



Published in final edited form as:

SIAM J Imaging Sci. 2015 ; 8(2): 1030–1069. doi:10.1137/140984002.

An Inexact Newton–Krylov Algorithm for Constrained Diffeomorphic Image Registration*

Andreas Mang[†] and George Biros[†]

Andreas Mang: andreas@ices.utexas.edu; George Biros: gbiros@acm.org

[†]Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, TX 78712-0027

Abstract

We propose numerical algorithms for solving large deformation diffeomorphic image registration problems. We formulate the nonrigid image registration problem as a problem of optimal control. This leads to an infinite-dimensional partial differential equation (PDE) constrained optimization problem. The PDE constraint consists, in its simplest form, of a hyperbolic transport equation for the evolution of the image intensity. The control variable is the velocity field. Tikhonov regularization on the control ensures well-posedness. We consider standard smoothness regularization based on H^1 - or H^2 -seminorms. We augment this regularization scheme with a constraint on the divergence of the velocity field (control variable) rendering the deformation incompressible (Stokes regularization scheme) and thus ensuring that the determinant of the deformation gradient is equal to one, up to the numerical error. We use a Fourier pseudospectral discretization in space and a Chebyshev pseudospectral discretization in time. The latter allows us to reduce the number of unknowns and enables the time-adaptive inversion for nonstationary velocity fields. We use a preconditioned, globalized, matrix-free, inexact Newton–Krylov method for numerical optimization. A parameter continuation is designed to estimate an optimal regularization parameter. Regularity is ensured by controlling the geometric properties of the deformation field. Overall, we arrive at a black-box solver that exploits computational tools that are precisely tailored for solving the optimality system. We study spectral properties of the Hessian, grid convergence, numerical accuracy, computational efficiency, and deformation regularity of our scheme. We compare the designed Newton–Krylov methods with a globalized Picard method (preconditioned gradient descent). We study the influence of a varying number of unknowns in time. The reported results demonstrate excellent numerical accuracy, guaranteed local deformation regularity, and computational efficiency with an optional control on local mass conservation. The Newton–Krylov methods clearly outperform the Picard method if high accuracy of the inversion is required. Our method provides equally good results for stationary and nonstationary velocity fields for two-image registration problems.

*This work was supported by AFOSR grants FA9550-12-10484 and FA9550-11-10339; by NSF grants CCF-1337393 OCI-1029022; by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under awards DE-SC0010518, DE-SC0009286, and DE-FG02-08ER2585; and by NIH grant 10042242.

AMS subject classifications. 68U10, 49J20, 35Q93, 65K10, 76D55, 90C20

Keywords

large deformation diffeomorphic image registration; optimal control; PDE constrained optimization; Stokes regularization; Newton–Krylov method; pseudospectral Galerkin method; Stokes solver

1. Introduction and motivation

Image registration has become a key area of research in computer vision and (medical) image analysis [54, 65]. The task is to establish spatial correspondence between two images, $m_R: \bar{\Omega} \rightarrow \mathbf{R}, \mathbf{x} \mapsto m_R(\mathbf{x})$, and $m_T: \bar{\Omega} \rightarrow \mathbf{R}, \mathbf{x} \mapsto m_T(\mathbf{x})$, with compact support on some domain $\Omega := (-\pi, \pi)^d \subset \mathbf{R}^d$ via a mapping $\mathbf{y}: \mathbf{R}^d \rightarrow \mathbf{R}^d, \mathbf{x} \mapsto \mathbf{y}(\mathbf{x})$, such that $m_T \circ \mathbf{y} \approx m_R$ [22, 27], where \circ is the function composition. Here, m_T is referred to as the *template image* (the image to be registered), m_R is referred to as the *reference image* (the fixed image), and $d \in \{2, 3\}$ is the data dimensionality. We limit ourselves to non-rigid image registration. The search for \mathbf{y} is typically formulated as a variational optimization problem [27, 54],

$$\min_{\mathbf{y} \in \mathcal{Y}} \left\{ \mathcal{J}[\mathbf{y}] := \frac{1}{2} \|m_R - m_T \circ \mathbf{y}\|_{L^2(\Omega)}^2 + \frac{\beta}{2} \mathcal{S}[\mathbf{y}] \right\}. \quad (1.1)$$

The proximity between m_R and $m_T \circ \mathbf{y}$ is measured on the basis of an L^2 -distance (other measures can be considered [54, 65]). The functional \mathcal{S} in (1.1) is a regularization model that is introduced to overcome ill-posedness. The regularization parameter $\beta > 0$ weights the contribution of \mathcal{S} . Various regularization models \mathcal{S} have been considered (see [13, 14, 19, 22, 25, 26, 28, 42, 63] for examples).

A key requirement in many image registration problems is that the mapping \mathbf{y} is a *diffeomorphism* [4, 14, 23, 68, 69, 70]. This translates into the necessary condition $\det(\nabla \mathbf{y}) > 0$, where $\nabla \mathbf{y} \in \mathbf{R}^{d \times d}$ is the *deformation gradient* (also referred to as the *Jacobian matrix*). An intuitive approach is to explicitly safeguard against nondiffeomorphic mappings \mathbf{y} by adding a constraint to (1.1) [39, 48, 61]. Another strategy is to perform the optimization on the manifold of diffeomorphic mappings [2, 3, 4, 44, 52, 53, 69, 70]. The latter models, in general, do not control geometric properties of the deformation field and may result in fields that are close to being nondiffeomorphic. Further, for certain image registration problems, restricting the search space to the manifold of diffeomorphisms does not necessarily guarantee that \mathbf{y} is *physically meaningful*. Some applications may benefit from extending these types of models by introducing additional constraints. One example for such a constraint is incompressibility (i.e., enforcing $\det(\nabla \mathbf{y}) = 1$; see also [10, 16, 52, 62]). Incompressibility is a requirement that might be of interest in medical image computing applications. If required, we can modify the incompressibility constraint to control the deviation of $\det(\nabla \mathbf{y})$ from identity. This will be the topic of a follow-up paper, in which we will extend our formulation. Here, we focus on the algorithmic issues of incorporating the incompressibility constraint. Furthermore, we remark that our optimal control formulation can naturally be extended to account for more complex constraints on the velocity field, for

example, constraints related to biophysical models (examples of such models can be found in [31, 40, 45, 51, 66]).

In what follows, we outline our method (see section 1.1), summarize the key contributions (see section 1.2), list the limitations of our method (see section 1.3), and relate our work to existing approaches (see section 1.4).

1.1. Outline of the method

We assume the images are compactly supported functions, defined on the open set $\Omega := (-\pi, \pi)^d \subset \mathbf{R}^d$ with boundary $\partial\Omega$ and closure $\bar{\Omega} := \Omega \cup \partial\Omega$. The deformation is modeled in an Eulerian frame of reference. We introduce a pseudotime variable $t > 0$ and solve on a unit time horizon for a velocity field $\mathbf{v} : \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}^d$, $(\mathbf{x}, t) \mapsto \mathbf{v}(\mathbf{x}, t)$, as follows:

$$\min_{\mathbf{v} \in \mathcal{V}} \left\{ \mathcal{J}[\mathbf{v}] := \frac{1}{2} \|m_R - m_1\|_{L^2(\Omega)}^2 + \int_0^1 \mathcal{S}[\mathbf{v}] dt \quad \text{subject to} \quad \mathcal{C}[m, \mathbf{v}] = 0 \right\}, \quad (1.2a)$$

where

$$\mathcal{C}[m, \mathbf{v}] := \begin{cases} \partial_t m + \nabla m \cdot \mathbf{v} & \text{in } \Omega \times (0, 1], \\ m - m_T & \text{in } \Omega \times \{0\}, \\ \gamma(\nabla \cdot \mathbf{v}) & \text{in } \Omega \times [0, 1]. \end{cases} \quad (1.2b)$$

with periodic boundary conditions on $\partial\Omega$. The parameter $\gamma \in \{0, 1\}$ in (1.2b) is introduced to enable or disable this constraint on the control \mathbf{v} . In PDE constrained optimization theory, m is referred to as the *state variable* and \mathbf{v} as the *control variable*. The first equation in (1.2b), in combination with its initial condition (second equation), models the flow of m_T subject to \mathbf{v} , where $m : \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}$, $(\mathbf{x}, t) \mapsto m(\mathbf{x}, t)$, represents the transported intensities of m_T . Accordingly, the *final state* $m_1 := m(\cdot, 1)$, $m_1 : \bar{\Omega} \rightarrow \mathbf{R}$, $\mathbf{x} \mapsto m_1(\mathbf{x})$, corresponds to $m_T \circ \mathbf{y}$ in (1.1). We measure the proximity between the deformed template image m_1 and the reference image m_R in terms of an L^2 -distance. Once we have found \mathbf{v} , we can compute \mathbf{y} from \mathbf{v} as a postprocessing step (this is also true for the deformation gradient $\nabla \mathbf{y}$; see Appendix D for details). The third equation in (1.2b) is a control on the divergence of \mathbf{v} and guarantees that the flow is incompressible (Stokes flow), i.e., the volume is conserved. This is equivalent to enforcing $\det(\nabla \mathbf{y}) = 1$ (see [33, p. 77ff.]).

We use either an H^1 - or an H^2 -seminorm for the smoothness regularization \mathcal{S} (resulting in Laplacian or biharmonic vector operators, respectively; see section 3.1). We report results for standard H^1 - and H^2 -regularization (neglecting $\nabla \cdot \mathbf{v} = 0$, i.e., $\gamma = 0$ in (1.2b)) and for a Stokes regularization scheme (incompressible flow; enforcing $\nabla \cdot \mathbf{v} = 0$, i.e., $\gamma = 1$ in (1.2b)).

In section 4 we will see that the optimality condition for (1.2) is a system of space-time nonlinear multicomponent PDEs for the transported image m , the velocity \mathbf{v} , and the adjoint variables for the transport and the divergence condition. Efficiently solving this system is a

significant challenge. For example, when we include the incompressibility constraint, the equation for the velocity ends up being a linear Stokes equation.

We solve for the first-order optimality conditions using a globalized, matrix-free, preconditioned Newton–Krylov method for the Schur complement of the velocity \mathbf{v} (a linearized Stokes problem driven by the image mismatch). We first derive the optimality conditions and then discretize using a pseudospectral discretization in space with a Fourier basis. We use a second-order Runge–Kutta method in time. The preconditioner for the Newton–Krylov schemes is based on the exact spectral inverse of the second variation of \mathcal{L} .

1.2. Contributions

The fundamental contributions are as follows:

- We design a numerical scheme for (1.2) with the following key features:
 - We use an adjoint-based Newton-type solver.
 - We provide a fast Stokes solver.
 - We introduce a spectral Galerkin method in time.
 - We design a parameter continuation method for automatically selecting the regularization parameter β .
 - Our framework guarantees deformation regularity.¹
- We provide a numerical study for the designed framework. We compare a globalized Picard method (preconditioned gradient descent) to an inexact Newton–Krylov and a Gauss–Newton–Krylov method. We report results for synthetic and real-world problems. We study spectral properties, grid convergence, and numerical accuracy of the proposed scheme. We study the effects of compressible (plain H^1 - and H^2 -regularization) and incompressible (Stokes regularization) deformation models. We report results for a varying number of unknowns in time (i.e., inverting for stationary and nonstationary velocity fields).

The numerical discretization (pseudospectral) allows for an efficient solution of the Stokes-like equations by eliminating the pressure (i.e., the adjoint variable for the incompressibility constraint $\nabla \cdot \mathbf{v} = 0$ in (1.2b)).

The inf-sup condition for pressure spaces [12, p. 200ff.] is not an issue with our scheme.² In fact, for smooth images, our scheme is spectrally accurate in space and second-order accurate in time. We will see that we can numerically enforce incompressibility up to almost

¹Note that controlling the magnitude of $\det(\nabla \mathbf{y})$ is not sufficient to guarantee that \mathbf{y} is locally well behaved. Therefore, our framework features geometric constraints that guarantee a nice, locally diffeomorphic mapping \mathbf{y} . In particular, we control the shear angle of the cells within the deformed grid during the parameter continuation in β .

²The inf-sup condition is an important requirement when solving Stokes-like problems via the finite element method (see [12, p. 200ff.]; examples for a finite element discretization of similar problem formulations can be found in [16, 62]). Satisfying the inf-sup condition ensures that the finite element solution exists and is stable and optimal. Essentially, we require two different finite element spaces for the discretization of the pressure and the velocity. The inf-sup condition is key for the decision on an adequate pair of spaces. For our scheme, we can use the same basis (Fourier) for the discretization of the pressure and the velocity. Also, it is very efficient since we can eliminate the pressure from the optimality system (see section 4.3.4).

machine precision. Also, our scheme allows for efficient preconditioning of the Hessian: at the cost of a diagonal scaling we obtain a problem with a bounded condition number.

Overall, we demonstrate that the designed framework(i) is efficient and accurate, (ii) features a precise control of the deformation regularity, and (iii) does not require manual tuning of parameters.

1.3. Limitations

The main limitations of our method are as follows:

- The considered model assumes a constant illumination of m_R and m_T (a consequence of the transport equation and the L^2 -distance in (1.2)). Therefore, it is (in its current state) not directly applicable to multimodal registration problems. Nevertheless, let us remark that the L^2 -distance is commonly used in practice [4, 16, 28, 41, 44, 48, 57, 71].
- The efficient use of a Fourier discretization for the PDEs requires periodic boundary conditions. If the images are not periodic, we artificially introduce periodic boundary conditions by mollification and zero padding.

1.4. Related work

Due to the vast body of literature, it is not possible to provide a comprehensive review of numerical methods for nonrigid image registration. Background on image registration formulations and numerics can be found in [27, 54, 65]. We limit the discussion to approaches that(i) model the deformation via a velocity field \mathbf{v} , (ii) view image registration as a problem of optimal control, and/or (iii) constrain \mathbf{v} to be divergence free (i.e., introduce a mass conservation equation as an additional constraint).

Fluid mechanical models have been introduced [18, 19] to overcome limitations of small deformation models [13, 25, 26]. The work in [18, 19] has been extended in [4, 23, 53, 68] using concepts from differential geometry. This class of approaches is referred to as *large deformation diffeomorphic metric mapping* (LDDMM). Under the assumption that \mathbf{v} is adequately smooth, it is guaranteed that \mathbf{y} is a diffeomorphism [23, 68]. The associated smoothness requirements are enforced by the regularization model \mathcal{S} (typically an H^p -norm) [3, 4]. The optimization is performed on the space of nonstationary velocity fields [4]. To reduce the number of unknowns, it has been suggested to perform the optimization either on the space of stationary velocity fields [1, 2, 44] or with respect to an initial momentum that entirely defines the flow of the map \mathbf{y} [3, 71, 74].

The idea of parameterizing a diffeomorphism \mathbf{y} via a stationary velocity field [1] has also been introduced to the demons registration framework [52, 69, 70]. Here, optimization is performed in a sequential fashion, alternating between updates resulting from the distance measure (forcing term) and the application of a smoothing operator to regularize the problem (typically through Gaussian smoothing [67, 69, 70]). This scheme is somewhat equivalent to the Picard scheme we discuss in our paper, but it is unclear how one couples it with line search or trust region techniques.

Approaches that more closely reflect an optimal control PDE constrained optimization formulation (1.2) are described in [10, 16, 41, 48, 49, 57, 62, 71]. The model in [16] is equivalent to (1.2). The model in [62] follows the traditional optical flow formulation [46]. The conceptual difference between our formulation and the latter is that in optical flow, the transport equation constraint appears in the objective (see, e.g., [46, 47, 62]) and is therefore only fulfilled approximately in an L^2 least squares sense. We treat it as a hard constraint instead. In [57] an optimal mass transport formulation is described, which is based on the Monge-Kantorovich problem. The formulations in [41, 71] do not account for incompressibility. The optical flow approach in [10] treats incompressibility as an L^2 -penalty. An optimal control formulation for a constant-in-time velocity field was proposed in [48, 49], in which the divergence of the velocity field is penalized along with smoothness constraints.

What sets our work apart are the numerical algorithms and the discretization scheme. Almost all existing efforts on large deformation diffeomorphic image registration that are closely related to our optimal control formulation exclusively use first-order information for numerical optimization [2, 4, 10, 16, 18, 19, 41, 44, 48, 49, 57, 62, 71, 74]. We use second-order information. The only work³ in the context of large deformation diffeomorphic image registration that to our knowledge uses second-order information is [3]. The model in [3] is based on the LDDMM framework [4, 23, 48, 49, 53, 44, 68]. The inversion is, likewise to [71], performed with respect to an initial momentum. No additional constraints on \mathbf{v} are considered. Another difference is that we use a Galerkin method in time to reduce the number of unknowns. This allows us to invert for stationary [1, 2, 44, 52, 69, 70] as well as time-varying [4, 10, 16, 41] velocity fields. Nothing changes in our formulation other than the number of unknowns. Furthermore, we globalize our methods with a line search strategy (i.e., we guarantee a sufficient decrease of the objective \mathcal{J}). This is a standard—yet important—ingredient for guaranteeing convergence, which is often not accounted for [2, 16, 17, 41, 52, 69, 70].

We, likewise to [10, 16, 48, 49, 52, 62], consider incompressibility as an optional constraint (see (1.2b)). Operating with divergence-free velocity fields is equivalent to enforcing $\det(\nabla \mathbf{y}) = 1$ up to numerical accuracy (see [33, p. 77ff.]; other formulations for controlling $\det(\nabla \mathbf{y})$ can be found in [14, 36, 37, 39, 50, 55, 58, 60, 61, 64, 73]). Unlike [10, 48, 49], which penalize the divergence of the velocity, we treat it exactly. We are not arguing that this approach is better per se. The use of penalties is adequate unless one has reasons to insist on an incompressible velocity field. In that case, a penalty method results in ill-conditioning.

Finally our pseudospectral formulation in space allows us to resolve several numerical difficulties related to the incompressibility constraint. For example, the inf-sup condition for pressure spaces is not an issue with our scheme. Regarding accuracy, for smooth images, our scheme is spectrally accurate in space and second-order accurate in time. We do not have to use different discretization models [16, 62] for solving the individual subsystems of the mixed-type (hyperbolic-elliptic) optimality conditions. Since we use second-order explicit time stepping in combination with Fourier spectral methods, we have at hand a scheme that

³After the submission of our work, another contribution on second-order numerical optimization for LD-DMM appeared [43].

displays minimal numerical diffusion and does not require flux-limiters [10, 16, 41, 71]. As we will see, the conditioning of the Hessian that appears in our Newton–Krylov scheme can be quite bad. Although the literature for preconditioners for PDE constrained optimization problem is quite rich (e.g., [6, 10, 8, 34]), none of these methods directly applies to our formulation. Developing effective preconditioning schemes for our formulation is ongoing work in our group.

