Published in final edited form as:

SIAM J Sci Comput. 2017 ; 39(6): B1064–B1101. doi:10.1137/16M1070475.

A SEMI-LAGRANGIAN TWO-LEVEL PRECONDITIONED NEWTON-KRYLOV SOLVER FOR CONSTRAINED DIFFEOMORPHIC IMAGE REGISTRATION

Andreas Mang* and George Biros*

*The Institute for Computational Engineering and Sciences, The University of Texas at Austin,
Austin, Texas, 78712-0027, US

Abstract

We propose an efficient numerical algorithm for the solution of diffeomorphic image registration problems. We use a variational formulation constrained by a partial differential equation (PDE), where the constraints are a scalar transport equation.

We use a pseudospectral discretization in space and second-order accurate semi-Lagrangian time stepping scheme for the transport equations. We solve for a stationary velocity field using a preconditioned, globalized, matrix-free Newton-Krylov scheme. We propose and test a two-level Hessian preconditioner. We consider two strategies for inverting the preconditioner on the coarse grid: a nested preconditioned conjugate gradient method (exact solve) and a nested Chebyshev iterative method (inexact solve) with a fixed number of iterations.

We test the performance of our solver in different synthetic and real-world two-dimensional application scenarios. We study grid convergence and computational efficiency of our new scheme. We compare the performance of our solver against our initial implementation that uses the same spatial discretization but a standard, explicit, second-order Runge-Kutta scheme for the numerical time integration of the transport equations and a single-level preconditioner. Our improved scheme delivers significant speedups over our original implementation. As a highlight, we observe a $20\times$ speedup for a two dimensional, real world multi-subject medical image registration problem.

Keywords

Newton; Krylov method; semi-Lagrangian formulation; KKT preconditioners; constrained diffeomorphic image registration; stationary velocity field registration; optimal control; PDE constrained optimization

AMS subject classifications

68U10; 49J20; 35Q93; 65K10; 65F08; 76D55

1. Introduction

Image registration finds numerous applications in image analysis and computer vision [45,67]. Image registration establishes *meaningful* spatial correspondence between two images $m_R: \bar{\Omega} \rightarrow \mathbf{R}$ (the “reference image”) and $m_T: \bar{\Omega} \rightarrow \mathbf{R}$ (the “template image”) of a scene such that the deformed template image $m_T \circ \mathbf{y}$ becomes similar to m_R , i.e., $m_T \circ \mathbf{y} \approx m_R$ [61]; the images are defined on an open set $\Omega \subset \mathbf{R}^d$, $d \in \{2, 3\}$, with closure $\bar{\Omega} := \Omega \cup \partial\Omega$ and boundary $\partial\Omega$. $\circ$ denotes the composition of two functions, and $\mathbf{y}: \bar{\Omega} \rightarrow \bar{\Omega}$ is the sought after deformation map. There exist various approaches to image registration; we refer to [33,61,67] for a lucid overview.

Image registration is typically formulated as a variational optimization problem with an objective functional that consists of a data fidelity term and a Tikhonov-type regularization norm [3]; the unregularized problem is ill-posed. Here, we follow up on our preceding work on *constrained diffeomorphic image registration* [58, 59]. In diffeomorphic image registration we require that the map $\mathbf{y}$ is a *diffeomorphism*, i.e., $\mathbf{y}$ is a bijection, continuously differentiable, and has a continuously differentiable inverse. Formally, we require that $\det \nabla \mathbf{y} > 0$, $\forall \mathbf{y} \in \mathbf{R}^{d \times d}$, $\forall \mathbf{x} \in \Omega$ and—under the assumption that $\mathbf{y}$ is orientation preserving— $\det \nabla \mathbf{y} > 0$, $\forall \mathbf{x} \in \Omega$.

Different approaches to guarantee a diffeomorphic $\mathbf{y}$ have appeared in the past. One approach is to penalize $\det \nabla \mathbf{y}$ as done in [22, 30, 42, 43, 65]. Another approach is to change the formulation; instead of inverting directly for the deformation map $\mathbf{y}$, we invert for its velocity $\mathbf{v} = d\mathbf{y}$. If $\mathbf{v}$ is sufficiently smooth it can be guaranteed that the resulting $\mathbf{y}$ is a diffeomorphism [10,31,70]. In our formulation, we augment this type of smoothness regularization by constraints on the divergence of $\mathbf{v}$ [58, 59]. For instance, for $\nabla \cdot \mathbf{v} = 0$ the flow becomes incompressible. This is equivalent to enforcing $\det \nabla \mathbf{y} = 1$ [40, pages 77ff.].

Velocity field formulations for diffeomorphic image registration can be distinguished between approaches that invert for a *time dependent* $\mathbf{v}$ [10, 31, 46, 58] and approaches that invert for a *stationary* $\mathbf{v}$ [48, 57, 59]. We invert for a stationary $\mathbf{v}$. We formulate the diffeomorphic image registration problem as a PDE constrained optimization problem, where the constraint is a transport equation for the scalar field $m: \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}$ (the image intensities). Due to ill-conditioning, non-convexity, large-problem size, infinite-dimensional structure, and the need for adjoint operators, such problems are challenging to solve. We use a reduced space Newton–Krylov method [58]. In reduced space methods we eliminate state variables (in our case the transported image) and iterate in the control variable space (in our case the velocity space). Newton methods typically display faster convergence than gradient-descent methods (see [58]). Using a Newton method, however, requires solving linear systems with the reduced space Hessian, which—upon discretization—is a large, dense, and ill-conditioned operator. Efficient preconditioning is critical for making our solver effective across a wide spectrum of image resolutions, regularization weights, and inversion tolerances. Standard preconditioning techniques like incomplete factorization cannot be applied since we do not have access to the matrix entries (too expensive to compute). Instead, we present a matrix-free, two-level preconditioner for the reduced space Hessian

that significantly improves performance. Another computational challenge of our formulation is that the reduced space formulation requires the exact solution of two hyperbolic transport equations—the state and adjoint equations of our problem—every time we evaluate the reduced gradient or apply the reduced space Hessian operator. We introduce a semi-Lagrangian formulation to further speed up our solver.

1.1. Outline of the Method

We are given two functions $m_R: \bar{\Omega} \rightarrow \mathbf{R}$ (fixed image) and $m_T: \bar{\Omega} \rightarrow \mathbf{R}$ (deformable image) compactly supported on an open set $\Omega := (-\pi, \pi)^d$, $d \in \{2, 3\}$, with boundary $\partial\Omega$, and closure $\bar{\Omega} := \Omega \cup \partial\Omega$. We solve for a stationary velocity field $\mathbf{v} \in \mathcal{U}$ and a mass source as follows [59]:

$$\min_{m, \mathbf{v}, w} \frac{1}{2} \|m_R - m_1\|_{L^2(\Omega)}^2 + \frac{\beta_v}{2} \|\mathbf{v}\|_{\mathcal{V}}^2 + \frac{\beta_w}{2} \|w\|_{\mathcal{W}}^2 \quad (1a)$$

subject to

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

$$m = m_T \quad \text{in } \Omega \times \{0\}, \quad (1c)$$

$$\nabla \cdot \mathbf{v} = w \quad \text{in } \Omega, \quad (1d)$$

and periodic boundary conditions on $\partial\Omega$. In our formulation $m_1(\mathbf{x}) := m(\mathbf{x}, t=1)$ —i.e., the solution of the hyperbolic transport equation (1b) with initial condition (1c)—is equivalent to $m_T \circ \mathbf{y}$; the deformation map $\mathbf{y}$ can be computed from $\mathbf{v}$ in a post-processing step (see, e.g., [58, 59]). The weights $\beta_v > 0$, and $\beta_w > 0$ control the regularity of $\mathbf{v}$.

The regularization norm for $\mathbf{v}$ not only alleviates issues related to the ill-posedness of our problem but also ensures the existence of a diffeomorphism $\mathbf{y}$ parameterized by $\mathbf{v}$ if chosen appropriately. The constraint in (1d) allows us to control volume change; setting $w = 0$ results in an incompressible diffeomorphism $\mathbf{y}$, i.e., the deformation gradient $\det \nabla \mathbf{y}$ is fixed to one for all $\mathbf{x} \in \Omega$. The deformation map $\mathbf{y}$ is no longer incompressible if we allow w to deviate from zero (this formulation has originally been introduced in [59]; a similar formulation can be found in [19]). We can control this deviation with β_w ; the regularization norm for w acts like a penalty on $\nabla \cdot \mathbf{v}$. We will specify and discuss the choices for the spaces $\mathcal{U}$, $\mathcal{V}$, and $\mathcal{W}$ in more detail §2 and §A.

