



Citation for published version:

Ehrhardt, MJ & Mauritz, M 2024 'Complex extension of optical flow and its practical evaluation for undersampled dynamic MRI' arXiv.

Publication date:
2024

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

COMPLEX EXTENSION OF OPTICAL FLOW AND ITS PRACTICAL EVALUATION FOR UNDERSAMPLED DYNAMIC MRI

M.J.EHRHARDT, M.MAURITZ

ABSTRACT. Reconstructing high-quality images from undersampled dynamic MRI data is a challenging task and important for the success of this imaging modality. To remedy the naturally occurring artifacts due to measurement undersampling, one can incorporate a motion model into the reconstruction so that information can propagate across time frames. Current models for MRI imaging are using the optical flow equation. However, they are based on real-valued images. Here, we generalise the optical flow equation to complex-valued images and demonstrate, based on two real cardiac MRI datasets, that the new model is capable of improving image quality.

1. INTRODUCTION

Magnetic resonance imaging (MRI) is a powerful medical imaging modality that allows for noninvasive visualisations of the human body. In particular, for dynamic MRI, there is a variety of applications, e.g. cardiac imaging [25, 2], dynamic contrast-enhanced magnetic resonance imaging [13], [9] and flow imaging [24]. However, the imaging times in MRI are inherently slow, leading to motion artifacts that degrade image quality. This issue can be addressed from a software perspective where one aims for acquisition speedups by decreasing the number of measurements and using mathematical modelling to improve image quality. In this article, we follow this approach and solve the ill-posed inverse problem $A_t \rho_t = y_t$ (the index t denotes the t -th motion state of the considered subject; the MRI forward operator A_t includes subsampling of measurements and hence may depend on t) of undersampled dynamic MRI by considering the variational problem

$$\min_{\rho, v} \sum_t \|A_t \rho_t - y_t\|^2 + \mathcal{R}(\rho, v),$$

where \mathcal{R} is a regularisation term that incorporates prior knowledge about the images into the reconstruction. More precisely, we aim to jointly reconstruct the temporal image sequence ρ and the velocity field v associated with the temporal evolution. To this end, we model the underlying motion using the optical flow equation [6]. This was already done for MRI and real-valued images [1]. Here, we generalise this idea by explicitly taking into account that images are complex-valued, i.e. we formulate an optical flow equation for complex-valued images. This is important since omitting the phase may lead to reconstruction artifacts. This problem can be seen in fig. 1. The figure shows reconstructions from fully sampled MRI measurements. The images on the right show a nonnegative real-valued reconstruction where some structures, marked by red boxes, of the complex-valued reconstructions on the left

are missing. The top row shows the magnitude of the reconstructions, while the bottom row depicts complex values: The hue represents the normalised phase and the brightness represents the magnitude. In particular, one can see that structures where the phase is different are missing in the real-valued images.

1.1. Related work. Dynamic inverse problems in general have been studied extensively. An overview of the topic can be found in [14, 23]. The particular topic of undersampled (dynamic) MRI has also been widely studied and many different approaches have been proposed to enhance reconstruction quality. Many works leveraged compressed sensing making use of the fact that MRI images can be sparsely represented in suitable domains [17, 16, 10]. In dynamic imaging, the temporal relation between motion states should be taken into account to improve the results. One approach is to decompose the (temporal) sequence of images into a low-rank part (background between time frames) and a sparse part (modelling dynamics between frames) [12, 18]. This way, the image sequence is represented suitably for the undersampled MRI problem. Another approach is taken in [26], where the motion between different motion states is estimated based on the frame-wise undersampled reconstruction of each state. Subsequently, this motion information is used to get a final reconstruction incorporating all motion states into a single reconstruction. Joint reconstruction and motion estimation approaches are also promising. This was investigated in [1, 29] where the underlying motion was modelled using (locally affine) optical flow. Besides classical image reconstruction, learning-based approaches have been considered in dynamic MRI. Those aim at learning spatiotemporal dependencies within the image sequence, and/or mimic the iterative nature of classical reconstruction algorithms [21, 20, 28].

Besides MRI, dynamic inverse problems arise in other fields of imaging. For instance, in dynamic PET imaging the continuity equation was used as the motion model to describe temporal changes between time points [22]. In tomography, an approach based on a template and its temporal deformation was taken [7].

When it comes to working with complex-valued images, two approaches are conceivable: First, one can parameterise the images using real and imaginary part. This is the standard approach that is basically always followed. However, in MRI imaging often the magnitude of the image is of most interest and one would like to pose prior assumptions on the magnitude. This is difficult to incorporate using real and imaginary part because calculating the magnitude is a nonlinear operation, hence the overall regularisation would involve a nonlinear operator. Therefore, one can work with a magnitude and phase parameterisation of the complex-valued images. This makes regularising the magnitude easier but comes at the cost that now the MRI forward operator is nonlinear. This was considered in [11, 27].