2. Outline

In section 3 the mathematical model is developed. The numerical strategies are described in section 4. In particular, we specify (i) the optimality conditions (see section 4.1), (ii) strategies for numerical optimization (see section 4.2), and (iii) implementation details (see section 4.3). Numerical experiments on synthetic and real-world data are reported in section 5. Final remarks can be found in section 6.

3. Continuous problem formulation

We provide a summary of the basic notation in Table 1. The original problem formulation is stated in section 1.1. The only missing building block is the considered choices for \mathcal{S} in (1.2a). This is what we discuss next. Note that we neglect any technicalities with respect to the associated function spaces; we assume that the considered functions are adequately smooth (i.e., sufficiently many derivatives exist and are bounded).

3.1. Regularization models

In contrast to [10] we do not explicitly enforce continuity in time. We relax the model to an L^2 -integrability instead (see (1.2a)). This relaxation still yields a velocity field that varies smoothly in time [16].

Quadratic smoothing regularization models are commonly used in nonrigid image registration [27, 54, 56, 65]. They can be defined as

$$\mathcal{S}[\mathbf{v}] := \frac{\beta}{2} \|\mathbf{v}\|_{\mathcal{W}}^2 = \frac{\beta}{2} \langle \mathcal{B}[\mathbf{v}], \mathcal{B}[\mathbf{v}] \rangle_{L^2(\Omega)}, \quad (3.1)$$

where \mathcal{B} is a differential operator that (together with its dual) defines the function space \mathcal{W} and $\beta > 0$ is a regularization parameter that balances the contribution of \mathcal{S} .

As images are functions of bounded variation, regularity requirements on $\mathbf{v} \in \mathcal{V}$, $\mathcal{V} := L^2([0, 1]; \mathcal{W})$ (i.e., the choice of \mathcal{W} in (3.1)) have to be considered with care (for an analytical result see [16]). Experimental analysis suggests that an H^1 -seminorm is appropriate if incompressibility is considered (i.e., $\gamma = 1$ in (1.2b); see also [16]). Thus,

$$\mathcal{B}[\mathbf{v}] = \nabla \mathbf{v}.$$

This choice is motivated from continuum mechanics and yields a viscous model of *linear Stokes flow* (see section 4.1; Stokes regularization). If we neglect the incompressibility constraint (i.e., $\gamma = 0$ in (1.2b)), we use a vectorial Laplacian operator,

$$\mathcal{B}[\mathbf{v}] = \Delta \mathbf{v},$$

instead. This choice is motivated by the fact that H^2 -norm-based quadratic regularization is commonly used in large deformation diffeomorphic image registration [4, 41, 44].

4. Numerics

We describe the numerical methods used to solve (1.2) next. Whenever discretized quantities are considered, a superscript h is added to the continuous variables and operators (i.e., the discretized representation of \mathbf{v} is denoted by \mathbf{v}^h). Likewise, if we refer to a discrete variable at a particular iteration, we will add the iteration index as a subscript (i.e., \mathbf{v}^h at iteration k is denoted by \mathbf{v}_k^h).

We discretize the data on a *nodal grid* in space and time. The number of spatial grid points is denoted by $\mathbf{n}_x := (n_x^1, \dots, n_x^d)^\top \in \mathbf{N}^d$ with spatial step size $\mathbf{h}_x := (h_x^1, \dots, h_x^d) \in \mathbf{R}_{>0}^d$. The number of time points is denoted by $n_t \in \mathbf{N}$ with step size $h_t = 1/n_t$, $h_t > 0$.

We use the *method of Lagrange multipliers* to numerically solve (1.2) with Lagrange multipliers $\lambda : \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}$, $(\mathbf{x}, t) \mapsto \lambda(\mathbf{x}, t)$ (for the hyperbolic transport equation in (1.2b)), and $p : \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}$, $(\mathbf{x}, t) \mapsto p(\mathbf{x}, t)$ (pressure; for the incompressibility constraint in (1.2b)). We use an *optimize-then-discretize* approach (for a discussion on advantages and disadvantages see [32, p. 57ff.]). The resulting optimality conditions are described next.

4.1. Optimality conditions

Computing variations of the Lagrangian with respect to perturbations of the state (m), adjoint (λ and p), and control (\mathbf{v}) variables, respectively, yields the (necessary) first-order optimality (KKT) conditions (in strong form)

$$\partial_t m + \nabla m \cdot \mathbf{v} = 0 \quad \text{in } \Omega \times (0, 1], \quad (4.1a)$$

$$-\partial_t \lambda - \nabla \cdot (\mathbf{v} \lambda) = 0 \quad \text{in } \Omega \times [0, 1], \quad (4.1b)$$

$$\gamma(\nabla \cdot \mathbf{v}) = 0 \quad \text{in } \Omega \times [0, 1]. \quad (4.1c)$$

$$\mathbf{g} := \beta \mathcal{A}[\mathbf{v}] + \gamma \nabla p + \mathbf{f} = 0 \quad \text{in } \Omega \times [0, 1], \quad (4.1d)$$

subject to the initial and terminal conditions

$$m = m_r \text{ in } \Omega \times \{0\} \quad \text{and} \quad \lambda = -(m_l - m_r) \text{ in } \Omega \times \{1\}$$

and periodic boundary conditions on $\partial\Omega$; (4.1d) is referred to as the *reduced gradient*, where $\mathbf{f} := \lambda \nabla m, \mathbf{f}: \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}^d, (\mathbf{x}, t) \mapsto \mathbf{f}(\mathbf{x}, t)$, is the applied *body force* and $\mathcal{A} = \mathcal{B}\mathcal{B}^H$ is the Gâteaux derivative of \mathcal{L} . In particular, we have

$$\mathcal{A}[\mathbf{v}] = -\Delta \mathbf{v} \quad (H^1\text{-regularization}; \gamma=1 \text{ in (4.1)}), \quad (4.2a)$$

$$\mathcal{A}[\mathbf{v}] = \Delta^2 \mathbf{v} \quad (H^2\text{-regularization}; \gamma=0 \text{ in (4.1)}), \quad (4.2b)$$

respectively. We refer to (4.1a) (hyperbolic initial value problem) as the *state equation*, to (4.1b) as the *adjoint equation* (hyperbolic final value problem) and to (4.1d) as the *control equation* (elliptic problem). Note that the adjoint equation (4.1b) is, likewise to (4.1a), a scalar conservation law that flows the mismatch between m_R and m_l backward in time. If we neglect the incompressibility constraint in (1.2b), γ in (4.1) is set to zero (i.e., (4.1) consists only of (4.1a), (4.1b), and (4.1d)).

Taking second variations of the Lagrangian yields the system

$$\partial_t \tilde{m} + \nabla \tilde{m} \cdot \mathbf{v} + \nabla m \cdot \tilde{\mathbf{v}} = 0 \quad \text{in } \Omega \times (0, 1], \quad (4.3a)$$

$$-\partial_t \tilde{\lambda} - \nabla \cdot (\mathbf{v} \tilde{\lambda}) - \nabla \cdot (\tilde{\mathbf{v}} \lambda) = 0 \quad \text{in } \Omega \times [0, 1], \quad (4.3b)$$

$$\gamma (\nabla \cdot \tilde{\mathbf{v}}) = 0 \quad \text{in } \Omega \times [0, 1], \quad (4.3c)$$

$$\mathcal{H}[\tilde{\mathbf{v}}] := \beta \mathcal{A}[\tilde{\mathbf{v}}] + \gamma \nabla \tilde{p} + \tilde{\mathbf{f}} = -\mathbf{g} \quad \text{in } \Omega \times [0, 1], \quad (4.3d)$$

subject to initial and terminal conditions $\tilde{m}_0 := \tilde{m}(\cdot, 0) = 0, \tilde{m}_0: \bar{\Omega} \rightarrow \mathbf{R}, \mathbf{x} \mapsto \tilde{m}_0(\mathbf{x})$, and $\tilde{\lambda}_1 := \tilde{\lambda}(\cdot, 1) = -\tilde{m}_1, \tilde{\lambda}_1: \bar{\Omega} \rightarrow \mathbf{R}, \mathbf{x} \mapsto \tilde{\lambda}_1(\mathbf{x}), \tilde{m}_1 := \tilde{m}(\cdot, 1), \tilde{m}_1: \bar{\Omega} \rightarrow \mathbf{R}, \mathbf{x} \mapsto \tilde{m}_1(\mathbf{x})$,

respectively, and periodic boundary conditions on $\partial\Omega$. Here, (4.3a), (4.3b), and (4.3d) are referred to as the *incremental* state, adjoint, and control equations, respectively; the incremental variables are denoted with a tilde. Further, \mathcal{H} in (4.3d) is referred to as the *reduced Hessian* and $\tilde{\mathbf{f}} := \tilde{\lambda} \nabla m + \lambda \nabla \tilde{m}, \tilde{\mathbf{f}}: \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}, (\mathbf{x}, t) \mapsto \tilde{\mathbf{f}}(\mathbf{x}, t)$, is the *incremental body force*. The operator \mathcal{A} in (4.3d) represents the second variation of \mathcal{S} with respect to the control \mathbf{v} . We use the same symbol as in (4.1), since the second variation of \mathcal{S} with respect to \mathbf{v} is identical to its first variation (the corresponding vectorial differential operators are given in (4.2a) and (4.2b), respectively).

4.2. Numerical optimization

We discuss strategies for numerical optimization next. We consider second-order Newton–Krylov methods (see section 4.2.1) and a first-order Picard method (see section 4.2.3).

We use a backtracking line search subject to the *Armijo condition* with search direction $\mathbf{s}_k \in \mathbf{R}^n$ and step size $\alpha_k > 0$ at (outer) iteration $k \in \mathbf{N}_0$ to ensure a sequence of monotonically decreasing objective values \mathcal{J}^h (we use default parameters; see [59, Algorithm 3.1, p. 37]). Note that each evaluation of \mathcal{J}^h requires a forward solve (i.e., the solution of (4.1a) to obtain m_1^h given some trial solution $\mathbf{v}_k^h \in \mathbf{R}^n$). Therefore, it is desirable to keep the number of line search steps at minimum.

4.2.1. Inexact Newton–Krylov method—Applying Newton’s method to (4.1) yields a large KKT system that has to be solved numerically at each *outer iteration* k . We will refer to the iterative solution of this system as *inner iterations*.⁴ In *reduced space methods*, incremental adjoint and state variables are eliminated from the system via block elimination (under the assumption that state and adjoint equations are fulfilled exactly) [8, 9]. We obtain the reduced KKT system

$$\mathcal{H}_k^h \tilde{\mathbf{v}}_k^h = -\mathbf{g}_k^h, \quad k \in \mathbf{N}, \quad (4.4)$$

where $\mathcal{H}_k^h \in \mathbf{R}^{n \times n}$ corresponds to the reduced Hessian in (4.3d) (i.e., the Schur complement of the full Hessian for the control variable \mathbf{v}^h) and $\tilde{\mathbf{v}}_k^h \in \mathbf{R}^n$ to the incremental control variable in (4.3) (which is nothing but the search direction \mathbf{s}_k mentioned earlier). Further, the right-hand side $\mathbf{g}_k^h \in \mathbf{R}^n$ corresponds to the reduced gradient in (4.1d).

The numerical scheme amounts to a sequential solution of the optimality conditions (4.1) and (4.3). Algorithm 1 illustrates a realization of an outer iteration.⁵ Note that we eliminate (4.1c) and (4.3c) from the optimality conditions (see section 4.3.4). The inner iteration (i.e., the solution of (4.4)) is what we discuss next.

⁴As opposed to the steps for updating \mathbf{v}_k^h , which we refer to as *outer iterations*; see Algorithm 1.

⁵Note that the scheme in Algorithm 1 also applies to the Picard method (see section 4.2.3). The only difference is the way we compute the search direction \mathbf{s}_k in line 5.

Forming or storing \mathcal{H}^h in (4.4) is computationally prohibitive. Therefore, it is desirable to use an iterative solver for which \mathcal{H}^h does not have to be assembled in practice. Krylov-subspace methods are a popular choice [5, 8, 15, 34], as they only require matrix-vector products. We use a PCG method, exploiting the fact that \mathcal{H}^h is positive definite (i.e., $\mathcal{H}^h \succ 0$; see section 4.2.2 for a discussion) and symmetric.

Solving (4.4) exactly can be prohibitively expensive and might not be justified if an iterate is far from the (true) solution [20]. A common strategy is to perform *inexact solves*. That is, starting with a large tolerance for the Krylov-subspace method we successively reduce the tolerance and by that solve more accurately for the search direction, as we approach a (local) minimizer [21, 24]. This can be achieved with the termination criterion

$$\|\mathcal{H}_\iota^h \tilde{\mathbf{v}}_\iota^h + \mathbf{g}_k^h\|_2 =: \|\mathbf{r}_\iota\|_2 \leq \eta_k \|\mathbf{g}_k^h\|_2 \quad (4.5)$$

for the Krylov-subspace method. Here, $\eta_k := \min(0.5, \sqrt{\|\mathbf{g}_k^h\|_2 / \|\mathbf{g}_0^h\|_2})$, $\eta_k \in [0, 1)$, is referred to as a *forcing sequence* (assuming superlinear convergence; details can be found in e.g., [59, p. 165ff.]); $\iota \in \mathbf{N}$ in (4.5) is the iteration index of the inner iteration (i.e., for the iterative solution of (4.4)) at a given outer iteration k .

Algorithm 1

Outer iteration of the designed inexact Newton–Krylov method.

```

1:  $\mathbf{v}_0^h \leftarrow 0$  compute  $m_0^h, \lambda_0^h, \mathcal{J}^h(\mathbf{v}_0^h)$ , and  $\mathbf{g}_0^h$ ;  $k \leftarrow 0$ 
2: while true do
3:   stop  $\leftarrow$  (4.9)
4:   if stop break
5:    $\mathbf{s}_k \leftarrow$  solve (4.4) given  $m_k^h, \lambda_k^h, \mathbf{v}_k^h$ , and  $\mathbf{g}_k^h$  ▷ Newton step
6:    $\alpha_k \leftarrow$  perform line search on  $\mathbf{s}_k$ 
7:    $\mathbf{v}_{k+1}^h \leftarrow \mathbf{v}_k^h + \alpha_k \mathbf{s}_k$ ,
8:    $m_{k+1}^h(t=0) \leftarrow m_t^h$ 
9:    $m_{k+1}^h \leftarrow$  solve (4.1a) forward in time given  $\mathbf{v}_{k+1}^h$  ▷ forward solve
10:   $\lambda_{k+1}^h(t=1) \leftarrow (m_R^h - m_{k+1}^h(t=1))$ 
11:   $\lambda_{k+1}^h \leftarrow$  solve (4.1b) backward in time given  $\mathbf{v}_{k+1}^h$  and  $m_{k+1}^h$  ▷ adjoint solve
12:  compute  $\mathcal{J}^h(\mathbf{v}_{k+1}^h)$  and  $\mathbf{g}_{k+1}^h$  given  $m_{k+1}^h, \lambda_{k+1}^h$ , and  $\mathbf{v}_{k+1}^h$ 
13:   $k \leftarrow k + 1$ 
14: end while

```

The course of an inner iteration follows the standard PCG steps (see, e.g., [59, p. 119, Algorithm 5.3]). During each inner iteration ι we have to apply \mathcal{H}^h in (4.3d) to a vector. We summarize this matrix-vector product in Algorithm 2. As can be seen, each application of \mathcal{H}^h requires an additional forward and adjoint solve (i.e., the solution of the incremental state and adjoint equations (4.3a) and (4.3b), respectively). This is a direct consequence of the block elimination in reduced space methods.

The number of inner iterations essentially depends on the spectrum of the operator \mathcal{H}^h . Typically, \mathcal{H}^h displays poor conditioning. An optimal preconditioner $\mathbf{P} \in \mathbf{R}^{n \times n}$ renders the number of iterations independent of n and β . The design of such a preconditioner is an open area of research [6, 7, 8, 34]. Standard techniques like incomplete factorizations or algebraic multigrid are not applicable, as they require the assembling of \mathcal{H}^h in (4.4). Geometric, matrix-free preconditioners are a valid option. This is something we will investigate in the future. Here, we consider a left preconditioner based on the exact spectral inverse of the regularization part of \mathcal{H}^h . That is, $\mathbf{P} := \mathcal{A}^h$ (implementation details can be found in section 4.3.3). Note that the PCG method only requires the action of \mathbf{P}^{-1} on a vector (i.e., a matrix-free implementation is in place). Since we use a Fourier spectral method, the cost of our preconditioning amounts to a spectral diagonal scaling. We will refer to this algorithm as the Newton-PCG (N-PCG) method.

Algorithm 2

Hessian matrix-vector product of the designed inexact Newton–Krylov algorithm at outer iteration $k \in \mathbf{N}$. We illustrate the computational steps required for applying \mathcal{H}^h in (4.3d) to the PCG search direction at inner iteration index $\iota \in \mathbf{N}$.