We use the method of Lagrange multipliers to solve (1). We first formally derive the optimality conditions and then discretize using a pseudospectral discretization in space with a Fourier basis (i.e., we use an *optimize-then-discretize* approach; see §3). We solve for the

first-order optimality conditions using a globalized, matrix-free, preconditioned, inexact Newton–Krylov algorithm for the velocity field $\mathbf{v}$ (see [58] for details). The hyperbolic transport equations are solved via a semi-Lagrangian method.

1.2. Contributions

Our Newton-Krylov scheme has originally been described in [58], in which we compared it to gradient-descent approach in the Sobolev space induced by the regularization operator (the latter approach is, e.g., used in [46]). The latter, as expected, is extremely slow and not competitive with (Gauss–)Newton schemes. In [59] we introduced and studied different regularization functionals, and compared the performance of our method against existing approaches for diffeomorphic image registration, in particular the Demons family of algorithms [71, 72]. Here we extend our preceding work in the following ways:

- We propose a semi-Lagrangian formulation for our entire optimality system, i.e., the state, adjoint, and incremental state and adjoint equations. We compare it with an stabilized Runge–Kutta method, which we also introduce here; we show that the semi-Lagrangian scheme has excellent stability properties.
- We introduce an improved preconditioner for the reduced Hessian system. It is a two-level pre-conditioner that uses spectral restriction and prolongation operators and a Chebyshev stationary iterative method for an approximate coarse grid solve.
- We provide an experimental study of the performance of our improved numerical scheme based on synthetic and real-world problems. We study self-convergence, grid convergence, numerical accuracy, convergence as a function of the regularization parameter, and the time to solution. We account for different constraints and regularization norms.

Taken together, the new algorithm results in order of magnitude speedups over the state-of-the-art. For example, for a magnetic resonance image of a brain with 512^2 resolution the new scheme is $18\times$ faster (see Tab. 8 in §4) than the scheme described in [59].

1.3. Limitations and Unresolved Issues

Several limitations and unresolved issues remain. We assume similar intensity statistics for m_R and m_T . This is a common assumption in many deformable registration algorithms. For multimodal registration problems we have to replace the squared L^2 -distance in (1a) with more involved distance measure; examples can be found in [61, 67]. We present results only for $d=2$. Nothing in our formulation and numerical approximation is specific to the two-dimensional case. In this work we discuss improvements of the algorithm used in our preceding work [58,59] en route to an effective three-dimensional solver. Once this three-dimensional solver is available, we will extend the study presented in [58], by providing a detailed comparison of our method against diffeomorphic image registration approaches of other groups in terms of efficiency and inversion accuracy.

1.4. Related Work

The body of literature on diffeomorphic image registration, numerical optimization in optimal control, preconditioning of KKT systems, and the effective solution of hyperbolic transport equations is extensive. We limit the discussion to work that is most relevant to ours.

1.4.1. Diffeomorphic Image Registration—Lucid overviews for image registration can be found in [33, 62,67]. Related work on velocity field based diffeomorphic image registration is discussed in [5,6,10,19,46, 55,58,59] and references therein. Related optimal control formulations for image registration are described in [9, 13, 19, 25, 55, 58, 59, 64, 66, 73]. Most work on velocity based diffeomorphic registration considers first order information for numerical optimization (see, e.g., [10,19,23,25,46,55,73]), with the exceptions of our own work [58, 59] and [6, 13, 47, 66]; only [13, 66] discuss preconditioning strategies (see also below). The application of a Newton–Krylov solver for incompressible and near-incompressible formulations (with an additional control on a mass-source term) for diffeomorphic image registration is, to the best of our knowledge, exclusive to our group [58,59].

1.4.2. PDE Constrained Optimization—There exists a huge body of literature for the numerical solution of PDE constrained optimization problems. The numerical implementation of an efficient solver is, in many cases, tailored towards the nature of the control problem, e.g., by accounting for the type and structure of the PDE constraints; see for instance [1, 15] (elliptic), [2, 37, 60, 69] (parabolic), or [13, 19, 55] (hyperbolic). We refer to [14, 21, 39, 49, 51] for an overview on theoretical and algorithmic developments in optimal control and PDE constrained optimization. A survey on strategies for preconditioning saddle point problems can be found in [12]. We refer to [20] for an overview on multigrid methods for optimal control problems.

Our preconditioner can be viewed as a simplified two level multigrid v-cycle with a smoother based on the inverse regularization operator and the coarse grid solve is inexact. We note, that more sophisticated multigrid preconditioners for the reduced Hessian exist [2,19]. Multigrid approaches have been considered in [19] for optical flow and in [13, 66] for the Monge-Kantorovich functional. The work of [19] is the most pertinent to our problem. It is a space-time multigrid in the full KKT conditions and the time discretization scheme is CFL restricted, and, thus, very expensive. The effectiveness of the smoother depends on the regularization functional—it is unclear how to generalize it to incompressible velocities. Our scheme is simpler to implement, supports general regularizations, and is compatible with our semi-Lagrangian time discretization. The preconditioner in [13] is a block triangular preconditioner based on a perturbed representation of the GN approximation of the full space KKT system. A similar preconditioner that operates on the reduced space Hessian is described in [66]. In some sense, we do not approximate the structure of our Hessian operator; we invert an exact representation. We amortize the associated costs as follows: (i) we solve for the action of the inverse inexact, and (ii) we invert the operator on a coarser grid.

1.4.3. The Semi-Lagrangian Method—We refer to [32] for a summary on solvers for advection dominated systems. Example implementations for the solution of hyperbolic transport equations that have been considered in the work cited above are implicit Lax-Friedrich schemes [13, 66], explicit high-order total variation diminishing schemes [19,25,46], or explicit, pseudospectral (in space) RK2 schemes [58,59]. These schemes suffer either from numerical diffusion and/or CFL time step restrictions. We use a high-order, unconditionally stable semi-Lagrangian formulation. Semi-Lagrangian methods are well established and have first been considered in numerical weather prediction [68]. The use of semi-Lagrangian schemes is not new in the context of diffeomorphic image registration. However, such schemes have only been used to solve for the deformation map and/or solve the forward problem [10, 23, 25, 48] and for the adjoint problem in the context of approximate gradient-descent methods.

1.5. Organization and Notation

We summarize our notation in Tab. 1. We summarize the optimal control formulation for diffeomorphic image registration in §2. We describe the solver in §3. We provide the optimality system and the Newton step in §3.1. We describe the discretization in §3.2. The schemes for integrating the hyperbolic PDEs that appear in our formulation are discussed in §3.3. We describe our Newton–Krylov solver in §3.4; this includes a discussion of the preconditioners for the solution of the reduced space KKT system. We provide numerical experiments in §4. We conclude with §5.