2. METHODS

2.1. MRI modelling. Solving the inverse problem in MRI imaging in its simplest form leads to inverting the discrete Fourier transform. Given enough measurements (i.e. enough coverage of Fourier space/k-space) this problem is well-posed by the stable invertibility of the Fourier transform. In practice, however, multiple receiver coils (here are the measurements detected) and undersampling lead to an ill-posed inverse problem. The MRI forward operator maps a complex-valued image ρ_t (which in our case is a time slice of a temporal sequence of images) to some complex space Y_t , where Y_t depends on the undersampling factor and the




FIGURE 1. Comparison between real-valued and complex-valued reconstructions in MRI. Left: Complex-valued reconstruction. Right: Real-valued reconstruction. Top row: Magnitude representation. The bottom row shows the same images but as a complex-valued image: The hue represents the normalised phase and the brightness represents the magnitude. The red boxes point to some missing structures in the real-valued image.

number of receiver coils. For instance, in a continuous setting one could choose $\rho \in L^p((0, T); BV(\Omega))$, hence $\rho_t \in BV(\Omega)$, or similar spaces with more regularity [6]. In a discrete setting (which will be considered in this article) we take $\rho \in X = \mathbb{C}^{N_t \times N_x \times N_y}$, where N_t , N_x and N_y denote the number of dimensions in the temporal, x and y direction, respectively. Therefore, it is $\rho_t \in X_t = \mathbb{C}^{N_x \times N_y}$ for each time slice. We next introduce the forward operator describing the measurements. Denoting by N_c the number of receiver coils, we have $Y_t = \mathbb{C}^{N_c \times N_{k,t}}$, $t = 1, \dots, N_t$, as the measurement space. Here, $N_{k,t}$ denotes the number of measured k-space points at time t after the undersampling has been applied. With the Fourier operator F , the coil sensitivity operators $C_{t,i}$ and the undersampling operators S_t (see section 4.4 for details) we can find for the i -th component (i.e. coil) of the MRI forward operator A_t at time point t

$$(1) \quad A_t: X_t \rightarrow Y_t, \quad (A_t \rho_t)_i = S_t F C_{t,i} \rho_t.$$

Hence, we model MRI measurements as

$$y_t = A_t \rho_t^\dagger + \varepsilon_t$$

where $\rho_t^\dagger \in X_t$ is the ground truth image that we want to reconstruct and ε_t is some additive noise.

2.2. Modelling motion for complex-valued images. In this work, we use the optical flow equation as a motion model to regularise the dynamic inverse problem. The optical flow equation is often used in computer vision and has proven valuable for 2D images [1, 5, 6]. Its defining property is constant image intensity along a trajectory with $\frac{d}{dt}\rho = v$ [14]. Thus, the underlying PDE reads

$$\partial_t \rho + \langle v, \nabla \rho \rangle_{\mathbb{R}} = 0$$

and during the reconstruction, both the image ρ as well as the velocity field v have to be reconstructed. Note, that since velocity and density are multiplied by each other, using the optical flow equation makes the problem nonconvex. To be able to apply the optical flow equation in MRI imaging (where complex-valued images are reconstructed), we propose to generalise this equation by allowing both the image ρ and the velocity v to be complex-valued. The inner product is therefore replaced by the standard inner product for complex vector spaces. This leads us to the model

$$(2) \quad \begin{aligned} 0 &= \partial_t(\rho^1 + i\rho^2) + \langle v^1 + iv^2, \nabla(\rho^1 + i\rho^2) \rangle_{\mathbb{C}} \\ &\iff \begin{cases} 0 = \partial_t \rho^1 + \langle v^1, \nabla \rho^1 \rangle_{\mathbb{R}} + \langle v^2, \nabla \rho^2 \rangle_{\mathbb{R}} \\ 0 = \partial_t \rho^2 + \langle v^2, \nabla \rho^1 \rangle_{\mathbb{R}} - \langle v^1, \nabla \rho^2 \rangle_{\mathbb{R}}. \end{cases} \end{aligned}$$

for $\rho = \rho^1 + i\rho^2 \in X$ and $v = v^1 + iv^2 \in V^2$ (note, that the velocity is a 2D vector). Here, V is a suitable space for the velocities such that the derivatives needed for the priors are well-defined. In the discrete setting, we may choose $V = X$. Associated with the above two equations we introduce the optical flow operator

$$(3) \quad M(\rho, v) = \begin{pmatrix} \partial_t \rho^1 + \langle v^1, \nabla \rho^1 \rangle_{\mathbb{R}} + \langle v^2, \nabla \rho^2 \rangle_{\mathbb{R}} \\ \partial_t \rho^2 + \langle v^2, \nabla \rho^1 \rangle_{\mathbb{R}} - \langle v^1, \nabla \rho^2 \rangle_{\mathbb{R}} \end{pmatrix}.$$

For our discrete space $X = \mathbb{C}^{N_t \times N_x \times N_y}$ the temporal derivative $\partial_t: X \rightarrow X$ and the inner product $\langle \cdot, \cdot \rangle_{\mathbb{R}}: X \times X \rightarrow \mathbb{R}$ are to be understood pointwise.

We use the same discrete versions of the differential operators as in [6]. This means we use forward differences in time and central differences in space with Neumann boundary conditions for the differential operators involved in the optical flow equation. For the differential operators of other parts of the optimization function we use forward differences with Neumann boundary conditions (see [6] for more details).

Remark (Choices on optical flow model). *There are other possible generalisations of the optical flow equation to complex-valued images. For instance, one could introduce separate velocities for real and imaginary part, hence decoupling both components. Also, a single velocity for both the real and imaginary part would be an option. Finally, only one of the two parts could be regularised by the optical flow equation and the other one evolves freely. However, these variants do not explicitly model interchange between real and imaginary part. To be able to model such an interchange, we used the above approach.*

Remark (Alternative choice for motion model). *An alternative motion model would be to use the continuity equation in density and momentum variables instead of density and velocity variables (see, for example, [4]). This way, the motion model does not incorporate a nonconvexity into the optimisation problem. Unfortunately, finding meaningful regularisers for the momentum variables is more challenging than for the velocities. For instance, an object moving along the x*

axis with constant velocity will have a zero gradient of the corresponding velocities. The momentum, however, is nonconstant as it also changes when the mass of the underlying object changes. Therefore, we decided to use the optical flow equation.

