	$\mathcal{N}_{0.35}^{33\%}$	$\mathcal{N}_{0.35}^{15\%}$	$\mathcal{N}_{0.15}^{33\%}$	$\mathcal{N}_{0.15}^{15\%}$	$\mathcal{S}^{33\%}$	$\mathcal{S}^{15\%}$	$\mathcal{L}_{0.35}^{33\%}$
BP		0.436	0.288	0.475	0.461	0.440	0.347	0.312
WAV	OV	0.457	0.268	0.514	0.475	0.454	0.312	0.281
WAV	DV	0.439	0.266	0.483	0.470	0.432	0.312	0.279
TV	OV	0.754	0.448	0.773	0.720	0.746	0.585	0.569
TV	DV	0.747	0.447	0.773	0.717	0.730	0.559	0.563
TGV ²	OV	0.616	0.279	0.665	0.544	0.586	0.433	0.443
TGV ²	DV	0.616	0.268	0.665	0.541	0.584	0.409	0.436

Table 4: SSIM values of the complex signal phase for various sampling patterns: optimal value over Bregman iterations (OV), and value at violation of discrepancy principle (DV). For $J = \text{TV}, \text{TGV}^2$ we use parameter $\alpha = 1.6$ and for $J = \text{WAV}$ we use $\alpha = 6$.

Method	Crite- rion	$\mathcal{N}_{0.35}^{33\%}$	$\mathcal{N}_{0.35}^{15\%}$	$\mathcal{N}_{0.15}^{33\%}$	$\mathcal{N}_{0.15}^{15\%}$	$\mathcal{S}^{33\%}$	$\mathcal{S}^{15\%}$	$\mathcal{L}_{0.35}^{33\%}$
BP		0.906	0.806	0.947	0.948	0.917	0.858	0.807
WAV	OV	0.914	0.771	0.937	0.923	0.910	0.823	0.760
WAV	DV	0.914	0.768	0.937	0.921	0.909	0.823	0.769
TV	OV	0.946	0.865	0.955	0.947	0.942	0.894	0.873
TV	DV	0.936	0.851	0.937	0.936	0.931	0.886	0.861
TGV ²	OV	0.947	0.755	0.959	0.951	0.942	0.886	0.847
TGV ²	DV	0.945	0.674	0.953	0.948	0.941	0.886	0.816

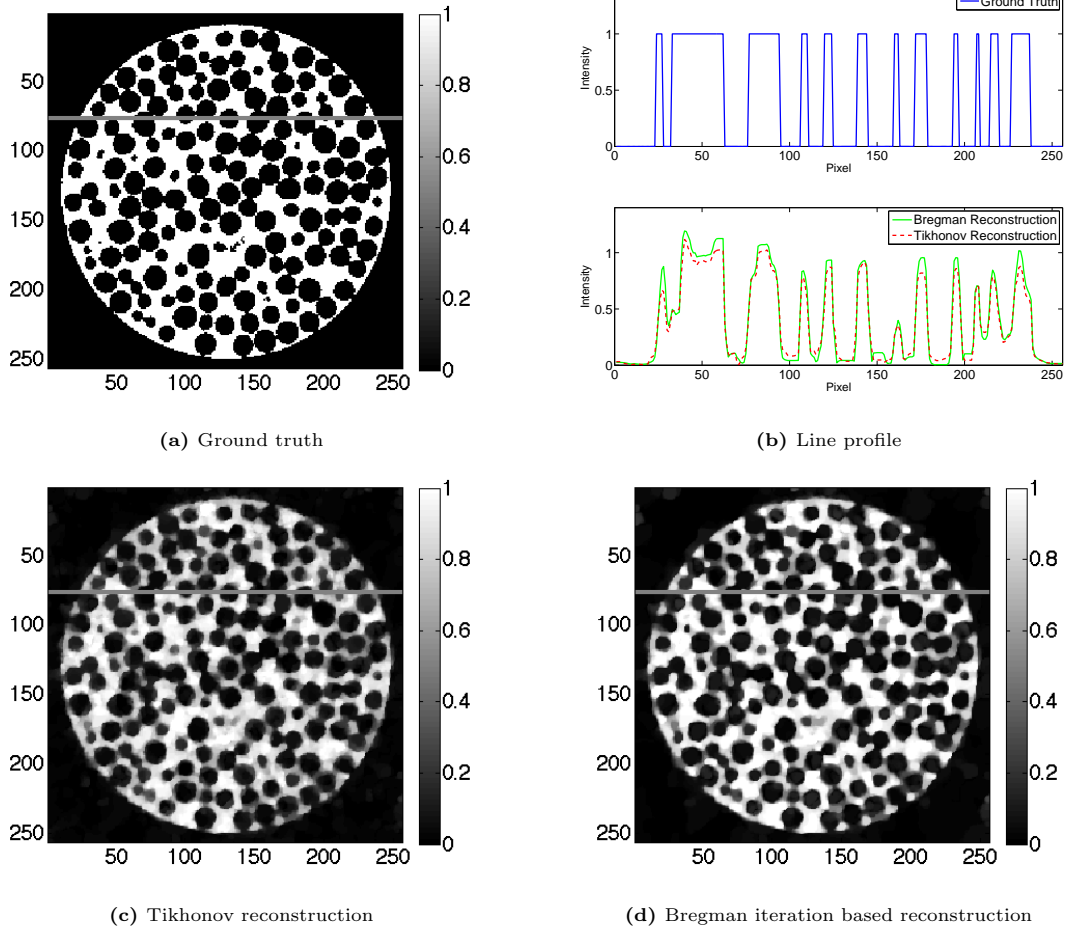


Figure 4: Difference between a solution u_1 of (3) and u_2 of (10), for $J(u) = \text{TV}(u)$ and the spiral sampling pattern $\mathcal{S}^{15\%}$. We denote the value of $\|f - \mathcal{S}\mathcal{F}u_2\|_2$ by $\tilde{\sigma}$ and choose α such that $\|f - \mathcal{S}\mathcal{F}u_1\|_2 = \|f - \mathcal{S}\mathcal{F}u_2\|_2 = \tilde{\sigma}$. Figure 4a shows the magnitude of the ground truth described in Section 3, and a horizontal line at pixel 76 indicating the profile that is being visualised in Figure 4b. Figure 4c and Figure 4d show the magnitude of u_1 and u_2 , respectively. It is clearly visible from the line profile, but also from the magnitude images, that the contrast of the Bregman iterated reconstruction is better preserved when compared to the ground truth.

