

Received October 15, 2020, accepted November 28, 2020, date of publication December 9, 2020, date of current version December 24, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3043638

Robust Image Reconstruction With Misaligned Structural Information

LEON BUNBERT¹ AND MATTHIAS J. EHRHARDT²

¹Department of Mathematics, University of Erlangen, 91058 Erlangen, Germany

²Institute for Mathematical Innovation, University of Bath, Bath BA2 7AY, U.K.

Corresponding author: Leon Bungert (leon.bungert@fau.de)

The work of Leon Bungert was supported by the European Union's Horizon 2020 Research and Innovation Programme through the Marie Skłodowska-Curie Grant Agreement 777826 (NoMADS). The work of Matthias J. Ehrhardt was supported in part by the Engineering and Physical Sciences Research Council (EPSRC) under Grant EP/S026045/1 and Grant EP/T026693/1, in part by the Faraday Institution under Grant EP/T007745/1, and in part by the Leverhulme Trust under Grant ECF-2019-478.

ABSTRACT Multi-modality (or multi-channel) imaging is becoming increasingly important and more widely available, e.g. hyperspectral imaging in remote sensing, spectral CT in material sciences as well as multi-contrast MRI and PET-MR in medicine. Research in the last decades resulted in a plethora of mathematical methods to combine data from several modalities. State-of-the-art methods, often formulated as variational regularization, have shown to significantly improve image reconstruction both quantitatively and qualitatively. Almost all of these models rely on the assumption that the modalities are perfectly registered, which is not the case in most real world applications. We propose a variational framework which jointly performs reconstruction and registration, thereby overcoming this hurdle. Our approach is the first to achieve this for different modalities and outranks established approaches in terms of accuracy of both reconstruction and registration. Numerical results on simulated and real data show the potential of the proposed strategy for various applications in multi-contrast MRI, PET-MR, and hyperspectral imaging: typical misalignments between modalities such as rotations, translations, zooms can be effectively corrected during the reconstruction process. Therefore the proposed framework allows the robust exploitation of shared information across multiple modalities under real conditions.

INDEX TERMS Image fusion, image reconstruction, image registration, multi-modality imaging, nonconvex and nonsmooth optimization, variational regularization.

I. INTRODUCTION

The aim of this paper is to find an approximate solution to an ill-posed inverse problem of the form $Au = f$, within the framework of structural regularization. Here A denotes the forward operator and f is measured data which is typically corrupted by noise. Examples include magnetic resonance imaging (MRI), positron emission tomography (PET), computed tomography (CT), image denoising, deblurring, or super-resolution, or possibly a combination of these tasks. A widely used approach of obtaining approximate solutions of inverse problems is *variational regularization* where prior knowledge, like for instance sparsity of the solution or its gradient, is enforced through a regularization functional, see [1]–[3] and references therein.

The associate editor coordinating the review of this manuscript and approving it for publication was Hengyong Yu.

In multi-modality imaging one is often in possession of a specific piece of prior knowledge, namely a side information v which is known to have some “common structure” with the true solution u of the inverse problem [4]. The literature on this topic is rich with some works as early as the 1990's so we only list a couple of key papers here. For instance, v could be a high-resolution photograph which assists the reconstruction of low-resolution hyperspectral images [5]–[9] or an anatomical (MRI or CT) image for the reconstruction of PET images [10]–[17], see Figure 1. This strategy has also been used for functional MRI (fMRI) [18], spectral CT [19], electrical impedance tomography (EIT) [20] and multi-contrast MRI [21]–[26].

While these approaches can greatly improve the image reconstruction, they all assume that the target solution u and the side information v are perfectly aligned. If they are badly aligned, most of these methods yield unsatisfactory




FIGURE 1. In many applications two images of different contrast and resolution are acquired. Images courtesy of D. Coomes, P. Markiewicz and J. Schott.




FIGURE 2. While reconstruction with structural side information (center) can yield very good reconstructions even for noisy data (left), the reconstruction is completely distorted when the side information is misaligned (right).

reconstructions, see Figure 2 for an illustration. In real-world applications such misalignment typically cannot be avoided since frequently the acquisitions of v and f happen at different times and through different modalities, see e.g. [27].

To mitigate this problem, we consider the following joint reconstruction and registration model

$$\min_{u, \phi} \mathcal{D}(A(u \circ \phi); f) + \mathcal{R}(u; v) + \mathcal{S}(\phi). \quad (1)$$

This model seeks to reconstruct an image u together with a deformation field ϕ , such that u has common structure with the side information v (enforced through the regularization term $\mathcal{R}(u; v)$) and the deformed image $u \circ \phi$ explains the data f well (enforced through the data fidelity term $\mathcal{D}(A(u \circ \phi); f)$). In addition, functional \mathcal{S} can be chosen to suitably regularize ϕ . Note that the minimization can be restricted to subsets of admissible images or deformation fields. In particular, one can use *parametric deformations* in order to only allow for rigid, affine, diffeomorphic, and other deformations.

The main advantage of our model over the ones discussed in the next section is its jointness which allows for more accurate estimations of the misalignment and leads to better image quality.

A. RELATED WORK

A common strategy in the literature is to perform a three-step strategy: first reconstruct u without the side information v , then to register u and v and finally to reconstruct again using a registered side information (cf. (22) below). This has been considered for PET-MR [28]. The registration of different modalities is often achieved through maximizing the mutual information [29], [30]. A slightly more advanced strategy is to alternate between reconstruction and registration. For instance, this strategy was used for PET-CT [16], [17],

super-resolution MRI with CT [31] and multi-contrast MRI [32]. However, it is not clear if this alternating procedure converges to a limiting reconstruction / registration and if so how the limit can be interpreted.

Related to our proposed approach is blind deconvolution [9], [33] which can correct for translations between side information and measured data. However, it cannot correct for more complicated transformations.

In [34] the authors propose the fusion of hyperspectral image data with a misaligned RGB image using a joint objective functional similar to (1) and optimize it using PALM [35]. In contrast to our work, they assume that both modalities are connected via a system of linear equations which is too much of an assumption for general modalities.

Similarly, in [36] a variational model for the fusion of misaligned multispectral and panchromatic images was proposed and solved with an alternating minimization approach.

In [37], [38] the related problem of simultaneously reconstructing an image while registering it to a template was studied. Models to jointly reconstruct and estimate the motion between a sequence of images observed through indirect measurements were proposed in [39], [40]. However, these methods can only deal with data from a single modality. Registering the reconstruction to a template of a different modality is not meaningful, hence we enforce registration indirectly through the regularization functional $\mathcal{R}(u; v)$.

There are only a few neural network based approaches in the literature for reconstruction with structural priors. We would like to refer to [41], [42] for pansharpening algorithms using deep neural nets. However, they do not consider misalignments between the panchromatic image and the multi-/hyperspectral data.