3. OPTIMISATION AND IMPLEMENTATION

Based on [1, 6] we minimise

$$(4) \quad F(\rho, v) = \sum_t \|A_t \rho_t - y_t\|^2 + \alpha_1 \mathcal{R}_1(\rho) + \alpha_2 \mathcal{R}_2(v) + \alpha_3 \mathcal{R}_3(M(\rho, v))$$

with respect to $\rho = \rho^1 + i\rho^2$, $v = (v_x^1, v_y^1) + i(v_x^2, v_y^2)$. Here, y_t is the measured MRI signal at time point t , and A_t and M are operators related to the MRI measurements and the optical flow equation, respectively (see (1) and (3)). The regularisers for the variables ρ and v are chosen to promote smoothness (the “level of smoothness” can be controlled via the ε_i parameters)

$$\begin{aligned} \mathcal{R}_1(\rho) &= \mathcal{H}_{\varepsilon_1}(\nabla \rho^1) + \mathcal{H}_{\varepsilon_1}(\nabla \rho^2) \\ \mathcal{R}_2(v) &= \mathcal{H}_{\varepsilon_2}(\nabla v_x^1) + \mathcal{H}_{\varepsilon_2}(\nabla v_y^1) + \mathcal{H}_{\varepsilon_2}(\nabla v_x^2) + \mathcal{H}_{\varepsilon_2}(\nabla v_y^2), \end{aligned}$$

where ∇ denotes the spatial gradient and \mathcal{H}_ε is the Huber loss with parameter ε . Defining for $x \in \mathbb{R}^2$

$$h_\varepsilon(x) = \begin{cases} \frac{\|x\|_2^2}{2\varepsilon} & \text{if } \|x\|_2 \leq \varepsilon, \\ \|x\|_2 - \frac{\varepsilon}{2} & \text{else,} \end{cases}$$

we have $\mathcal{H}_\varepsilon(x) = \sum_{i=1}^N h_\varepsilon(x_i)$ for $x \in (\mathbb{R}^2)^N$. Finally, we use a soft constraint

$$\mathcal{R}_3(M(\rho, v)) = \mathcal{H}_{\varepsilon_3}(M(\rho, v))$$

since we cannot expect the optical flow equation to be satisfied perfectly. Compared to the classical L^1 loss the Huber loss is differentiable, and the parameter ε controls the extent of smoothness within the Huber loss, i.e. it controls whether the Huber loss is more similar to the norm $\|\cdot\|_2$ or the squared norm $\|\cdot\|_2^2$. Tuning this parameter allows us to adapt to the considered problem and find a better model.

The considered optimisation problem (4) is nonconvex due to the optical flow equation. To deal with the nonconvexity we follow the same strategy as in [1, 6] and apply a block coordinate descent (see algorithm 2). This means we iteratively optimise the minimisation function only for ρ or v while keeping the other variable fixed. Each resulting subproblem is solved using the FISTA algorithm [3] (applied to $f + g$ with f smooth and $g = 0$, see algorithm 1) with early stopping (iterations where stopped if the relative L^2 error was below $\delta = 10^{-5}$). Furthermore, motivated by the coarse-to-fine approach in image registration, we smoothed the density ρ and the velocity v after each sub-optimisation using a Gaussian smoothing with standard deviation σ/i_{outer} . Here, σ is again a tunable parameter, and $i_{\text{outer}} \in \{1, \dots, n_{\text{outer}}\}$ denotes the current outer iteration, i.e. the smoothing decays with growing iterations. For the maximal number of iterations we used $n_{\text{outer}} = 200$, $n_\rho = 1400$ and $n_v = 3200$. The optimisation algorithm was implemented in Python using SIRF [19] for reading the MR data and applying the associated measurement operators, and CIL [15] to implement the actual optimisation algorithm. **The code will be made available upon publication and is attached for the review process.**

Algorithm 1 FISTA with early stopping for minimizing smooth function f

```

1: Inputs:
   smooth function  $f$  with Lipschitz constant  $L$ , initialisation
    $x_{\text{init}}$ 
   Algorithm parameters  $n, \delta$ 
2: Initialize:
    $x^0 \leftarrow x_{\text{init}}, \hat{x}^1 \leftarrow x_{\text{init}}, t^1 \leftarrow 1$ 
3: for  $j = 1$  to  $n$  do
4:    $x^j \leftarrow \hat{x}^j - \frac{1}{L} \nabla f(\hat{x}^j)$ 
5:    $t^{j+1} \leftarrow \frac{1 + \sqrt{1 + 4(t^j)^2}}{2}$ 
6:    $\hat{x}^{j+1} \leftarrow x^j + \left(\frac{t^j - 1}{t^{j+1}}\right)(x^j - x^{j-1})$ 
7:   if  $\|x^j - x^{j-1}\|_2 / \|x^{j-1}\|_2 < \delta$  then
8:     break
9: return  $x^j$ 