2. Optimal Control Formulation

We consider a PDE constrained formulation, where the constraints consist of a scalar transport equation for the image intensities. We solve for a stationary velocity field $\boldsymbol{\nu} \in \mathcal{U}$ and a mass-source $w \in \mathcal{W}$ as follows [59]:

$$\min_{m, \boldsymbol{\nu}, w} \mathcal{J}[\boldsymbol{\nu}, w] = \frac{1}{2} \|m_1 - m_R\|_{L^2(\Omega)}^2 + \frac{\beta_v}{2} \|\boldsymbol{\nu}\|_{\mathcal{V}}^2 + \frac{\beta_w}{2} \|w\|_{\mathcal{W}}^2 \quad (2a)$$

subject to

$$\partial_t m + \nabla m \cdot \boldsymbol{\nu} = 0 \quad \text{in } \Omega \times (0, 1], \quad (2b)$$

$$m = m_T \quad \text{in } \Omega \times \{0\}, \quad (2c)$$

$$\nabla \cdot \boldsymbol{\nu} = w \quad \text{in } \Omega \quad (2d)$$

and periodic boundary conditions on Ω . We measure the distance between the reference image m_R and the deformed template image m_1 via a squared L^2 -distance. The contributions

of the regularization models for w and $\mathbf{v}$ are controlled by the weights $\beta_v > 0$ and $\beta_w > 0$, respectively. We consider an H^1 -regularization norm for w , i.e.,

$$\|w\|_{H^1(\Omega)}^2 := \int_{\Omega} \nabla w \cdot \nabla w + w^2 dx. \quad (3)$$

We consider three quadratic regularization models for $\mathbf{v}$; an H^1 -, an H^2 -, and an H^3 -seminorm:

$$|\mathbf{v}|_{H^1(\Omega)^d}^2 := \int_{\Omega} \nabla \mathbf{v} : \nabla \mathbf{v} dx, \quad |\mathbf{v}|_{H^2(\Omega)^d}^2 := \int_{\Omega} \Delta \mathbf{v} \cdot \Delta \mathbf{v} dx, \quad \text{and} \quad |\mathbf{v}|_{H^3(\Omega)^d}^2 := \int_{\Omega} \nabla \Delta \mathbf{v} : \nabla \Delta \mathbf{v} dx.$$

(4)

The use of an H^1 -seminorm is motivated by related work in computational fluid dynamics; we will see that the first order variations of our formulation will result in a system that reflects a linear Stokes model under the assumption that we enforce $\nabla \cdot \mathbf{v} = 0$ [25, 58, 59, 64]. We use an H^2 -seminorm if we neglect the incompressibility constraint (2d). This establishes a connection to related formulations for diffeomorphic image registration [10,46,48]; an H^2 -norm is the paramount model in many algorithms (or its approximation via its Green's function; a Gaussian kernel) [10].

Remark 1

The norm on w acts like a penalty on $\nabla \cdot \mathbf{v}$. In fact, we can eliminate (2d) from (2) by inserting $\nabla \cdot \mathbf{v}$ for w into the regularization norm in (2a). If we neglect the incompressibility constraint (2d) the space $\mathcal{U}$ for $\mathbf{v}$ is given by the Sobolev space $\mathcal{V}$ (this formulation is, e.g., used in [46] for a non-stationary velocity with H^2 -regularity in space and L^2 regularity in time). If we set w in (2d) to zero, the computed velocity will be in the space of divergence free velocity fields with Sobolev regularity in space, as defined by $\mathcal{V}$ (examples for this formulation can be found in [25, 58, 59, 64]). For a non-zero w we additionally require that the divergence of $\mathbf{v}$ is in $\mathcal{W}$. An equivalent formulation is, e.g., presented in [18,19]. They use H^1 -regularity for $\mathbf{v}$ and stipulate L^2 -regularity for its divergence, and proof existence of the state and adjoint variables for smooth images [19]. In particular, they provide existence results for a unique, H^1 -regular solution of the forward problem under the assumption of H^1 -regularity for the template image. The same regularity requirements hold true for the adjoint equation. In our formulation, we not only require $\mathbf{v}$ to be an H^1 -function, but also that its divergence is in H^1 (according to (3)). Another approach to impose regularity on $\mathbf{v}$ is to not only control the divergence but also control its curl (see, e.g., [4,55]). We provide additional remarks in §A.

3. Numerics and Solver

In what follows, we describe our numerical solver for computing a discrete approximation to the continuous problem. We use a globalized, preconditioned, inexact, reduced space¹

(Gauss–)Newton–Krylov method. Our scheme is described in detail in [58]. We will briefly recapitulate the key ideas and main building blocks.

We use the (formal) Lagrangian method [56] to solve (2); the Lagrangian functional $\mathcal{L}$ is given by

$$\mathcal{L}[\phi] := \mathcal{J}[\mathbf{v}, w] + \int_0^1 \langle \partial_t m + \mathbf{v} \cdot \nabla m, \lambda \rangle_{L^2(\Omega)} dt + \langle m_0 - m_T, \mathbf{v} \rangle_{L^2(\Omega)} - \langle \nabla \cdot \mathbf{v} - w, p \rangle_{L^2(\Omega)} \quad (5)$$

with $\phi := (m, \lambda, p, w, \mathbf{v})$ and Lagrange multipliers $\lambda \times \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}$ for the hyperbolic transport equation (2b), $v: \bar{\Omega} \rightarrow \mathbf{R}$ for the initial condition (2c), and $p: \bar{\Omega} \rightarrow \mathbf{R}$ for the incompressibility constraint (2d) (we neglect the periodic boundary conditions for simplicity). The Lagrange multiplier functions inherit the boundary conditions of the forward operator.

Remark 2

We can consider two numerical strategies to tackle (2). We can either use an *optimize-then-discretize* approach or a *discretize-then-optimize* approach. We choose the former, i.e., we compute variations of the continuous problem and then discretize the optimality system. In general, this approach does not guarantee that the discretization of the gradient is consistent with the discretized objective. Further, it is not guaranteed that the discretized forward and adjoint operators are transposes of one another. Likewise, it is not guaranteed that the discretized Hessian is a symmetric operator. We report numerical experiments to quantify these errors; we will see that they are below the tolerances we target for the inversion. By using a discretize-then-optimize approach one can (by construction) guarantee that the derived operators are consistent. However, it is, e.g., not guaranteed that the forward and adjoint operators (in the transposed sense) yield the same numerical accuracy (see, e.g., [29, 44]). We refer, e.g., to [21, 39] for additional remarks on the discretization of optimization and control problems.

3.1. Optimality Conditions and Newton Step

From Lagrange multiplier theory we know that we require vanishing variations of $\mathcal{L}$ in (5) with respect to the state, adjoint, and control variables ϕ for an admissible solution to (2). We present the steps necessary to evaluate the reduced gradient and Hessian matvec. The associated PDE operators are derived using calculus of variations, and invoking Green's identities. We will see that the optimality conditions of our problem form a system of PDEs. This system needs to be solved to find a solution of (2). We will only present the strong form of our reduced space formulation.² Note, that we also eliminate the incompressibility constraint from the optimality system (see [58, 59] for details; we comment on this in more

¹By *reduced space* we mean that we will only iterate on the reduced space of the velocity $\mathbf{v}$; we assume that the state and adjoint equations are fulfilled exactly. This is different to *all-at-once* or *full space approaches*, in which one iterates on all unknown variables simultaneously (see §3.1 and §A).

²We refer to [16, 17] for more details on reduced-space methods.

detail in §A); we only iterate on the reduced space for the velocity field $\mathbf{v}$. The *control/decision equation* (reduced gradient) for our problem is given by

$$\mathbf{g}(\mathbf{v}) := \beta_v \mathcal{A}[\mathbf{v}] + \mathcal{K}[\mathbf{b}] = \beta_v \mathcal{A}[\mathbf{v}] + \mathcal{K}\left[\int_0^1 \lambda \nabla m \, dt\right] \quad (6)$$

with (pseudo-)differential operators $\mathcal{A}$ (regularization) and $\mathcal{K}$ (projection); the definitions are given below. Formally, we require $\mathbf{g}(\mathbf{v}^*) = \mathbf{0}$ for an admissible solution $\mathbf{v}^*$ to (2). We can compute this minimizer iteratively using $\mathbf{g}$ in a gradient descent scheme. To evaluate $\mathbf{g}$ we need to find the space-time fields m and λ given a candidate $\mathbf{v}$. We can compute m by solving the *state equation* (primal)

$$\partial_t m + \mathbf{v} \cdot \nabla m = 0 \quad \text{in } \Omega \times (0, 1], \quad (7a)$$

$$m = m_T \quad \text{in } \Omega \times \{0\}, \quad (7b)$$

with periodic boundary conditions on Ω *forward* in time. Once we have found m at $t = 1$ we can compute λ by solving the *adjoint* or *costate equation* (dual)

$$-\partial_t \lambda - \nabla \cdot \lambda \mathbf{v} = 0 \quad \text{in } \Omega \times [0, 1), \quad (8a)$$

$$\lambda = -(m - m_R) \quad \text{in } \Omega \times \{1\}, \quad (8b)$$

with periodic boundary conditions on Ω *backward* in time; notice, that for vanishing $\nabla \cdot \mathbf{v}$ (8a) will also be a transport equation.