II. MATHEMATICAL MODEL

A. NOTATION

In this section we thoroughly define all quantities which are involved in model (1) above. The optimization takes place over images $u \in \mathcal{U}$ and deformation fields $\phi \in \mathcal{V}$. We assume that an image $u \in \mathcal{U}$ has its pixels located at some fixed positions x_i in \mathbb{R}^d for $i = 1, \dots, n$, where $d \in \mathbb{N}$ denotes the image dimension (typically $d \in \{2, 3\}$), which allows us to identify $\mathcal{U} \cong \mathbb{R}^n$ where $u_i := u(x_i)$. Furthermore, this lets us identify a deformation field $\phi \in \mathcal{V}$, which we denote as mapping $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $x \mapsto \phi(x)$, with an array in $\mathbb{R}^{n \times d}$ by letting the i -th row be given by $\phi_i := \phi(x_i) \in \mathbb{R}^d$ for $i = 1, \dots, n$. For $u \in \mathcal{U}$ and $\phi \in \mathcal{V}$ we define their composition as

$$u \circ \phi := J(u; \phi), \quad (2)$$

where $J : \mathcal{U} \times \mathcal{V} \rightarrow \mathcal{U}$ denotes an interpolation operator. Mathematically speaking, being an interpolation operator requires two properties: firstly, it shall hold $J(u; \text{id}) = u$ for all $u \in \mathcal{U}$ where $\text{id} \in \mathcal{V}$ denotes the identity deformation field $\text{id}(x) = x$. Secondly, for fixed $\phi \in \mathcal{V}$ the map $u \mapsto J(u; \phi)$ shall be linear. In contrast, for fixed $u \in \mathcal{U}$ the map $\phi \mapsto J(u; \phi)$ is nonlinear, in general. However, we

assume that J is continuously differentiable in its second argument, which is the case for biquadratic (-cubic, -quartic, etc.) spline interpolation, for instance. Note that linear interpolation is not differentiable and is hence not considered here. The differentiability assumption is not only of mathematical relevance but is also important from an algorithmic point of view (cf. [43]) since it assures that optimization algorithms based on derivatives are well-defined.

The forward operator in (1) is a linear map $A : \mathcal{U} \rightarrow \mathbb{R}^m$ and $f \in \mathbb{R}^m$ is the measured data. In some applications, e.g. MRI, one needs a complex-valued data space. This can be incorporated in our setting by identifying the complex numbers with \mathbb{R}^2 . For defining the regularization functional \mathcal{R} in (1) we will also need the notion of a gradient $\nabla u \in \mathbb{V}$ of an image $u \in \mathcal{U}$. To this end we use forward finite differences, as detailed for instance [9]. Throughout the paper, we let $\|x\|$ denote the standard Euclidean norm of a vector x .

B. DATA FIDELITIES

The choice of the data fidelity \mathcal{D} in (1) typically depends on the statistical noise modelling of the data acquisition, cf. [1] and references therein. For Gaussian noise—which models the noise in MRI or hyperspectral imaging, for instance—one uses the squared Euclidean norm

$$\mathcal{D}(u; f) = \frac{1}{2} \|u - f\|^2. \tag{3}$$

For Poisson distributed noise—which occurs during PET measurements, for instance—one utilizes the Kullback–Leibler distance

$$\mathcal{D}(u; f) = \int_{\Omega} u - f + f \log\left(\frac{f}{u}\right) dx. \tag{4}$$

There are also relevant non-differentiable fidelities, as for instance, the 1-norm $\mathcal{D}(u; f) = \|u - f\|_1$ which is used for the removal of impulse noise. However, they do not fit into our optimization framework (cf. Section III-A) which requires a differentiable fidelity term.

C. STRUCTURE-PROMOTING REGULARIZATION

The main challenge of our model is that it combines data from different modalities. In particular, the image u , observed through the data f , and the available side information v typically have completely different contrasts. This makes the methods [37]–[40] inapplicable. We only assume structural similarities on u and v . While there are also other approaches [12], we model this similarity in terms of shared image edges. Mathematically, this means that u and v ought to have parallel gradients, which can be expressed as

$$\nabla u_i - \frac{\nabla v_i \nabla v_i^T}{\|\nabla v_i\|^2} \nabla u_i = 0 \tag{5}$$

for all $i = 1, \dots, n$. Since gradients are orthogonal to level sets, this condition is also referred to as *parallel level sets*, see [4], [44], [45]. Note that the left hand side in (5) vanishes whenever ∇u and ∇v are collinear and is equal to ∇u if they

are orthogonal. To enforce this constraint, one defines the *directional total variation* (cf. [22]) with respect to $v \in \mathcal{U}$ as

$$dTV(u; v) = \sum_{i=1}^n \|P_i \nabla u_i\|, \tag{6}$$

$$P_i = \mathbb{1} - \xi_i \xi_i^T, \tag{7}$$

$$\xi_i = \gamma \frac{\nabla v_i}{\sqrt{\|\nabla v_i\|^2 + \varepsilon^2}}, \tag{8}$$

where $\mathbb{1}$ denotes the $d \times d$ identity matrix. The parameter $\gamma \in [0, 1)$ steers the influence of the side information from no ($\gamma = 0$) to high ($\gamma \approx 1$), where $\gamma = 1$ should be avoided for theoretical reasons, see [9], [46]. For $\gamma = 0$ one observes that dTV reduces to the standard total variation TV . Furthermore, $\varepsilon > 0$ is a small parameter which assures that (8) is well-defined even if $\nabla v_i = 0$. Note that if $\gamma = 1$ and $\varepsilon = 0$, then (7) is the projection onto the orthogonal complement of $\text{span}(\nabla v_i)$ and in this case (5) holds for all $i = 1, \dots, n$ if and only if $dTV(u; v) = 0$. Conversely, choosing $\varepsilon > 0$ and $\gamma < 1$ implies that $\nabla u_i = 0$ is energetically cheaper for dTV than $\nabla u_i = \nabla v_i$. This ensures that no artificial edges of v are introduced.

For many modalities the physical assumption that images are nonnegative makes sense. In these cases, we enforce this via the characteristic function

$$\iota_+(u) = \begin{cases} 0, & \text{if } u_i \geq 0 \quad \forall i = 1, \dots, n, \\ \infty, & \text{else,} \end{cases} \tag{9}$$

and set the regularization function \mathcal{R} in (1) as

$$\mathcal{R}(u; v) := \alpha dTV(u; v) + \iota_+(u). \tag{10}$$

Here, $\alpha > 0$ denotes a regularization parameter which steers the amount of regularization used.

We would like to stress that the overarching framework of this contribution is not limited to this particular regularizer and can be used in conjunction with other models, too, for example [4]–[6], [10], [12], [14], [18], [21], [23], [24].

D. PARAMETRIC DEFORMATIONS

In practice, frequently the kind of deformation is approximately known. For instance, a CT and an MRI image of the brain of the same patient can be assumed to be connected by an affine transform of the form

$$\phi(x) = Mx + b, \quad x \in \mathbb{R}^d. \tag{11}$$

where $M \in \mathbb{R}^{d \times d}$ and $b \in \mathbb{R}^d$. Special cases include *rigid motions* in the plane with rotation matrix

$$M = R_\theta := \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}, \quad \theta \in (0, 2\pi] \tag{12}$$

and *shears* with the shear matrix

$$M = S_a := \begin{pmatrix} 1 & a \\ 0 & 1 \end{pmatrix}, \quad a \in \mathbb{R}. \tag{13}$$

However, one can also think of higher-order parametrizations of the form $\phi(x) = H(x, x) + Mx + b$ for $x \in \mathbb{R}^d$ where $H : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ denotes a bilinear form parametrized by a tensor with d^3 degrees of freedom. For more examples of parametric deformations we refer to [43]. For a review on different methods in image registration we refer to [43], [47]. An advantage of using parametrized deformations is that the number of variables in the optimization problem (1) is dramatically reduced since not a whole deformation field with $n \times d$ entries has to be estimated but only the parameters of the parametrization. Remember that n denotes the number of pixel which can be very large, and d denotes the dimension of the physical space which is typically 2 or 3. In the case of affine transformations (11) the number of parameters is $d^2 + d = 6$ for $d = 2$ and for rigid motion in the plane it is only 3. Hence, in these cases the number of parameters is significantly smaller than $n \times d$. Moreover, parametric deformations do not require sophisticated regularization, as used for instance in [39], [40].