```

Algorithm 2 Optimization of $F(\rho, v)$

```

1: Inputs:
   Measurement  $y$ 
   Model parameters  $\alpha_1, \alpha_2, \alpha_3, \varepsilon_1, \varepsilon_2, \varepsilon_3$ 
   Algorithm parameters  $\sigma, n_{\text{outer}}, n_\rho, n_v, \delta$ 
2: Initialize:
    $v^1 \leftarrow 0, \hat{v}^1 \leftarrow 0, \rho^1 \leftarrow 0, \hat{\rho}^1 \leftarrow 0$ 
3: for  $i = 1$  to  $n_{\text{outer}}$  do
4:    $\rho^{i+1} \leftarrow \text{FISTA}(\rho \mapsto F(\rho, \hat{v}^i), \hat{\rho}^i)$  with  $n_\rho$  iterations
5:    $\hat{\rho}^{i+1} \leftarrow \text{smooth}_{\sigma/i}(\rho^{i+1})$ 
6:    $v^{i+1} \leftarrow \text{FISTA}(v \mapsto F(\hat{\rho}^{i+1}, v), \hat{v}^i)$  with  $n_v$  iterations
7:    $\hat{v}^{i+1} \leftarrow \text{smooth}_{\sigma/i}(v^{i+1})$ 
8:   if  $\frac{1}{2} \|v^{i+1} - v^i\|_2 / \|v^i\|_2 + \frac{1}{2} \|\rho^{i+1} - \rho^i\|_2 / \|\rho^i\|_2 < \delta$  then
9:     break
10: return  $\rho^{i+1}, v^{i+1}$ 

```

3.1. Models for comparison. We compare our proposed method (4) (we call this model OF and label corresponding reconstructions OF, too) to three slightly different variants:

- (a) Frame-wise model (FW): For each time point a separate reconstruction is computed and no correlation between time points is assumed. This corresponds to optimising (4) for the density ρ only and setting α_2 and α_3 to zero. Thus, we minimise

$$F(\rho) = \sum_t \|A_t \rho_t - y_t\|^2 + \alpha_1 \mathcal{R}_1(\rho)$$

with respect to $\rho = \rho^1 + i\rho^2$ applying algorithm 1 with n_ρ iterations, where

$$\mathcal{R}_1(\rho) = \mathcal{H}_{\varepsilon_1}(\nabla \rho^1) + \mathcal{H}_{\varepsilon_1}(\nabla \rho^2)$$

as above.

- (b) Time derivative model (DT): The term containing the optical flow equation is replaced by a term containing time derivatives only. This means we fix

the velocities to zero and optimise (4) for the density ρ only. Thus, we minimise

$$F(\rho) = \sum_t \|A_t \rho_t - y_t\|^2 + \alpha_1 \mathcal{R}_1(\rho) + \alpha_3 \mathcal{R}_3(M(\rho, 0))$$

with respect to $\rho^1 + i\rho^2$ applying algorithm 1 with n_ρ iterations, where

$$\mathcal{R}_1(\rho) = \mathcal{H}_{\varepsilon_1}(\nabla \rho^1) + \mathcal{H}_{\varepsilon_1}(\nabla \rho^2).$$

and

$$\mathcal{R}_3(M(\rho, 0)) = \mathcal{H}_{\varepsilon_3}([\partial_t \rho^1, \partial_t \rho^2]).$$

as above promote spatial and temporal regularity, respectively.

- (c) Optical flow using ground truth velocities (**Cheat-OF**): With this model we want to investigate how our method would perform if we had access to the velocities underlying the motion, i.e. this model gives us the best possible reconstruction within our framework. To this end, we use our model (4) with fully sampled data to determine ground truth velocities v_{GT} . Then, we optimise (4) for the density ρ only while fixing the velocity to be v_{GT} . Thus, we minimise

$$F(\rho) = \sum_t \|A_t \rho_t - y_t\|^2 + \alpha_1 \mathcal{R}_1(\rho) + \alpha_3 \mathcal{R}_3(M(\rho, v_{\text{GT}}))$$

with respect to $\rho = \rho^1 + i\rho^2$ applying algorithm 1 with n_ρ iterations, where

$$\mathcal{R}_1(\rho) = \mathcal{H}_{\varepsilon_1}(\nabla \rho^1) + \mathcal{H}_{\varepsilon_1}(\nabla \rho^2)$$

and

$$\mathcal{R}_3(M(\rho, v_{\text{GT}})) = \mathcal{H}_{\varepsilon_3}(M(\rho, v_{\text{GT}}))$$

as above.

3.2. Metrics. We consider two metrics to compare the quality of two images: PSNR (which is also used as the metric for the hyperparameter optimisation) and SSIM (which is the secondary alternative metric to compare image quality). To compute them, we use the respective functions provided by *scikit-image* where the "data_range" parameter is obtained from the ground truth image. Also, we compute the metrics for each time point and compute the final metric as the average over all time points.

3.3. Hyperparameter optimisation. Hyperparameter optimisation was carried out using *scikit-optimize* to find good sets of parameters for each model. We chose PSNR (computed on the dynamic mask, see section 4.2) as the metric to be minimised by the Bayesian optimisation. For the Bayesian optimisation we used the Optimizer class with a Gaussian process as the base estimator and expected improvement as the acquisition function. For the model in (4) and its optimisation we used the following hyperparameters. The regularisation parameters $\alpha_1, \alpha_2, \alpha_3$, the smoothness parameters $\varepsilon_1, \varepsilon_2, \varepsilon_3$ and the algorithm parameter σ which gives the initial amount of smoothing of the density and velocity after each outer optimisation step. The number of tested parameters depends on the model. The more model parameters exist, the more combinations of parameters were considered within the Bayesian optimisation. More details can be found in the code.