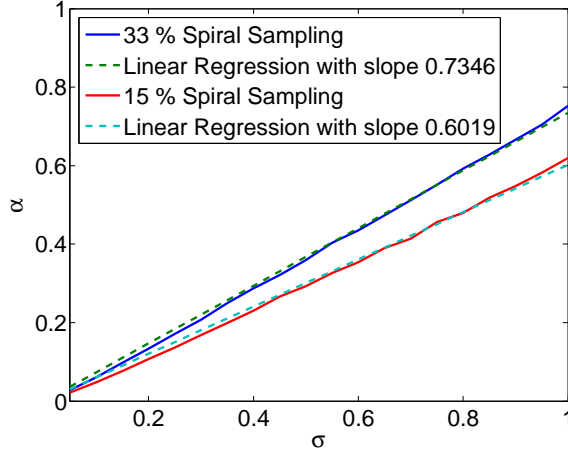


Figure 5: A plot of $\alpha(\sigma)$ against the noiselevels $\sigma = \|\mathcal{S}\mathcal{F}u_\alpha - y_\sigma\|_2$. Here, u_α is the solution of (4) for input data $y_\sigma = y + n_\sigma$, with n_σ being normal-distributed with zero mean and standard deviation σ . The clean data y is produced via $y = \mathcal{S}\mathcal{F}u$, where u is the velocity phantom described above, and \mathcal{S} is the sampling operator derived from either $\mathcal{S}^{33\%}$ (blue line) or $\mathcal{S}^{15\%}$ (red line).

What remains is to further quantify the choice of the regularisation parameter α . In the absence of noise we want to point out that the amount of samples basically affects the scales of features that can be recovered, but does not seem to affect the scaling of the regularisation parameter (unless we miss the centre frequency y_0). That is due to the fact that sub-sampling only alters the number of k -space elements, but not their scaling. However, if we change either mean (and therefore the scaling) or the number of grid points of the image, this clearly affects the scaling of the data in k -space. For the normalized discrete Fourier-transform that satisfies $\mathcal{F}^*\mathcal{F} = I$ the centre k -space frequency equals

$$y_0 = \frac{1}{\sqrt{n_1 n_2}} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} u_{i,j} = \sqrt{n_1 n_2} \text{mean}(u).$$

Thus, either rescaling u by a factor c or by regridding u with a factor of c (leading to a multiplication of the product $n_1 n_2$ with c^2) yields to a scaling cy in k -space, and therefore requires a rescaling of α by the same factor in order to obtain the similar results. Thus, if we have a reference choice $\hat{\alpha}$ for the Bregman iteration for a given k -space centre frequency \hat{y}_0 and number of pixels $\hat{n}_1 \hat{n}_2$, we can basically rescale $\hat{\alpha}$ for a k -space centre frequency \tilde{y}_0 , and number of pixels $\tilde{n}_1 \tilde{n}_2$ via

$$\tilde{\alpha} = \sqrt{\frac{\tilde{n}_1 \tilde{n}_2 |\tilde{y}_0|}{\hat{n}_1 \hat{n}_2 |\hat{y}_0|}} \hat{\alpha}. \quad (12)$$

The scale of the magnitude image of the simulated data described in this section is $[0, 1]$, whereas the mean of the absolute image is about 0.316. In order to obtain similar results on any other dataset as in this paper, the data should be scaled accordingly.

However, in the presence of noise it can be observed that the number of samples m affects the scaling of the regularization parameter. For instance, for the spiral sampling schemes $\mathcal{S}^{33\%}$ and $\mathcal{S}^{15\%}$ we observe a behaviour as described in Figure 5. The plot in this figure shows the regularization parameters α that correspond to values $\sigma = \|\mathcal{S}\mathcal{F}u_\alpha - y\|_2$. There is a linear relation between the choice of α and the values of σ , but we also discover that the slope of these linear relations differs for the two sampling schemes $\mathcal{S}^{15\%}$ and $\mathcal{S}^{33\%}$.

Note however, that in order to find a suitable α for the Bregman iteration we basically just need to make sure that α is sufficiently large. From Figure 5 we observe that a larger m yields a larger α , and thus, a large choice of α for a fully sampled reconstruction will work for an under-sampled one as well (at the cost of slightly more Bregman iterations). Nevertheless, from experimental validation we usually found

$$\alpha \sim c \|\mathcal{F}^{-1} \mathcal{S}^T y\|_\infty, \quad (13)$$

to be a good choice to create an over-regularized initial Bregman iterate, with c depending on the scaling of the regularisation operator in J .

3.4 Discrepancy Principle and Image Quality Measures

We now justify the use of the discrepancy principle to identify the point at which to stop the Bregman iterations. For the sake of brevity, we concentrate on $J = \text{TV}$, remarking that the observations derived are generally

comparable for $J = \text{TGV}^2$ and $J = \text{WAV}$. As described, we compute the solutions of (10) for a range of parameters α and iterations $k = 1, \dots, 15$. This is done for all the sampling patterns with the exception of line sampling. We plot in Figure 6 and Figure 7, respectively, colour-coded SSIM and PSNR values of α and k . In this, we follow the visualisation approach employed in [42]. Superimposed over the colour-coded results is the choice of the Bregman iterate by the discrepancy principle (white line) as well as by the optimal SSIM or PSNR (dashed black line). We also plot the optimal (dashed line) and discrepancy principle (solid line) SSIM or PSNR values against α below the colour-coded graph along with the value at the initial iterate (dash-dotted line).