We incorporate the parametrization through a map P from the parameter space \mathbb{R}^p to the space of deformation fields \mathbb{V} . In general, P can be nonlinear—as it is the case for rigid motion, for instance—but for affine deformations (11) it is linear and given by $P : \mathbb{R}^6 \rightarrow \mathbb{V}$, defined as

$$P(\varphi) = \left[x \mapsto \begin{pmatrix} 1 + \varphi_1 & \varphi_2 \\ \varphi_3 & 1 + \varphi_4 \end{pmatrix} x + \begin{pmatrix} \varphi_5 \\ \varphi_6 \end{pmatrix} \right] \quad (14)$$

in two spatial dimensions. Note that the parameters φ model the *deviation from the identity* such that $P(0) = \text{id}$. For a more compact notation we suppress the dependency on the parametrization P and write

$$u_\varphi := u \circ P(\varphi) = J(u; P(\varphi)), \quad (15)$$

where we used the composition operator (2). Furthermore, incorporating parametric deformations into (1) we obtain

$$\min_{u \in \mathbb{U}, \varphi \in \mathbb{R}^p} \mathcal{D}(Au_\varphi; f) + \mathcal{R}(u; v) + \mathcal{S}(\varphi). \quad (16)$$

III. ALGORITHMIC FRAMEWORK

A. OPTIMIZATION

Algorithm 1 Proximal Alternating Linearized Minimization (PALM) to Solve (17)

Input: $u, \varphi, \sigma > 0, \tau > 0, K \in \mathbb{N}$

Output: u, φ

- 1: **for** $k = 1, \dots, K$ **do**
- 2: $u = \text{prox}_{\sigma \mathcal{R}}(u - \sigma \partial_u \mathcal{H}(u, \varphi))$
- 3: $\varphi = \text{prox}_{\tau \mathcal{S}}(\varphi - \tau \partial_\varphi \mathcal{H}(u, \varphi))$
- 4: **end for**

In this section, we detail the optimization scheme that we use to solve problem (16) numerically. Before we formulate the algorithm, we cast problem (16) into the form

$$\min_{u, \varphi} \mathcal{H}(u, \varphi) + \mathcal{R}(u) + \mathcal{S}(\varphi), \quad (17)$$

where $\mathcal{H}(u, \varphi) = \mathcal{D}(Au_\varphi; f)$ denotes the data fidelity term and we abbreviated $\mathcal{R}(u) = \mathcal{R}(u; v)$. Note that—due to the smoothness assumption on the interpolation operator in (2)—the function \mathcal{H} is smooth in both variables if the fidelity \mathcal{D} is. Furthermore, the objective function in (17) is non-convex in the joint variable (u, φ) and also non-convex in φ . However, it is convex in the variable u . Still, due to the overall non-convexity one can in general only expect to find critical points of (17) using gradient-based algorithms and the outcome depends on the initialization of the numerical optimization algorithm. Another difficulty arises from the non-smoothness of the regularization functionals in (17) which excludes many gradient-based algorithms. Here we employ the proximal alternating linearized minimization (PALM) algorithm [35] which amounts to alternating forward-backward splitting, see Algorithm 1.

The proximal operator [48]–[50] of a proper function $\mathcal{J} : \mathbb{R}^k \rightarrow \mathbb{R} \cup \{\infty\}$ is defined as

$$\text{prox}_{\mathcal{J}}(x) := \arg \min_{y \in \mathbb{R}^k} \left(\frac{1}{2} \|x - y\|^2 + \mathcal{J}(y) \right). \quad (18)$$

The proximal operator for dTV with nonnegativity constraint (10) can be computed using the fast gradient projection algorithm (also known as FISTA) [51], see [22] for details.

Let us now study the gradients of the data fidelity used in Algorithm 1. With the chain rule and (15), these are given by

$$\partial_u \mathcal{H}(u, \varphi) = \partial_u J(u; P(\varphi))^* A^* \partial \mathcal{D}(Au_\varphi; f) \quad (19)$$

$$\partial_\varphi \mathcal{H}(u, \varphi) = \partial P(\varphi)^* \partial_\phi J(u; P(\varphi))^* A^* \partial \mathcal{D}(Au_\varphi; f) \quad (20)$$

where the asterisk denotes the adjoint / transpose of a linear map. Note that since the interpolation operator J is linear in the first argument, it holds that $\partial_u J = J$ and its adjoint—which involves inverting the deformation—can be explicitly calculated for affine deformations. In our implementation we use the `scipy` [52] routine `RectBivariateSpline` which allows to evaluate bivariate splines of order two or higher together with their derivatives $\partial_\phi J$. Furthermore, the linear operator $\partial P(\varphi)^*$ can be explicitly calculated in dependency on the type of parametrization used. For affine deformations, for instance, map P in (14) is linear; hence, it holds $\partial P = P$ and the adjoint map $P^* : \mathbb{V} \rightarrow \mathbb{R}^6$ can be computed easily.

Finally, we address the choice of step sizes σ and τ appearing in Algorithm 1. From a theoretical point of view it suffices to choose them smaller than the reciprocal of the Lipschitz constants of the gradient maps in (19), (20) to obtain convergence of the algorithm. However, in practice those Lipschitz constants are hard to compute analytically and typically one can only find pessimistic upper bounds for them. Since this slows down convergence, we employ a backtracking scheme instead, see [9] for details.

B. INITIALIZATION STRATEGY

Since image registration alone is already a severely non-convex problem, one can expect that the degree of non-convexity of our joint model (1) is even higher. In particular,

the outcome of Algorithm 1 will depend heavily on its initialization, in general. Hence, in order to be able to detect large deformations between the side information v and the image u such that $Au \approx f$, we employ a scale-space strategy commonly used in image registration [43]. To overcome large deformations we reduce the resolution of the problem such that these correspond to a few pixels only. We also employ an over-smoothing strategy based on the regularization parameter α in (10). It is well-known for a large class of regularization functionals of the form $\mathcal{R} = \alpha \mathcal{J}$, that solutions of the associated variational regularization method converge to an element in the null-space of \mathcal{J} as $\alpha \rightarrow \infty$ (see [53] for total variation and [54] for the general case). For the case of directional total variation, i.e. $\mathcal{J} = \text{dTV}$, this means that for high regularization parameters \mathcal{R} the reconstructed images become less oscillatory and better align with the side information. However, to avoid a strong loss of contrast in the reconstructions [53], one would like to choose the regularization parameter as small as possible. In this work we propose a combination of resolution-based and over-smoothing strategy to solve (1) which is detailed in Algorithm 2. In a nutshell, the algorithm first reconstructs low-resolution images with high regularization parameters first and then successively decreases the regularization parameter while reconstructing better resolved images.

Algorithm 2 Scale-Space Strategy to Solve (17)

Input: u, φ , resolutions $n_1 < \dots < n_M = n$, regularization parameters $\alpha_1 > \dots > \alpha_M$ in (10)

Output: u, φ

- 1: **for** $i = 1, \dots, M$ **do**
 - 2: down-sample v to resolution n_i
 - 3: up-sample u to resolution n_i
 - 4: compute $(u, \varphi) = \text{PALM}(u, \varphi)$
 - 5: **end for**
-

IV. NUMERICAL RESULTS

In all numerical experiments we will use affine parametric deformations (cf. Section II-D), since they model many realistic deformations and are computationally efficient due to the low number of free parameters. Furthermore, we do not regularize the affine deformation fields, meaning that we choose $\mathcal{S}(\varphi) = 0$ for all $\varphi \in \mathbb{R}^p$ in (16). Correspondingly, the second proximal operator in Algorithm 1 reduces to the identity operator. The details of the scale-space Algorithm 2 are specified below but at this point we already remark that we use $K = 500$ iterations of the PALM algorithm in each loop of Algorithm 2. The model parameters in (8) were chosen similar to [9], [22] as $\gamma = 0.9995$, $\varepsilon = 0.01 \cdot \max_{i=1}^n \|\nabla v_i\|$. The numerical experiments have been carried out in Python using ODL [55]. The line integrals in the PET experiment were computed using the Python module `scikit-image` [56] and for interpolation we use biquadratic splines from the `scipy` [52] routine `RectBivariateSpline`. The source

code which reproduces all experiments can be found on https://github.com/mehrhhardt/robust_guided_reconstruction.