1:	$\tilde{m}_\iota^h(t=0) \leftarrow 0$	
2:	$\tilde{m}_\iota^h \leftarrow \text{solve (4.3a) forward in time given } m_k^h, \mathbf{v}_k^h, \text{ and } \tilde{\mathbf{v}}_\iota^h$	▷ incremental forward solve
3:	$\tilde{\lambda}_\iota^h(t=1) \leftarrow -\tilde{m}_\iota^h(t=1)$	
4:	$\tilde{\lambda}_\iota^h \leftarrow \text{solve (4.3b) backward in time given } \lambda_k^h, \mathbf{v}_k^h, \text{ and } \tilde{\mathbf{v}}_\iota^h$	▷ incremental adjoint solve
5:	apply \mathcal{H}_ι^h in (4.3d) to the PCG search direction given $\lambda_k^h, \tilde{\lambda}_\iota^h, m_k^h, \text{ and } \tilde{m}_\iota^h$	

4.2.2. GN approximation—Even though \mathcal{H}^h is in the proximity of a (local) minimum by construction positive semidefinite (i.e., $\mathcal{H}^h \geq 0$) it can be indefinite or singular far away from the solution. Accordingly, the search direction is not guaranteed to be a descent direction. One remedy is to terminate the inner iteration whenever negative curvature occurs [21]. Another approach is to use a quasi-Newton approximation. We consider a GN approximation $\mathcal{H}_{\text{GN}}^h$ instead. Here, we drop certain expressions of \mathcal{H}^h , which in turn guarantees that $\mathcal{H}_{\text{GN}}^h \succ 0$. In particular, we drop all expressions in (4.3) in which λ appears. Accordingly, we obtain the (continuous) system

$$\partial_t \tilde{m} + \nabla \tilde{m} \cdot \mathbf{v} + \nabla m \cdot \tilde{\mathbf{v}} = 0 \quad \text{in } \Omega \times (0, 1], \quad (4.6a)$$

$$-\partial_t \tilde{\lambda} - \nabla \cdot (\mathbf{v} \tilde{\lambda}) = 0 \quad \text{in } \Omega \times [0, 1], \quad (4.6b)$$

$$\gamma(\nabla \cdot \tilde{\mathbf{v}}) = 0 \quad \text{in } \Omega \times [0, 1], \quad (4.6c)$$

$$\mathcal{H}_{\text{GN}}[\tilde{\mathbf{v}}] := \beta \mathcal{A}[\tilde{\mathbf{v}}] + \gamma \nabla p + \tilde{\lambda} \nabla m = -\mathbf{g} \quad \text{in } \Omega \times [0, 1]. \quad (4.6d)$$

We expect the rate of convergence to drop from quadratic to (super-)linear when turning to (4.6). However, if the L^2 -distance can be driven to zero, we recover fast local convergence close to the true solution \mathbf{v}^\star , even if the adjoint variable is neglected. This is due to the fact that (4.1b) models the flow of the mismatch backward in time, such that $\lambda \rightarrow 0$ for $\mathbf{v} \rightarrow \mathbf{v}^\star$. We refer to this method as the (*inexact*) GN-PCG method [8, 9]. We remark that all algorithmic details described in this note apply to both Newton–Krylov methods.

4.2.3. Picard method—We consider a globalized Picard iteration (fixed point iteration) in addition to the described Newton–Krylov methods. Based on (4.1d) we have

$$\mathbf{v}_{k+1}^h = -(\beta \mathcal{A}^h)^{-1} [\gamma \nabla^h p_k^h + \mathbf{f}_k^h]. \quad (4.7)$$

Since we use Fourier spectral methods, the inversion of \mathcal{A}^h in (4.7) comes at the cost of a diagonal scaling (implementation details can be found in section 4.3.3). Accordingly, this scheme does not require the (iterative) solution of a linear system. However, it potentially results in a larger number of outer iterations until convergence as we expect the optimization problem to be poorly conditioned.

We do not directly use the solution of (4.7) as a new iterate but compute a search direction \mathbf{s}_k instead. This in turn allows us to perform a line search on \mathbf{s}_k . That is, we subtract the new from the former iterate. This scheme can be viewed as a gradient descent in the function space induced by \mathcal{W} (i.e., a *preconditioned gradient descent scheme*; see Appendix C).

Note that \mathbf{s}_k is, in contrast to Newton methods, arbitrarily scaled. Therefore, we provide an augmented implementation that tries to estimate an optimal scaling during the course of optimization. Details can be found in section 4.3.5.

4.2.4. Termination criteria—The termination criteria are in accordance with [56] (see [29, p. 305 ff.] for a discussion) given by

$$\begin{aligned}
\text{(C1)} \quad \mathcal{J}^h(\mathbf{v}_{k-1}^h) - \mathcal{J}^h(\mathbf{v}_k^h) &< \tau_{\mathcal{J}} (1 + \mathcal{J}^h(\mathbf{v}_0^h)), \\
\text{(C2)} \quad \|\mathbf{v}_{k-1}^h - \mathbf{v}_k^h\|_{\infty} &< \sqrt{\tau_{\mathcal{J}}} (1 + \|\mathbf{v}_k^h\|_{\infty}), \\
\text{(C3)} \quad \|\mathbf{g}_k^h\|_{\infty} &< \sqrt[3]{\tau_{\mathcal{J}}} (1 + \mathcal{J}^h(\mathbf{v}_0^h)), \\
\text{(C4)} \quad \|\mathbf{g}_k^h\|_{\infty} &< 1\text{E}3 \varepsilon_{\text{mach}}, \\
\text{(C5)} \quad k &> n_{\text{opt}}.
\end{aligned} \tag{4.8}$$

Here, $\tau_{\mathcal{J}} > 0$ is a user defined tolerance, $\varepsilon_{\text{mach}} > 0$ is the machine precision, and $n_{\text{opt}} \in \mathbf{N}$ is the maximal number of outer iterations. The algorithm is terminated if

$$\{(\text{C1}) \wedge (\text{C2}) \wedge (\text{C3})\} \vee (\text{C4}) \vee (\text{C5}), \tag{4.9}$$

where \wedge denotes the *logical or* and \vee the *logical and* operator, respectively.

4.3. Algorithmic details

This section provides additional specifics on the implementation. In particular, we describe (i) the numerical discretization (see section 4.3.1), (ii) the parameterization in time (see section 4.3.2), (iii) the inversion of the operator \mathcal{A}^h (see section 4.3.3), and (iv) strategies for the parameter selection (see section 4.3.5).

4.3.1. Numerical discretization—We use a (regular) nodal grid for the discretization in space and time. The problem is defined on the space-time interval $\Omega \times [0, 1]$, where $\Omega := (-\pi, \pi)^d$. Accordingly, we obtain the time step size via $h_t = 1/n_t$. The cell size (pixel or voxel size) $\mathbf{h}_x := (h_x^1, \dots, h_x^d) \in \mathbf{R}_{\geq 0}^d$ for a spatial grid cell can be computed via $h_x^i = 2\pi/n_x^i$, $i = 1, \dots, d$, where n_x^i is the number of grid points along the i th spatial direction x^i .

The derivative operators are discretized via Fourier spectral methods [11]. The time integrator for the forward and adjoint solves is an explicit second-order Runge–Kutta method, which, in connection with Fourier spectral methods, displays minimal numerical diffusion.

Following standard numerical theory for hyperbolic equations, the step size $h_t > 0$ is bounded from above by $h_{t,\text{max}} := \varepsilon_{\text{CFL}} / \max(\|\mathbf{v}^h\|_{\infty} \oslash \mathbf{h}_x)$, $h_{t,\text{max}} > 0$ (Courant–Friedrich–Lewy (CFL) condition). Here, \oslash denotes a Hadamard division and $\varepsilon_{\text{CFL}} > 0$ is the CFL number. The theoretical bound for $h_{t,\text{max}}$ is attained for $\varepsilon_{\text{CFL}} = 1$. We use $\varepsilon_{\text{CFL}} = 0.2$ for all experiments. Since we use a spectral Galerkin method in time (see section 4.3.2), we can adaptively adjust n_t (and therefore h_t) for the forward and adjoint solves as required by the CFL condition.

4.3.2. Spectral Galerkin method—To reduce the number of unknowns, \mathbf{v} is expanded *in time* in terms of basis functions $b_l: [0, 1] \rightarrow \mathbf{R}$, $t \mapsto b_l(t)$, $l = 1, \dots, n_{\oslash}$.

$$\mathbf{v}(\mathbf{x}, t) = \sum_{l=1}^{n_c} b_l(t) \mathbf{v}_l(\mathbf{x}), \quad (4.10)$$

where $\mathbf{v}_l: \mathbf{R}^d \rightarrow \mathbf{R}^d, \mathbf{x} \mapsto \mathbf{v}_l(\mathbf{x})$, is a coefficient field. The coefficients \mathbf{v}_l are the new unknowns of our problem. This reduces the number of unknowns in time from n_t to n_c where $n_c \ll n_t$. Thus, we can invert for a stationary ($n_c = 1$) or a nonstationary velocity field as required. Nothing changes in our formulation—just the number of unknowns.

We use *Chebyshev polynomials* as basis functions b_l on account of their excellent approximation properties as well as their orthogonality (see Appendix A for details). The expansion (4.10) solely affects \mathcal{S} and the (incremental) control equation (i.e., (4.1d) and (4.3d)); \mathbf{v} is computed from the coefficient fields $\mathbf{v}_l, l = 1, \dots, n_c$, during the forward and adjoint solves according to (4.10).

4.3.3. Inversion: Regularization operators—The Picard iteration in (4.7) as well as the preconditioning of (4.4) require the inversion of the differential operator \mathcal{A}^h . Since we use Fourier spectral methods this inversion can be accomplished at the cost of a spectral diagonal scaling. However, \mathcal{A}^h has a nontrivial kernel (which only includes constant functions due to the periodic boundary conditions). We make \mathcal{A}^h invertible by setting the base frequency of the inverse of \mathcal{A}^h (including the scaling by β) to one. This ensures not only invertibility, but also that the constant part of the (incremental) body force \mathbf{f} (or $\tilde{\mathbf{f}}$, respectively) remains in the kernel of our regularization scheme. This in turn allows us to invert for constant velocity fields.

4.3.4. Elimination of ρ and $\boldsymbol{\rho}$ —In our numerical scheme, we eliminate ρ and by that (4.1c) from (4.1). Details on the derivation can be found in Appendix B. We obtain

$$\tilde{\mathbf{g}} := \beta \mathcal{A}[\mathbf{v}] + \mathcal{K}[\mathbf{f}] = 0, \quad \mathcal{K}[\mathbf{f}] := -\nabla(\Delta^{-1}(\nabla \cdot \mathbf{f})) + \mathbf{f}. \quad (4.11)$$

to replace (4.1d), where \mathbf{f} is the body force as defined in section 4.1 and \mathcal{A} the first variation of \mathcal{S} with respect to \mathbf{v} (see (4.2a) and (4.2b), respectively). It immediately follows that we obtain

$$\hat{\mathcal{H}}[\tilde{\mathbf{v}}] := \beta \mathcal{A}[\tilde{\mathbf{v}}] + \mathcal{K}[\tilde{\mathbf{f}}] = -\tilde{\mathbf{g}} \quad (4.12)$$

to replace (4.3d); $\tilde{\mathbf{f}}$ denotes the incremental body force defined in section 4.1.

4.3.5. Parameter selection—To the extent possible, it is desirable to design a numerical scheme that does not require a selection of parameters (black-box solver). This is challenging for previously unseen data. In general, the user should only be required to decide on the following:

- The desired accuracy of the inversion (controlled by the tolerance $\tau_{\mathcal{J}}$; see section 4.2.4).
- The desired properties of the mapping y (controlled by ε_{θ} or $\varepsilon_{\bar{F}}$ respectively; see below).
- The budget we are willing to assign to the computation (controlled by n_{opt} ; see section 4.2.4).

For the purpose of this numerical study we proceed as follows.

Optimization: We set the maximum number of iterations n_{opt} (see (4.8)) to 1E6, as we do not want our algorithm to terminate early (i.e., we make sure that we terminate only if either we reach the defined tolerances or we no longer observe a decrease in \mathcal{J}^h). For the convergence study in section 5, we use the relative change of the ℓ^{∞} -norm of the gradient \mathbf{g}^h as a stopping criterion, as we are interested in studying convergence properties. This enables an unbiased comparison in terms of the required work to solve an optimization problem up to a desired accuracy. In particular, we terminate the optimization if the relative change of the reduced gradient \mathbf{g}^h is larger than or equal to three orders of magnitude.

Following standard textbook literature [29, 56] we use the stopping criteria in (4.8) for the remainder of the experiments. We set the tolerance to $\tau_{\mathcal{J}} = 1\text{E-}3$. We *qualitatively* did not observe significant differences in the final results for the experiments performed in this study, when turning to smaller tolerances. We will further elaborate on the required accuracy for the inversion (i.e., the registration quality) in a follow-up paper.

The tolerance of the PCG method is set as discussed in section 4.2.1 (see (4.5)). The maximal number of iterations for the PCG method is set to n (order of the reduced KKT system in (4.4)). In theory, this guarantees that the PCG method converges to a solution. This choice not only ensures that we provide an unbiased study (i.e., we do not terminate early) but also makes sure that we do not miss any issues in the implementation or parameter selection. We converged for all experiments conducted in this study after only a fraction of n inner iterations. This statement is confirmed by the reported number of PDE solves.⁶

For all our experiments we initialized the line search with a factor of $\alpha_k = 1$ (see section 4.2). This is a sensible choice, as search directions obtained from second-order methods are nicely scaled (i.e., we expect α_k to be 1). However, this is not the case for the Picard scheme (i.e., the preconditioned gradient descent). Our implementation features an option to memorize the scaling of s_k for the next outer iteration. That is, we introduce an additional scaling factor $\tilde{\alpha}_k > 0$ that is applied to s_k before entering the line search (initialized with $\tilde{\alpha}_k = 1$). If the line search kicks in, we downscale $\tilde{\alpha}_k$ by α_k . On the contrary, we upscale $\tilde{\alpha}_k$ by a factor of two if $\alpha_k = 1$.

⁶“PDE solve” refers to the solution of one of the hyperbolic PDEs that appear in the optimality system (4.1) and (4.3).

PDE solver: The number of time steps n_t is bounded from below due to stability requirements (see section 4.3.1). Since we use an expansion in time (see section 4.3.2), it is possible to adaptively adjust n_t , so that numerical stability is attained.

However, we fix n_t for the numerical experiments in section 5.3 as we are interested in studying the convergence behavior with respect to the employed grid size. We set n_t to $4 \max(\mathbf{n}_x)$. This is a pessimistic choice. If we still encounter instabilities (as judged by monitoring the CFL condition (see section 4.3.1)), s_k is scaled by a factor of 0.5 until numerical stability is attained, *before* entering the line search. For all numerical experiments conducted in this study, we did not observe any instabilities for the Newton–Krylov methods. However, for the Picard method we observed instabilities in the case when we did not consider the rescaling procedure detailed above. This is due to the fact that s_k is arbitrarily scaled for first-order methods (as opposed to second-order methods). By introducing the additional scaling parameter $\tilde{\alpha}_k$ we could stabilize the Picard method—we did not observe a violation of the CFL condition for any of the conducted experiments (for n_t fixed).

Regularization: Estimating an optimal value for β is an area of research by itself. A variety of methods has been designed (see, e.g., [72]). A key difficulty is computational complexity. Methods based on the assumption that differences between model output and observed data are associated with random noise (such as generalized cross validation) might not be reliable in the context of nonrigid image registration. This is due to the fact that the noise in the images is likely to be highly structured [38]. Another possibility is to estimate the regularization parameter on the basis of the spectral properties of the Hessian (see section 5.3.1). That is, we can estimate the condition number of the problem during the PCG solves for the unregularized problem using the Lanczos algorithm (see [30, p. 528]). We can do this very efficiently by initializing the problem with a zero velocity field. Given \mathbf{v} is zero, the application of the Hessian within the PCG is computationally inexpensive, as a lot of the terms in the optimality systems drop (see section 4.1). However, the level of regularization depends not only on properties of the data, but also on regularity requirements on \mathbf{y} .

Another common strategy is to perform a *parameter continuation* in β (see, e.g., [35, 38]). In [38] it has been suggested to inform the algorithm about the required regularity of a solution on the basis of a lower bound on the L^2 -distance between the reference and the deformed template image. The decision on such bound, however, might not be intuitive for practitioners. Further, one is ultimately interested not only in a small residual but also in a bounded determinant of the deformation gradient. Therefore, we propose to inform the algorithm on regularity requirements in terms of a lower bound $\varepsilon_F \in (0, 1)$ on $\det(\mathbf{F}_1^h)$ (i.e., a bound on the tolerable compression of a volume element). If the Stokes regularization scheme ($\gamma = 1$ in (1.2b) and \mathcal{S} is an H^1 -seminorm) is considered, bounds on geometric constraints of the deformation of a volume element can be used. In particular, we use a lower bound $\varepsilon_\theta > 0$ on the acute angle of a grid element. The upper bound on the obtuse angle is given by $2\pi - \varepsilon_\theta$. Note that it is actually necessary to monitor geometric properties to guarantee a local diffeomorphism; a lower bound on $\det(\mathbf{F}_1^h)$ is not sufficient.

Our algorithm proceeds as follows. In the first step, the registration problem is solved for a large value of β ($\beta = 1$ in our experiments) so that we underfit the data.⁷ Subsequently, β is reduced by one order of magnitude until we reach ϵ_F (or ϵ_θ). From there on, a binary search is performed. The algorithm is terminated if the change in β is below 5% of the value for β , for which ϵ_F (or ϵ_θ) was breached. We add a lower bound of $1E-6$ on β as well as a lower bound for the relative change of the L_2 -distance of $1E-2$ to ensure that we do not perform unnecessary work. We never reached these bounds for the experiments conducted in this study.

Presmoothing: A numerical challenge in image computing is that images are functions of bounded variation. Therefore, an accurate computation of the derivatives becomes more involved. A common approach to ensure numerical stability and avoid the Gibbs phenomena is to reduce high-frequency information in the data. We use a Gaussian smoothing, which is parametrized by a user-defined standard deviation $\sigma > 0$. We experimentally found a value of $\sigma = 2\pi/\min(n_x)$ to be adequate for the problems at hand. However, we note that we increased σ by a factor of 2 for one set of experiments in section 5.3.2. We also note that we implemented a method for grid and scale continuation for the images. This avoids the problem of deciding on σ . We will investigate an automatic selection strategy for σ in a follow-up paper.

It is important to note that the sensitivity of second-order derivatives to noise in the data is problematic. Therefore, we refrain from applying the N-PCG method to nonsmooth images.

5. Numerical experiments

We report results only in two dimensions. We test the algorithm on real-world and synthetic registration problems (see section 5.1). The measures to analyze the registration results are summarized in section 5.2. We conduct a numerical study (see section 5.3), which includes an analysis of (i) the spectral properties of the Hessian (see section 5.3.1), (ii) grid convergence (see section 5.3.2), and (iii) the effects of varying the number of the unknowns in time (see section 5.3.3). We additionally report results for a fully automatic registration on high-resolution images based on the designed parameter continuation in β (see section 5.4).