Our principal observations from the analysis of Figures 6 and 7 are as follows. First of all, the lines indicating, respectively, the violation of the discrepancy principle and the optimal PSNR or SSIM value, both follow the same trend. Typically for the magnitude, the two lines are off by just one or two iterates. For PSNR the violation of the discrepancy principle is reached slightly earlier than the optimal quality measure value, while for SSIM the situation is reversed. For the phase the picture looks similar for both SSIM and PSNR, with the discrepancy principle being reached earlier than the optimal quality measure value. Here, in contrast to the magnitude, a larger number of iterations may be needed after the violation of the discrepancy principle to reach an optimal result. Nevertheless, the exact difference in value at both of the iterations is small.

We now verify that this difference for the two sampling patterns $\mathcal{N}_{0.15}^{15\%}$ and $\mathcal{S}^{33\%}$ is indeed small. The corresponding reconstructions for $\mathcal{S}^{33\%}$ and $\alpha = 1.5$ are plotted in Figure 8. Visually, we in fact prefer the overall appearance of the reconstructions obtained at the violation of the discrepancy principle. The reconstructions obtained at optimal SSIM or PSNR value reintroduce more wanted small-scale structures at the cost of unwanted noise. It is also worth mentioning that it is difficult to find significant differences between the considered sampling schemes with regard to the performance of the discrepancy principle.

As a short remark on the case $J = \text{WAV}$, we have observed that the agreement between the discrepancy principle and optimal iterate is even better. This, however, does not imply the superiority of wavelets in terms of quality of reconstruction. We will discuss this in the next section.

Returning to Figures 6 and 7, we observe that the plots support the view that $1/\alpha$ behaves as a step size parameter for the approximate solution of the constrained problem (9). Indeed, larger choices of α correspond to smaller differences of SSIM and PSNR between subsequent iterates. This leads to the question of how to choose α initially. Smaller values lead to an earlier violation of the discrepancy principle. However, for such a choice we may miss an optimal solution, because the accuracy of the iterations is low for small α (i.e., large step size). The goal is then to choose α such that the correct balance between accuracy and computational cost (number of iterations) is obtained. Our results indicate that a good choice for the present data set and $J = \text{TV}$ is $\alpha = 1.5$. As the regularisation parameter is primarily influenced by the domain Ω , the noise level σ , and the sampling pattern, a similar choice should be reasonable for other problems with similar experimental setup. For $J = \text{TGV}^2$ we find that the same (large) choice of α works as for TV, since the regularisation functionals are closely related. For $J = \text{WAV}$ the parameter α in general lies in a different range, due to the different scale of the operator. For the present setting we find the choice $\alpha = 6$ reasonable.

3.5 Comparison of TV, TGV and WAV

We are now going to compare computational results of (10) between $J = \text{WAV}$, $J = \text{TV}$ and $J = \text{TGV}$. We consider two noise levels, $\sigma = 0.2$ and $\sigma = 0.5$, and the two sampling patterns $\mathcal{N}_{0.15}^{15\%}$ and $\mathcal{S}^{33\%}$ that we have chosen above.

For $J = \text{TV}$ and $J = \text{TGV}^2$ we use the choice $\alpha = 1.5$ of the first-order regularisation parameter, that we have justified above. Moreover, for TGV^2 we pick the second order regularisation parameter $\beta = 1.1\alpha = 1.65$. The study of SSIM/PSNR values (from similar plots as Figures 6 and 7 for TV) and visual inspection of the results indicate that this is a good choice for all the different sampling patterns and methods. It will be observed to provide improved phase reconstruction in comparison to TV at the cost of some loss in the quality of magnitude reconstruction. The choice $\beta = 2\alpha$ would provide results closer to TV in both respects – in particular with phase reconstruction worse than with the present choice of $\beta = 1.1\alpha$, which provides a good balance between the reconstruction quality of magnitude and phase. This choice of β is in the usual range $[1, 10]\alpha$ of sensible values for β on $\Omega = [1, 256]^2$, cf. [43, 51]. The exact choice depends on the qualities desired from the solution.

Low-Noise Results For these results we apply additional Gaussian noise of standard deviation $\sigma = 0.2$ to the original noise-free computer-generated data. The quality measures of the results for all methods are displayed in Table 5 and Table 6, respectively, for the two different sampling patterns. The corresponding reconstructions are visualised in Figure 9 and Figure 10, respectively. Some of the artefacts that we discuss in the next paragraph are emphasised as an example with arrows in Figure 9.

Our main observation is that $J = \text{TGV}^2$ performs best with regard to the reconstruction of the phase, while $J = \text{TV}$ can be considered to be somewhat better in reconstructing the magnitude. The choice $J = \text{WAV}$




Figure 6: SSIM of magnitude (ρ , top row) and phase (ϕ , bottom row) for Bregmanised TV reconstruction of the computer-generated test data with various sampling patterns and noise $\sigma = 0.2$. The parameter α is on the horizontal and the Bregman iteration on the vertical axis. The colours code the SSIM value, also shown in the small lower graph. The continuous line corresponds to violation of the discrepancy principle, and the dashed line to the optimal SSIM. The dash-dotted line in the small graph indicates the SSIM for the first iteration.




Figure 7: PSNR of magnitude (ρ , top row) and phase (ϕ , bottom row) for Bregmanised TV reconstruction of the computer-generated test data with various sampling patterns and noise $\sigma = 0.2$. The parameter α is on the horizontal and the Bregman iteration on the vertical axis. The colours code the PSNR value, also shown in the small lower graph. The continuous line corresponds to violation of the discrepancy principle, and the dashed line to the optimal PSNR. The dash-dotted line in the small graph indicates the PSNR for the first iteration.