What is missing to complete the picture for $\mathbf{g}$ is a specification of the operators $\mathcal{A}$ and $\mathcal{K}$. The differential operator $\mathcal{A}$ in (6) corresponds to the first variation of the seminorms in (4). We have

$$\mathcal{A}[\mathbf{v}] = -\Delta \mathbf{v}, \quad \mathcal{A}[\mathbf{v}] = \Delta^2 \mathbf{v}, \quad \text{and} \quad \mathcal{A}[\mathbf{v}] = \Delta^3 \mathbf{v} \quad (9)$$

for the H^1 , H^2 , and H^3 case, respectively, resulting in an *elliptic*, *biharmonic*, or *triharmonic* integro-differential control equation for $\mathbf{v}$, respectively. The pseudo-differential operator $\mathcal{K}$ in (6) originates from the elimination of p and (1d). For instance, if we set $w = 0$ we obtain the Leray operator $\mathcal{K}[\mathbf{b}] := -\nabla \Delta^{-1} \nabla \cdot \mathbf{b} + \mathbf{b}$; for non-zero w this operator becomes more complicated (see [58, 59] for details on the derivation of this operator). Combining the state (primal) (7), the adjoint (8), and the control equation (6) provides the formal optimality conditions (see §A).

A common strategy to compute a minimizer for (2) is to use

$$\tilde{\nu} = \nu + (\beta_{\nu} \mathcal{A})^{-1} \mathcal{K} [\int_0^1 \lambda \nabla m \, dt]$$

as a search direction (see, e.g., [46]). We opt for a (Gauss–)Newton–Krylov method instead, due to its superior rate of convergence (see [58] for a comparison). Formally, this requires second variations of $\mathcal{L}$. The expression for the action of the reduced space Hessian $\mathcal{H}$ on a vector $\tilde{\nu}$ (i.e., the *incremental control/decision equation*) is given by

$$\mathcal{H}[\tilde{\nu}](\nu) := \beta_{\nu} \mathcal{A}[\tilde{\nu}] + \mathcal{K}[\tilde{b}] = \beta_{\nu} \mathcal{A}[\tilde{\nu}] + \mathcal{K}[\int_0^1 \tilde{\lambda} \nabla m + \lambda \nabla \tilde{m} \, dt]. \quad (10)$$

The operators $\mathcal{A}$ and $\mathcal{K}$ are as defined above. We, likewise to the reduced gradient $\mathbf{g}$ in (6), need to find two space-time fields $\tilde{m}$ and $\tilde{\lambda}$. We can find the *incremental state variable* $\tilde{m}$ by solving

$$\partial_t \tilde{m} + \nu \cdot \nabla \tilde{m} + \tilde{\nu} \cdot \nabla m = 0 \quad \text{in } \Omega \times (0, 1], \quad (11a)$$

$$\tilde{m} = 0 \quad \text{in } \Omega \times \{0\}, \quad (11b)$$

with periodic boundary conditions on Ω , *forward in time*. Once we have found $\tilde{m}$ we can compute the *incremental adjoint variable* $\tilde{\lambda}$ by solving

$$-\partial_t \tilde{\lambda} - \nabla \cdot (\tilde{\lambda} \nu + \lambda \tilde{\nu}) = 0 \quad \text{in } \Omega \times [0, 1), \quad (12a)$$

$$\tilde{\lambda} = -\tilde{m} \quad \text{in } \Omega \times \{1\}, \quad (12b)$$

with periodic boundary conditions on Ω *backward in time*. Thus, each time we apply the Hessian to a vector we have to solve two PDEs—(11a) and (12a).

3.2. Discretization

We subdivide the time interval $[0, 1]$ into $n_t \in \mathbf{N}$ uniform steps $\ell^j, j = 0, \dots, n_t$ of size $h_t = 1/n_t$. We discretize $\Omega := (-\pi, \pi)^d$ via a regular grid with cell size $\mathbf{h}_x = (h_x^1, \dots, h_x^d)^\top \in \mathbf{R}_{>0}^d$, $\mathbf{h}_x = 2\pi \oslash \mathbf{n}_x$, $\mathbf{n}_x = (n_x^1, \dots, n_x^d)^\top \in \mathbf{N}^d$; we use a pseudospectral discretization with a Fourier basis. We discretize the integral operators based on a midpoint rule. We use cubic splines as a basis function for our interpolation model.

3.3. Numerical Time Integration

An efficient, accurate, and stable time integration of the hyperbolic PDEs that appear in our optimality system is critical for our solver to be effective. Each evaluation of the objective functional $\mathcal{J}$ in (2a) requires the solution of (7a) (forward in time). The evaluation of the reduced gradient $\mathbf{g}$ in (6) requires an additional solution of (8a) (backward in time). Applying the reduced space Hessian $\mathcal{H}$ (Hessian matvec) in (10) necessitates the solution of (11a) (forward in time) and (12a) (backward in time).

3.3.1. Second order Runge-Kutta Schemes—In our original work [58, 59] we solved the transport equations based on an RK2 scheme (in particular, Heun's method). This method—in combination with a pseudospectral discretization in space—offers high accuracy solutions, minimal numerical diffusion, and spectral convergence for smooth problems at the cost of having to use a rather small time step due to its conditional stability; the time step size h_t has to be chosen according to considerations of stability rather than accuracy. This scheme can become unstable, even if we adhere to the conditional stability (see §4.1 for examples). One strategy to stabilize our solver is to rewrite the transport equations in antisymmetric form [35, 53]. Here we extend this stable scheme to the adjoint problem and the Hessian operator. We do so by deriving the antisymmetric form of the forward operator and then formally computing its variations. It is relatively straightforward but we have not seen this in the literature related to inverse transport problems. We present the associated PDE operators in §B. We refer to this solver as RK2A scheme. It is evident that the discretization in antisymmetric form requires more work (see §B). We provide estimates in terms of the number of FFTs we have to perform in Tab. 12 in §C.

3.3.2. Semi-Lagrangian Formulation—Next, we describe our semi-Lagrangian formulation. To be able to apply the semi-Lagrangian method to the transport equations appearing in our optimality systems, we have to reformulate them. Using the identity $\nabla \cdot u\mathbf{v} = u\nabla \cdot \mathbf{v} + \nabla u \cdot \mathbf{v}$ for some arbitrary scalar function $u: \bar{\Omega} \rightarrow \mathbf{R}$, we obtain

$$\partial_t m + \mathbf{v} \cdot \nabla m = 0 \quad \text{in } \Omega \times (0, 1], \quad (13a)$$

$$-\partial_t \lambda - \mathbf{v} \cdot \nabla \lambda - \lambda \nabla \cdot \mathbf{v} = 0 \quad \text{in } \Omega \times [0, 1), \quad (13b)$$

$$\partial_t \tilde{m} + \mathbf{v} \cdot \nabla \tilde{m} + \tilde{\mathbf{v}} \cdot \nabla m = 0 \quad \text{in } \Omega \times (0, 1], \quad (13c)$$

$$-\partial_t \tilde{\lambda} - \mathbf{v} \cdot \nabla \tilde{\lambda} - \tilde{\lambda} \nabla \cdot \mathbf{v} - \nabla \cdot \lambda \tilde{\mathbf{v}} = 0 \quad \text{in } \Omega \times [0, 1). \quad (13d)$$

These equations are all of the general form $d_t u = \mathbf{v} \cdot \nabla u = f(u, \mathbf{v})$, where $u: \bar{\Omega} \times [0, 1] \rightarrow \mathbf{R}$ is some arbitrary scalar function and $d_t := \partial_t + \mathbf{v} \cdot \nabla$. If the Lagrangian

derivative vanishes, i.e., $d_t u = 0$, u is constant along the characteristics $\mathbf{X} : [\tau_0, \tau_1] \rightarrow \mathbf{R}^d$ of the flow, where $[\tau_0, \tau_1] \subseteq [0, 1]$. We can compute $\mathbf{X}$ by solving the ODE

$$d_t \mathbf{X}(t) = \mathbf{v}(\mathbf{X}(t)) \quad \text{in } (\tau_0, \tau_1] \quad (14a)$$

$$\mathbf{X}(t) = \mathbf{x} \quad \text{at } \{\tau_0\} \quad (14b)$$

The solution of (14) requires the knowledge of the velocity field $\mathbf{v}$ at points that do not coincide with the computational grid; we have to interpolate $\mathbf{v}$ in space.³

The idea of pure Lagrangian schemes is to solve $d_t u = f(u, \mathbf{v})$ along the characteristic lines (14). The key advantage of these methods is that they are essentially unconditionally stable [68]; i.e., the time step h_t may be chosen according to accuracy considerations rather than stability considerations.⁴ On the downside the solution will no longer live on a regular grid; the grid changes over time and eventually might become highly irregular. Semi-Lagrangian methods can be viewed as a hybrid between Lagrangian and Eulerian methods; they combine the best from both worlds—they operate on a regular grid and are unconditionally stable.