5.1. Data

We consider synthetic and real-world registration problems.⁸ These are illustrated in Figure 1. All images have been normalized to an intensity range of $[0, 1]$. The synthetic problems are constructed by solving the forward problem to create an artificial template image m_T given some image m_R and some velocity field \mathbf{v}^\star (sinusoidal images and UT images in Figure 1). Here, \mathbf{v}^\star is chosen to live on the manifold of divergence-free velocity fields to provide a testing environment for the Stokes regularization scheme. Further, \mathbf{v}^\star is by construction assumed to be constant in time (i.e., $n_c = 1$). In particular, we have

⁷Note that for large β the optimization problem is almost quadratic, so that Newton-Krylov methods converge quickly.

⁸The hand images are taken from [56].

$$v^{i,\star}(x^i, t) = 0.5 \sin(x^i) \cos(x^i) \quad \forall t \in [0, 1], \quad i=1, \dots, d. \quad (5.1)$$

5.2. Measures of performance

We report the number of the (outer) iterations and the number of the hyperbolic PDE solves to assess the computational work load. The latter is a good proxy for the wall clock time. It provides a transparent comparison between the designed first- and second-order methods, given that the number of the hyperbolic PDE solves varies between these methods: Two hyperbolic PDE solves are required during each iteration of the Picard method (we have to solve (4.1a) and (4.1b) to evaluate the reduced gradient in (4.1d); see also Algorithm 1). For the Newton–Krylov methods, we require an additional two hyperbolic PDE solves per inner iteration (we have to solve (4.3a) and (4.3b) to compute the Hessian matrix-vector product given in (4.3d); see also Algorithm 2). Each evaluation of \mathcal{J}^h in (1.2) (i.e., each line search step) requires an additional hyperbolic PDE solve (we have to compute the deformed template image at $t=1$ by solving (4.1a)). Note that the solution of the forward and adjoint problems is the key bottleneck of our algorithm. We will report the wall clock times in a follow-up paper, in which we extend the current framework to a three-dimensional implementation. This study focuses on algorithmic features.

We report the relative change of (i) the L^2 -distance, (ii) the objective functional \mathcal{J}^h , and (iii) the ℓ^∞ -norm of the (reduced) gradient \mathbf{g}^h to assess the quality of the inversion. We additionally report values for the determinant of the deformation gradient to study local deformation properties. These measures are defined more explicitly in Table 2.

We visually support this quantitative analysis on the basis of snapshots of the registration results. Information on the reconstruction accuracy can be obtained from pointwise maps of the residual difference between m_R and m_I . The deformation regularity and the mass conservation can be assessed via images of the pointwise determinant of the deformation gradient and/or of a deformed grid overlaid onto m_I . Details on how these are obtained and on how to interpret them can be found in Appendix D.

5.3. Numerical study

We study the spectral properties of the Hessian (see section 5.3.1), grid convergence (see section 5.3.2), as well as the influence of an increase in the number of the unknowns in time (see section 5.3.3).

5.3.1. Spectral analysis

Purpose: We study the ill-posedness and the ill-conditioning of the problem at hand. We report spectral properties of \mathcal{H}^h . We study the eigenvalues and the eigenvectors with respect to different choices for β . We study the differences between plain H^2 -regularization and the Stokes regularization scheme (H^1 -regularization).

Setup: This study is based on the UT images (the true solution $\mathbf{v}^{h,\star}$ is divergence free; see section 5.1 for more details on the construction of this synthetic registration problem; $\mathbf{n}_X =$

$(64, 64)^T$ and $n_c = 1$ so that $n = 8192$). The eigendecomposition $\mathbf{V} \mathbf{\Lambda} \mathbf{V}^{-1}$, $\mathbf{V} = (\mathbf{v}_i)_{i=1}^n$, $\mathbf{v}_i \in \mathbf{R}^n$, $\|\mathbf{v}_i\|_2 = 1$, $\mathbf{\Lambda} = \text{diag}(\Lambda_{11}, \dots, \Lambda_{nn})$, $\Lambda_{ii} > 0$, is computed at the true solution \mathbf{v}^h, \star to guarantee that $\mathcal{H}^h > 0$. The spectrum is computed for three different choices of β : for the unregularized problem ($\beta = 0$), an empirically determined (moderate) value ($\beta = 1\text{E-}3$), and solely for the regularization model ($\beta = 1\text{E}6$).

Results: Figure 2 displays the trend of the absolute value of the eigenvalues Λ_{ii} , $i = 1, \dots, n$. They are sorted in *descending* order for $\beta = 0$ and in *ascending* order otherwise. If an eigenvalue drops below machine precision (i.e., $1\text{E-}16$), it is set to $1\text{E-}16$ (only for visualization purposes). The extremal real and imaginary part of the eigenvalues is summarized in Table 3. Figure 3 provides the spatial variation of the eigenvectors $\mathbf{v}_i \in \mathbf{R}^n$ that correspond to the eigenvalues Λ_{ii} , $i \in \{1, 5, 20, 100, 1000\}$, in Figure 2 with respect to different choices for β and different regularization schemes. We only display the first component \mathbf{v}_i^1 of the coefficient field $\mathbf{v}_i := (\mathbf{v}_i^1, \mathbf{v}_i^2)$. The pattern for the second component is (qualitatively) alike.

Observations: The most important observations are that (i) the regularization schemes display a very similar behavior (as judged by the clustering of the eigenvalues as well as the spatial variation of the eigenvectors for $\beta = 1\text{E}6$), (ii) the smoothness of the eigenvectors *decreases* with a *decreasing* regularization parameter β and *increasing* eigenvalues (for $\beta = 1\text{E-}3$ and $\beta = 1\text{E}6$) for both regularization schemes, and (iii) it is less clear how to identify the smooth eigenvectors within the eigenspace of the Stokes regularization scheme.

The eigenvalues Λ_{ii} , $i = 1, \dots, 8192$, drop rapidly for the unregularized problem, approaching almost machine precision for $i \approx 4000$ (see Figure 2). This demonstrates ill-conditioning and ill-posedness. The eigenvalues are bounded away from zero for the regularized problem. Increasing β shifts the trend of $|\Lambda_{ii}|$ to larger numbers. The values in Table 3 confirm that $\mathcal{H}^h > 0$ (up to almost machine precision) at the true solution \mathbf{v}^h, \star .

Turning to the eigenvector plots, we can see that the first eigenvector displays a delta peak like structure for both regularization schemes, since there is no local coupling of the spatial information. For the regularized problem we can observe a smooth spatial variation for the eigenvectors associated with large eigenvalues for both regularization schemes. The first eigenvector plot is almost constant for $\beta = 1\text{E}6$ (bottom row of each block in Figure 3). The structure of the pattern for $\beta = 1\text{E}6$ is analogue for both schemes, which indicates similarities in the behavior of both schemes. As the index i increases, the eigenvectors become more oscillatory. We can observe strong differences between the two schemes for a moderate regularization ($\beta = 1\text{E-}3$; middle row of each block in Figure 3). Also, we can observe a more complex structure for the Stokes regularization scheme for small eigenvalues (i.e., it is difficult to identify where the smoothest eigenvectors are located within the eigenspace).

Conclusion: We conclude that the Hessian operator is singular if we do not include a smoothness regularization model for the control variable. For practical values of the regularization parameter, the Hessian behaves as a compact operator; larger eigenvalues are

associated with smooth eigenvectors. It is well known that designing a preconditioner for such operators is challenging.

5.3.2. Convergence study—We study the grid convergence of the considered iterative optimization methods on the basis of synthetic registration problems. We use a rigid setting to prevent bias originating from adaptive changes during the computations. That is, the results are computed on a single resolution level. No grid, scale, or parameter continuation is applied. The number of the time points is fixed to $n_t = 4\max(n_x)$ for all experiments. Further, we use empirically determined values for the regularization parameter β , namely, $\beta \in \{1\text{E}-2, 1\text{E}-3\}$. Since we are interested in studying the convergence properties of our method, we consider the relative change of the ℓ^∞ -norm of the reduced gradient \mathbf{g}^h as a stopping criterion. This yields a fair comparison between the different optimization methods, as a reduction in the norm of \mathbf{g}^h directly reflects how well an optimization problem is solved (i.e., we exploit that $\mathbf{g}^h = 0$ is a necessary condition for a minimizer). We terminate if the relative change of the ℓ^∞ -norm of \mathbf{g}^h is at least three orders of magnitude. However, since the Picard method tends to converge slowly for low tolerances with respect to the gradient, we stop if we detect a stagnation in the objective. In particular, we terminate the optimization if the change in the objective in ten consecutive iterations was equal or below $1\text{E}-6$. We solve for a stationary velocity field (i.e., $n_c = 1$).

C^∞ registration problem

Purpose: We study the numerical behavior for smooth registration problems. We report results for grid convergence and deformation regularity. We compare the Picard, GN-PCG, and N-PCG methods.

Setup: This experiment is based on the sinusoidal images (see section 5.1 for more details on the construction of this synthetic registration problem). Therefore, $m_T, m_R \in C^\infty(\Omega)$, and $\mathbf{v}^\star \in L^2([0, 1]; C^\infty(\Omega)^d)$ so that the excellent convergence properties of Fourier spectral methods are expected to pay off. Additionally, it is not problematic to apply the N-PCG method. We report results for different grid sizes $\mathbf{n}_x = (n_x^1, n_x^2)^\top$, $n_x^i \in \{64, 128, 256\}$, $i = 1, 2$, $n_t = 4\max(n_x)$. No presmoothing is applied. We use an experimentally determined value of $\beta = 1\text{E}-3$ for all experiments. The remainder of the parameters is chosen as stated in the introduction to this section as well as in section 4.3.5.

Results: The grid convergence results are summarized in Table 4. Values derived from the deformation gradient \mathbf{F}_1^h are reported in Table 5. Exemplary results for the plain H^2 -regularization ($\gamma = 0$) and the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization) are displayed in Figure 4. The definitions of the quantitative measures reported in Table 4 and Table 5 can be found in Table 2.

Observations: The most important observations are that (i) there are significant differences in computational work between the Picard and Newton–Krylov methods with the latter being much more efficient, (ii) the differences between the Newton–Krylov methods are insignificant, (iii) the rate of convergence is independent of the grid resolution, and (iv) the numerical accuracy is almost at the order of machine precision.

The registered images are quantitatively (see Table 4) and qualitatively (see Figure 4) in excellent agreement. For the considered tolerance (reduction of the L^∞ -norm of the reduced gradient by three orders of magnitude) we can reduce the L^2 -distance between three (compressible deformation) and four (incompressible deformation) orders of magnitude (see Table 4). The search direction of the Newton–Krylov methods is nicely scaled. No additional line search steps are necessary. We require 1.57 to 1.71 line search steps for the Picard iteration (on average). Note that we prescale the search direction of the Picard method by an additional parameter $\tilde{\alpha}_k$, which is estimated during the computation (see section 4.3.5 for details). Otherwise, the number of the line search steps would be seven to eight on average for the Picard method. The Picard method did stagnate during the computations. This is why the gradient has not been reduced by three orders of magnitude for the Picard method. However, it is in general possible to reduce the gradient accordingly. We decided to report only until stagnation for the Picard method as the number of the iterations would significantly increase without making any real progress.

The Newton–Krylov methods display quick convergence. Only five outer iterations are necessary to reduce the gradient by more than four orders of magnitude. The results demonstrate a significant difference in the computational work between first- and second-order methods for the considered tolerance.

The reconstruction quality improves by approximately one order of magnitude when switching from plain H^2 -regularization ($\gamma = 0$) to a Stokes regularization scheme ($\gamma = 1$; H^1 -regularization), as judged by the relative change in the L^2 -distance. This is expected, since the synthetic problem has been created under the assumption of mass conservation (i.e., $\nabla \cdot \mathbf{v}^\star = 0$). Second, we expect a smaller contribution of the H^1 -regularization model on the solution for the same values of β .

From a theoretical point of view, we expect N-PCG to outperform GN-PCG (quadratic versus superlinear convergence). The reported results demonstrate an almost identical performance. This is due to the fact that we can drive the residual almost to zero, such that we can recover fast local convergence for the GN-PCG method (see section 4.2.2 for details).

The Picard method converges faster for the Stokes regularization scheme. However, the differences between the Picard and the Newton–Krylov methods are still significant with an approximately four-fold difference in n_{PDE} . For Newton–Krylov methods, we can globally observe a slight increase in the number of the inner iterations when switching from plain H^2 -regularization to the Stokes regularization scheme. These differences have to be attributed to a varying relative change in the reduced gradient.⁹

The results reported in Table 5 demonstrate an excellent numerical accuracy for the mass conservation for all numerical schemes. The error in the determinant of the deformation gradient is $\mathcal{O}(1\text{E-}12)$, i.e., we achieve an accuracy that is almost at the order of machine precision for a grid resolution of $\mathbf{n}_x = (256, 256)^\top$ and $n_t = 1024$.

⁹Note that the tolerance of the Krylov-subspace method and therefore the number of the inner iterations depends on the gradient (see (4.5)).

Conclusion: We conclude that we can interchangeably use the Newton–Krylov methods. Therefore, given that N-PCG is more sensitive to noise and discontinuities in the data, we will exclusively consider GN-PCG for the remainder of the experiments. Also, if we require an inversion with high accuracy, the Newton–Krylov methods clearly outperform the Picard method (i.e., the preconditioned gradient descent).

Images with sharp features

Purpose: We study the grid convergence and deformation regularity for an image with sharp features. We compare the Picard and the GN-PCG method.

Setup: We consider the UT images (see section 5.1 for details on the construction of this synthetic registration problem). We report results for experimentally determined values of $\beta \in \{1\text{E-}2, 1\text{E-}3\}$ with respect to different grid resolution levels

$\mathbf{n}_x = (n_x^1, n_x^2)^\top$, $n_x^i \in \{64, 128, 256\}$, $i = 1, 2$, $n_t = 4\max(\mathbf{n}_x)$. The remainder of the parameters are chosen as stated in the introduction of this section and in section 4.3.5. Both plain H^2 -regularization ($\gamma = 0$) as well as the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization) are considered.

For images of size $\mathbf{n}_x = (256, 256)^\top$ and a Stokes regularization scheme, we observed difficulties in the inversion (only the number of the outer and the inner iterations increased; the algorithm still converges to the same solution), due to a strong forcing (i.e., the sharp features pushed the solver at an early stage to a solution that was far away from the final minimizer). We increased the smoothing by a factor of two as a remedy. This is not an issue for the practical application of our algorithm, as our framework features a method for performing a scale continuation as well as a continuation in the regularization parameter. Therefore, the user does not have to decide on σ or on β . In addition to that, we currently investigate adaptive approaches to automatically detect insufficient smoothness during the course of the optimization to prevent a deterioration in the convergence behavior.

The remainder of the parameters are chosen as stated in the introduction of this section as well as in section 4.3.5.

Results: Table 6 summarizes the results of the convergence study. We illustrate intermediate results with respect to the first 13 (outer) iterations k in Figure 5 (plain H^2 -regularization; $\gamma = 0$). We report the trend of the individual building blocks of \mathcal{J}^h (contribution of the L^2 -distance and the regularization model \mathcal{S}^h) in Figure 6. We report measures of deformation regularity in Table 7.

Observations: The most important observations are (i) the GN-PCG method displays a quicker convergence than the Picard method, (ii) we cannot achieve the same inversion accuracy with the Picard method as compared to the GN-PCG method, and (iii) the number of the (inner and outer) iterations increases and is no longer independent of the resolution level.

The rate of convergence decreases compared to the results reported in the former section (see Table 4). Overall, we require more outer and inner iterations to solve the registration problem.

The residual differences between m_R and m_1 clearly depend on the choice of β (see Table 4 and Figure 6). We achieve a similar reduction in the L^2 -distance for both the Picard and the GN-PCG method (two to four orders of magnitude). The residual differences are less pronounced when switching from plain H^2 -regularization to the Stokes regularization scheme as compared to the results reported in section 5.3.2.

We cannot guarantee that it is possible to reduce the gradient by three orders of magnitude if we use the Picard method. Even if we do not include a condition to terminate if we observe stagnation (i.e., the change in \mathcal{J}^h is below or equal to $1E-6$ for 10 consecutive iterations), it is for some of the experiments not possible to reduce the gradient by three orders of magnitude as the changes of the objective hit our numerical accuracy (which causes the line search to fail). We do not observe this issue when considering the GN-PCG method. Further, there are significant differences in terms of the computational work. If we do not account for the stagnation of the Picard method we have observed a number of hyperbolic PDE solves that is well above $\mathcal{O}(1E4)$. Clearly, in a practical application we terminate the Picard method at an earlier stage, as we no longer make significant progress. However, in this part of the study we are interested in the convergence properties. This experiment demonstrates that we cannot guarantee a high inversion accuracy (i.e., a significant reduction in the gradient) when turning to first-order methods. Note that we have stabilized the Picard method by introducing an additional scaling parameter for the search direction that prevents additional line search steps (see section 4.3.5). If we neglect this scaling, we observe seven to nine line search steps on average (results not included in this study) for the considered problem; also, the optimization fails at an early stage. The search direction obtained via the GN-PCG method is nicely scaled; no additional line search steps are necessary.

The trend of the \mathcal{J}^h , the L^2 -distance and \mathcal{S}^h in Figure 6 confirm these observations. The plots in Figure 6 illustrate that the Picard and the GN-PCG method perform very similarly during the first few outer iterations. However, after about four outer iterations the differences between the methods manifest, in particular with respect to the reduction of the L^2 -distance. This observation confirms standard numerical optimization theory on convergence properties of the Picard and the inexact Newton–Krylov methods.

Focusing on the GN-PCG method we can observe that the number of the outer iterations is almost constant across different grid sizes. However, the effectiveness of the spectral preconditioner decreases with an increasing grid size as well as with a reduction of the regularization parameter (as judged by an increase in the number of the inner iterations). This demonstrates that the preconditioner is not optimal. A similar behavior can be observed for the Picard method.¹⁰

¹⁰Note that the Picard method is a gradient descent scheme in the function space induced by \mathcal{W} . We can interpret the inverse of \mathcal{A}^h as a preconditioner acting on the body force f . This operator is exactly the spectral preconditioner we use for the Newton–Krylov methods, which explains the similar behavior.

The numerical accuracy of the incompressibility constraint deteriorates (slightly, but not significantly) as compared to the results reported in the former section. In particular, we obtain a numerical accuracy of $\mathcal{O}(1\text{E-}5)$ for the GN-PCG method (see Table 7).