4. DATA

We work with two datasets taken from the OCMR dataset [8] containing cardiac cine series (long-axis view). The data consists of multi-coil k-space data acquired with Cartesian sampling on a Siemens MAGNETOM Avanto scanner (1.5T) (see [8] for more details). We use fully sampled data and apply the subsampling manually. A fully sampled reconstruction of both datasets can be found in fig. 2. Based on subject 1, we additionally consider a simulated dataset that resembles a real-world MRI image but satisfies the optical flow equation perfectly. This way, we can assess the model capacities and our optimisation procedure in an idealised setting.

4.1. Ground truth images. To assess the reconstruction quality, we need a suitable “ground truth” image to compare the results to. One possibility would be to use the fully sampled measurements to get reconstructions for each time frame separately using the FW model. In doing so, however, one gets time trajectories that are slightly irregular and do not resemble realistic motion (see fig. 2 second column on the right) due to, e.g., in-frame motion. Therefore, we use a “dynamic ground truth” that results from minimising (4) with fully sampled data and high dynamic regularisation parameters. This results in reconstructions that are more regular in time (see fig. 2 first column on the right). This approach also has the advantage that we have “ground truth” velocities available that describe the dynamics within the ground truth image sequence. These can be used to assess the quality of our reconstructed velocities.

In the discussion section, we will briefly comment on the results being similar when using the frame-wise “ground-truth”.

4.2. Mask focusing on dynamic parts in images. We are primarily interested in the dynamic parts of the image and want to improve image quality in these regions particularly. Since large parts of the images are basically constant in time, computing loss metrics between ground truth images and reconstructions over the whole image would consequently result in focusing on the wrong parts. Therefore, we computed masks focusing on the dynamic parts of the images (see fig. 2 right column), and computed the image metrics on these masks only.

4.3. Simulated data. We test our proposed model on simulated data. To this end, we used a dataset from [8] (subject 1 in fig. 2) and solved (4) for fully sampled data. Then, we smoothed the obtained velocities in space and time and used the resulting velocities to simulate a sequence of images starting with the image at the first time point from solving (4). The thus obtained images and velocities are therefore guaranteed to satisfy the optical flow equation and the velocities are smooth enough to be reconstructable. Based on these simulated ground truth images we generated measurements by applying the MRI forward operator A_t . Finally, Gaussian noise (zero mean) was added where the standard deviation was chosen such that unregularised reconstructions yielded qualitatively similar images to unregularised reconstructions based on the original measured data.

4.4. Undersampling. We next describe the used undersampling. The data we are working with has been acquired with a standard Cartesian sampling scheme, and we consider four times undersampling in the phase encoding direction with full sampling of the rows in the readout direction. More precisely, we use the same $x = 15$ percent central rows in each time frame and randomly select $25 - x$ percent

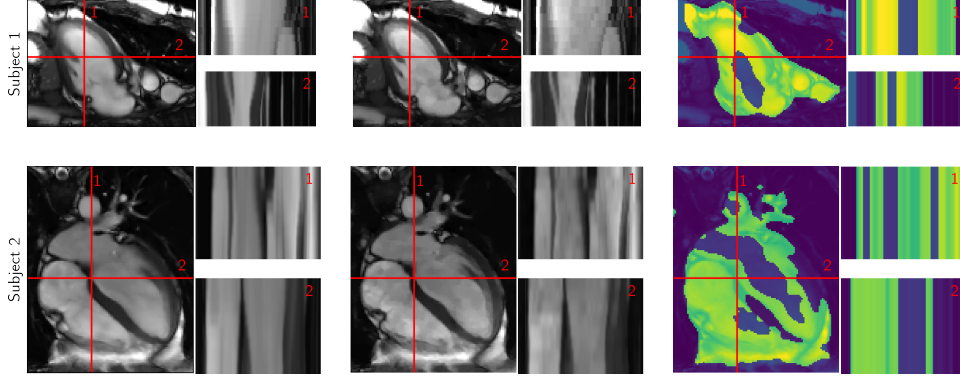


FIGURE 2. Each row shows a reconstruction for a different data set. The first column shows a dynamic ground truth (obtained as a reconstruction using (4) and fully-sampled data). The second column shows the same for a frame-wise ground truth (obtained from reconstructing each time frame separately (FW model) using fully-sampled data). The last column shows an overlay of the reconstructions with the dynamic masks. Within each column, the left side shows one time slice of the reconstructions, the right side shows time-space images along the red lines.

of the outer rows (same number of rows on each side of the central rows) in such a way that the selected outer rows between consecutive time frames do not overlap.

5. EXPERIMENTS AND RESULTS

5.1. Simulated data. The results for the simulated data are summarised in table 1 and can visually be inspected in fig. 3. The top row in that figure shows the magnitude of the simulated images. In the second row, the same images are represented as complex-valued images where the hue represents the normalised phase and the brightness represents the magnitude. From this row of images, we see that the phase varies spatially, but its temporal changes are small.

The other rows depict the difference between the reconstructions using the respective model and the simulated images (magnitude images are compared). We can clearly see that DT improves over the frame-wise approach FW, and that OF improves over DT. The improved reconstruction qualities can also be seen in PSNR values in table 1 (first column), where OF improves over FW and DT by 10dB and 4dB, respectively. Moreover, we see, that **Cheat-OF**, for which ground truth velocities are available in the reconstruction, is basically not improving over OF. Similar observations can be made by comparing the SSIM values in table 1.