The semi-Lagrangian scheme involves two steps: For each time step t^j we have to compute the departure point $\mathbf{X}_D := \mathbf{X}(t = t^{j-1})$ of a fluid parcel by solving the characteristic equation (14) backward in time, with initial condition $\mathbf{X}(t = t^j) = \mathbf{x}$.⁵ We revert to a uniform grid by interpolation. The second step is to compute the transported quantity along the characteristic $\mathbf{X}$. The accuracy of the semi-Lagrangian method is sensitive to the time integrator for solving (14) as well as the interpolation scheme used to evaluate the departure points $\mathbf{X}_D$. We discuss the individual building blocks of our solver next.

Tracing the Characteristic: For each time step t^j of the integration of a given transport equation we have to trace the characteristic $\mathbf{X}$ backward in time in an interval $[t^{j-1}, t^j] \subset [0, 1]$. Since we invert for a stationary velocity field $\mathbf{v}$ we have to trace $\mathbf{X}$ (i.e., compute the departure points $\mathbf{X}_D$) only once in every Newton iterations used for all time steps.⁶ We use an explicit RK2 scheme (Heun's method) to do so [68]. We illustrate the computation of the characteristic in Fig. 1. Each evaluation of the right hand side of (14) requires interpolation.

Interpolation: We use a cubic spline interpolation model to evaluate the transported quantities along the characteristic $\mathbf{X}$. We pad the imaging data to account for the periodic boundary conditions. The size of the padding zone is computed at every iteration based on

³Notice that the scheme becomes more complicated if $\mathbf{v}$ is non-stationary; we have to interpolate in time and space.

⁴For rapidly varying velocity fields, instabilities may still occur.

⁵The direction of time integration depends on the transport equation. For simplicity, we will limit the description of the semi-Lagrangian method to transport equations that are solved forward in time. Notice that (13) also contains equations that have to be solved backward in time.

⁶In total, we actually need compute two characteristics, one for the forward (state or primal) equations and one for the backward (adjoint or dual) equations; one defined forward in time and one defined backward in time.

the maximal displacement between the original grid nodes and the departure points $\mathbf{X}_D$; we also account for the support of the basis functions of the interpolation model.⁷

Transport: To transport the quantity of interest we have to solve equations of the form

$$d_t u(\mathbf{X}(t), t) = f(u(\mathbf{X}(t)), t, \boldsymbol{\nu}(\mathbf{X}(t))) \quad \text{in } [t^{j-1}, t^j] \quad (15)$$

along the characteristic $\mathbf{X}$. We use an explicit RK2 scheme to numerically solve (15). Since u will be needed along the characteristic $\mathbf{X}$ we have to interpolate u at the computed departure points $\mathbf{X}_D$.

3.4. Numerical Optimization

We use a globalized, inexact, matrix-free (Gauss–)Newton–Krylov method for numerical optimization. Our solver has been described and tested in [58]. In what follows, we will briefly revisit this solver and from thereon design a nested, two-level preconditioner for the reduced space optimality conditions.

3.4.1. Newton–Krylov Solver—The Newton step for updating $\boldsymbol{\nu}_k^h \in \mathbf{R}^n$, $n = d \prod_{i=1}^d n_x^i$, is in general format given by

$$\mathcal{H}^h \tilde{\boldsymbol{\nu}}_k^h = -g_k^h, \quad \boldsymbol{\nu}_{k+1}^h = \boldsymbol{\nu}_k^h + \alpha_k \tilde{\boldsymbol{\nu}}_k^h, \quad (16)$$

where $\mathcal{H}^h \in \mathbf{R}^{n,n}$, is the reduced space Hessian operator, $\tilde{\boldsymbol{\nu}}_k^h \in \mathbf{R}^n$ the search direction, and $g_k^h \in \mathbf{R}^n$ the reduced gradient.⁸ We globalize our iterative scheme on the basis of a backtracking line search subject to the Armijo–Goldstein condition with step size $\alpha_k > 0$ at iteration $k \in \mathbf{N}$ (see, e.g., [63, page 37]). We keep iterating until the relative change of the gradient $\|g_k^h\|_{\text{rel}} := \|g_k^h\|_{\infty} / \|g_0^h\|_{\infty}$ is smaller or equal to $1\text{E-}2$ or $\|g_k^h\|_{\infty} \leq 1\text{E-}5$ (other stopping conditions can be used; see e.g. [38, pages 305 ff.]). We refer to the steps necessary for updating $\boldsymbol{\nu}_k^h$ as *outer iterations* and to the steps necessary for “inverting the reduced Hessian” in (16) (i.e., the steps necessary to solve for the search direction $\tilde{\boldsymbol{\nu}}_k^h$) as *inner iterations* (see [58,59] for more details).

We use a Krylov iterative solver to compute $\tilde{\boldsymbol{\nu}}_k^h$. To evaluate the reduced gradient $g_k^h \in \mathbf{R}^n$ on the right hand side of (16) (see (6)) we have to solve (7a) *forward in time* and (8a) *backward in time* for a given iterate $\boldsymbol{\nu}_k^h$.⁹ Once we have found the gradient, we can solve (16). The reduced space Hessian in (16) is a large, dense, ill-conditioned operator. Solving this system is a significant challenge; we use a PCG method [50]. Indefiniteness of $\mathcal{H}^h$ can

⁷We have tested a more accurate implementation that applies the interpolation step on a grid of half the cell size $\mathbf{h}_x$ to minimize the interpolation error as well as the numerical diffusion. We prolong and restrict the data in the Fourier domain. The gain in numerical accuracy (one to two digits) did not justify the significant increase in CPU time.

⁸The Hessian matvec is given by (10); the expression for the reduced gradient is given in (6).

⁹In general we associate the cost for solving (7a) to the evaluation of the objective in (2a), i.e., to the line search.

be avoided by using a GN approximation to the true Hessian¹⁰ or by terminating the PCG solve in case negative curvature occurs. By using a GN approximation, we sacrifice speed of convergence; quadratic convergence drops to superlinear convergence; we locally recover quadratic convergence as λ^h tends to zero.

An important property of Krylov subspace methods is that we do not have to store or form the reduced space Hessian $\mathcal{H}^h$; we merely need an expression for the action of $\mathcal{H}^h$ on a vector; this is exactly what (10) provides. Each application of $\mathcal{H}^h$ (i.e., each PCG iteration) requires the solution of (11a) and (12a). This results in high computational costs. We use inexact solves (see [63, pages 165ff.] and references therein) to reduce these costs. Another key ingredient to keep the number of PCG iterations small is an effective preconditioner. This is what we discuss next.

3.4.2. Preconditioner—The design of an optimal preconditioner for KKT systems arising in large-scale inverse problems is an active area of research [13, 15–17, 41].¹¹ Standard techniques, like incomplete factorizations, are not applicable as they require the assembling of $\mathcal{H}^h$. We provide two matrix-free strategies below.

Given the reduced gradient $\mathbf{g}$, the Newton step in the reduced space is given by

$$\mathcal{H}[\tilde{\mathbf{v}}](\mathbf{v}) = \beta_v \mathcal{A}[\tilde{\mathbf{v}}] + \mathcal{K} \left[\int_0^1 \tilde{\lambda} \nabla m + \lambda \nabla \tilde{m} \, dt \right] = \beta_v \mathcal{A}[\tilde{\mathbf{v}}] + \mathcal{Q}[\tilde{\mathbf{v}}] = -\mathbf{g}. \quad (17)$$

We have introduced the operator $\mathcal{Q}[\tilde{\mathbf{v}}] := \mathcal{Q}[\tilde{\lambda}, m, \lambda, \tilde{m}](\mathbf{v}, \tilde{\mathbf{v}})$ in (17) for notational convenience and to better illustrate its dependence on $\tilde{\mathbf{v}}$; the incremental state and adjoint variables, $\tilde{m}$ and $\tilde{\lambda}$, are functions of $\tilde{\mathbf{v}}$ through (11a) and (12a), respectively.