Figure 8: Optimal reconstructions of (10) with $J = \text{TV}$ and $\alpha = 1.5$ for the sampling pattern $\mathcal{S}^{33\%}$ in terms of discrepancy, and SSIM and PSNR of the phase of the complex image.

performs clearly worst in our experiments, introducing significant artefacts. The rationale for the differences between TV and TGV^2 is that the small choice of β in TGV^2 allows good reconstruction of the smooth higher-order features in the phase, while TV, closely corresponding to a large choice of β in TGV^2 , is good at reconstructing piecewise-constant functions. This is the situation with the magnitude picture. As indicated in Figure 9, TGV^2 introduces some waviness and slanting in the magnitude, too closely modelling the higher-order features introduced by noise and sparse sampling. On the other hand, TV tries to approximate the originally smooth phase by a piecewise constant function, thus exhibiting the well-known stair-casing effect. This prohibits the reconstruction of small-scale features. TGV^2 performs better at reconstructing small-scale features of the phase, and preserving its smoothness.

High-Noise Results For these results we apply additional Gaussian noise of standard deviation $\sigma = 0.5$ to the original noise-free computer-generated data. The quality measure values of all the methods are displayed in Table 7 and Table 8 for two different sampling patterns. The corresponding reconstructions are visualised in Figure 11 and Figure 12, respectively. Our conclusions from observing these results are more or less the same as in the low noise case. The stair-casing of TV is even more prominent now, as is the over-smoothing of the magnitude by TGV^2 . However, the phase reconstruction by TGV^2 is not significantly degraded from the low-noise test case.

Proposed Guideline The analysis presented suggests the following:

1. Pick a regularisation functional J with desirable characteristics.
2. Select parameter α according to (12) and (13), with c chosen according to the scaling of the regulariser, and to balance between the requirement of efficient computation (small α) and high precision (the larger α , the better).
3. In the case of TGV^2 regularisation, select the 2nd parameter β . The object here is to balance between the requirement of good recovery of smooth features ($\beta \sim \alpha$), and the recovery of flat areas ($\beta \gg \alpha$). We stress that these ratios hold for spatial discretisation with step size $h = 1$, as is the case with ∇ defined as in Section 2.1, and typical image sizes. In general the ratio between β/α is multiplied by h .
4. Estimate noise level σ for data; ideally it would be known or is obtained with an additional experiment.
5. Iterate (10) until the discrepancy principle (11) is violated.

3.6 Application to Real Data

In the previous sections, we have analysed extensive computations on computer-generated test data. We have sought to justify the use of the discrepancy principle as a stopping criterion for Bregman iterations, and to find a good initial regularisation parameter α . Moreover, we have evaluated a range of sampling patterns for the




Figure 9: Results of Bregmanised TV, TGV², and WAV reconstruction with the discrepancy principle as stopping criterion, along with the backprojection, of the computer-generated test data with noise $\sigma = 0.2$ for the sampling pattern $\mathcal{N}_{0.15}^{15\%}$. The top picture presents the magnitude and the bottom picture the phase. The arrows in (d) indicate areas where the stair-casing of TV is particularly noticeable. In (e) the arrow points to the “slanting” of TGV², noticeable around other pebbles as well. The “blockiness” of WAV is visible overall in (c).




Figure 10: Results of Bregmanised TV, TGV², and WAV reconstruction with discrepancy principle as stopping criterion, along with the backprojection, of the computer-generated test data with noise $\sigma = 0.2$ for the sampling pattern $\mathcal{S}^{33\%}$. The top picture presents the magnitude and the bottom picture the phase.




Figure 11: Results of Bregmanised TV, TGV^2 , and WAV reconstruction with discrepancy principle as stopping criterion, along with the backprojection, of the computer-generated test data with noise $\sigma = 0.5$ for the sampling pattern $\mathcal{N}_{0.15}^{15\%}$. The top picture presents the magnitude and the bottom picture the phase.




Figure 12: Results of Bregmanised TV, TGV^2 , and WAV reconstruction with discrepancy principle as stopping criterion, along with the the backprojection, of the computer-generated test data with noise $\sigma = 0.5$ for the sampling pattern $\mathcal{S}^{15\%}$. The top picture presents the magnitude and the bottom picture the phase.

Table 5: SSIM and PSNR values of both magnitude (ρ) and phase (ϕ) for Bregmanised TV, TGV² ($\alpha = 1.5$) and WAV ($\alpha = 6$) reconstruction, along with the backprojection, of the computer-generated test data with noise $\sigma = 0.2$, and sampling pattern $\mathcal{N}_{0.15}^{15\%}$. Both the discrepancy principle value (DV) and optimal value (OV) by corresponding SSIM/PSNR are indicated, along with the Bregman iteration number on which this is reached.

Method	Crite- rion	SSIM(ρ)		SSIM(ϕ)		PSNR(ρ)		PSNR(ϕ)	
		value	iter.	value	iter.	value	iter.	value	iter.
BP		0.461		0.948		20.3		31.5	
WAV	OV	0.475	11	0.923	18	19.6	13	29.3	18
WAV	DV	0.470	12	0.921	12	19.6	12	29.2	12
TV	OV	0.720	4	0.947	12	22.2	6	31.1	13
TV	DV	0.717	5	0.936	5	22.1	5	30.1	5
TGV ²	OV	0.544	7	0.951	9	21.5	10	31.3	14
TGV ²	DV	0.541	6	0.948	6	21.1	6	30.8	6

Table 6: SSIM and PSNR values of both magnitude (ρ) and phase (ϕ) for Bregmanised TV, TGV² ($\alpha = 1.5$) and WAV ($\alpha = 6$) reconstruction, along with the backprojection, of the computer-generated test data with noise $\sigma = 0.2$, and sampling pattern $\mathcal{S}^{33\%}$. Both the discrepancy principle value (DV) and optimal value (OV) by corresponding SSIM/PSNR are indicated, along with the Bregman iteration number on which this is reached.