This good reconstruction result of the OF model is backed by the reconstructed velocities depicted in fig. 4 (bottom row). We see that the general form of the velocities has been recovered compared to the ground truth velocities (top row), hence a good reconstruction quality is expected.

5.2. Subject 1. The summarised results of subject 1 can be found in table 1 and fig. 5. We see that our proposed method OF performs better than FW or DT. The




FIGURE 3. Results for the simulated data. The top row shows the evolution (magnitude) of the simulated ground truth images for the time points t_0 , t_2 , t_4 , t_6 (out of 8 time points). The second row shows the same images but as a complex-valued image: The hue represents the normalised phase and the brightness represents the magnitude. The other rows depict the difference between the magnitudes of the reconstructions and the magnitude of the simulated images on the dynamic mask. Rows 3-5: FW, DT and OF.

comparison between OF and the frame-wise approach yields PSNR improvements (a similar pattern can be observed in the SSIM values) of almost 3dB, showing that the complex optical flow equation helps the reconstruction by describing the underlying motion. The gains of OF over DT are not very pronounced. However, they improve to over 1dB when looking at finer structures in more detail as can be seen in fig. 5 in the third row. Further, the results in table 1 and fig. 5 also show that the method **Cheat-OF** including ground truth velocities outperforms all the other methods by a lot (which is no surprise since this method basically has access to the fully sampled measurements via the ground truth velocities obtained from a fully sampled reconstruction using (4)). This improvement can be explained by comparing the reconstructed to the “ground truth” velocities in fig. 7. We see that the reconstruction quality is decent for the real part, whereas the reconstructed imaginary part differs significantly from the “ground truth” velocities’ imaginary part. Overall, the velocities are not as well reconstructed as for the simulated data, which explains the improved results using **Cheat-OF**.




FIGURE 4. Reconstructed velocities of one time frame for the simulated data using our proposed complex-valued optical flow method. First row: Ground truth velocities in x and y direction. Second row: Reconstructed velocities in x and y direction

5.3. Subject 2. The summarised results of subject 2 can be found in table 1 and fig. 6. We see that OF improves PSNR over FW and DT by 2dB and 1.4dB, respectively. Therefore, the proposed optical flow model helps with reconstructing dynamic MRI images. Again, a similar conclusion can be drawn from the SSIM values. Looking at the reconstructed phase (fig. 6 rows 3 and 4), we see that all methods reconstruct the phase well on an overall level (row 3). However, the magnified area in row 4 shows differences in the reconstructed phases in the detailed structures. Here, OF gives a better result over FW and DT. This improved phase reconstruction is associated with a better reconstruction of details, as can be seen in the second row. Again, **Cheat-OF**, as expected, improves over all other methods.

6. DISCUSSION

Overall, our experiments demonstrated that explicitly considering the complex-valued optical flow equation (2) is useful for modelling motion between complex-valued images. In particular, our simulated experiment showed large reconstruction improvements. This, of course, is to be expected as the underlying ground truth motion perfectly satisfies the optical flow equation and the corresponding velocities follow the posed prior assumptions. Still, this simplified example shows that the optimisation procedure is suitable for the problem, and meaningful results are obtained. Moreover, we have mentioned already in section 5.1 that **Cheat-OF** is basically not improving over OF in the simulation, suggesting that OF is able to produce nearly optimal reconstructions in case the underlying images satisfy the considered prior assumptions.

Comparing the ground truth velocities of the simulation (fig. 4) and of the measured subject 1 (fig. 7) we see that they are of different quality. In particular, the simulated velocity appears to be much smoother in space, hence they are more in accordance with our prior modelling. On the other hand, the ground truth velocities of subject 1 are much rougher in space which is not reflected in the reconstructed velocity. This of course influences reconstruction quality and partly explains the

		simulation	subject 1 GT(dyn)	subject 2 GT(dyn)	subject 1 GT(FW)	subject 2 GT(FW)
PSNR	FW	34.1 ± 1.2	27.3 ± 0.7	25.6 ± 0.6	27.7 ± 0.8	26.4 ± 0.9
	DT	40.1 ± 0.9	29.6 ± 0.3	26.2 ± 0.3	29.8 ± 0.4	27.5 ± 0.5
	OF	44.1 ± 0.6	30.2 ± 0.5	27.6 ± 0.4	30.2 ± 0.5	28.0 ± 0.5
	Cheat-OF	44.2 ± 0.3	34.4 ± 0.7	32.0 ± 0.7	33.1 ± 0.4	28.7 ± 1.0
SSIM	FW	0.943 ± 0.010	0.828 ± 0.014	0.812 ± 0.015	0.847 ± 0.014	0.852 ± 0.018
	DT	0.978 ± 0.004	0.869 ± 0.009	0.819 ± 0.012	0.880 ± 0.009	0.869 ± 0.011
	OF	0.991 ± 0.001	0.892 ± 0.010	0.873 ± 0.012	0.888 ± 0.011	0.880 ± 0.011
	Cheat-OF	0.992 ± 0.001	0.959 ± 0.009	0.942 ± 0.009	0.934 ± 0.004	0.872 ± 0.018

TABLE 1. Summary of the results. The columns show the metrics for the respective data sets (with dynamic ground truth GT(dyn) and frame-wise ground truth GT(FW)). The displayed values are (average) \pm (standard deviation) taken over the considered time points. The metric for each time point is computed on a dynamic mask (see section 4.2, fig. 2).