The error metrics we use in this section are structural similarity (SSIM) [57] for the comparison of reconstructed and ground truth images, and the relative difference (RD), given by

$$\text{RD} = \frac{\|\varphi - \varphi_{\text{gt}}\|}{\|\varphi_{\text{gt}}\|} \times 100 [\%], \quad (21)$$

for the comparison of reconstructed motion parameters φ with ground truth parameters φ_{gt} . Whereas more advanced evaluation metrics (cf. e.g. [58] for semantic evaluation of neural network based MR reconstruction) would be desirable, they are beyond the scope of this paper and the combination of SSIM, RD, and visual assessment appears to be sufficient for our applications.

In our numerical experiments we compare our method with a three-step method which computes an initial reconstruction, registers the side information, and subsequently reconstructs again (cf. the references in Section I-A). Based on the structural side information v and data f it computes

$$\tilde{u} \in \arg \min_{u \in \mathbb{U}} \mathcal{D}(Au; f) + \alpha \text{TV}(u), \quad (22a)$$

$$\varphi^* \in \arg \max_{\varphi \in \mathbb{R}^p} \mathcal{MI}(v, \tilde{u}_\varphi), \quad (22b)$$

$$u^* \in \arg \min_{u \in \mathbb{U}} \mathcal{D}(Au_{\varphi^*}; f) + \alpha \text{dTV}(u; v). \quad (22c)$$

The maximum is taken over all parametrized affine deformation fields $\varphi \in \mathbb{R}^p$. Furthermore, $\mathcal{MI}(\cdot, \cdot)$ denotes the mutual information [59] of images which is a popular distance measure for the registration of multi-modal images. The affine registration step (22b) was performed using the MATLAB routine `imregtform`.

A. UNDERSAMPLED MULTI-CONTRAST MRI RECONSTRUCTION

In our first experiment we perform multi-contrast MRI reconstruction from undersampled data. Since all contrasts are likely to share a structural information, less data is needed if the expected redundancy is exploited. Here we assume that a T2-weighted has been acquired using “full data” and reconstructed without artefacts. We then use this image to reconstruct a T1-weighted image using only a fraction of the data. In this experiment we simulate the k -space data based on reconstructions of fully sampled clinical data.

The forward operator is a discrete Fourier transform defined on complex-valued images of size 256×256 , followed by a sampling operator corresponding to 15 equidistant spokes and a low-pass sampling of the 10×10 center of the k -space. This yields a sampling of approximately 3% of the k -space. This sampling is easy to approximate on any clinical MRI scanner. The noise modelling here is Gaussian which is why we use the squared Euclidean norm (3) as data fidelity. The image domain is $[-1, 1]^2$ which corresponds to a pixel width of approximately 0.0078. The deformation field in this




FIGURE 3. Multi-contrast MRI reconstruction. Pseudo-inverse, TV, and misaligned dTV do not correct motion. The proposed method corrects the deformation and the reconstruction satisfyingly agrees with the ground truth. We give (SSIM, RD) values of the reconstructed images and deformation parameters. For pseudo-inverse and TV no deformation was corrected and SSIM is given with respect to the deformed ground truth.

simulation is a rigid transform followed by zooming, given by

$$\phi_{\text{zoom}}(x) = \alpha R_{\theta} x + b$$

where R_{θ} is a rotation matrix (12) with angle $\theta = 0.1$, $b = (-0.02, -0.08)^T$ is a translation vector, and $\alpha = 0.85$ is a zoom factor. The latter means that the size of the side information image is only 85% of the size of the image which underlies the k -space data. Since nonnegativity is not a meaningful assumption for complex-valued images, we drop this constraint in (10) for this test case. The resolutions and regularization parameters in Algorithm 2 were chosen as $n_k = (32^2, 62^2, 128^2, 256^2)$ and $\alpha_k = 5 \cdot 10^{-3} \cdot (5^3, 5^2, 5, 1)$. For the TV reconstruction we used $\alpha_k = 5 \cdot 10^{-4} \cdot (5^3, 5^2, 5, 1)$. Here, and also for the following experiments, the resolutions were determined by successively dividing the resolution of the side information by a factor of two. The regularization parameters at the target resolution were chosen experimentally in order to maximize SSIM. Typically, the regularization parameter of TV has to be chosen one or two orders of magnitude smaller than the one for dTV.

The first row of Figure 3 shows the sampling pattern and the deformed image which was used to generate the data. Furthermore, it shows the side information and the ground truth image which we would like to reconstruct.

In the second row of Figure 3 we show the zero-filled reconstruction (pseudo-inverse), a standard TV reconstruction, the result of the three-step method (22), and our proposed method. Obviously the pseudo-inverse and TV method yield unsatisfactory results since they do not utilize the side information. Furthermore, they visualize the strong degree of ill-posedness of the problem where the sampled Fourier

data alone cannot be used for a meaningful reconstruction. In contrast, both the three-step method (22) and our proposed method successfully correct the deformation between side information and data and yield good reconstructions despite the high degree of undersampling in the data. Notably, the proposed method yields better reconstructions of fine structures as can be seen in the zooms of Figure 3.

The previous observations are also supported by quantitative metrics. The SSIM values between the pseudo-inverse and TV reconstruction and the deformed ground truth image which underlies the data are very low which is due to the strong undersampling. In contrast, the three-step method and the proposed method have high similarity values with the ground truth image, with slightly better numbers of the proposed method. Also the relative difference between the estimated affine deformation parameters and the ground truth parameters is lower for the proposed method than for the three-step method which underlines the superiority of our joint approach over the three-step method (22) in this application. We suspect that worse estimation of the deformation field of the three-step method is due to the fact that it must utilize the blurry and blocky TV reconstruction for registration whereas the proposed method iteratively improves registration and reconstruction.

We visualize our algorithm in Figure 4 where we plot the reconstructions and deformations fields of the scale-space Algorithm 2, respectively. Since our reconstructed motion parameters correspond to the deviation from the identity (see (14)), we let the yellow arrows indicate the field $\phi_{\text{zoom}} - \text{id}$ and the red arrows the corresponding expression for the reconstructed fields. Already at the second resolution stage 60×60 the estimated deformation field is visually




FIGURE 4. From left to right: MRI reconstructions of increasing resolutions generated by Algorithm 2. The deviation of the ground truth field ϕ_{zoom} and the reconstructed fields from the identity are shown in yellow and red, respectively. The ground truth field is only visible in the first image where the estimated deformation field is still inaccurate. Relative differences (RD) between reconstructed and ground truth motion parameters are given, as well.