Method	Crite- rion	SSIM(ρ)		SSIM(ϕ)		PSNR(ρ)		PSNR(ϕ)	
		value	iter.	value	iter.	value	iter.	value	iter.
BP		0.440		0.917		19.4		29.7	
WAV	OV	0.454	9	0.910	13	19.5	12	28.7	12
WAV	DV	0.432	12	0.909	14	19.5	14	28.7	14
TV	OV	0.746	4	0.942	8	23.3	6	31.0	9
TV	DV	0.730	5	0.931	5	23.1	5	30.0	5
TGV ²	OV	0.586	7	0.942	8	22.8	8	30.9	9
TGV ²	DV	0.584	6	0.941	6	22.2	6	30.5	6

regularisation functionals $J = \text{TV}$, $J = \text{TGV}^2$ and $J = \text{WAV}$. We have concluded that TGV² performs overall best for phase reconstruction, and picked two sampling patterns of primary interest. These were the centrally concentrated 15% normal-distributed sampling patterns $\mathcal{N}_{0.15}^{15\%}$, and the 33% spiral $\mathcal{S}^{33\%}$.

We now plan to verify our conclusions by applying the methods to a real data set obtained from the literature [5]. All experiments were performed on a Bruker DMX 200 spectrometer with a vertical 4.7 T superconducting magnet. The experiments were performed using a 64 mm diameter radiofrequency (r.f.) coil operating at 199.7 MHz for proton (¹H). The magnet was equipped with a three-axis shielded gradient producing a maximum gradient strength of 0.136 T m⁻¹ in the x , y , and z directions.

Velocity images were acquired of water flowing through a packed bed. The bed used was a cylindrical column of inner diameter 27 mm, randomly packed with 5 mm diameter glass spheres. The column was of length 1 m. Deionised water was pumped in a closed circuit by a Verder VG330-10 gear pump and controlled by a PC-operated Bronkhorst Cori-flow (model M55C4-AAD-11- K-C) mass flow controller; the flow rate of water was varied from 0 to 60 kg h⁻¹. The T₁ relaxation time constant of the water was reduced to 50 ms by adding gadolinium chloride to the water at a concentration of 0.49 g L⁻¹. Velocity images were obtained using a spin-echo sequence that was designed to minimise the total echo time [4, 7]. The flow-encoding gradients were applied for duration $\delta = 1.19$ ms and separated by an observation time (Δ) of 3.3 ms. Velocity images were obtained with a field-of-view of 30 mm \times 30 mm at an in-plane resolution of 178 μm \times 178 μm with a 1.5 mm slice thickness. The repetition time of the experiment was 300 ms and a 4 step phase-cycle was used, giving a total acquisition time of about 7 min. Data were acquired for the full k-space matrix to enable comparison with a conventional fully sampled acquisition.

The methods are used as a guideline parametrised according to the good choices found above, including the parameter choice $\alpha = 1.5$ for $J = \text{TV}$ and $J = \text{TGV}^2$. This is justified, because the experimental setup of the real data set matches the computer-generated data in terms of dynamic range, domain $\Omega = \{1, \dots, 256\}^2$, and roughly the noise level. Regarding the noise level σ , we do not know it exactly in advance. In order to apply the discrepancy principle, we therefore need to estimate it. We do this by calculating the standard deviation of the data within the small top-left rectangle $\{1, \dots, 48\}^2$ of Ω . This rectangle consists of only background, so we may assume that the noise-free signal should be zero there. This gives $\sigma \approx 0.45$.

The reconstructions are visualised in Figure 14 for $\mathcal{N}_{0.15}^{15\%}$. We include the results for $J = \text{TV}$ and $J = \text{TGV}^2$,

Table 7: SSIM and PSNR values of both magnitude (ρ) and phase (ϕ) for Bregmanised TV, TGV² and WAVreconstruction, along with the backprojection, of the computer-generated test data with noise $\sigma = 0.5$, and sampling pattern $\mathcal{N}_{0.15}^{15\%}$. Both the discrepancy principle value (DV) and optimal value (OV) by corresponding SSIM/PSNR are indicated, along with the Bregman iteration number on which this is reached.

Method	Crite- rion	SSIM(ρ)		SSIM(ϕ)		PSNR(ρ)		PSNR(ϕ)	
		value	iter.	value	iter.	value	iter.	value	iter.
BP		0.350		0.837		16.6		26.3	
WAV	OV	0.329	5	0.824	5	17.2	5	25.0	5
WAV	DV	0.321	6	0.822	6	17.0	6	24.9	6
TV	OV	0.533	2	0.867	3	19.6	3	27.0	4
TV	DV	0.488	3	0.867	3	19.6	3	26.9	3
TGV ²	OV	0.393	4	0.886	3	18.6	3	27.6	3
TGV ²	DV	0.385	3	0.886	3	18.6	3	27.6	3

Table 8: SSIM and PSNR values of both magnitude (ρ) and phase (ϕ) for Bregmanised TV, TGV² and WAVreconstruction, along with the backprojection, of the computer-generated test data with noise $\sigma = 0.5$, and sampling pattern $\mathcal{S}^{33\%}$. Both the discrepancy principle value (DV) and optimal value (OV) by corresponding SSIM/PSNR are indicated, along with the Bregman iteration number on which this is reached.