FIGURE 5. Reconstructions for subject 1. Left to right columns: Ground truth image, FW, DT, OF and Cheat-OF. First row shows magnitude of reconstructions. Second row: Difference of reconstruction to ground truth on the mask. Third row: Magnified area and corresponding metrics on magnification.




FIGURE 6. Reconstructions for subject 2. Left to right columns: Ground truth image, FW, DT, OF and Cheat-OF. First and second row show magnitude of reconstructions with respective metrics. Third and fourth row show the phase of the reconstructions. Second and last row are the indicated magnified areas.




FIGURE 7. Comparison between ground truth velocities (top row) and reconstructed velocities (bottom row) of subject 1 for a single time frame.

differences in reconstruction quality compared to the simulation. In future work, it will be interesting to investigate different priors that better resemble the underlying velocities’ structure. For instance, one could use Haar wavelets, include regularisation in time, or follow a learning-based approach that can derive very specialised and problem-specific priors.

For our “ground truth” images we argued in section 4.2 that using “dynamic ground truth” images is more appropriate in our setting. However, we repeated the procedures for subjects 1 and 2 with “frame-wise ground truth” images (for the simulation both versions of ground truth images coincide) with similar results, where PSNR improvements are slightly less pronounced (2.5dB over FW for subject 1 and 1.6dB over FW for subject 2, see table 1 right two columns). This suggests that the overall results do not depend on the choice of “ground truth” images.

7. CONCLUSIONS

Dynamic MRI (like cardiac MRI) is an important clinical imaging modality for which measurement undersampling is vital to get reasonably fast acquisitions. To remedy the image reconstruction artifacts naturally appearing due to this undersampling, motion models such as the optical flow equation can be used to propagate information across time frames. Current models have followed this approach for real-valued images. In this work, we have generalised the procedure to complex-valued images by proposing a complex version of the optical flow equation that can account for phase changes between time frames. We tested the new model on two real cardiac datasets and saw reconstruction improvements of the complex optical flow model over the comparison methods.

We used central differences for spatial derivatives in the optical flow equation. In future work, one could analyse which discrete differentiation approach works best in the context of the complex optical flow equation. Related is the question of whether a CFL condition changes when transitioning from the real to the complex optical flow equation.

We also saw that including “ground truth” velocities into the model further improved reconstruction quality implying that the currently chosen regularisation for the velocities is not optimal. In future work, it will be interesting to further investigate how to best choose regularisers for the velocities to achieve better performance (such an optimal choice is not even clear for the standard real-valued optical flow motion model).

Overall, we observed that the complex optical flow model is well-suited for cardiac MRI data. Of course, other options are conceivable. For instance, the continuity equation is a typical way of describing motion (especially in three dimensions). Also, a learning-based approach could be vital to adapt the model to the underlying problem so that problem-specific features can be taken into account. It will be interesting to compare different approaches with each other.

ACKNOWLEDGEMENTS

MM would like to express gratitude to Carola-Bibiane Schönlieb for generously hosting a research stay and supporting this work. MM acknowledges support from

the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2044 –, Mathematics Münster: Dynamics – Geometry – Structure, and under the Collaborative Research Centre 1450–431460824, InSight, University of Münster.

MJE acknowledges support from the EPSRC (EP/S026045/1, EP/T026693/1, EP/V026259/1, EP/Y037286/1) and the European Union Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement RE-MODEL.

Calculations for this publication were performed on the HPC cluster PALMA II of the University of Münster, subsidised by the DFG (INST 211/667-1).

Special thanks to Christoph Kolbitsch and Kris Thielemans for their insightful discussions on dynamic MR, which greatly enriched the research process.

REFERENCES

- [1] Angelica I Aviles-Rivero, Noémie Debroux, Guy Williams, Martin J Graves, and Carola-Bibiane Schönlieb. Compressed sensing plus motion (cs + m): A new perspective for improving undersampled mr image reconstruction. *Med Image Anal.*, 2021.
- [2] Elwin C. Bassett, Eugene G. Kholmovski, Brent D. Wilson, Edward V. R. DiBella, Derek J. Dossdall, Ravi Ranjan, Christopher J. McGann, and Daniel Kim. Evaluation of highly accelerated real-time cardiac cine mri in tachycardia. *NMR in Biomedicine*, 27(2):175–182, 2014.
- [3] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [4] Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numer. Math.*, 84(3):375–393, 2000.
- [5] 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.
- [6] Martin Burger, Hendrik Dirks, and Carola-Bibiane Schönlieb. A variational model for joint motion estimation and image reconstruction. *SIAM Journal on Imaging Sciences*, 11(1):94–128, 2018.
- [7] Chong Chen, Barbara Gris, and Ozan Öktem. A new variational model for joint image reconstruction and motion estimation in spatiotemporal imaging. *SIAM Journal on Imaging Sciences*, 12(4):1686–1719, 2019.
- [8] Chong Chen, Yingmin Liu, Philip Schniter, Matthew Tong, Karolina Zareba, Orlando Simonetti, Lee Potter, and Rizwan Ahmad. Ocmr (v1.0)–open-access multi-coil k-space dataset for cardiovascular magnetic resonance imaging. 2020. arXiv:2008.03410.
- [9] Richard Kinh Gian Do, Henry Rusinek, and Bachir Taouli. Dynamic contrast-enhanced mr imaging of the liver: current status and future directions. *Magnetic resonance imaging clinics of North America*, 17(2):339–49, 2009.
- [10] Li Feng, Monvadi B. Srichai, Ruth P. Lim, Alexis Harrison, Wilson King, Ganesh Adluru, Edward V. R. Dibella, Daniel K. Sodickson, Ricardo Otazo, and Daniel Kim. Highly accelerated real-time cardiac cine mri using k-t sparse-sense. *Magnetic Resonance in Medicine*, 70(1):64–74, 2013.
- [11] J.A. Fessler and D.C. Noll. Iterative image reconstruction in mri with separate magnitude and phase regularization. In *2004 2nd IEEE International Symposium on Biomedical Imaging: Nano to Macro (IEEE Cat No. 04EX821)*, pages 209–212 Vol. 1, 2004.
- [12] Hao Gao, Stanislas Rapacchi, Da Wang, John Moriarty, Conor Meehan, James Sayre, Gerhard Laub, Paul Finn, and Peng Hu. Compressed sensing using prior rank, intensity and sparsity model (prism): applications in cardiac cine mri. In *Proceedings of the 20th Annual Meeting of ISMRM, Melbourne, Australia*, volume 2242, 2012.
- [13] Yaron Gordon, Sasan Partovi, and Matthias et al. Müller-Eschner. Dynamic contrast-enhanced magnetic resonance imaging: fundamentals and application to the evaluation of the peripheral perfusion. *Cardiovascular diagnosis and therapy*, 4(2):147–164, 2014.