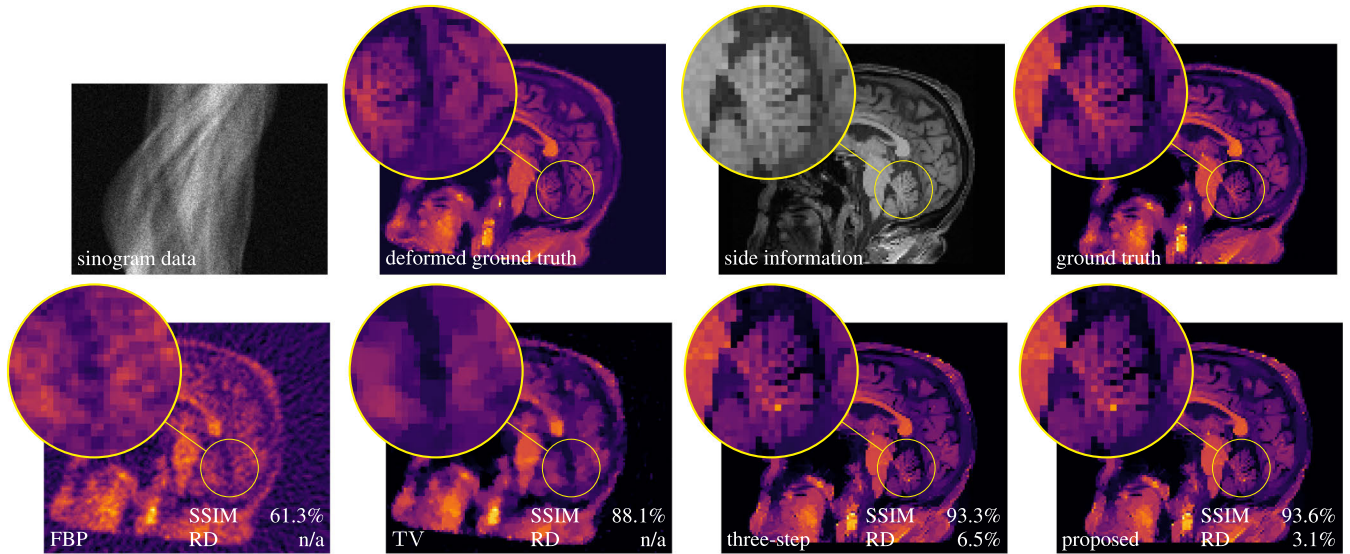


FIGURE 5. PET reconstructions with structural MR side information. Filtered back-projection (FBP) and TV do not correct motion and yield poor reconstructions. Both the three-step and the proposed method correct the deformation and the reconstruction satisfyingly agrees with the ground truth. We give SSIM and RD values of the reconstructed images and deformation parameters. For FBP and TV no deformation was corrected and SSIM is given with respect to the deformed ground truth.

not distinguishable from the ground truth field which is also supported by the decreasing relative differences between the estimated and ground truth motion parameters.

B. PET-MR RECONSTRUCTION

In this experiment we consider PET-MR, where we aim to reconstruct a tracer distribution using a fully sampled T1-weighted MR image of size 144×144 as side information. The forward operator is modelled by a parallel beam X-ray transform with 200 angles equispaced in $(0, \pi]$ and 192 bins. The sinogram data were simulated using a ground truth image deformed with respect to the side information through the rigid deformation

$$\phi_{\text{rigid}}(x) = R_{\theta}x + b, \tag{23}$$

where R_{θ} is a rotation matrix (12) with angle $\theta = 0.1 \approx 5.7^{\circ}$ and $b = (0.02, 0.08)^T$ is a translation vector.

In this experiment we simulate data based on a dTV-regularized reconstruction of clinical data, see [28]. The data is an instance of a Poisson random variable with parameter $Ax + r$, where the background r is chosen as constant 7 and the forward operator is scaled to about $1.3 \cdot 10^6$ expected counts in the data. Correspondingly, the data fidelity used is the Kullback–Leibler distance (4).

Again, images are in $[-1, 1]^2$ and hence the pixel width of the side information is $0.013\bar{8}$. The resolutions and regularization parameters in Algorithm 2 were chosen as $n_k = (9^2, 18^2, 36^2, 72^2, 144^2)$ and $\alpha_k = 4 \cdot 10^{-1} \cdot (10^4, 10^3, 10^2, 10, 1)$. For the TV experiment we used $\alpha_k = 4 \cdot 10^{-2} \cdot (10^4, 10^3, 10^2, 10, 1)$.

The sinogram data and the deformed image which was used to generate the data are shown in the top row of Figure 5. Furthermore, we show the side information and the ground truth image. In the second row of Figure 5 we show four different reconstructions: the first one obtained through filtered




FIGURE 6. From left to right: PET reconstructions of increasing resolutions generated by Algorithm 2. The deviation of the ground truth field ϕ_{rigid} and the reconstructed fields from the identity are shown in yellow and red, respectively. The ground truth field is only visible in the first three images where the estimated deformation field is still inaccurate. In the first image the reconstructed field is zero and thus not plotted. Relative differences (RD) between reconstructed and ground truth motion parameters are given, as well.

back-projection, the second one utilizing TV regularization, the third one using the three-step method (22), and the fourth one being the proposed method. The first two methods, which do not use the side information or correct any motion, exhibit poor image quality due to strong noise in the sinogram. On the other hand, both the three-step and the proposed method correct the deformation and the reconstructions are in very good agreement with the ground truth image.

Quantitatively, the SSIM values for filtered back-projection and TV are relatively low whereas they are comparably high for both the three-step and the proposed method, with slightly better values for the proposed method. The same is also true for the relative errors of the computed deformation fields.

The scale-space generated by Algorithm 2 is again visualized in Figure 6 with observations similar to Section IV-A.

We also test the limits of the three-step method and our method by performing PET reconstructions where the measured sinogram data and the side information at hand are offset by the rigid deformation ϕ_{rigid} from (23) with $\theta \in \{0.2, 0.4, \dots, 1.2\}$. The first row of Figure 7 shows the deformed ground truth images which were used to generate the PET data. In the second row we show the reconstructions generated by the three-step method (22) together with SSIM and RD values for the images and deformation parameters, respectively.

The remaining five rows of Figure 7 show the results of our proposed method for different sizes $M \in \{1, 2, 3, 4, 5\}$ of the scale space in Algorithm 2. The method behaves very reliably in the sense that a larger scale space implies enables the method to estimate larger deformations correctly. In particular, for scale space of size 3 or larger our method breaks down much later than the three-step method (22) which internally uses state-of-the-art multimodal registration algorithms. This suggests a high degree of robustness of our method and also shows that the scale space algorithm 2 is inevitable for joint reconstruction / registration problems.

Also the error metrics SSIM and RD are better for the proposed method than for the three-step method, as it was the case in the previous experiments.

C. HYPERSPECTRAL SUPER-RESOLUTION

In this section we fuse a 100×100 spectral data channel with a highly resolved panchromatic side information of size

400×400 . We assume Gaussian noise and hence utilize the Euclidean fidelity (3).

We consider three test cases *rigid*, *shear*, and *nonlinear* where we artificially introduce different deformations between data and side information. For *rigid* we utilize the deformation ϕ_{rigid} in (23) with $\theta = 0.1 \approx 5.7^\circ$ and $b = (0.06, -0.04)^T$. The pixels of the images lie in $[-1, 1]^2$ and have width 0.02 (data image) and 0.005 (side information). Hence, we try to correct an effective off-set of approximately 12 pixels. The other deformations are given by $\phi_{\text{shear}}(x) = S_a x + b$ —where S_a is the shear matrix (13) with $a = 0.08$ —and

$$\phi_{\text{nonlinear}}(x) = \phi_{\text{rigid}}(x) + 0.05 \begin{pmatrix} x_2^2 \\ -x_1^3 \end{pmatrix}.$$

We notice that $\phi_{\text{nonlinear}}$ is a nonlinear perturbation of ϕ_{rigid} and, in particular, is *not affine*. In the scale-space Algorithm 2 we chose five different resolutions $n_k = (25^2, 50^2, 100^2, 200^2, 400^2)$ and corresponding regularization parameters $\alpha_k = 10^{-3} \cdot (10^4, 10^3, 10^2, 10, 1)$. These parameters were used for all reconstructions involving dTV and for all test cases. In the experiment with standard total variation TV as the regularizer we used the parameters $\alpha_k = 10^{-5} \cdot (10^4, 10^3, 10^2, 10, 1)$ instead.