Conclusion: We conclude that the GN-PCG is less sensitive, provides a better inversion accuracy, and overall displays quicker convergence if a high accuracy of the inversion is required and, therefore it is to be preferred.

5.3.3. Number of unknowns in time

Purpose: It is not immediately evident how the number of the coefficient fields

$\mathbf{v}_l^h: \mathbf{R}^d \rightarrow \mathbf{R}^d, l = 1, \dots, n_c$ affects the registration quality. We study the effects of varying n_c on the reconstruction quality and the rate of convergence. We also provide advice on how to decide on n_c .

Setup: We report results for registration problems of varying complexity. The analysis is limited to the GN-PCG method. We consider the hand images ($\mathbf{n}_x = (128, 128)^\top$) and the brain images ($\mathbf{n}_x = (200, 200)^\top$). The number of the time steps is fixed to $n_t = 2\max(\mathbf{n}_x)$. The regularization parameter is empirically set to $\beta = 1\text{E-}3$ and $\beta = 2\text{E-}2$, respectively. We consider the full set of stopping conditions in (4.8) with $\tau_{\mathcal{J}} = 1\text{E-}3$, as we no longer compare different methods. The remainder of the parameters are set as stated in section 4.3.5.

One possibility to estimate an adequate number of coefficients for the registration of unseen images m_R and m_T is to compute the relative spectral power (see Table 2) of an individual coefficient field \mathbf{v}_l^h for different choices of n_c . If only a small number of coefficients is necessary to recover the deformation, this energy should decrease rapidly with an increasing l . The problem is stationary for $n_c = 1$.

Results: The trend of the relative ℓ^2 -norm (i.e., the spectral power) of an individual coefficient field \mathbf{v}_l^h for different choices of $n_c \in \{1, 2, 4, 8, 16\}$ is plotted in Figure 7. Convergence results are reported in Table 8. A qualitative comparison of the registration results for different choices of n_c can be found in Figure 8.

Observations: The most important observation is that we obtain the same results for stationary as well as time varying velocity fields for two-image registration problems. Qualitatively, we cannot observe any differences for a varying number of coefficient fields (see Figure 8). This observation is confirmed by the values for the relative reduction in the L^2 -distance in Table 8. Increasing the number of the coefficients slightly reduces the L^2 -distance. These differences, however, are practically insignificant. In particular, we (on average) observe a relative change in the L^2 -distance of $6.50\text{E-}2 \pm 1.20\text{E-}3$ (hand images, plain H^2 -regularization) and $5.37\text{E-}1 \pm 9.75\text{E-}3$ (brain images, plain H^2 -regularization). Also, we obtain identical deformation patterns as judged by careful visual inspection (see Figure 8) and the variations in the determinant of the deformation gradient. We obtain identical results for the UT images (for plain H^2 -regularization and the Stokes regularization scheme; results are not included in this study).

Turning to the required work load, we observe that the differences are also insignificant. The number of the outer iterations is almost constant; just the number of the inner iterations varies. In particular, we require 7 outer iterations with (on average) ≈ 280 inner iterations (hand images; plain H^2 -regularization; $\gamma = 0$) and 20–21 outer iterations with (on average) ≈ 693 inner iterations (brain images; plain H^2 -regularization; $\gamma = 0$). However, we have to keep in mind that each application of the reduced Hessian is slightly more expensive and we require more memory as n_c increases. That is, we have to store more coefficient fields v_i^h (to all of which the regularization operator has to be applied). The cost of the forward and adjoint solves (which is the key bottleneck), however, is (almost) the same, since we expand v^h (note that this expansion is not necessary for $n_c = 1$; see section 4.3.2).

The power spectrum of the coefficient fields drops quickly (see Figure 7). This also indicates that only a small number of coefficients is required to obtain an excellent agreement between the images. However, we expect the differences to manifest, when registering time series of images (multiple time frames). Here, we might benefit from being able to invert for a time varying velocity field.

Conclusion: We conclude that it is sufficient to use stationary velocity fields for two-image registration problems.

5.4. Parameter continuation to estimate β

Purpose: We study the stability and accuracy of the designed parameter continuation method (see section 4.3.5) and the associated control over the properties of the mapping. That is, we study how the quantities of interest (determinant of the deformation gradient and L^2 -distance) behave during the course of the parameter continuation and how close we actually approach the given bounds.

Setup: The registration problems are solved on images with a grid size of $n_x = (512, 512)^\top$. The number of time points is adapted as required by monitoring the CFL condition (see section 4.3.1). We use the full set of stopping conditions (see section 4.2.4) with a tolerance of $\tau_{\mathcal{J}} = 1\text{E-}3$. We consider the hand images and the brain images (see Figure 1). We invert for a stationary velocity field (i.e., $n_c = 1$). In case we consider a plain H^1 - and H^2 -regularization (smoothness regularization; $\gamma = 0$), we set the lower bound on $\det(F_1^h)$ to $1\text{E-}1$ (hand images) and $5\text{E-}2$ (brain images), respectively. For the case of the Stokes regularization ($\gamma = 1$; H^1 -regularization) we set the bound on the grid angle to $\varepsilon_\theta = \pi/16$ (11.25°). The remainder of the parameters are set as described in section 4.3.5.

Results: We report the obtained estimates for β as well as results for the reconstruction quality and deformation regularity in Table 9. We provide an exemplary illustration of the obtained registration results in Figure 9. We report results for the course of the parameter continuation in Figure 10.

Observations: The most important observation is that we can precisely control the properties of our mapping without having to manually tune any parameters. We only have to decide on

geometric bounds (the smallest tolerable deformation of a grid element or a bound on the shear angle of the grid cell), the decision on which is intuitive for practitioners.

The accuracy of our method (in space) is only limited by the grid resolution (i.e., how many frequencies we can resolve; this statement is confirmed by the experiments conducted in section 5.3.2) as well as the defined bounds on the binary search used to estimate β (see section 4.3.5). Clearly, the desired level of accuracy competes with the computational work load we are willing to invest.

For plain H^1 - and H^2 -regularization, we achieve an excellent agreement between m_R and m_I (see Figure 9) with a reduction of the L^2 -distance by approximately half an order of magnitude for the brain images and 1.5 orders of magnitude for the hand images (see Table 9). The discrepancy between the lower bound e_F and $\min(\det(\mathbf{F}_h^1))$ for the obtained optimal value of β is small. In particular, we are, e.g., bounded from above by an absolute difference of $1.13\text{E-}3$ for the brain images (H^2 -regularization) and $5.10\text{E-}3$ for the hand images (H^2 -regularization). These values are well above the attainable accuracy reported in section 5.3.2.

For the results reported for the Stokes regularization scheme we can qualitatively (see Figure 9) and quantitatively (see Table 9) observe that enforcing incompressibility up to numerical accuracy is a too-strong prior for the considered problem. However, the key observation and intention of this experiment is to demonstrate that we attain a deformation that is very well behaved (with $\det(\mathbf{F}_1^h)=1$). A direct comparison to the result obtained for the incompressible case reveals that the mapping is diffeomorphic but displays a large variation in the magnitude of the determinant of the deformation gradient (see the leftmost image in the middle row and bottom row as well as the corresponding maps for $\det(\mathbf{F}_1^h)$ in Figure 9). If we further decrease the bound on $\det(\mathbf{F}_1^h)$ we will lose control and generate a mapping that locally is close to being nondiffeomorphic. We again emphasize that the intention of this work is the study of algorithmic properties. We will address the practical benefit of exploiting a model of (near-)incompressible flow in a follow-up paper and refer to [10, 16, 52, 62] for potential applications. This exemplary result on real-world data demonstrates that it might be beneficial to consider a relaxation of the incompressibility constraint in order to improve the mismatch between the considered images while maintaining as much control on the deformation regularity as possible.

In future work, we will focus on improvements of the computational efficiency for estimating β . We have tested combining it with a grid continuation but could not observe strong improvements. We will also investigate the idea of providing a coarse estimation of β via the spectral properties of the Hessian and from there on do a parameter continuation (see section 4.3.5 for additional comments).

Conclusion: We conclude that the designed framework is highly accurate, is stable, guarantees deformation regularity (assuming that the user-defined tolerance is sufficiently bounded away from irregularities), and does not require any additional tuning of parameters. The user merely has to provide a lower bound on an acceptable volume change or a bound

on the distortion of a volume element (shear angle), the decision on which is intuitive for practitioners.

6. Conclusions

We have presented numerical methods for large deformation diffeomorphic nonrigid image registration that (i) operate in a black-box fashion, (ii) introduce novel algorithmic features (including a second-order Newton–Krylov method, spectral preconditioning, an efficient solver for Stokes problems, and a spectral Galerkin method in time), (iii) is stable and efficient and (iv) guarantees deformation regularity with an explicit control on the quality of the deformation.

In addition, we have conducted a detailed numerical study to demonstrate computational performance and numerical behavior on synthetic and real-world problems. The most important observations of our study are the following:

- The Newton–Krylov methods outperform the globalized Picard method (see section 5.3.2) as we increase the image size and the registration fidelity.
- We can enforce incompressibility with high accuracy. The numerical accuracy (in space) is only limited by the resolution of the data (see section 5.3.2).
- We can compute deformations that are guaranteed to be regular (i.e., a local diffeomorphism) up to user specifications. Controlling the magnitude of $\det(\mathbf{F}_1)$ is not sufficient, as volume elements might still collapse (deformation field with strong shear). Therefore, we introduced a parameter continuation in β that can be interfaced not only with lower bounds on the determinant of the deformation gradient but also with bounds on the geometric properties of the grid cells (in particular, the shear angle of a grid cell; see section 5.4).
- The experiments reported in this study demonstrate that it is adequate to limit the inversion to stationary velocity fields when considering two-image registration problems. This observation is in accordance with results reported for other classes of large deformation image registration algorithms [1, 3, 44, 52, 69, 70]. We have additionally provided advice on how to decide on the number of unknowns in time (see section 5.3.3).

The control equation for the velocity is a space-time nonlinear elliptic system. But the main cost in our formulation is the solution of transport problems to compute the image transformation and the adjoint variables. That is, we require two hyperbolic PDE solves for computing the gradient (which essentially corresponds to one Picard iteration or one outer iteration of a Newton–Krylov method; see Algorithm 1) and an additional two hyperbolic PDE solves for evaluating the incremental control equation (Hessian matrix-vector product in Newton–Krylov methods; see Algorithm 2) in each inner iteration of the Krylov-subspace method. Because we use a pseudospectral discretization in space, the elliptic solve for the

Picard iteration is (for quadratic regularization models) only at the cost of a spectral diagonal scaling. For the Newton–Krylov methods, we have to solve a linear system using an iterative solver. The Picard scheme has a lower cost per iteration but requires more iterations than the Newton–Krylov scheme.

Our results demonstrate that there is a significant difference in stability, computational work, and accuracy between the Picard and Newton–Krylov methods, especially when we require a high accuracy of the inversion. If we require an inaccurate solution or use a strong regularization, the differences between Picard and Newton–Krylov methods are less pronounced. Better preconditioning of the Hessian would make the Newton–Krylov approach preferable across the spectrum of accuracy requirements. The Newton–Krylov approach is not significantly more complex, since we essentially use the same numerical tools that have been used for the solution of the first-order optimality conditions. Also, the individual building blocks ((incremental) forcing term, regularization operator \mathcal{A} and the projection operator \mathcal{K}) that appear in the first- and second-order optimality system are very similar. Therefore, the difference of solving the first- or the second-order optimality conditions essentially amounts to interfacing a Krylov-subspace method to solve the saddle point problem.

By formulating the nonrigid image registration as a problem of optimal control, we target the design of a generic, biophysically constrained framework for large deformation diffeomorphic image registration. Further, there are many applications that do require incompressible or near-incompressible deformations, for example, in medical image analysis. Our framework provides such a technology.

We report results only in two dimensions. Nothing in our formulation and numerical approximation is specific to the two-dimensional case. The next steps will be its extension to three-dimensions and to problems that have time sequences of images. For such cases, we expect to have to invert for a nonstationary velocity field. In addition, we aim at designing a framework that allows for a relaxation of the incompressibility constraint, as we observed in this study that incompressibility might be a too-strong prior. We also observed (results not included in this study) that the use of an incompressibility constraint can promote shear. In a follow-up paper, we will target this problem by introducing a novel continuum mechanical model that allows us to control the shear inside the deformation field.

Table 4.

Acknowledgments

We would like to thank Amir Gholaminejad, Georg Stadler, and Bryan Quaife for helpful discussions and comments. We would like to thank Florian Tramnitzke for his initiative work on this project. Any opinions, findings, and conclusions or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the AFOSR or the NSF.

References

1. Arsigny, V.; Commowick, O.; Pennec, X.; Ayache, N. A Log-Euclidean framework for statistics on diffeomorphisms. Springer; New York: 2006. Medical Image Computing and Computer-Assisted Intervention, Lect. Notes in Comput. Sci 4190; p. 924-931.

2. Ashburner J. A fast diffeomorphic image registration algorithm. *Neuro Image*. 2007; 38:95–113. [PubMed: 17761438]
3. Ashburner J, Friston KJ. Diffeomorphic registration using geodesic shooting and Gauss-Newton optimisation. *Neuro Image*. 2011; 55:954–967. [PubMed: 21216294]
4. Beg MF, Miller MI, Trouné A, Younes L. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int J Comput Vis*. 2005; 61:139–157.
5. Benzi M, Golub GH, Liesen J. Numerical solution of saddle point problems. *Acta Numer*. 2005; 14:1–137.
6. Benzi M, Haber E, Taralli L. A preconditioning technique for a class of PDE-constrained optimization problems. *Adv Comput Math*. 2011; 35:149–173.
7. Biros G, Doğan G. A multilevel algorithm for inverse problems with PDE constraints. *Inverse Problems*. 2008; 24
8. Biros G, Ghattas O. Parallel Lagrange-Newton-Krylov-Schur methods for PDE-constrained optimization. Part I. The Krylov–Schur Solver, SIAM. *J Sci Comput*. 2005; 27:687–713.
9. Biros G, Ghattas O. Parallel Lagrange-Newton-Krylov-Schur methods for PDE-constrained optimization. Part II. The Lagrange–Newton Solver and its application to Optimal Control of Steady Viscous Flows, SIAM. *J Sci Comput*. 2005; 27:714–739.
10. Borzi A, Ito K, Kunisch K. Optimal control formulation for determining optical flow. *SIAM J Sci Comput*. 2002; 24:818–847.
11. Boyd, JP. *Chebyshev and Fourier Spectral Methods*. Dover; Mineola, NY: 2000.
12. Brezzi, F.; Fortin, M., editors. *Mixed and Hybrid Finite Element Methods*. Springer; New York: 1991.
13. Broit, C. PhD thesis. Computer and Information Science, University of Pennsylvania; Philadelphia: 1981. *Optimal Registration of Deformed Images*.
14. Burger M, Modersitzki J, Ruthotto L. A hyperelastic regularization energy for image registration. *SIAM J Sci Comput*. 2013; 35:B132–B148.
15. Byrd RH, Curtis FE, Nocedal J. An inexact SQP method for equality constrained optimization. *SIAM J Optim*. 2008; 19:351–369.
16. Chen K, Lorenz DA. Image sequence interpolation using optimal control. *J Math Imaging Vision*. 2011; 41:222–238.
17. Chen K, Lorenz DA. Image sequence interpolation based on optical flow, segmentation and optimal control. *IEEE Trans Image Process*. 2012; 21:1020–1030. [PubMed: 22167632]
18. Christensen GE, Rabbitt RD, Miller MI. 3D brain mapping using a deformable neuroanatomy. *Phys Med Biol*. 1994; 39:609–618. [PubMed: 15551602]
19. Christensen GE, Rabbitt RD, Miller MI. Deformable templates using large deformation kinematics. *IEEE Trans Image Process*. 1996; 5:1435–1447. [PubMed: 18290061]
20. Dembo RS, Eisenstat SC, Steihaug T. Inexact Newton methods. *SIAM J Numer Anal*. 1982; 19:400–408.
21. Dembo RS, Steihaug T. Truncated-Newton algorithms for large-scale unconstrained optimization. *Math Program*. 1983; 26:190–212.
22. Droske M, Rumpf M. A variational approach to non-rigid morphological registration. *SIAM Appl Math*. 2003; 64:668–687.
23. Dupuis P, Germander U, Miller MI. Variational problems on flows of diffeomorphisms for image matching. *Quart Appl Math*. 1998; 56:587–600.
24. Eisenstat SC, Walker HF. Choosing the forcing terms in an inexact Newton method. *SIAM J Sci Comput*. 1996; 17:16–32.
25. Fischer B, Modersitzki J. Fast diffusion registration. *Contemp Math*. 2002; 313:117–129.
26. Fischer B, Modersitzki J. Curvature based image registration. *J Math Imaging Vision*. 2003; 18:81–85.
27. Fischer B, Modersitzki J. Ill-posed medicine—an introduction to image registration. *Inverse Problems*. 2008; 24:1–16.
28. Frohn-Schauf C, Henn S, Witsch K. Multigrid based total variation image registration. *Comput Vis Sci*. 2008; 11:101–113.