Method	Crite- rion	SSIM(ρ)		SSIM(ϕ)		PSNR(ρ)		PSNR(ϕ)	
		value	iter.	value	iter.	value	iter.	value	iter.
BP		0.289		0.759		14.4		23.6	
WAV	OV	0.314	5	0.817	5	17.1	5	24.7	5
WAV	DV	0.308	6	0.808	6	16.7	6	24.5	6
TV	OV	0.516	2	0.857	3	19.5	3	26.5	3
TV	DV	0.475	3	0.857	3	19.5	3	26.5	3
TGV ²	OV	0.399	4	0.880	3	18.6	3	27.3	3
TGV ²	DV	0.389	3	0.880	3	18.6	3	27.3	3




Figure 13: Magnitude (left) and phase (right) of one frame of the real data set. The square marks the region displayed in the computational results of Figure 14.

but exclude $J = \text{WAV}$ due to unsatisfactory performance. We have no quality measures to report, since we do not have a ground-truth to compare against. Our analysis of the results is therefore entirely visual. As can be observed, the performance of the methods is very good indeed, for both TV and especially TGV^2 . Some stair-casing artefacts can be observed with $J = \text{TV}$, but this is minor, and avoided by $J = \text{TGV}^2$. Moreover, the reconstruction of edges of the magnitude image, and the smoothness of the phase image is good for TGV^2 . Further, not only TV, but also TGV^2 manages to perform a flat reconstruction of the background (blue), where there is mainly noise in the signal. In regions of the image that contain signal, the magnitude will vary owing to partial volume effects and attenuation of the signal in regions of high shear. The TV and TGV^2 reconstructions both yield high quality reconstructions of the magnitude, although there is some evidence of unrealistic stair-casing artefacts in the TV reconstruction.

4 Conclusion

In this paper, by analysing computations on computer-generated test data, we have derived a guideline for the reconstruction of sub-sampled MRI phase data. Our approach promotes sparsity of the reconstruction through a regularisation functional, as well as contrast enhancement through an iterative reconstruction approach called the Bregman iteration. We have verified our approach by application of the guideline to a real data set whose experimental setup matches the computer-generated data in the essential parameters: domain, sampling patterns, dynamic range, and noise level. These data parametrise the guideline, primarily the choice of the initial parameter α .

Considering sampling schemes, our findings verify existing literature. The scheme should be a good mixture of incoherent samples covering the most important parts of the k -space (which tends to be the centre). In the present application, out of the schemes considered, a widely dispersed Gaussian pattern and a spiral satisfy these requirements. Line sampling, as was used in earlier works, performs relatively poorly.

With regard to the different regularisers evaluated, we can say that Daubechies-3 wavelets perform unsatisfactorily for the low resolutions inherent in our application area. Total variation (TV) performs reasonably, but suffers from stair-casing artefacts, which is in the nature of the TV model. Second-order Total Generalised Variation (TGV^2) significantly improves upon TV in the phase reconstruction that we are interested in; in case of the computer-generated test data, better magnitude reconstruction cannot be expected, as the original data set is piecewise constant. This is where TV performs the best. Best reconstruction for this data set with the regularisers discussed in this paper should be achievable with split TV regularisation of the magnitude and TGV^2 regularisation of the phase. In our real data experiment also the magnitude is smooth aside from discontinuities arising from solid particles. TGV^2 arguably offers overall best performance, providing a good compromise between sharp boundaries and smoothness.

Finally, we can attest that Bregman iteration combined with the Morozov discrepancy principle as a stopping criterion provides a good strategy in the absence of knowledge of ground truth. Knowing only the noise level, it gives good quality reconstruction with much better preservation of contrast than with standard Tikhonov-type reconstruction.

5 Acknowledgements

This work has been financially supported by the King Abdullah University of Science and Technology (KAUST) Award No. KUK-I1-007-43, and the EPSRC / Isaac Newton Trust Small Grant 2012/13 “Non-smooth geometric reconstruction for high resolution MRI imaging of fluid transport in bed reactors”, EPSRC grant EP/F047991/1, and Microsoft Research Connections.

References

- [1] P. Gatehouse, J. Keegan, L. Crowe, S. Masood, R. Mohiaddin, K.-F. Kreitner and D. Firmin, *Applications of phase-contrast flow and velocity imaging in cardiovascular MRI*, *European Radiology* **15** (2005), 2172–2184.
- [2] P. Callaghan, *Rheo-NMR: nuclear magnetic resonance and the rheology of complex fluids*, *Reports on Progress in Physics* **62** (1999), 599–670.
- [3] A. Sederman, M. Johns, P. Alexander and L. Gladden, *Structure-flow correlations in packed beds*, *Chemical engineering Science* **53** (1998), 2117–2128.
- [4] M. H. Sankey, D. J. Holland, L. F. Gladden and A. J. Sederman, *Magnetic resonance velocity imaging of liquid and gas two-phase flow in packed beds*, *Journal of Magnetic Resonance* **196** (2009), 142–148.