The first row of Figure 8 shows the data *rigid* together with the side information misaligned through a rigid transform. Below we show four different reconstructions, including a target reconstruction which was computed using aligned data and serves as substitute for a ground truth solution for this real data set. The second one is a standard TV reconstruction, which does not utilize the side information at all and hence yields the expected poor and blurry super-resolution result. Instead of the three-step method—which we compared with our method in Sections IV-A and IV-B—the third reconstruction is now a dTV reconstruction which uses misaligned data and side information and does not correct the deformation. This result corresponds to a naive super-resolution which does not take into account the misalignment between the two modalities and hence produces an erroneous image with many artefacts. The last reconstruction is obtained using our proposed method which corrects for the misalignment between data and side information and almost perfectly agrees with the target reconstruction which is also reflected by its SSIM value.




FIGURE 7. Testing the limits of our method: first row: rotated images which were used to create the PET data, second row: results of the three-step method, other rows: proposed reconstructions for difference sizes M of the scale space. Failed registrations are marked with a red frame. Note that relative differences are larger for small displacements since these have a small norm (cf. (14) and (21)).

In Figure 9 we apply our method also to the other data sets shear and nonlinear. Also here our algorithm was able to correct the deformation between data and side information sufficiently to achieve reconstructions which are very similar to the target reconstruction, as also shown by SSIM. Note that even though the deformation field in the data set

nonlinear is not affine, the reconstruction is almost as good as those for the affine data sets rigid and shear and the SSIM to the target reconstruction is only slightly lower than for rigid or shear. Since the ground truth deformation field is non-affine for this data set, it is not meaningful to provide relative distances of the estimated parameters here.

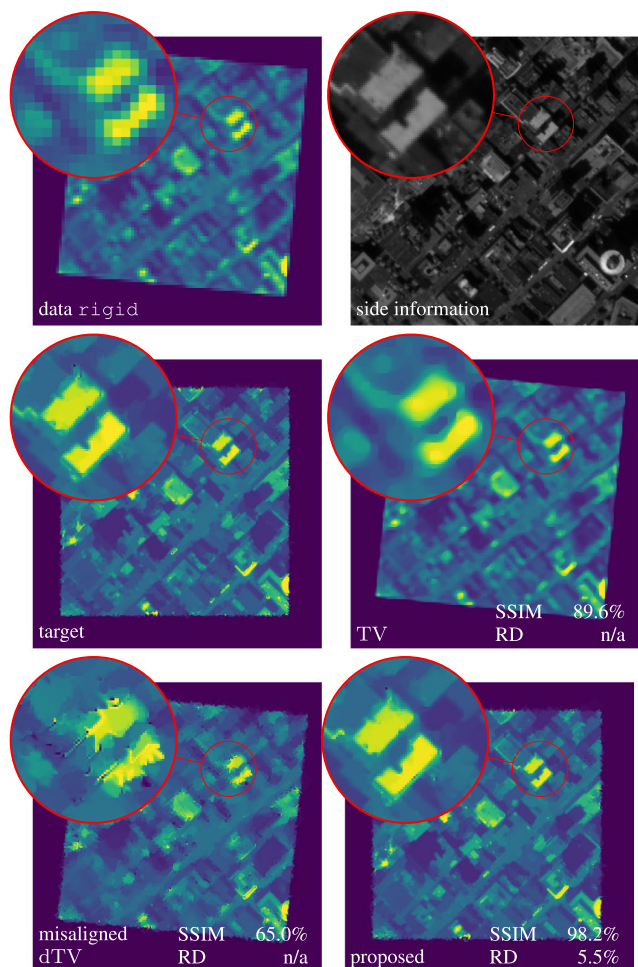


FIGURE 8. Hyperspectral super-resolution. Reconstruction of data set *rigid* using three different methods. TV and misaligned dTV do not correct deformations and yield poor results. The proposed method corrects the deformation and the reconstruction is in very good agreement with the target. We give SSIM and RD values of the reconstructed images and deformation parameters. For TV and misaligned dTV no deformation was corrected and for TV SSIM is given with respect to a deformed target image.

V. DISCUSSION

In this work we have proposed a variational framework for jointly reconstructing an image from indirect measurements while registering it to a structural side information from a potentially different modality. We have then specialized our model to the case of parametric deformations. In numerical experiments on multi-contrast MR, PET-MR, and hyperspectral super-resolution we have compared our method with standard reconstruction methods (pseudo inversion, total variation regularization) and with a three-step method which first reconstructs an auxiliary image, then registers the side information, and finally reconstructs again.

While the standard methods, which do not use any structural prior, do not produce satisfactory results for the type of inverse problems we are considering, both the three-step method and the proposed method yield comparable results and their pros and cons are discussed in more detail in the following.




FIGURE 9. Reconstructions of data sets *shear* and *nonlinear*. The proposed method corrects the deformations and yields very good reconstructions. We give SSIM values of the reconstructed images and deformation parameters for both data sets and RD values for the affine data set *shear*.

The clear benefit of the three-step method is that it is straightforward to implement since it can be constructed by combining any two (black-box) reconstruction and multi-modal registration algorithms. Furthermore, by using state-of-the-art solvers of the sub-problems one can create a very efficient method. However, the drawback of this strategy becomes evident in our numerical experiments on MR and PET-MR where—due to the ill-posedness of the problems—the initial reconstruction does not suffice to estimate an accurate deformation field. Correspondingly, the final reconstructions of the three-step method is not as accurate as our proposed approach. Furthermore, as seen in Section IV-B the three-step methods seems to be less robust to large deformations since it utilizes a possibly poor initial reconstruction which makes the registration hard.

Thanks to its joint approach, which *simultaneously* estimates the deformation field and the reconstruction, our proposed method is able to estimate deformations more accurately which reflected in better image quality and higher robustness to larger deformations. Another benefit of our method is that, despite its non-convexity, it can be solved with optimization algorithms which come with convergence guarantees.

Incorporating more general deformation models than the affine one considered here can be achieved by suitably choosing the regularization functionals \mathcal{R} and \mathcal{S} in (1) or even choosing a joint regularization. For diffeomorphic deformations this can be done similarly to [40] and for optical flow approaches see [39]. In any case, judging from what the authors report in these works, the scale space algorithm 2 is necessary also for more complicated types of deformations.

Finally, it is also worth discussing the limitations of the assumption that side information and ground truth solution have structural similarity in terms of parallel gradients (cf. (5)). For the scenario of an aligned side information it was demonstrated in [13], [15], [45] that features which are only present in PET data but not in the structural MR prior can be recovered in the reconstruction. Similarly, excess edges in the MR image were shown not to propagate into the reconstruction. A comprehensive overview of different MR-informed PET reconstruction methods is given in [60]. While these observations were made for registered modalities, one should mention that in our case the registration (enforced through the regularization functional $\mathcal{R}(u; v)$) heavily depends on the side information. In particular, a successful registration can only be obtained if most of the edges in the data are also present in the side information. In the extreme case where the side information is a constant image and no registration can be performed, this condition is of course violated.

VI. CONCLUSION

Multi-modality imaging is becoming ever more important in many disciplines from medicine to material sciences. Mathematical approaches to combine data from multiple modalities exist but are prone to imperfect misalignment. In this work we proposed a generic framework to reconstruct an image from indirect measurements while registering it to a structural side information from a potentially different modality. Numerical experiments for a specific regularizer (directional total variation) and a variety of applications including hyperspectral imaging in remote sensing, PET-MR and multi-contrast MRI underpinned the aptness of the proposed approach to correct misalignments between data and structural side information which cause existing algorithms to fail. Missing robustness to misalignment has been the biggest hurdle for integrating structure promoting regularizers into routine use, for instance in the clinic. Thus, the proposed framework paves the way for their translation into various applications.