- [14] Andreas Hauptmann, Ozan Öktem, and Carola Schönlieb. *Image Reconstruction in Dynamic Inverse Problems with Temporal Models*, pages 1–31. Springer International Publishing, Cham, 2021.
- [15] J. S. Jørgensen, E. Ametova, G. Burca, G. Fardell, E. Papoutsellis, E. Pasca, K. Thielemans, M. Turner, R. Warr, W. R. B. Lionheart, and P. J. Withers. Core imaging library - part i: a versatile python framework for tomographic imaging. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2204):20200192, 2021.
- [16] Dong Liang, Bo Liu, JiunJie Wang, and Leslie Ying. Accelerating sense using compressed sensing. *Magnetic Resonance in Medicine*, 62(6):1574–1584, 2009.
- [17] 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.
- [18] 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.
- [19] Evgueni Ovtchinnikov, Richard Brown, Christoph Kolbitsch, Edoardo Pasca, Casper da Costa-Luis, Ashley G. Gillman, Benjamin A. Thomas, Nikos Efthimiou, Johannes Mayer, Palak Wadhwa, Matthias J. Ehrhardt, Sam Ellis, Jakob S. Jørgensen, Julian Matthews, Claudia Prieto, Andrew J. Reader, Charalampos Tsoumpas, Martin Turner, David Atkinson, and Kris Thielemans. Sirt: Synergistic image reconstruction framework. *Computer Physics Communications*, 249:107087, 2020.
- [20] Chen Qin, Jo Schlemper, Jose Caballero, Anthony N. Price, Joseph V. Hajnal, and Daniel Rueckert. Convolutional recurrent neural networks for dynamic mr image reconstruction. *IEEE Transactions on Medical Imaging*, 38(1):280–290, 2019.
- [21] Jo Schlemper, Jose Caballero, Joseph V. Hajnal, Anthony N. Price, and Daniel Rueckert. A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Trans. Medical Imaging*, 37(2):491–503, 2018.
- [22] Bernhard Schmitzer, Klaus P. Schäfers, and Benedikt Wirth. Dynamic cell imaging in pet with optimal transport regularization. *IEEE Transactions on Medical Imaging*, 39(5):1626–1635, 2020.
- [23] Thomas Schuster, Bernadette Hahn, and Martin Burger. Dynamic inverse problems: modelling—regularization—numerics. *Inverse Problems*, 34(4):040301, mar 2018.
- [24] Zoran Stankovic, Bradley D. Allen, Julio Garcia, Kelly B. Jarvis, and Michael Markl. 4d flow imaging with mri. *Cardiovascular Diagnosis and Therapy*, 4(2), 2014.
- [25] Sonja Sudarski, Thomas Henzler, and Holger et al. Haubenreisser. Free-breathing sparse sampling cine mr imaging with iterative reconstruction for the assessment of left ventricular function and mass at 3.0 t. *Radiology*, 282(1):74–83, 2017.
- [26] Muhammad Usman, David Atkinson, Freddy Odille, Christoph Kolbitsch, Ghislain Vaillant, Tobias Schaeffter, Philip G. Batchelor, and Claudia Prieto. Motion corrected compressed sensing for free-breathing dynamic cardiac mri. *Magnetic Resonance in Medicine*, 70(2):504–516, 2013.
- [27] Tuomo Valkonen. A primal–dual hybrid gradient method for nonlinear operators with applications to mri. *Inverse Problems*, 30(5):055012, may 2014.
- [28] Junwei Yang, Thomas Küstner, and Peng et al. Hu. End-to-end deep learning of non-rigid groupwise registration and reconstruction of dynamic mri. *Frontiers in Cardiovascular Medicine*, 9, 2022.
- [29] Ningning Zhao, Daniel O’Connor, and Basarab Adrian et al. Motion compensated dynamic mri reconstruction with local affine optical flow estimation. *EEE transactions on bio-medical engineering*, 66(11):3050–3059, 2019.