- [5] D. J. Holland, D. M. Malioutov, A. Blake, A. J. Sederman and L. F. Gladden, *Reducing data acquisition times in phase-encoded velocity imaging using compressed sensing.*, Journal of Magnetic Resonance **203** (2010), 236–46.
- [6] E. Fukushima, *Nuclear magnetic resonance as a tool to study flow*, Annual Review of Fluid Mechanics **31** (1999), 95–123.
- [7] D. J. Holland, L. F. Gladden, C. R. Müller, J. S. Dennis and A. J. Sederman, *Spatially resolved measurement of anisotropic granular temperature in gas-fluidized beds*, Powder Technology **182** (2008), 171–181.
- [8] A. B. Tayler, D. J. Holland, A. J. Sederman and L. F. Gladden, *Exploring the origins of turbulence in multiphase flow using compressed sensing MRI*, Physical Review Letters **108** (2012), 264505.
- [9] K. Kose, *Instantaneous flow-distribution measurements of the equilibrium turbulent region in a circular pipe using ultrafast NMR imaging*, Physical Review A **44** (1991), 2495–2504.
- [10] A. J. Sederman, M. D. Mantle, C. Buckley and L. F. Gladden, *MRI technique for measurement of velocity vectors, acceleration, and autocorrelation functions in turbulent flow*, Journal of Magnetic Resonance **166** (2004), 182–189.
- [11] P. Galvosas and P. T. Callaghan, *Fast magnetic resonance imaging and velocimetry for liquids under high flow rates.*, Journal of Magnetic Resonance **181** (2006), 119–25.
- [12] E. J. Candès, J. Romberg and T. Tao, *Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information*, Information Theory, IEEE Transactions on **52** (2006), 489–509.
- [13] D. L. Donoho, *Compressed sensing*, Information Theory, IEEE Transactions on **52** (2006), 1289–1306.
- [14] M. Lustig, D. Donoho and J. M. Pauly, *Sparse MRI: The application of compressed sensing for rapid MR imaging*, Magn. Reson. Med. **58** (2007), 1182–1195, doi:[10.1002/mrm.21391](https://doi.org/10.1002/mrm.21391).
- [15] B. Jung, M. Honal, P. Ullmann, J. Hennig and M. Markl, *Highly k-t-space-accelerated phase-contrast MRI*, Magnetic Resonance in Medicine **60** (2008), 1169–1177.
- [16] E. Paciok and B. Blümich, *Ultrafast microscopy of microfluidics: compressed sensing and remote detection*, Angewandte Chemie (International Edition) **50** (2011), 5258–60.
- [17] J. Paulsen, V. S. Bajaj and A. Pines, *Compressed sensing of remotely detected MRI velocimetry in microfluidics.*, Journal of Magnetic Resonance **205** (2010), 196–201.
- [18] L. Rudin, S. Osher and E. Fatemi, *Nonlinear total variation based noise removal algorithms*, Physica D **60** (1992), 259–268.
- [19] K. Bredies, K. Kunisch and T. Pock, *Total generalized variation*, SIAM Journal on Imaging Sciences **3** (2011), 492–526, doi:[10.1137/090769521](https://doi.org/10.1137/090769521).
- [20] J.-B. Hiriart-Urruty and C. Lemaréchal, *Convex analysis and minimization algorithms I*, Springer, 1993.
- [21] H. W. Engl, M. Hanke and A. Neubauer, *Regularization of inverse problems*, volume 375, Springer, 1996.
- [22] Y. Meyer, *Oscillating patterns in image processing and nonlinear evolution equations*, The Fifteenth Dean Jacqueline B. Lewis Memorial Lectures, American Mathematical Society, 2001.
- [23] M. Benning and M. Burger, *Ground States and Singular Vectors of Convex Variational Regularization Methods*, Methods and Applications of Analysis, special issue in honor of S. Oshers 70th birthday (2012). Accepted.
URL <http://arxiv.org/abs/1211.2057>
- [24] L. Bregman, *The relaxation method for finding the common point of convex sets and its application to the solution of problems in convex programming*, USSR Comp. Math. Math. Phys. **7** (1967), 200–217.
- [25] K. Kiwiel, *Proximal minimization methods with generalized bregman functions*, SIAM J. Cont. Optim. **35** (1997), 1142–1168.
- [26] S. Osher, M. Burger, D. Goldfarb, J. Xu and W. Yin, *An iterative regularization method for total variation-based image restoration*, SIAM Multiscale Model. Simul. **4** (2005), 460–489.
- [27] V. A. Morozov, *On the solution of functional equations by the method of regularization*, Soviet Math. Doklady **7** (1966), 414–417.

- [28] M. Benning, C. Brune, M. Burger and J. Müller, *Higher-order tv methods – enhancement via bregman iteration*, Journal of Scientific Computing **54** (2013), 269–310, doi:[10.1007/s10915-012-9650-3](https://doi.org/10.1007/s10915-012-9650-3).
- [29] M. Guerquin-Kern, M. Haberlin, K. Pruessmann and M. Unser, *A fast wavelet-based reconstruction method for magnetic resonance imaging*, Medical Imaging, IEEE Transactions on **30** (2011), 1649–1660.
- [30] I. Daubechies, *Ten lectures on wavelets*, Society for Industrial and Applied Mathematics, 1992, 1 edition.
- [31] K. Bredies, K. Kunisch and T. Valkonen, *Properties of L^1 -TGV 2 : The one-dimensional case*, Journal of Mathematical Analysis and Applications **398** (2013), 438–454, doi:[10.1016/j.jmaa.2012.08.053](https://doi.org/10.1016/j.jmaa.2012.08.053).
URL <http://math.uni-graz.at/mobis/publications/SFB-Report-2011-006.pdf>
- [32] K. Bredies and T. Valkonen, *Inverse problems with second-order total generalized variation constraints*, in: *Proceedings of SampTA 2011 – 9th International Conference on Sampling Theory and Applications, Singapore*, 2011.
- [33] C. M. Tsai and D. G. Nishimura, *Reduced aliasing artifacts using variable-density k-space sampling trajectories.*, Magnetic Resonance in Medicine **43** (2000), 452–8.
- [34] W. Yin, S. Osher, D. Goldfarb and J. Darbon, *Bregman iterative algorithms for ℓ^1 -minimization with applications to compressed sensing*, SIAM J. Imaging Sci **1** (2008), 143–168.
- [35] W. Yin and S. Osher, *Error forgetting of bregman iteration*, Journal of Scientific Computing **54** (2013), 684–695.
- [36] O. Scherzer, *The use of Morozov’s discrepancy principle for Tikhonov regularization for solving nonlinear ill-posed problems*, Computing **51** (1993), 45–60.
- [37] M. Bostock, D. Holland and D. Nietlispach, *Compressed sensing reconstruction of undersampled 3D NOESY spectra: application to large membrane proteins*, Journal of Biomolecular NMR **54** (2012), 15–32.
- [38] Z. 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** (2004), 600–612, doi:[10.1109/TIP.2003.819861](https://doi.org/10.1109/TIP.2003.819861).
- [39] G. Strang and T. Nguyen, *Wavelets and filter banks*, Wellesley Cambridge Press, 1996.
- [40] A. Chambolle and P.-L. Lions, *Image recovery via total variation minimization and related problems*, Numerische Mathematik **76** (1997), 167–188, doi:[10.1007/s002110050258](https://doi.org/10.1007/s002110050258).
- [41] F. Knoll, K. Bredies, T. Pock and R. Stollberger, *Second order total generalized variation (TGV) for MRI*, Magnetic Resonance in Medicine **65** (2011), 480–491.
- [42] K. Papafitsoros and C.-B. Schönlieb, *A combined first and second order variational approach for image reconstruction*, Journal of Mathematical Imaging and Vision (2013), doi:[10.1007/s10851-013-0445-4](https://doi.org/10.1007/s10851-013-0445-4). Published online.
URL <http://www.maths.cam.ac.uk/postgrad/cca/files/TVBHpaper.pdf>
- [43] T. Valkonen, K. Bredies and F. Knoll, *Total generalised variation in diffusion tensor imaging*, SIAM Journal on Imaging Sciences **6** (2013), 487–525, doi:[10.1137/120867172](https://doi.org/10.1137/120867172).
URL <http://tuomov.iki.fi/mathematics/dtireg.pdf>
- [44] M. Mantle, A. Sederman and L. Gladden, *Single-and two-phase flow in fixed-bed reactors: MRI flow visualisation and lattice-Boltzmann simulations*, Chemical Engineering Science **56** (2001), 523–529.
- [45] B. Manz, L. F. Gladden and P. B. Warren, *Flow and dispersion in porous media: Lattice-Boltzmann and NMR studies*, AIChE Journal **45** (1999), 1845–1854.
- [46] D. Donoho, M. Duncan, X. Huo and O. Levi, *Wavelab* (2000). Software.
URL <http://www-stat.stanford.edu/~wavelab>
- [47] D. Gabay, *Applications of the method of multipliers to variational inequalities*, in: *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems, Studies in Mathematics and its Applications*, volume 15, Edited by M. Fortin and R. Glowinski, North-Holland, Amsterdam, 1983, 299–331.
- [48] T. Goldstein and S. Osher, *The split bregman method for $L1$ -regularized problems*, SIAM Journal on Imaging Sciences **2** (2009), 323–343.