REFERENCES

- [1] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen, *Variational Methods in Image Processing* (Applied Mathematical Sciences), vol. 167. Cham, Switzerland: Springer, 2008.
- [2] K. Ito and B. Jin, *Inverse Problems: Tikhonov Theory and Algorithms*. Singapore: World Scientific, 2014.
- [3] M. Benning and M. Burger, "Modern regularization methods for inverse problems," *Acta Numerica*, vol. 27, pp. 1–111, May 2018.
- [4] M. J. Ehrhardt, "Multi-modality imaging with structure-promoting regularisers," 2020, *arXiv:2007.11689*. [Online]. Available: <http://arxiv.org/abs/2007.11689>
- [5] C. Ballester, V. Caselles, L. Igual, J. Verdera, and B. Rougé, "A variational model for P+XS image fusion," *Int. J. Comput. Vis.*, vol. 69, no. 1, pp. 43–58, Aug. 2006.
- [6] M. Möller, T. Wittman, A. L. Bertozzi, and M. Burger, "A variational approach for sharpening high dimensional images," *SIAM J. Imag. Sci.*, vol. 5, no. 1, pp. 150–178, Jan. 2012.
- [7] L. Loncan, L. B. De Almeida, J. M. Biucas-Dias, X. Briottet, J. Chanussot, N. Dobigeon, S. Fabre, W. Liao, G. A. Licciardi, M. Simoes, J. Y. Tourmeret, M. A. Vezanzones, G. Vivone, Q. Wei, and N. Yokoya, "Hyperspectral pansharpening: A review," *IEEE Trans. Geosci. Remote Sens.*, vol. 3, no. 3, pp. 27–46, Sep. 2015.
- [8] J. Duran, A. Buades, B. Coll, C. Sbert, and G. Blanchet, "A survey of pansharpening methods with a new band-decoupled variational model," *ISPRS J. Photogramm. Remote Sens.*, vol. 125, pp. 78–105, Mar. 2017.
- [9] L. Bungert, D. A. Coomes, M. J. Ehrhardt, J. Rasch, R. Reisenhofer, and C.-B. Schönlieb, "Blind image fusion for hyperspectral imaging with the directional total variation," *Inverse Problems*, vol. 34, no. 4, Apr. 2018, Art. no. 044003.
- [10] R. M. Leahy and X. Yan, "Incorporation of anatomical MR data for improved functional imaging with PET," in *Proc. Int. Process. Med. Imag. Cham, Switzerland: Springer*, 1991, pp. 105–120.
- [11] G. Gindi, M. Lee, A. Rangarajan, and I. G. Zubal, "Bayesian reconstruction of functional images using anatomical information as priors," *IEEE Trans. Med. Imag.*, vol. 12, no. 4, pp. 80–670, Jan. 1993.
- [12] J. E. Bowsher, H. Yuan, L. W. Hedlund, T. G. Turkington, G. Akabani, A. Badea, W. C. Kurylo, C. T. Wheeler, G. P. Cofer, M. W. Dewhirst, and G. A. Johnson, "Utilizing MRI information to estimate F18-FDG distributions in rat flank tumors," in *Proc. IEEE Symp. Conf. Rec. Nucl. Sci.*, 2004, pp. 2488–2492.
- [13] M. J. Ehrhardt, P. Markiewicz, M. Liljeroth, A. Barnes, V. Kolehmainen, J. S. Duncan, L. Pizarro, D. Atkinson, B. F. Hutton, S. Ourselin, K. Thielemans, and S. R. Arridge, "PET reconstruction with an anatomical MRI prior using parallel level sets," *IEEE Trans. Med. Imag.*, vol. 35, no. 9, pp. 2189–2199, Sep. 2016.
- [14] A. Mehranian, M. A. Belzunce, F. Niccolini, M. Politis, C. Prieto, F. Turkheimer, A. Hammers, and A. J. Reader, "PET image reconstruction using multi-parametric anato-functional priors," *Phys. Med. Biol.*, vol. 62, no. 15, pp. 5975–6007, Jul. 2017.
- [15] J. Rasch, E.-M. Brinkmann, and M. Burger, "Joint reconstruction via coupled bregman iterations with applications to PET-MR imaging," *Inverse Problems*, vol. 34, no. 1, Jan. 2018, Art. no. 014001.
- [16] Y.-J. Tsai, A. Bousse, S. Ahn, C. W. Stearns, S. Arridge, B. F. Hutton, and K. Thielemans, "Algorithms for solving misalignment issues in penalized PET/CT reconstruction using anatomical priors," in *Proc. IEEE Nucl. Sci. Symp. Med. Imag. Conf. Proc. (NSSMIC)*, Nov. 2018, pp. 1–3.
- [17] Y.-J. Tsai, A. Bousse, S. Arridge, C. W. Stearns, B. F. Hutton, and K. Thielemans, "Penalized PET/CT reconstruction algorithms with automatic realignment for anatomical priors," 2019, *arXiv:1911.08012*. [Online]. Available: <http://arxiv.org/abs/1911.08012>
- [18] J. Rasch, V. Kolehmainen, R. Nivajärvi, M. Kettunen, O. Gröhn, M. Burger, and E.-M. Brinkmann, "Dynamic MRI reconstruction from undersampled data with an anatomical prescan," *Inverse Problems*, vol. 34, no. 7, Jul. 2018, Art. no. 074001.
- [19] D. Kazantsev, J. S. Jørgensen, M. S. Andersen, W. R. B. Lionheart, P. D. Lee, and P. J. Withers, "Joint image reconstruction method with correlative multi-channel prior for X-ray spectral computed tomography," *Inverse Problems*, vol. 34, no. 6, Jun. 2018, Art. no. 064001.
- [20] V. Kolehmainen, M. J. Ehrhardt, and S. R. Arridge, "Incorporating structural prior information and sparsity into EIT using parallel level sets," *Inverse Problems Imag.*, vol. 13, no. 2, pp. 285–307, 2019.
- [21] B. Bilgic, V. K. Goyal, and E. Adalsteinsson, "Multi-contrast reconstruction with Bayesian compressed sensing," *Magn. Reson. Med.*, vol. 66, no. 6, pp. 1601–1615, 2011.
- [22] M. J. Ehrhardt and M. M. Betcke, "Multicontrast MRI reconstruction with structure-guided total variation," *SIAM J. Imag. Sci.*, vol. 9, no. 3, pp. 1084–1106, Jan. 2016.
- [23] J. Huang, C. Chen, and L. Axel, "Fast Multi-contrast MRI Reconstruction," *Magn. Reson. Imag.*, vol. 32, no. 10, pp. 52–1344, 2014.
- [24] P. Song, L. Weizman, J. F. C. Mota, Y. C. Eldar, and M. R. D. Rodrigues, "Coupled dictionary learning for multi-contrast MRI reconstruction," *IEEE Trans. Med. Imag.*, vol. 39, no. 3, pp. 621–633, Mar. 2020.
- [25] A. J. Obert et al., "¹H-guided reconstruction of ¹⁹F gas MRI in COPD patients," *Magn. Reson. Med.* vol. 84, pp. 1336–1346, 2020, doi: [10.1002/mrm.28209](https://doi.org/10.1002/mrm.28209).
- [26] I. Teh, D. McClymont, E. Carruth, J. Omens, A. McCulloch, and J. E. Schneider, "Improved compressed sensing and super-resolution of cardiac diffusion MRI with structure-guided total variation," *Magn. Reson. Med.*, vol. 84, pp. 1868–1880, 2020, doi: [10.1002/mrm.28245](https://doi.org/10.1002/mrm.28245).
- [27] N. Yokoya, C. Grohnfeldt, and J. Chanussot, "Hyperspectral and multi-spectral data fusion: A comparative review of the recent literature," *IEEE Geosci. Remote Sens. Mag.*, vol. 5, no. 2, pp. 29–56, Jun. 2017.
- [28] M. J. Ehrhardt, P. Markiewicz, and C.-B. Schönlieb, "Faster PET reconstruction with non-smooth priors by randomization and preconditioning," *Phys. Med. Biol.*, vol. 64, no. 22, Nov. 2019, Art. no. 225019.