29. Gill, PE.; Murray, W.; Wright, MH. Practical Optimization. Academic Press; Waltham, MA: 1981.
30. Golub, GH.; Van Loan, CF. Matrix Computations. 3. Johns Hopkins University Press; Baltimore: 1996.
31. Gooya A, Pohl KM, Bilello M, Cirillo L, Biros G, Melhem ER, Davatzikos C. GLISTR: Glioma image segmentation and registration. *IEEE Trans Med Imaging*. 2013; 31:1941–1954. [PubMed: 22907965]
32. Gunzburger, MD. Perspectives in flow control and optimization. SIAM; Philadelphia: 2003.
33. Gurtin, ME. *Math Sci Engrg*. Academic Press; New York: 1981. An Introduction to Continuum Mechanics; p. 158
34. Haber E, Ascher UM. Preconditioned all-at-once methods for large, sparse parameter estimation problems. *Inverse Problems*. 2001; 17:1847–1864.
35. Haber E, Ascher UM, Oldenburg D. On optimization techniques for solving nonlinear inverse problems. *Inverse Problems*. 2000; 16:1263–1280.
36. Haber E, Horesh R, Modersitzki J. Numerical optimization for constrained image registration. *Numer Linear Algebra*. 2010; 17:343–359.
37. Haber E, Modersitzki J. Numerical methods for volume preserving image registration. *Inverse Problems*. 2004; 20:1621–1638.
38. Haber E, Modersitzki J. A multilevel method for image registration. *SIAM J Sci Comput*. 2006; 27:1594–1607.
39. Haber E, Modersitzki J. Image registration with guaranteed displacement regularity. *Int J Comput Vis*. 2007; 71:361–372.
40. Hand L, Hipwell JH, Eiben B, Barratt D, Modat M, Ourselin S, Hawkes DJ. A nonlinear biomechanical model based registration method for aligning prone and supine MR breast images. *IEEE Trans Med Imaging*. 2014; 33:682–694. [PubMed: 24595342]
41. Hart, GL.; Zach, C.; Niethammer, M. An optimal control approach for deformable registration. *Proceedings of CVPR, IEEE*; 2009; p. 9-16.
42. Henn S. A multigrid method for a fourth-order diffusion equation with application to image processing. *SIAM J Sci Comput*. 2005; 27:831–849.
43. Hernandez H. Gauss-Newton inspired preconditioned optimization in large deformation diffeomorphic metric mapping. *Phys Med Biol*. 2014; 59:6085–6115. [PubMed: 25254606]
44. Hernandez M, Bossa MN, Olmos S. Registration of anatomical images using paths of diffeomorphisms parameterized with stationary vector field flows. *Int J Comput Vis*. 2009; 85:291–306.
45. Hoge C, Davatzikos C, Biros G. Brain-tumor interaction biophysical models for medical image registration. *SIAM J Sci Comput*. 2008; 30:3050–3072.
46. Horn BKP, Shunck BG. Determining optical flow. *Artificial Intelligence*. 1981; 17:185–203.
47. Kalmoun EM, Garrido L, Caselles V. Line search multilevel optimization as computational methods for dense optical flow. *SIAM J Imaging Sci*. 2011; 4:695–722.
48. Lee E, Gunzburger M. An optimal control formulation of an image registration problem. *J Math Imaging Vision*. 2010; 36:69–80.
49. Lee E, Gunzburger M. Analysis of finite element discretization of an optimal control fomulation of the image registration problem. *SIAM J Numer Anal*. 2011; 49:1321–1349.
50. Loeckx, D.; Maes, F.; Vandermeulen, D.; Suetens, P. Nonrigid image registration using free-form deformations with a local rigidity constraint. Springer; New York: 2004. *Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Comput. Sci* 3216; p. 639-646.
51. Mang A, Toma A, Schuetz TA, Becker S, Eckey T, Mohr C, Petersen D, Buzug TM. Biophysical modeling of brain tumor progression: From unconditionally stable explicit time integration to an inverse problem with parabolic PDE constraints for model calibration. *Med Phys*. 2012; 39:4444–4459. [PubMed: 22830777]
52. Mansi T, Pennec X, Sermesant M, Delingette H, Ayache N. iLogDemons: A Demonsbased registration algorithm for tracking incompressible elastic biological tissues. *Int J Comput Vis*. 2011; 92:92–111.

53. Miller MI. Computational anatomy: Shape, growth and atrophy comparison via diffeomorphisms. *Neuro Image*. 2004; 23:S19–S33. [PubMed: 15501089]
54. Modersitzki, J. *Numerical Methods for Image Registration*. Oxford University Press; New York: 2004.
55. Modersitzki J. FLIRT with rigidity—image registration with a local non-rigidity penalty. *Int J Comput Vis*. 2008; 76:153–163.
56. Modersitzki, J. *FAIR: Flexible Algorithms for Image Registration*. SIAM; Philadelphia, Pennsylvania, US: 2009.
57. Museyko O, Stiglmayr M, Klamroth K, Leugering G. On the application of the Monge-Kantorovich problem to image registration. *SIAM J Imaging Sci*. 2009; 2:1068–1097.
58. Nielsen, M.; Johansen, P.; Jackson, AD.; Laurrup, B. *Brownian warps: A least committed prior for non-rigid registration*. Springer; New York: 2002. *Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Comput. Sci* 2489; p. 557-564.
59. Nocedal, J.; Wright, SJ. *Numerical Optimization*. Springer; New York: 2006.
60. Pennec, X.; Stefanescu, R.; Arsigny, V.; Fillard, P.; Ayache, N. Riemannian elasticity: A statistical regularization framework for non-linear registration. Springer; New York: 2005. *Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Comput. Sci* 3750; p. 943-950.
61. Rohlfing T, Maurer CR, Bluemke DA, Jacobs MA. Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *IEEE Trans Medical Imaging*. 2003; 22:730–741. [PubMed: 12872948]
62. Ruhnau P, Schnörr C. Optical Stokes flow estimation: An imaging-based control approach. *Exp Fluids*. 2007; 42:61–78.
63. Rumpf M, Wirth B. A nonlinear elastic shape averaging approach. *SIAM J Imaging Sci*. 2009; 2:800–833.
64. Sdika M. A fast nonrigid image registration with constraints on the Jacobian using large scale constrained optimization. *IEEE Trans Medical Imaging*. 2008; 27:271–281. [PubMed: 18334448]
65. Sotiras A, Davatzikos C, Paragios N. Deformable medical image registration: A survey. *IEEE Trans Medical Imaging*. 2013; 32:1153–1190. [PubMed: 23739795]
66. Sundar, H.; Davatzikos, C.; Biros, G. *Biomechanically constrained 4D estimation of myocardial motion*. Springer; New York: 2009. *Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Comput. Sci* 5762; p. 257-265.
67. Thirion JP. Image matching as a diffusion process: An analogy with Maxwell’s demons. *Med Image Anal*. 1998; 2:243–260. [PubMed: 9873902]
68. Trouvé A. Diffeomorphism groups and pattern matching in image analysis. *Int J Comput Vis*. 1998; 28:213–221.
69. Vercauteren, T.; Pennec, X.; Perchant, A.; Ayache, N. *Symmetric log-domain diffeomorphic registration: A demons-based approach*. Springer; New York: 2008. *Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Comput. Sci* 5241; p. 754-761.
70. Vercauteren T, Pennec X, Perchant A, Ayache N. Diffeomorphic demons: Efficient non-parametric image registration. *Neuro Image*. 2009; 45:S61–S72. [PubMed: 19041946]
71. Vialard FX, Risser L, Rueckert D, Cotter CJ. Diffeomorphic 3D. image registration via geodesic shooting using an efficient adjoint calculation. *Int J Comput Vis*. 2012; 97:229–241.
72. Vogel, CR. *Computational Methods for Inverse Problems*. SIAM; Philadelphia: 2002.
73. Yanovsky, I.; Thompson, PM.; Osher, S.; Loew, AD. *Topology preserving log-unbiased nonlinear image registration: Theory and implementation*. *Proceedings CVPR, IEEE*; 2007; p. 1-8.
74. Younes L. Jacobi fields in groups of diffeomorphisms and applications. *Quart Appl Math*. 2007; 650:113–134.

Appendix A. Expansion in time: Derivation

This section summarizes modifications of the regularization operator as well as the (incremental) control equation on account of the expansion in time (see section 4.3.2). Inserting (4.10) into (3.1) yields

$$\int_0^1 \mathcal{S}[\mathbf{v}] dt = \frac{\beta}{2} \sum_{l=1}^{n_c} \sum_{l'=1}^{n_c} c_{ll'} \int_{\Omega} \langle \mathcal{B}[\mathbf{v}_l], \mathcal{B}[\mathbf{v}_{l'}] \rangle d\mathbf{x}, \quad \text{where } c_{ll'} := \int_0^1 b_l b_{l'} dt \forall l, l' \in \mathbf{N}. \quad (\text{A.1})$$

Taking first and second variations with respect to the l th expansion coefficient \mathbf{v}_l yields the control equation

$$\beta \sum_{l'=1}^{n_c} c_{ll'} \mathcal{A}[\mathbf{v}_{l'}] + \int_0^1 b_l \mathcal{K}[\mathbf{f}] dt = 0, \quad l=1, \dots, n_c. \quad (\text{A.2})$$

and the incremental control equation

$$\mathcal{H}_l[\tilde{\mathbf{v}}_l] := \beta \sum_{l'=1}^{n_c} c_{ll'} \mathcal{A}[\tilde{\mathbf{v}}_{l'}] + \int_0^1 b_l \mathcal{K}[\tilde{\mathbf{f}}] dt, \quad l=1, \dots, n_c. \quad (\text{A.3})$$

respectively. Accordingly, the operators \mathcal{B} and \mathcal{A} simply act on \mathbf{v}_l instead of \mathbf{v} . The definition for these operators can be found in section 4.1.

We use a global basis on the unit time horizon for the expansion (see section 4.3.2). We use *Chebyshev polynomials* as basis functions in (4.10) on account of their excellent approximation properties as well as their orthogonality. The latter property considerably reduces the computational complexity, since $c_{ll'} = 0$ for all $l, l', l' \neq l$, and $c_{ll} = 1$ (see (A.1), (A.2) and (A.3)). To avoid Runge's phenomenon (see, e.g., [11, p. 82ff.]) *Chebyshev–Gauss–Labatto nodes* are used.

Appendix B. Incompressibility constraint: Elimination

Here, we derive the elimination of p and \bar{p} from the optimality systems for the Stokes regularization scheme. We only consider a quadratic H^1 -regularization for the velocity \mathbf{v} . However, the same line of arguments applies to the H^2 -regularization model.

Applying the divergence to (4.1d) results in¹¹

¹¹Since we discuss the implementation of the incompressibility constraint we set $\gamma = 1$.

$$-\nabla \cdot \beta \Delta \mathbf{v} + \Delta p + \nabla \cdot \mathbf{f} = 0 \quad \text{in } \Omega \times [0, 1].$$

Under the optimality assumption $\nabla \cdot \mathbf{v} = 0$ it follows from the definition of the vectorial Laplacian that $p = -\Delta^{-1}(\nabla \cdot \mathbf{f})$. Inserting this expression into (4.1d) projects \mathbf{v} onto the manifold of divergence-free velocity fields and as such eliminates (4.1c) (assuming that the initial \mathbf{v} is divergence free). Accordingly, we obtain the control equation (reduced gradient)

$$-\beta \Delta \mathbf{v} + \mathcal{K}[\mathbf{f}] = 0, \quad \text{where } \mathcal{K}[\mathbf{f}] := -\nabla(\Delta^{-1}(\nabla \cdot \mathbf{f})) + \mathbf{f}. \quad (\text{B.1})$$

to replace (4.1d). This expression is equivalent to (4.11) in the case when an H^1 -regularization model is used (i.e., $\mathcal{A} = -\Delta$). As stated above, the derivation also holds for the H^2 -regularization operator. We only have to replace $-\beta \Delta \mathbf{v}$ by $\beta \Delta^2 \mathbf{v}$. Computing the second variations of the weak form of the eliminated system yields the incremental control equation

$$-\beta \mathcal{A}[\tilde{\mathbf{v}}] + \mathcal{K}[\tilde{\mathbf{f}}] = -\tilde{\mathbf{g}},$$

where $\tilde{\mathbf{g}}$ is the reduced gradient in (B.1).

Appendix C. Relation to LDDMM

In this section we relate our work to [41] and by that to approaches based on LDDMM [3, 4, 23, 53, 68]. Since the work in [4, 41] is based on first-order information, we only consider the reduced gradient in (4.1d) (setting $\gamma = 0$). In weak form we have

$$\int_0^1 \langle \mathbf{g}, \tilde{\mathbf{v}} \rangle_{L^2(\Omega)} dt = \int_0^1 \langle \beta (\mathcal{B} \mathcal{B}^H)[\mathbf{v}] + \mathbf{f}, \tilde{\mathbf{v}} \rangle_{L^2(\Omega)} dt = \int_0^1 \langle \beta \mathbf{v} + (\mathcal{B} \mathcal{B}^H)^{-1}[\mathbf{f}], \tilde{\mathbf{v}} \rangle_{\mathcal{W}} dt.$$

The expression $\beta \mathbf{v} + (\mathcal{B} \mathcal{B}^H)^{-1}[\mathbf{f}] = \mathbf{v} + (\beta \mathcal{B} \mathcal{B}^H)^{-1}[\mathbf{f}]$ is exactly the gradient in the function space \mathcal{W} that has been used in [41]. This expression yields the *preconditioned gradient descent scheme*

$$\mathbf{v}_{k+1}^h = \mathbf{v}_k^h - \alpha_k ((\beta \mathcal{B}^h \mathcal{B}^{h,H})^{-1}[\mathbf{f}_k^h] - \mathbf{v}_k^h),$$

where $(\beta \mathcal{B}^h \mathcal{B}^{h,H})^{-1}[\mathbf{f}^h]$ is nothing but a Picard iterate (see (4.7)). Subtracting \mathbf{v}_k^h translates this iterate into an update. This is exactly the formulation we have used in this work (see section 4.2.3) so that the considered first-order method is equivalent to the solver used in [41] (under the assumption that $\alpha_k = 1$, i.e., if we neglect the line search). Accordingly, the same line of arguments used in [41] to relate their work to LDDMM apply to our numerical framework.

Appendix D. Measures of deformation regularity

D.1. Deformation map

To visualize the deformation pattern, \mathbf{y} has to be inferred from \mathbf{v} . This can be done by solving

$$\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{v} = \mathbf{v} \text{ in } \Omega \times (0, 1], \quad \mathbf{u} = 0 \text{ in } \Omega \times \{0\}, \quad (\text{D.1})$$

with periodic boundary conditions on $\partial\Omega$. Here, $\mathbf{u} : \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}^d$, $(\mathbf{x}, t) \mapsto \mathbf{u}(\mathbf{x}, t)$, is a displacement field and $\mathbf{y} := \mathbf{x} - \mathbf{u}_1$, $\mathbf{y} : \bar{\Omega} \rightarrow \mathbf{R}^d$, where $\mathbf{u}_1 := \mathbf{u}(\cdot, t=1)$, $\mathbf{u}_1 : \bar{\Omega} \rightarrow \mathbf{R}^d$, $\mathbf{x} \mapsto \mathbf{u}_1(\mathbf{x})$.

Visualization

As can be seen in the visualization of the deformed grids, the mapping \mathbf{y} actually corresponds to the inverse of the deformation map applied to an image. This reflects the fact that our model is formulated in an Eulerian frame of reference. Note that all images reported are high-resolution vector graphics. Zooming in on the digital version of the paper will reveal local properties of the deformation map.

D.2. Deformation gradient

It is well known from calculus that the determinant of the Jacobian matrix $\det(\nabla \mathbf{y})$ can be used to assess invertibility of \mathbf{y} as well as local volume change, provided that $\mathbf{y} \in C^2(\Omega)^d$. In the framework of continuum mechanics, we can obtain this information from the deformation tensor field $\mathbf{F} : \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}^{d \times d}$, where \mathbf{F} is related to \mathbf{v} by

$$\partial_t \mathbf{F} + (\mathbf{v} \cdot \nabla) \mathbf{F} = (\nabla \mathbf{v}) \mathbf{F} \text{ in } \Omega \times (0, 1], \quad \mathbf{F} = \mathbf{I} \text{ in } \Omega \times \{0\}. \quad (\text{D.2})$$

with periodic boundary conditions on $\partial\Omega$. Here, $\mathbf{I} = \text{diag}(1, \dots, 1) \in \mathbf{R}^{d \times d}$; $\det(\mathbf{F}_1)$ is equivalent to $\det(\nabla \mathbf{y})$, where $\mathbf{F}_1 := \mathbf{F}(\cdot, t=1)$, $\mathbf{F}_1 : \bar{\Omega} \rightarrow \mathbf{R}^{d \times d}$, $\mathbf{x} \mapsto \mathbf{F}_1(\mathbf{x})$.

Visualization

We limit the color map for the display of $\det(\mathbf{F}_1)$ to $[0, 2]$. In particular, the color map ranges from black (compression: $\det(\mathbf{F}_1) \in (0, 1)$; black corresponds to values of 0 or below (due to clipping), which represents a singularity or the loss of mass, respectively) to orange (mass conservation: $\det(\mathbf{F}_1) = 1$) to white (expansion: $\det(\mathbf{F}_1) > 1$; white represents values of 2 or greater (due to clipping)).




Figure 1.

Synthetic and real-world registration problems. All images have been normalized to an intensity range of $[0, 1]$. The registration problems are referred to as sinusoidal images (top row, left), UT images (top row, right), hand images (bottom row, left), and brain images (bottom row, right). Each row displays the reference image m_R , the template (deformable) image m_T , and a map of their pointwise difference (from left to right as identified by the inset in the images). We provide an illustration of the deformation pattern \mathbf{y} (overlaid onto m_T) for the synthetic problems. This mapping is computed from \mathbf{v}^\star in (5.1) via (D.1).

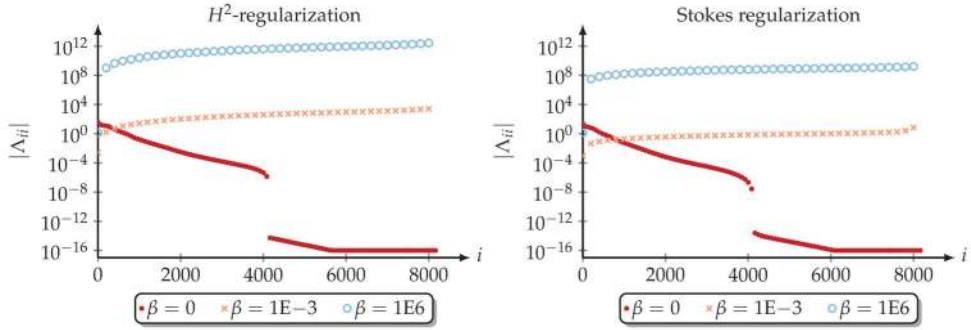


Figure 2.

Trend of the absolute value of the eigenvalues Λ_{ii} , $i = 1, \dots, 8192$, of the reduced Hessian \mathcal{H}^h for plain H^2 -regularization ($\gamma = 0$; left) and the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization; right) for $\beta \in \{0, 1E-2, 1E6\}$ (as indicated in the legend of the plots). We report the trend of the entire set of 8192 eigenvalues. The test problem is the UT images (see section 5.1 for details on the construction of this synthetic registration problem; $\mathbf{n}_x = (64, 64)^\top$ and $n_c = 1$). The Hessian is computed at the true solution \mathbf{v}^h, \star to ensure that $\mathcal{H}^h > 0$ (this statement is confirmed by the values reported in Table 3). The eigenvalues (absolute value) are sorted in descending order for the unregularized problem (i.e., for $\beta = 0$) and in ascending order otherwise (i.e., for $\beta = 1E-3$ and $\beta = 1E6$).