We use a left preconditioner $\mathbf{P}^{-1}$; our solver will see the system $\mathbf{P}^{-1} \mathcal{H}^h \tilde{\mathbf{v}}_k^h = -\mathbf{P}^{-1} \mathbf{g}_k^h$. Ideally the preconditioned matrix will have a much better spectral condition number and/or eigenvalues that are clustered around one. An ideal preconditioner is one that has vanishing costs for its construction and application and at the same time represents an excellent approximation to the Hessian operator $\mathcal{H}^h$ so that $\mathbf{P}^{-1} \mathcal{H}^h \approx \mathbf{I}_n$ [11]. These are in general competing goals. Since we use a PCG method to iteratively solve (17), we only require the action of $\mathbf{P}^{-1}$ on a vector.

Regularization Preconditioner: In our original work [58,59], we use a preconditioner that is based on the exact, spectral inverse of the regularization operator $\mathcal{A}^h$, i.e.,

$$\mathbf{P}_{\text{REG}} = \beta_v \mathcal{A}^h = \beta_v \mathbf{W} \mathbf{\Gamma} \mathbf{W}^{-1}, \quad \mathbf{P}_{\text{REG}} \in \mathbf{R}^{n,n}, \quad (18)$$

¹⁰This corresponds to dropping all terms with λ in (10) and (12a) (see [58,59]).

¹¹We study the spectral properties of $\mathcal{H}^h$ for the compressible and incompressible case in [58].

where $W^{-1} = I_d \otimes \hat{W} \in C^{m,n}$, $I_d = \text{diag}(1, \dots, 1) \in \mathbf{R}^{d,d}$, $\hat{W}$ is a DFT matrix and $\Gamma = I_d \otimes \hat{\Gamma} \in \mathbf{R}^{n,n}$ are the spectral weights for the Laplacian, biharmonic, or triharmonic differential operators in (9). The operator $\mathcal{A}^h$ has a non-trivial kernel; to be able to invert this operator analytically we replace the zero entries in Γ by one. If we apply P_{REG}^{-1} to the reduced Hessian in (17) the system we are effectively solving is a low-rank, compact perturbation of the identity:

$$\tilde{\mathbf{v}}^h + (\beta_v \mathcal{A}^h)^{-1} \mathcal{Q}^h [\tilde{\mathbf{v}}^h] = (\mathbf{I}_n + (\beta_v \mathcal{A}^h)^{-1} \mathcal{Q}^h) \tilde{\mathbf{v}}^h. \quad (19)$$

Notice that the operator P_{REG} acts as a smoother on $\mathcal{Q}^h$. Applying and inverting this preconditioner has vanishing computational costs (due to our pseudospectral discretization). This preconditioner becomes ineffective for small regularization parameters β_v and a high inversion accuracy (i.e., small tolerances for the relative reduction of the reduced gradient; see, e.g., [59]).

Nested Preconditioner: We use a coarse grid correction by an inexact solve to provide an improved preconditioner. This corresponds to a simplified two-level multigrid v-cycle, where the smoother is the inverse of the regularization operator and the coarse grid solve is inexact. We introduce spectral restriction and prolongation operators to change from the fine to the coarse grid and vice versa. The action of the preconditioner, i.e., of the action of the reduced space Hessian in (17), is computed on the coarse grid. This preconditioner only operates on the low frequency modes due to the restriction to a coarser grid. In our implementation, we treat the high and the low frequency components separately; we apply the nested preconditioner to the low frequency modes and leave the high frequency modes untouched. We separate the frequency components by applying an ideal low- and high-pass filter to the vector the preconditioner is applied to. As we will see below, we will actually treat the high frequency components with a smoother that is based on the inverse of our regularization operator, i.e., the Hessian to be preconditioned does not correspond to the reduced space Hessian in (17) but the preconditioned Hessian in (19). We refer to this preconditioner as P_{2L} .

The effectiveness of this scheme is dictated by the computational costs associated with the inversion of the (coarse grid) Hessian operator $\mathcal{H}^{2h} \in \mathbf{R}^{n/2, n/2}$. One strategy for applying this preconditioner is to compute the action of the inverse of $\mathcal{H}^{2h}$ using a nested PCG method. From the theory of Krylov subspace methods we know that we have to solve for the action of this inverse with a tolerance that is smaller than the one we use to solve (17) (exact solve; we refer to this approach as PCG(ϵ), where $\epsilon \in (0, 1)$ is the scaling for the tolerance used to solve (17)) for the outer PCG method to not break down. This increased accuracy may lead to impractical computational costs, especially since each application of $\mathcal{H}^{2h}$ requires two PDE solves, and we expect this preconditioner to have a very similar conditioning as $\mathcal{H}^h$. Another strategy is to solve the system inexactly. This requires the use of flexible Krylov subspace methods or a Chebyshev (CHEB) method (see e.g. [8, pages 179ff.]) with a fixed number of iterations (we refer to this strategy as CHEB(k), where k is

the number of iterations). This makes the work spent on inverting the preconditioner constant but the inexactness might lead to a less effective preconditioner. Another bottleneck is the fact that the CHEB method requires estimates of the spectral properties of the operator we try to invert; estimating the eigenvalues is expensive and can lead to excessive computational costs. We provide implementation details next, some of which are intended to speed up the formation and application of our nested preconditioner.

- **Spectral Preconditioning of $\mathcal{H}^h$:** Since the application of the inverse of the regularization operator $\mathcal{A}^h$ comes at almost no cost, we decided to use the spectrally preconditioned Hessian operator in (19) within our two-level scheme, with a small technical modification. The left preconditioned Hessian in (19) is not symmetric. We can either opt for Krylov methods that do not require the operator we try to invert to be symmetric, or we employ a spectral split preconditioner. We opt for the latter approach to be able to use a PCG method, attributed to its efficiency. The split preconditioned system is given by

$$(\mathbf{I}_n + (\beta_v \mathcal{A})^{-1/2} \mathcal{Q}^h (\beta_v \mathcal{A})^{-1/2}) \mathbf{s} = - (\beta_v \mathcal{A})^{-1/2} \mathbf{g},$$

where $\mathbf{s} := (\beta_v \mathcal{A})^{-1/2} \tilde{\mathbf{v}}$. Notice, that the inverse of the regularization operator can be viewed as a smoother, which establishes a connection of our scheme to more sophisticated multigrid strategies [2,19].

- **Eigenvalue Estimates for the CHEB Method:** The computational costs for estimating eigenvalues of $\mathbf{P}_{2L}$ are significant. Our assumption is that we have to estimate the extremal eigenvalues only once for the registration for a given set of images (we will experimentally verify this assumption; see §4.2.1); if we change the regularization parameter we simply have to scale the estimated eigenvalues. Notice that we can efficiently estimate the eigenvalues for a zero velocity field since a lot of the terms drop in the optimality systems. We compute an estimate for the largest eigenvalue $e_{\max}$ based on an implicitly restarted Lanczos algorithm. We approximate the smallest eigenvalue analytically under the assumption that $\mathcal{Q}^h$ is a low-rank operator of order $\mathcal{O}(1)$; $e_{\min} \approx \min(\mathbf{I}_n + (\beta_v \mathbf{\Gamma})^{-1})$.
- **Hyperbolic PDE Solves:** Each matvec with $\mathcal{H}^{2h}$ requires the solution of (11a) and (12a) on the coarse grid. We exclusively consider the semi-Lagrangian formulation to speed up the computation. In general we assume that we do not need high accuracy solutions for our preconditioner. This might even be true for the PDE solves within each Hessian matvec.¹²
- **Restriction/Prolongation:** We use spectral restriction and prolongation operators. We do not apply an additional smoothing step after or before we restrict the data to the coarser grid. We actually observed that applying an

¹²We have also tested an approximation of the forcing term by dropping all second order terms of the RK2 scheme for numerically integrating (11a) and (12a). Since we have observed instabilities in the RK2 schemes and due to the effectiveness of the semi-Lagrangian method (see §4.1) we do not report results for this preconditioner.

additional Gaussian smoothing with a standard deviation of $2h_x$ (i.e., one grid point on the coarser grid) significantly deteriorates the performance of our preconditioner for small grid sizes (e.g., 64×64). A more detailed study on how the choices for the restriction and prolongation operators affect the performance of our solver with respect to changes in the regularity of the underlying objects remains for future work.