- [49] S. Boyd, N. Parikh, E. Chu, B. Peleato and J. Eckstein, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Foundations and Trends in Machine Learning **3** (2011), 1–122.
- [50] A. Chambolle and T. Pock, *A first-order primal-dual algorithm for convex problems with applications to imaging*, Journal of Mathematical Imaging and Vision **40** (2011), 120–145, doi:[10.1007/s10851-010-0251-1](https://doi.org/10.1007/s10851-010-0251-1).
- [51] T. Valkonen, F. Knoll and K. Bredies, *TGV for diffusion tensors: A comparison of fidelity functions*, in: *Journal of Inverse and Ill-posed problems special issue for IP:MEIS 2012*, Antalya, Turkey, 2012, doi:[10.1515/jip-2013-0005](https://doi.org/10.1515/jip-2013-0005). Published online.
URL <http://tuomov.iki.fi/mathematics/ipmsproc.pdf>
- [52] T. Valkonen and M. Liebmann, *GPU-accelerated regularisation of large diffusion-tensor volumes*, in: *Computing special issue for ESCO 2012*, Pilsen, Czech Republic, 2012. To appear.
URL <http://tuomov.iki.fi/mathematics/escoproc.pdf>
- [53] E. Esser, X. Zhang and T. F. Chan, *A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science*, SIAM Journal on Imaging Sciences **3** (2010), 1015–1046.

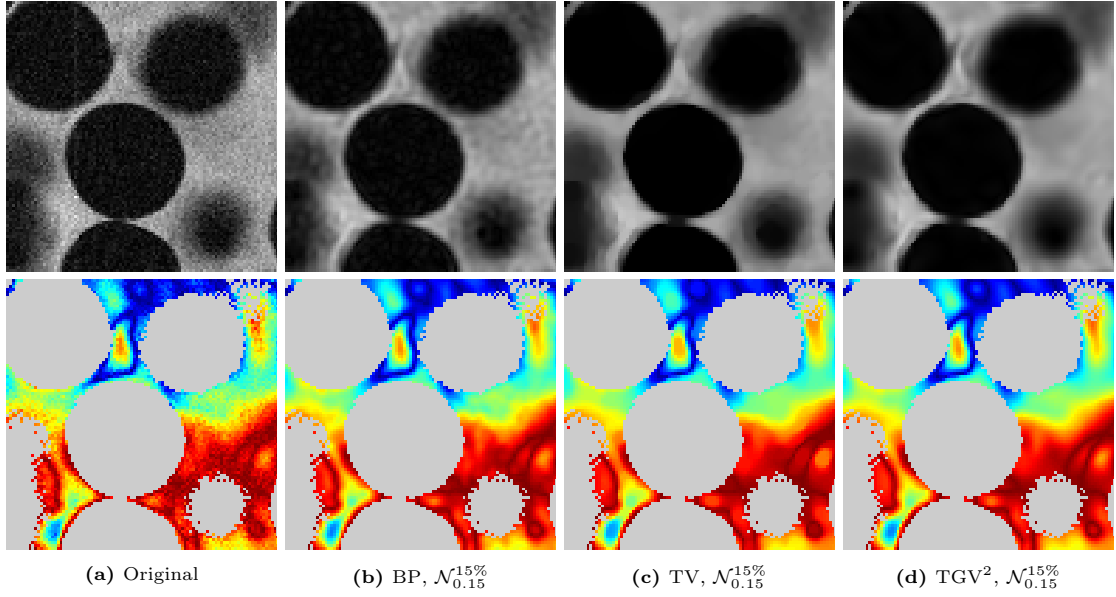


Figure 14: Results of Bregmanised TV and TGV² reconstruction with discrepancy principle as the stopping criterion, along with the backprojection, of the real measurement data for the sampling pattern $\mathcal{N}_{0.15}^{15\%}$. The top picture presents the magnitude and the bottom picture the phase.




Figure 15: Unwrapped velocity phase difference $\phi_1 - \phi_2$ based on the phase ϕ_1 in Figure 14 and corresponding phase ϕ_2 for a second measurement.