Figure 3. Eigenvector plots of the reduced Hessian $\mathcal{H}^h \in \mathbf{R}^{n \times n}$, $n = 8192$, for $\beta \in \{0, 1E-3, 1E6\}$ for plain H^2 -regularization ($\gamma = 0$; top) and for the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization; bottom). The results correspond to the eigenvalue plots reported in Figure 2. We refer to Figure 2 and the text for details on the experimental setup. Each plot provides the spatial variation of the portion of an eigenvector $\mathbf{v}_i \in \mathbf{R}^n$ associated with the first component of the coefficient field ν_i^h , $I = n_c$, $n_c = 1$. The individual plots correspond to the eigenvalues $\Lambda_{ii} > 0$, $i = 1, 5, 20, 100, 1000$ in Figure 2. The range of the values for ν_i^1 is provided below each plot.




Figure 4. Qualitative comparison of exemplary registration results of the convergence study reported in Table 4. In particular, we display the results for the N-PCG method for a grid size of $\mathbf{n}_x = (256, 256)^T$. We refer to Table 4 and the text for details on the experimental setup. We report results for plain H^2 -regularization ($\gamma = 0$; images to the left) and the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization; images to the right). We display the deformed template image m_1 , a pointwise map of the residual differences between m_R and m_1 (which appears completely white, as the residual differences are extremely small), as well as a pointwise map of the determinant of the deformation gradient $\det(F_1)$ (from left to right as identified by the inset in the images). The values for the $\det(F_1)$ are reported in Table 5. Information on how to interpret these images can be found in Appendix D.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

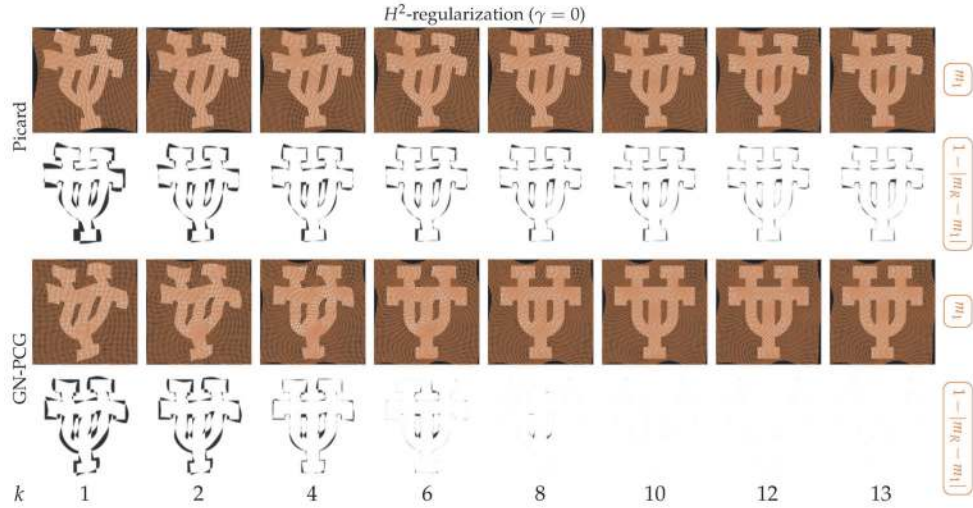


Figure 5. Illustration of the course of the optimization for the Picard (top block) and the GN-PCG (bottom block) methods with respect to the (outer) iteration index k for exemplary results of the convergence study reported in Table 6. We refer to Table 6 and the text for details on the experimental setup. We report results for plain H^2 -regularization ($\gamma = 0$), with an empirically chosen regularization parameter of $\beta = 1E-3$ for images of grid size $\mathbf{n}_X = (256, 256)^T$. We report results until convergence of the GN-PCG method ($k^\star = 13$). We display the deformed template m_1 (top row) and a map of the pointwise difference between m_R and m_1 (bottom row) for both iterative optimization methods (as identified by the inset on the right of the images). Information on how to interpret these images can be found in Appendix D.

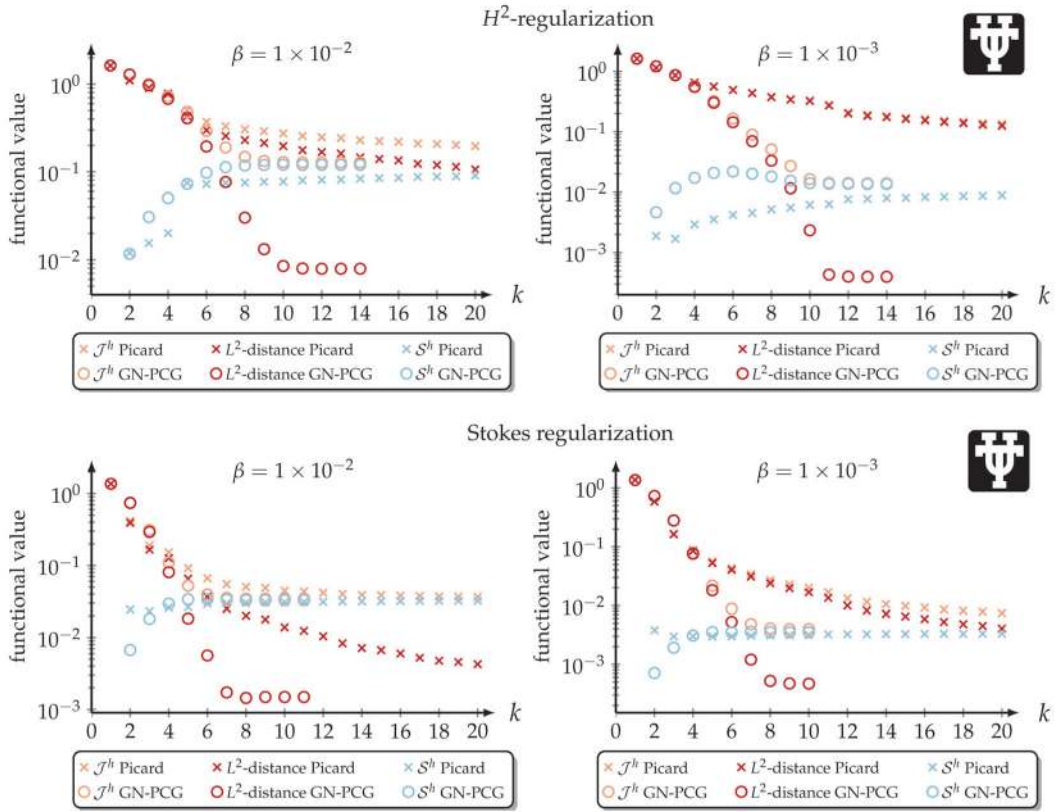


Figure 6.

Trend of the objective \mathcal{J}^h , the L^2 -distance, and the regularization model \mathcal{S}^h (logarithmic scale) for the Picard and the GN-PCG method with respect to the (outer) iteration index k for exemplary results of the convergence study reported in Table 6. We refer to Table 6 and the text for details on the experimental setup. The trend of the functionals is plotted for different (empirically determined) choices of β (left column: $\beta = 1\text{E-}2$; right column: $\beta = 1\text{E-}3$) and a grid size of $\mathbf{n}_x = (256, 256)^\top$. We report results for plain H^2 -regularization ($\gamma = 0$; top row) and the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization; bottom row).




Figure 7.

Relative power spectrum of the individual coefficient fields $v_l^h: \mathbf{R}^d \rightarrow \mathbf{R}^d$, $l = 1, \dots, n_c$, for different choices of n_c used to solve the considered registration problems. We report results for plain H^2 -regularization. The reported results correspond to Table 8. We refer to Table 8 and the text for details on the experimental setup. We report exemplary results for the brain images ($n_x = (200, 200)^T$; left) and the hand images ($n_x = (128, 128)^T$; right). We choose n_c to be in $\{1, 2, 4, 8, 16\}$ as indicated in the legend of each plot. The definition of the relative ℓ^2 -norm (relative power spectrum) of v_l^h can be found in Table 2.

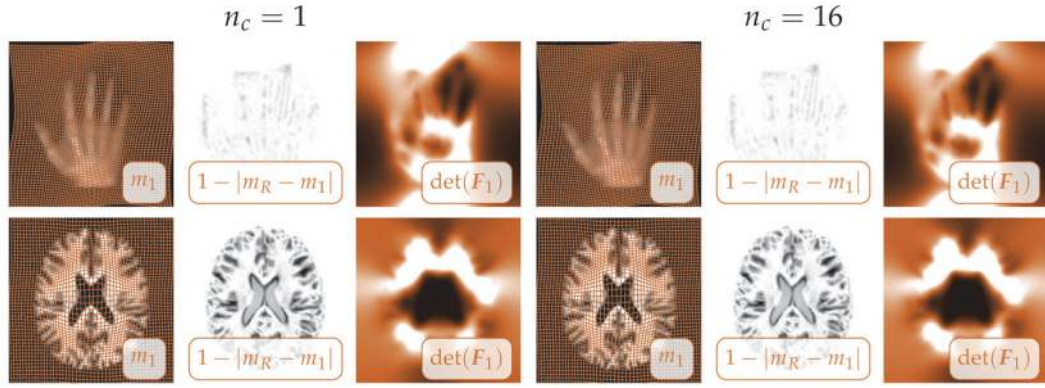


Figure 8.

Qualitative comparison of exemplary registration results for $n_c = 1$ (images to the left) and $n_c = 16$ (images to the right) of the study reported in Table 8. We refer to Table 8 and the text for details on the experimental setup. We report results for the hand images ($\mathbf{n}_x = (128, 128)^\top$) and the brain images ($\mathbf{n}_x = (200, 200)^\top$) for plain H^2 -regularization. We display (for each experiment) the deformed template image m_1 , a pointwise map of the absolute difference between m_R and m_1 , and a map of the determinant of the deformation gradient $\det(\mathbf{F}_1)$ (from left to right as identified by the inset in the images). Information on how to interpret these images can be found in Appendix D.

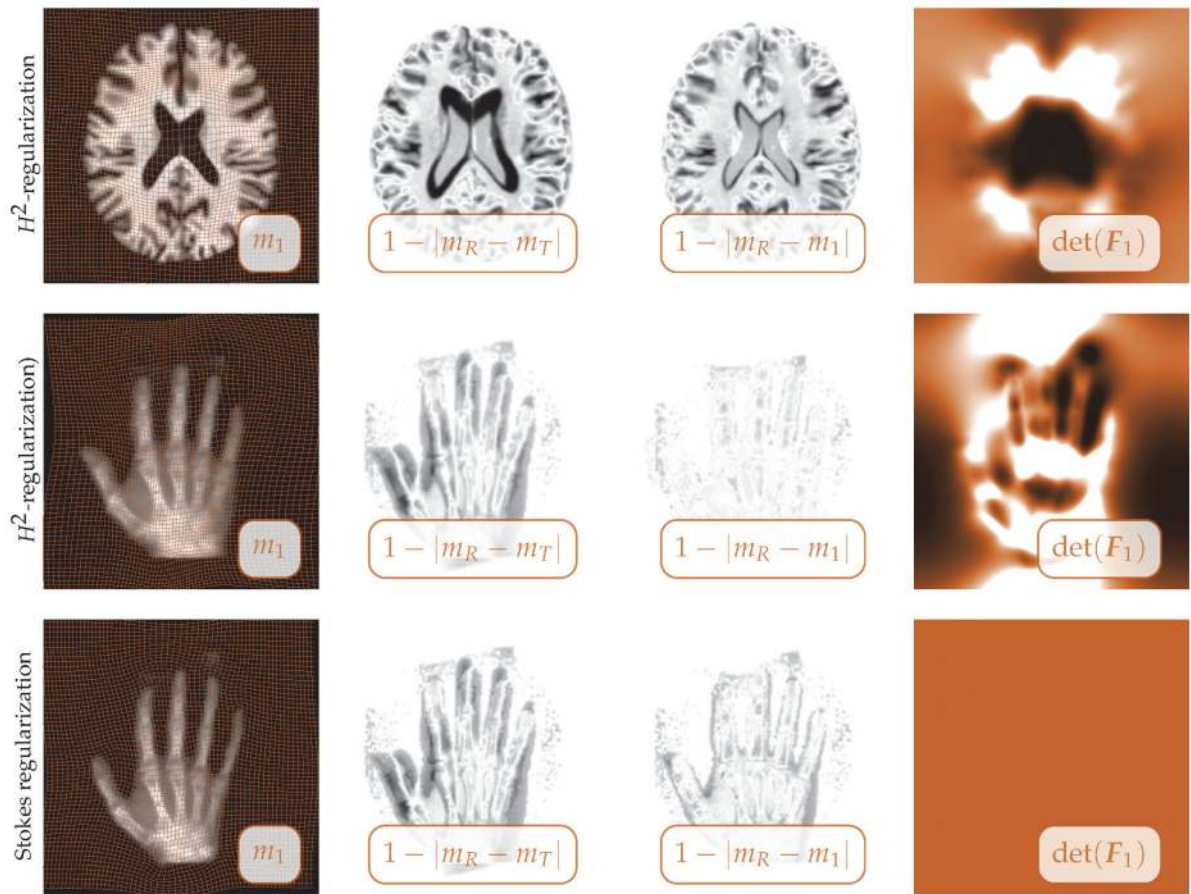


Figure 9.

Qualitative illustration of exemplary registration results of the results for the parameter continuation in β reported in Table 9. We refer to Table 9 and the text for details on the experimental setup. We report results for the brain images (top row; plain H^2 -regularization; $\gamma = 0$) and the hand images (middle row: plain H^2 -regularization ($\gamma = 0$); bottom row: Stokes regularization scheme ($\gamma = 1$; H^1 -regularization)). We display the deformed template image m_1 , a map of the absolute difference between m_R and m_T and between m_R and m_1 and a map of the determinant of the deformation gradient $\det(F_1)$ (from left to right as indicated in the inset in the images). Information on how to interpret these images can be found in Appendix D.

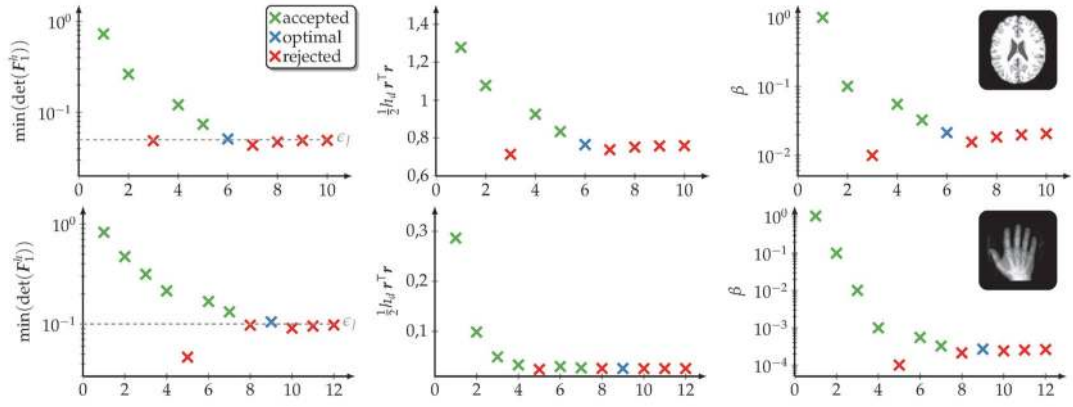


Figure 10. Exemplary illustration of the course of the parameter continuation in β for the quantitative results reported in Table 9. We refer to Table 9 and the text for details on the experimental setup. We report results for the brain images (top row) and the hand images (bottom row) for plain H^2 -regularization. For each experiment, we display (from left to right) (i) the trend of the minimal value of the determinant of the deformation gradient (the dashed line indicates the user-defined lower bound on $\det(\mathbf{F}_1^h)$), (ii) the trend of the L^2 -distance (h_d is the grid cell volume and $\mathbf{r} := m_R^h - m_T^h$, $\mathbf{r} \in \mathbf{R}^{\tilde{n}}$, $\tilde{n} = \prod_{i=1}^2 n_x^i$) and (iii) the trend of β , all with respect to the parameter continuation step. We indicate our judgment on the results in color. That is, if a result is accepted (i.e., $\min(\det(\mathbf{F}_1^h)) \geq \varepsilon_p$) we plot the marker in green, and if a result is rejected (i.e., $\min(\det(\mathbf{F}_1^h)) < \varepsilon_p$) we plot the marker in red. The optimal value is plotted in blue. The plots correspond to the results reported in Table 9.

Table 1

Notation (frequently used acronyms and symbols).

Notation	Description
GN	Gauss–Newton
KKT system	Karush–Kuhn–Tucker system
PCG	preconditioned conjugate gradient method
PDE	partial differential equation
PDE solve	solution of hyperbolic PDEs of optimality systems (4.1) and (4.3)
m_R	reference (fixed) image ($m_R : \mathbf{R}^d \rightarrow \mathbf{R}$)
m_T	template image (image to be registered; $m_T : \mathbf{R}^d \rightarrow \mathbf{R}$)
y	mapping (deformation; $y : \mathbf{R}^d \rightarrow \mathbf{R}^d$)
v	velocity field (control variable; $v : \mathbf{R}^d \times [0, 1] \rightarrow \mathbf{R}^d$)
m	state variable (transported image; $m : \mathbf{R}^d \times [0, 1] \rightarrow \mathbf{R}$)
m_1	state variable at $t = 1$ (deformed template image; $m_1 : \mathbf{R}^d \rightarrow \mathbf{R}$)
λ	adjoint variable (transported mismatch; $\lambda : \mathbf{R}^d \times [0, 1] \rightarrow \mathbf{R}$)
f	body force (drives the registration; $f : \mathbf{R}^d \times [0, 1] \rightarrow \mathbf{R}^d$)
F_1	deformation gradient (tensor field) at $t = 1$ ($F_1 : \mathbf{R}^d \rightarrow \mathbf{R}^{d \times d}$)
\mathcal{J}	objective functional
\mathcal{S}	regularization functional
\mathcal{A}	differential operator (first and second variation of \mathcal{S})
β	regularization parameter
γ	parameter that enables ($\gamma = 1$) or disables ($\gamma = 0$) the incompressibility constraint
n_t	number of time points (discretization)
n_c	number of coefficient fields (spectral Galerkin method in time)
n_x	number of grid points (discretization; $\mathbf{n}_x = (n_x^1, \dots, n_x^d)^\top$)
g	reduced gradient (first variation of Lagrangian with respect to v)
\mathcal{H}	reduced Hessian (second variation of Lagrangian with respect to v)

Table 2

Overview of the quantitative measures that are used to assess the registration performance. We report the number of outer iterations (steps for updating the control variable \mathbf{v}^h) and the number of hyperbolic PDE solves (i.e., how often we have to solve one of the hyperbolic PDEs (4.1a), (4.1b), (4.3a), and (4.3b) that appear in the optimality systems) to assess the work load. We report the relative change of the L^2 -distance, the objective, and the reduced gradient to assess the quality of the inversion. We report values for the determinant of the deformation gradient to assess the regularity of the computed deformation map. We report the relative power spectrum of the coefficients \mathbf{v}_l^h to assess which of the coefficients of the expansion in (4.10) are significant.