- **Filters:** We use simple cut-off filters before applying the restriction and prolongation operators, with a cut-off frequency of half the frequency that can be represented on the finer grid.

3.5. Implementation Details and Parameter Settings

Here, we briefly summarize some of the implementation details and parameter choices.

- **Image Data:** Our solver can not handle images with discontinuities. We ensure that the images are adequately smooth by applying a Gaussian smoothing kernel with an empirically selected standard deviation of one grid point in each spatial direction. We normalize the intensities of the images to $[0, 1]$ prior to registration.
- **PDE Solves:** We use a CFL number of 0.2 for the explicit RK2 schemes; we observed instabilities for some of the test cases for a CFL number of 0.5. The semi-Lagrangian method is unconditionally stable; we test different CFL numbers.
- **Restriction/Prolongation:** We use spectral prolongation and restriction operators within our pre-conditioner (more implementation details for our preconditioner can be found in the former section). We do not perform any other grid, scale, or parameter continuation to speed up our computations.
- **Interpolation:** We consider a C^2 -continuous cubic spline interpolation model. We extend our data periodically to account for the boundary conditions.
- **Regularization:** Since we study the behavior of our solver as a function of the regularization parameters, we will set their value empirically. For practical applications, we have designed a strategy that allows us to probe for an ideal regularization parameter; we perform a parameter continuation that is based on a binary search and considers bounds on the determinant of the deformation gradient as a criterion; see [58,59].
- **Globalization:** We use a backtracking line search subject to the Armijo–Goldstein condition to globalize our Newton–Krylov scheme (see, e.g., [63, page 37]).
- **Stopping Criteria:** We terminate the inversion if the relative change of the gradient is smaller or equal to $1E-2$ or $\|g_k^h\|_\infty \leq 1E-5$ (other stopping conditions can be used; see, e.g., [38, pages 305 ff.]).

- **Hessian:** We use a GN approximation to the reduced space Hessian $\mathcal{H}^h$ to avoid indefiniteness. This corresponds to dropping all expressions with λ in (10) and (12a) (see [58] for more details); we recover quadratic convergence for $\lambda \rightarrow 0$.
- **KKT solve:** If not noted otherwise, we will solve the reduced space KKT system in (16) inexactly, with a forcing sequence that assumes quadratic convergence (see [63, pages 165ff.] and references therein); we use a PCG method to iteratively solve (16).
- **PC solve:** We compute the action of the inverse of the 2-level preconditioner either exactly using a nested PCG method or inexactly based on a nested CHEB method with a fixed number of iterations.

4. Numerical Experiments

We report numerical experiments next. The error of our discrete approximation to the control problem depends on the smoothness of the solution, the smoothness of the data, and the numerical errors/order-of-accuracy of our scheme. We perform a detailed numerical study to quantify these errors experimentally. We start with a comparison of the numerical schemes for solving the hyperbolic PDEs that appear in the optimality system and the Newton step (see §4.1). The second set of experiments analyzes the effectiveness of our schemes for preconditioning the reduced space KKT system (see §4.2).

All experiments are carried out for $d = 2$ using Matlab R2013a on a Linux cluster with Intel Xeon X5650 Westmere EP 6-core processors at 2.67GHz with 24GB DDR3-1333 memory. We illustrate the synthetic and real world data used for the experiments in Fig. 2 and Fig. 3, respectively.¹³

4.1. Hyperbolic PDE solver

We study the performance of the time integrators for the hyperbolic transport equations. We only consider the problems SMOOTH A and SMOOTH B in Fig. 2 as these are constructed to be initially resolved on the considered grids Ω^h . This allows us to study grid convergence without mixing in any additional problems due to potential sharp transitions in the intensity values of the image data. We will see that these simple test cases can already break standard numerical schemes.

4.1.1. Self-Convergence: State and Adjoint Equation

Purpose: To study the numerical stability and accuracy of the considered schemes for integrating the hyperbolic transport equations that appear in our optimality system.

Setup: We study the self-convergence of the considered numerical time integrators. We consider the RK2 scheme (pseudospectral discretization in space), the stabilized RK2A scheme (pseudospectral discretization in space), and the SL method (cubic interpolation combined with a pseudospectral discretization; see §3.3 for details). We test these schemes

¹³The HAND images in Fig. 3 are taken from [62]. The BRAIN images in Fig. 3 are taken from the “Nonrigid Image Registration Evaluation Project” (NIREP) available at <http://nirep.org> (data sets na01 and na02) [26].

for the synthetic problems SMOOTH A and SMOOTH B in Fig. 2. We consider the state and the adjoint equation. We compute the relative ℓ_2 -error between the solution of the transport equations (state equation (7a) and adjoint equation (8a)) obtained on a spatial grid of size $\mathbf{n}_x$ and the solution obtained on a spatial grid of size $\tilde{\mathbf{n}}_x = 2\mathbf{n}_x$. We compute this error in the Fourier domain; formally, the error is given by

$$\|\delta u^h\|_{\text{rel}} := \|M[\mathbf{W}^{-1}u^h]_{\mathbf{n}_x} - [\mathbf{W}^{-1}u^h]_{\tilde{\mathbf{n}}_x}\|_2 / \|[\mathbf{W}^{-1}u^h]_{\tilde{\mathbf{n}}_x}\|_2$$

for a given numerical solution u^h . Here, $[\cdot]_{\mathbf{n}}$ indicates that the data is represented on a grid of size $\mathbf{n}$; M is a prolongation operator that maps the data from a grid of size $\mathbf{n}_x$ to a grid of size $\tilde{\mathbf{n}}_x$; and $\mathbf{W}^{-1}$ represents the forward Fourier operator. We use a CFL number of 0.2 to compute the number of time steps n_t for the RK2 and the RK2A method. For the SL method we use the CFL numbers 0.2, 1, and 5. We expect the error to tend to zero for an increasing number of discretization points.

Results: We report results for the self-convergence of our numerical schemes in Tab. 2. We illustrate a subset of these results in Fig. 4.

Observations: The most important observations are that (i) our SL scheme delivers an accuracy that is at the order of the RK2A and the RK2 scheme with a speed up of one order of magnitude,¹⁴ and (ii) that our standard RK2 scheme can become unstable if we combine it with a spectral discretization—even for smooth initial data and a smooth velocity field (run #31, run #36, run #71, and run #76 in Tab. 2). This instability is a consequence of the absence of numerical diffusion; it is completely unrelated to the CFL condition. The RK2A and the SL method remain stable across all considered test cases with a similar performance. The rate of convergence for the RK2 and the RK2A scheme are excellent; we expect second order convergence in time and spectral convergence in space (this has been verified; results not reported here). The self-convergence for the SL(0.2) method is at the order, but overall slightly better, than the one observed for the RK2 and the RK2A scheme. The error for the self-convergence increases by one order of magnitude if we increase the CFL number for the SL method to 1 or 5, respectively. Switching from test problem SMOOTH A to SMOOTH B the self-convergence deteriorates for both methods. We can observe that we can not fully resolve the problem SMOOTH B if we solve the equations on a spatial grid with less than 128 nodes along each spatial direction; the errors range between $\mathcal{O}(1\text{E} - 1)$ and $\mathcal{O}(1\text{E} - 2)$ (run #21 through run #25 for the state equation and run #61 through run #70 for the adjoint equation; see also Fig. 4). Notice that we can fully resolve the initial data and the velocity field for smaller grid sizes.

We can also observe that we loose about one order of magnitude in the rate of convergence if we switch from the state to the adjoint equation—even for the mild case SMOOTH A. This observation is consistent across all solvers. This demonstrates that the adjoint equation is in

¹⁴We have also tested an implementation of the SL method that delivers more accurate solutions (less numerical diffusion) at the expense of a significant increase in time to solution. In this scheme we upsampled the data to a grid of size $2\mathbf{n}_x$ whenever we had to interpolate. The associated gain in accuracy did not justify the increase in computational cost.

general more difficult to solve than the state equation. This can be attributed to the fact that the adjoint equation is a transport equation for the residual; the residual has, in general, less regularity than the original images (see also [73]).