- [29] W. M. Wells, III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multi-modal volume registration by maximization of mutual information," *Med. Image Anal.*, vol. 1, no. 1, pp. 35–51, 1996.
- [30] J. Čížek, K. Herholz, S. Vollmar, R. Schrader, J. Klein, and W.-D. Heiss, "Fast and robust registration of PET and MR images of human brain," *NeuroImage*, vol. 22, no. 1, pp. 434–442, May 2004.
- [31] M. P. Heinrich, M. Jenkinson, M. Brady, and J. A. Schnabel, "Robust super-resolution reconstruction with multi-modal registration and guidance," *Proc. Med. Image Understand. Anal.*, pp. 81–86, 2012.
- [32] Q. Zhu, Y. Ren, Z. Qiu, and W. Wang, "Robust MR image super-resolution reconstruction with cross-modal edge-preserving regularization," *Int. J. Imag. Syst. Technol.*, vol. 29, no. 4, pp. 491–500, Dec. 2019.
- [33] L. Bungert, M. J. Ehrhardt, and R. Reisenhofer, "Robust blind image fusion for misaligned hyperspectral imaging data," *PAMM*, vol. 18, no. 1, Dec. 2018, Art. no. e201800033.
- [34] Y. Fu, Y. Zheng, L. Zhang, Y. Zheng, and H. Huang, "Simultaneous hyper-spectral image super-resolution and geometric alignment with a hybrid camera system," *Neurocomputing*, vol. 384, pp. 282–294, Apr. 2020.
- [35] J. Bolte, S. Sabach, and M. Teboulle, "Proximal alternating linearized minimization for nonconvex and nonsmooth problems," *Math. Program.*, vol. 146, nos. 1–2, pp. 459–494, Aug. 2014.
- [36] C. Chen, Y. Li, W. Liu, and J. Huang, "SIRF: Simultaneous satellite image registration and fusion in a unified framework," *IEEE Trans. Image Process.*, vol. 24, no. 11, pp. 4213–4224, Nov. 2015.
- [37] C. Chen and O. Öktem, "Indirect image registration with large diffeomorphic deformations," *SIAM J. Imag. Sci.*, vol. 11, no. 1, pp. 575–617, Jan. 2018.
- [38] L. F. Lang, S. Neumayer, O. Öktem, and C.-B. Schönlieb, "Template-based image reconstruction from sparse tomographic data," *Appl. Math. Optim.*, vol. 82, no. 3, pp. 1081–1109, 2020.
- [39] M. Burger, H. Dirks, and C.-B. Schönlieb, "A variational model for joint motion estimation and image reconstruction," *SIAM J. Imag. Sci.*, vol. 11, no. 1, pp. 94–128, Jan. 2018.
- [40] V. Corona, A. I. Aviles-Rivero, N. Debrox, C. Le Guyader, and C.-B. Schönlieb, "Variational multi-task MRI reconstruction: Joint reconstruction, registration and super-resolution," 2019, *arXiv:1908.05911*. [Online]. Available: <http://arxiv.org/abs/1908.05911>
- [41] J. Yang, X. Fu, Y. Hu, Y. Huang, X. Ding, and J. Paisley, "PanNet: A deep network architecture for pan-sharpening," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 5449–5457.
- [42] Q. Yuan, Y. Wei, X. Meng, H. Shen, and L. Zhang, "A multiscale and multidepth convolutional neural network for remote sensing imagery pan-sharpening," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 11, no. 3, pp. 978–989, Mar. 2018.
- [43] J. Modersitzki, *FAIR: Flexible Algorithms for Image Registration*, vol. 6. Philadelphia, PA, USA: SIAM, 2009.
- [44] M. J. Ehrhardt and S. R. Arridge, "Vector-valued image processing by parallel level sets," *IEEE Trans. Image Process.*, vol. 23, no. 1, pp. 9–18, Jan. 2014.
- [45] M. J. Ehrhardt, K. Thielemans, L. Pizarro, D. Atkinson, S. Ourselin, B. F. Hutton, and S. R. Arridge, "Joint reconstruction of PET-MRI by exploiting structural similarity," *Inverse Problems*, vol. 31, no. 1, Jan. 2015, Art. no. 015001.
- [46] M. Hintermüller, M. Holler, and K. Papafitsoros, "A function space framework for structural total variation regularization with applications in inverse problems," *Inverse Problems*, vol. 34, no. 6, Jun. 2018, Art. no. 064002.
- [47] F. P. M. Oliveira and J. M. R. S. Tavares, "Medical image registration: A review," *Comput. Methods Biomech. Biomed. Eng.*, vol. 17, no. 2, pp. 73–93, 2014.
- [48] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Cham, Switzerland: Springer, 2011.
- [49] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 123–231, 2013.
- [50] A. Chambolle and T. Pock, "An introduction to continuous optimization for imaging," *Acta Numerica*, vol. 25, pp. 161–319, May 2016.
- [51] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2419–2434, Nov. 2009.
- [52] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, and S. J. van der Walt, "SciPy 1.0: Fundamental algorithms for scientific computing in Python," *Nature Methods*, vol. 17, no. 3, pp. 261–272, 2020.
- [53] M. Burger and S. Osher, "A guide to the TV zoo," in *Level Set and PDE Based Reconstruction Methods in Imaging*. Cham, Switzerland: Springer, 2013, pp. 1–70.
- [54] L. Bungert and M. Burger, "Solution paths of variational regularization methods for inverse problems," *Inverse Problems*, vol. 35, no. 10, 2019, Art. no. 105012.
- [55] J. Adler, H. Kohr, and O. Öktem. (2017). *Operator Discretization Library (ODL)*. [Online]. Available: <https://github.com/odlgroup/odl>
- [56] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu, "Scikit-image: Image processing in Python," *PeerJ*, vol. 2, p. e453, Jun. 2014.
- [57] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [58] M. Seitzer, G. Yang, J. Schlemper, O. Oktay, T. Würfl, V. Christlein, T. Wong, R. Mohiaddin, D. Firmin, J. Keegan, D. Rueckert, and A. Maier, "Adversarial and perceptual refinement for compressed sensing MRI reconstruction," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2018, pp. 232–240.
- [59] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Mutual-information-based registration of medical images: A survey," *IEEE Trans. Med. Imag.*, vol. 22, no. 8, pp. 986–1004, Aug. 2003.
- [60] J. Bland, A. Mehranian, M. A. Belzunce, S. Ellis, C. Costa-Luis, C. J. McGinnity, A. Hammers, and A. J. Reader, "Intercomparison of MR-infused PET image reconstruction methods," *Med. Phys.*, vol. 46, no. 11, pp. 5055–5074, Nov. 2019.



LEON BUNGERT received the master's degree in mathematics and the Ph.D. degree (*summa cum laude*) from the University of Erlangen, Germany, in 2017 and 2020, respectively. His research interests include nonlinear eigenvalue problems, partial differential equations (on graphs), inverse problems, and image reconstruction.



MATTHIAS J. EHRHARDT received the Diploma degree (Hons.) in industrial mathematics from the University of Bremen, Germany, in 2011, and the Ph.D. degree in medical imaging from University College London, U.K., in 2015. He held a post-doctoral position with the Cambridge Image Analysis group, Department for Applied Mathematics and Theoretical Physics, University of Cambridge, U.K., from 2016 to 2018. He is currently a Prize Fellow at the Institute for Mathematical Innovation, University of Bath, U.K. His research interests include optimization, inverse problems, computational imaging, and machine learning.

• • •