Description	Definition
# of required outer iterations	k^*
# of required hyperbolic PDE solves	n_{PDE}
Relative change of L^2 -distance	$\ m_1^h - m_R^h\ _{2,\text{rel}} := \ m_1^h - m_R^h\ _2^2 / \ m_T^h - m_R^h\ _2^2$
Relative change of objective value	$\delta \mathcal{J}_{\text{rel}}^h := \mathcal{J}^h(\mathbf{v}_{k^*}^h) / \mathcal{J}^h(\mathbf{v}_0^h)$
Relative change of reduced gradient	$\ \mathbf{g}^h\ _{\infty,\text{rel}} := \ \mathbf{g}_{k^*}^h\ _{\infty} / \ \mathbf{g}_0^h\ _{\infty}$
Determinant of deformation gradient	$\det(\mathbf{F}_1^h)$
Relative power spectrum of \mathbf{v}_l^h	$\ \mathbf{v}_l^h\ _{2,\text{rel}} := \ \mathbf{v}_l^h\ _2 / \ \{\mathbf{v}_{l'}^h\}_{l'=1}^{n_c}\ _2$

Table 3

Extrema of the eigenvalues Λ_{ji} , $i = 1, \dots, 8192$, of the reduced Hessian reported in Figure 2. We report values for plain H^2 -regularization ($\gamma = 0$; top block) and the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization; bottom block). We refer to Figure 2 and the text for details on the experimental setup. We report the smallest and the largest real part as well as the largest absolute value of the imaginary part of the eigenvalues Λ_{ji} with respect to different choices of the regularization parameter $\beta \in \{0, 1E-2, 1E6\}$.

H^2 -regularization ($\gamma = 0$)			
β	$\min\{(\operatorname{Re}(\Lambda_{ii}))_{i=1}^n\}$	$\max\{(\operatorname{Re}(\Lambda_{ii}))_{i=1}^n\}$	$\max\{(\operatorname{Im}(\Lambda_{ii}))_{i=1}^n\}$
0	-8.35E-15	3.72E1	3.26E-7
1E-3	2.59E-3	4.19E3	0.
1E6	1.30	4.19E12	0.

Stokes regularization ($\gamma = 1$)			
β	$\min\{(\operatorname{Re}(\Lambda_{ii}))_{i=1}^n\}$	$\max\{(\operatorname{Re}(\Lambda_{ii}))_{i=1}^n\}$	$\max\{(\operatorname{Im}(\Lambda_{ii}))_{i=1}^n\}$
0	-3.74E-13	2.48E1	5.92E-6
1E-3	1.00E-3	2.56E1	2.59E-6
1E6	1.29	2.05E9	1.43E-7

Table 4

Quantitative analysis of the convergence of the Picard, N-PCG, and GN-PCG methods. The test problem is the sinusoidal images (see section 5.1 for details on the construction of this synthetic registration problem). We compare convergence results for plain H^2 -regularization ($\gamma = 0$; top block) and the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization; bottom block). We report results for different grid sizes $\mathbf{n}_x = (\mathbf{n}_x^1, \mathbf{n}_x^2)^\top$, $\mathbf{n}_x^i \in \{64, 128, 256\}$, $i = 1, 2$. We invert for a stationary velocity field (i.e., $n_c = 1$). We terminate the optimization if the relative change of the ℓ^∞ -norm of the reduced gradient \mathbf{g}^h is at least three orders of magnitude or if the change in \mathcal{J}^h between 10 successive iterations is below or equal to $1\text{E-}6$ (i.e., the algorithm stagnates). The regularization parameter is empirically set to $\beta = 1\text{E-}3$. The number of the (outer) iterations (k^\star), the number of the hyperbolic PDE solves (n_{PDE}), and the relative change of (i) the L^2 -distance ($\|m_R^h - m_1^h\|_{2,\text{rel}}$), (ii) the objective ($\delta \mathcal{J}_{\text{rel}}^h$), and (iii) the (reduced) gradient ($\|\mathbf{g}^h\|_{\infty,\text{rel}}$) and the average number of the required line search steps $\bar{\alpha}$ are reported. Note that we introduced a memory for the step size into the Picard method (preconditioned gradient descent) to stabilize the optimization (see section 4.3.5 and the description of the results). The definitions for the reported measures are summarized in Table 2.

	\mathbf{n}_x^i	k^\star	n_{PDE}	$\ m_R^h - m_1^h\ _{2,\text{rel}}$	$\delta \mathcal{J}_{\text{rel}}^h$	$\ \mathbf{g}^h\ _{\infty,\text{rel}}$	$\bar{\alpha}$	
H^2 -regularization	Picard	64	30	420	4.78E-3	6.05E-2	7.86E-3	1.72
		128	34	414	4.69E-3	6.04E-2	6.48E-3	1.82
		256	33	414	4.68E-3	6.03E-2	7.42E-3	1.75
	GN-PCG	64	5	78	4.63E-3	6.05E-2	5.57E-5	1.00
		128	4	45	4.59E-3	6.04E-2	5.60E-4	1.00
		256	4	45	4.58E-3	6.03E-2	5.10E-4	1.00
	N-PCG	64	5	63	4.63E-3	6.05E-2	1.94E-4	1.00
		128	5	61	4.57E-3	6.04E-2	2.83E-4	1.00
		256	5	61	4.57E-3	6.03E-2	2.86E-4	1.00
Stokes regularization	Picard	64	57	269	6.61E-4	1.56E-2	2.55E-3	1.68
		128	52	250	5.79E-4	1.55E-2	3.55E-3	1.65
		256	50	233	5.67E-4	1.55E-2	2.31E-3	1.71
	GN-PCG	64	5	88	5.86E-4	1.56E-2	7.30E-5	1.00
		128	5	86	4.87E-4	1.55E-2	8.09E-5	1.00
		256	5	86	4.86E-4	1.55E-2	8.33E-5	1.00

α	$\ g\ _{\text{rel}}$	δ	$\ m_R - m_q\ _{2,\text{rel}}$	n_{PDE}	$\star k$	n_x
1.00	2.60E-4	1.56E-2	5.86E-4	75	5	64
1.00	2.69E-4	1.55E-2	4.87E-4	75	5	128
1.00	2.51E-4	1.55E-2	4.87E-4	75	5	256

N-PCG

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Obtained values for $\det(\mathbf{F}_1^b)$ of exemplary registration results of the convergence study reported in Table 2 with respect to different iterative optimization methods (Picard, N-PCG, and GN-PCG). We refer to Table 4 and the text for details on the experimental setup. We report results for plain H^2 -regularization ($\gamma = 0$; top block) and for the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization; bottom block) for a grid size of $\mathbf{n}_x = (256, 256)^T$.

Table 5

	$\min(\det(\mathbf{F}_1^b))$	$\max(\det(\mathbf{F}_1^b))$	$\text{mean}(\det(\mathbf{F}_1^b))$	$\text{std}(\det(\mathbf{F}_1^b))$
Picard	8.81E-1	1.19	1.00	7.58E-2
GN-PCG	8.83E-1	1.19	1.00	7.56E-2
N-PCG	8.83E-1	1.19	1.00	7.54E-2
Picard	1.00	1.00	1.00	4.55E-12
GN-PCG	1.00	1.00	1.00	4.52E-12
N-PCG	1.00	1.00	1.00	4.52E-12

Table 6

Quantitative analysis of the convergence for the Picard and the GN-PCG method. The test problem is the UT images (see section 5.1 for more details on the construction of this synthetic registration problem). We compare convergence results for plain H^{β} -regularization ($\gamma = 0$; top block) and the Stokes regularization scheme ($\gamma = 1$; bottom block; H^{β} -regularization) for empirically chosen regularization parameters $\beta \in \{1E-2, 1E-3\}$. We report results for different grid sizes $\mathbf{n}_x = (n_x^1, n_x^2)^T$, $n_x^i \in \{64, 128, 256\}$, $i = 1, 2$, $n_t = 4$ $\max(\mathbf{n}_x)$. We invert for a stationary velocity field (i.e., $n_c = 1$). We terminate the optimization if the relative change of the ℓ^{∞} -norm of the reduced gradient \mathbf{g}^h is larger than or equal to three orders of magnitude or if the change in \mathcal{J}^h between 10 successive iterations is below or equal to $1E-6$ (i.e., the algorithm stagnates). We report the number of the (outer) iterations (k^{\star}), the number of the hyperbolic PDE solves (n_{PDE}) and the relative change of (i) the L^2 -distance ($\|m_R^h - m_1^h\|_{2,\text{rel}}$), (ii) the objective ($\delta \mathcal{J}_{\text{rel}}^h$), and (iii) the (reduced) gradient ($\|\mathbf{g}^h\|_{\infty,\text{rel}}$) as well as the average number of the line search steps $\bar{\alpha}$. Note that we introduced a memory for the step size into the Picard method to stabilize the optimization (see section 4.3.5 and the description of the results). The definitions for the reported measures can be found in Table 2. This study directly relates to the results for the smooth registration problem (see section 5.3.2, in particular Table 4).

	n_x^i	k^{\star}	n_{PDE}	$\ m_R^h - m_1^h\ _{2,\text{rel}}$	$\delta \mathcal{J}_{\text{rel}}^h$	$\ \mathbf{g}^h\ _{\infty,\text{rel}}$	$\bar{\alpha}$		
H^{β} -regularization	$\beta = 1E-2$	Picard	64	130	752	1.99E-2	1.18E-1	2.06E-2	1.68
			128	231	1589	8.29E-3	8.47E-2	4.03E-2	1.72
			256	388	3022	5.05E-3	7.92E-2	9.86E-2	1.68
		GN-PCG	64	7	282	1.94E-2	1.16E-1	3.30E-4	1.00
			128	9	450	8.09E-3	8.37E-2	6.07E-4	1.00
			256	13	789	4.87E-3	7.85E-2	7.28E-4	1.00
	$\beta = 1E-3$	Picard	64	339	4410	1.84E-3	1.54E-2	2.46E-2	1.56
			128	466	9671	7.19E-4	9.84E-3	5.98E-2	1.64
			256	632	8690	5.35E-4	8.81E-3	1.08E-1	1.64
		GN-PCG	64	7	670	1.44E-3	1.50E-2	2.59E-4	1.00
			128	8	929	4.47E-4	9.60E-3	4.76E-4	1.00
			256	13	1744	2.45E-4	8.55E-3	4.04E-4	1.00
Stokes regularization	$\beta = 1E-2$	64	92	448	2.73E-3	3.62E-2	7.54E-3	1.64	
		128	136	958	8.92E-4	2.38E-2	1.63E-2	1.72	
		256	143	1004	1.09E-3	2.53E-2	1.75E-2	1.68	

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

	n_x	k^*	n_{PDE}	$\ m_R^h - m_{I_1}^h\ _{2,rel}$	δ	$\ g^h\ _{\infty,rel}$	$\bar{\alpha}$
GN-PCG	64	7	313	2.72E-3	3.59E-2	4.70E-4	1.00
	128	9	437	8.88E-4	2.37E-2	5.59E-4	1.00
	256	10	514	1.09E-3	2.52E-2	5.49E-4	1.00
Picard	64	216	2823	9.99E-4	4.69E-3	1.34E-2	1.56
	128	179	5465	3.57E-4	2.77E-3	1.79E-2	1.69
	256	175	5569	5.27E-4	3.07E-3	2.22E-2	1.68
GN-PCG	64	8	1162	7.57E-4	4.55E-3	8.23E-4	1.00
	128	9	1069	2.15E-4	2.65E-3	5.67E-4	1.00
	256	9	769	3.43E-4	2.92E-3	6.44E-4	1.00

$\beta = 1E-3$

Table 7

Values for the determinant of the deformation gradient $\det(\mathbf{F}_1^h)$ for exemplary results of the convergence study reported in Table 6. We report results for the Picard and the GN-PCG method. We refer to Table 6 and the text for details on the experimental setup. We report results for the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization). The regularization parameter is set to $\beta = 1\text{E-}3$. The grid size is $\mathbf{n}_X = (256, 256)^\top$. These results directly relate to those reported for the smooth registration problem (see section 5.3.2, in particular Table 5).

	$\min(\det(\mathbf{F}_1^h))$	$\max(\det(\mathbf{F}_1^h))$	$\text{mean}(\det(\mathbf{F}_1^h))$	$\text{std}(\det(\mathbf{F}_1^h))$
Picard	10.00E-1	1.00	1.00	2.03E-5
GN-PCG	10.00E-1	1.00	1.00	2.29E-5

Table 8

Comparison of the inversion results for the GN-PCG method for a varying number of the spatial coefficient fields $v_l^h; \mathbf{R}^d \rightarrow \mathbf{R}^d, l = 1, \dots, n_c$ (i.e., we change the number of the unknowns in time). We report results for plain \mathcal{H}^2 -regularization. We consider the hand images ($n_x = (128, 128)^\top$; top block) and the brain images ($n_x = (200, 200)^\top$; bottom block). We consider the full set of stopping conditions in (4.8) with $\tau \mathcal{J} = 1\text{E-}3$, as we no longer study grid convergence and/or compare different optimization methods. We report the number of the (outer) iterations (K^*), the number of the hyperbolic PDE solves (n_{PDE}) and the relative change of (i) the L^2 -distance ($\|m_R^h - m_1^h\|_{2,\text{rcd}}$), (ii) the objective ($\delta \mathcal{J}_{\text{rcd}}^h$), and (iii) the (reduced) gradient ($\|g^h\|_{\infty,\text{red}}$) as well as the minimal and maximal values of the determinant of the deformation gradient. The definitions of these measures can be found in Table 2. The number of the coefficient fields n_c used to solve the individual registration problems is chosen to be in $\{1, 2, 4, 8, 16\}$.

	n_c	k	n_{PDE}	$\ m_R^h - m_1^h\ _{2,\text{rcd}}$	$\delta \mathcal{J}_{\text{rcd}}^h$	$\ g^h\ _{\infty,\text{red}}$	$\min(\det(\mathbf{F}_1^h))$	$\max(\det(\mathbf{F}_1^h))$
Hand images	1	7	279	6.65E-2	8.52E-2	2.27E-2	2.14E-1	6.60
	2	7	279	6.60E-2	8.48E-2	2.29E-2	2.15E-1	6.44
	4	7	283	6.41E-2	8.24E-2	2.29E-2	2.07E-1	6.49
	8	7	277	6.41E-2	8.24E-2	2.51E-2	2.08E-1	6.46
	16	7	281	6.41E-2	8.24E-2	2.56E-2	2.08E-1	6.45
Brain images	1	20	669	5.49E-1	6.68E-1	3.71E-2	4.93E-2	6.47
	2	20	667	5.47E-1	6.66E-1	3.72E-2	4.96E-2	6.45
	4	21	710	5.31E-1	6.51E-1	3.66E-2	4.55E-2	7.33
	8	21	710	5.30E-1	6.51E-1	3.52E-2	4.53E-2	7.37
	16	21	708	5.30E-1	6.51E-1	3.54E-2	4.53E-2	7.37

Table 9

Quantitative analysis of the parameter continuation in β . We report results for the hand images and the brain images for different regularization schemes (see Figure 1). The spatial grid size for the images is $\mathbf{n}_x = (512, 512)^\top$. The number of the time points n_t is chosen adaptively (see section 4.3.5 for details). We use the full set of stopping conditions (see section 4.2.4) with a tolerance of $\tau_{\mathcal{J}} = 1\text{E-}3$. We report results for plain H^1 - and H^2 -regularization ($\gamma = 0$; top block) as well as for the Stokes regularization scheme ($\gamma = 1$; H^1 -regularization; bottom block). We invert for a stationary velocity field (i.e., $n_c = 1$). We report (i) the considered lower bound on the deformation gradient ($\varepsilon_{\mathcal{F}}$) or the grid angle (ε_{θ}), (ii) the number of the required estimation steps, (iii) the minimal value for the deformation gradient for the optimal regularization parameter, (iv) the computed optimal value for $\beta(\beta^*)$, (v) the minimal change in $\beta(\delta\beta_{\min})$, as well as (vi) the relative change in the L^2 -distance ($\|m_R^h - m_1^h\|_{2,\text{rel}}$).

Smoothness regularization							
Data	\mathcal{S}	ε_r	steps	$\min(\det(\mathbf{F}_1^h))$	β^*	$\delta\beta_{\min}$	$\ m_R^h - m_1^h\ _{2,\text{rel}}$
Brain images	H^1	5E-2	9	5.09E-2	3.11E-1	5.00E-3	7.01E-1
	H^2	5E-2	10	5.11E-2	2.13E-2	5.00E-4	5.52E-1
Hand images	H^1	1E-1	10	1.13E-1	3.39E-2	5.00E-4	8.74E-2
	H^2	1E-1	12	1.05E-1	2.69E-4	5.00E-6	6.83E-2
Stokes regularization							
Data	\mathcal{S}	ε_{θ}	steps	$\min(\det(\mathbf{F}_1^h))$	β^*	$\delta\beta_{\min}$	$\ m_R^h - m_1^h\ _{2,\text{rel}}$
Hand images	H^1	$\pi/16$	10	10.00E-1	2.13E-2	5.00E-4	1.17E-1