As for the time to solution we can observe that the SL(0.2) scheme delivers a performance that is at the order of the RK2A(0.2) scheme (slightly worse). We have to switch to use a CFL number of 1 to be competitive with the RK2(0.2) scheme. For a CFL number of 5 the SL scheme outperforms the RK2 and RK2A scheme by about one order of magnitude in terms of time to solution. Intuitively, one would expect that the SL method delivers much more pronounced speedup due to the unconditional stability. However, the discrepancy is due to the fact that we essentially replace a large number of highly optimized FFT operations with cubic spline interpolation operations. We report estimates for computational complexity in terms of FFTs and IPs in Tab. 12 in §C. We can see in Tab. 13 in §C that the differences in CPU time between these two operations are significant.

Conclusions: We can not guarantee convergence to a valid solution if we use a standard RK2 scheme in combination with a spectral discretization, even for smooth initial data; we have to use more sophisticated schemes. We provide two alternatives: a stabilized RK2 scheme (RK2A) and an SL scheme (see §3.3 for details). Both schemes remained stable across all experiments. The SL scheme delivers a performance that is very similar to the RK2A scheme with a speedup of one order of magnitude—even for our non-optimized implementation.¹⁵

4.1.2. Convergence to RK2A

Purpose: To assess (i) the convergence of the SL method to the solution of the RK2A scheme and by that (ii) the numerical errors that might affect the overall convergence of our Newton–Krylov solver.

Setup: We assess the convergence of the SL method to a solution computed on the basis of the RK2A scheme for the state and the adjoint equation (7a) and (8a), respectively. Based on our past experiments (see [58, 59]) we assume that the solution of the RK2A scheme is a silver standard. We compute the reference solution on a grid of size $\mathbf{n}_x = (512, 512)^\top$ with a CFL number of 0.2 (RK2A(0.2)). Likewise to the former experiment we compute the discrepancy between the numerical solutions in the Fourier domain; i.e., we report relative errors $\|\delta m_1^h\|_{\text{rel}}$ and $\|\delta \lambda_0^h\|_{\text{rel}}$, where we have $\tilde{\mathbf{n}}_x = (512, 512)^\top$ for the RK2A(0.2) reference solution. We report results for different discretization levels (varying number of grid points $\mathbf{n}_x$ and $\mathbf{n}_p$; the CFL numbers for the SL method are 0.2, 1, 2, 5, and 10). As a reference, we also compute errors for the RK2A scheme for a CFL number of 0.2. We also compute convergence errors for the gradient; the setup is the same as for the experiment for the adjoint and state equation.

¹⁵Matlab’s FFT library is based on the highly optimized FFTW library (see <http://www.fftw.org>; [36]). For interpolation we use Matlab’s built in interp2 routine (Matlab R2013a; we report timings in Tab. 13 in §C). We expect an additional speedup if we switch to an optimized, three-dimensional C++ implementation, something we will investigate in future work.

Results: We report the relative error between the solution computed based on our SL formulation and the RK2A scheme in Tab. 3. The error estimates for the reduced gradient can be found in Tab. 4.

Observations: The most important observation is that the SL scheme converges to the RK2A(0.2) reference solution with a similar rate than the RK2A scheme itself. The SL scheme delivers an equivalent or even better rate of convergence than the RK2A scheme for a CFL number of 0.2. We lose one to two digits if we switch to higher CFL numbers; this loss in accuracy might still be acceptable for our Newton–Krylov solver to converge to almost identical solutions, something we will investigate below. Likewise to the former experiment we can again observe that the error for the adjoint equation are overall about one order of magnitude larger than those obtained for the state equation; this observation is again consistent for both schemes—the RK2A scheme and the SL scheme (see for instance run #41 and run #45; and run #42 and run #46 in Tab. 3).

Conclusions: Our SL scheme behaves very similar than the RK2A scheme with the benefit of an orders of magnitude reduction in computational work load due to the unconditional stability. We expect significant savings, especially for evaluating the Hessian, as accuracy requirements for the Hessian and its preconditioner are less significant than those for the reduced gradient for our Newton–Krylov solver to still converge. If high accuracy solutions are required, we can simply increase the number of time points to match the accuracy obtained for the RK2A scheme at the expense of an increase in CPU time.

4.1.3. Adjoint Error

Purpose: To assess the numerical errors of the discretized forward and adjoint operator.

Setup: We solve the state equation (7a) and the adjoint equation (8a) on a grid of size $\mathbf{n}_x = (256, 256)^T$ for a varying number of time points n_t . We consider the problem SMOOTH A in Fig. 2 to setup the equations. We report the relative error between the discretized forward operator $\mathcal{E}^h$ and the discretized adjoint operator

$(\mathcal{E}_h)^T : \delta_{\text{ADJ}} := |\langle \mathcal{E}^h m_0^h, \mathcal{E}^h m_0^h \rangle - \langle (\mathcal{E}^h)^T \mathcal{E}^h m_0^h, m_0^h \rangle| / |\langle \mathcal{E}^h m_0^h, \mathcal{E}^h m_0^h \rangle|$. The continuous forward operator is self-adjoint, i.e., $\mathcal{E} = \mathcal{E}^T$; the error should tend to zero if our numerical scheme preserves this property.¹⁶

Results: We report the relative adjoint errors in Tab. 5.

Observations: The most important observation is that the RK2A scheme is self-adjoint (up to machine precision) whereas the error for the SL method ranges between $\mathcal{O}(1\text{E} - 6)$ and $\mathcal{O}(1\text{E} - 3)$ as a function of n_t . If we solve the problem with a CFL number of 2 or smaller, the adjoint error is below or at the order of the accuracy we typically solve the inverse

¹⁶Our solver is based on an optimize-then-discretize approach (see §3). We can not guarantee that the properties of the continuous operators of our constrained formulation and its variations are preserved after discretization. In a discretize-then-optimize approach the discretization is differentiated, which will result in consistent operators; we refer to [39, pages 57ff.] for a more detailed discussion on the pros and cons; we also discuss this in §3.

problem with in practical applications (relative change of the gradient of $1\text{E}-2$ or $1\text{E}-3$ and an absolute tolerance for the ℓ^∞ -norm of the reduced gradient of $1\text{E}-5$).

Conclusions: Our SL scheme is not self-adjoint. The numerical errors are acceptable for the tolerances we use in practical applications—even for moderate CFL numbers. If we intend to solve the problem with a higher accuracy, we might have to either use a larger number of time steps or switch to the RK2A scheme to guarantee convergence. We already note that we have not observed any problems in terms of the convergence (failure to converge) nor the necessity for any additional line search steps in our solver, even if we considered a CFL number of 10.

4.2. Preconditioner

Next, we analyze the performance of our preconditioners (see §3.4).

4.2.1. Eigenvalue Estimation—We need to estimate the extremal eigenvalues of P_{2L} if we use the CHEB method to compute the action of its inverse. This estimation results in a significant amount of computational work if we have to do it frequently (about 30 matvecs for the estimation of $e_{\max}$). We estimate the smallest eigenvalue based on an analytical approximation; we estimate the largest eigenvalue numerically (see §3.4.2 for details).

Purpose: To assess if the estimates for the largest eigenvalue vary significantly during the course of on inverse solve.

Setup: We solve the inverse problem for different sets of images; we consider the UT, HAND, HEART, and BRAIN images in Fig. 3. We terminate the inversion if the gradient is reduced by three orders of magnitude or if $\|g_k^h\|_\infty \leq 1\text{E}-5$, $k = 0, 1, 2, \dots$. We estimate the largest eigenvalue every time the preconditioner is applied. We consider a compressible diffeomorphism (H^2 -regularization). The solution is computed using a GN approximation. We report results for different regularization weights β_v .

Results: We summarize the estimates for the largest eigenvalue $e_{\max}$ in Tab. 6.

Observations: The most important observation is that the estimates for the largest eigenvalue do not vary significantly during the course of the iterations for most of the considered test cases. We have verified this for different reference and template images and as such for varying velocity fields. Our results suggest that we might have to only estimate the eigenvalues once for the initial guess—a zero velocity field. The costs for applying the Hessian for a zero velocity field are small—several expressions in (10), (11a), and (12a) drop or are constant. Our results suggest that the changes in the eigenvalues are a function of the changes in the magnitude of the velocity field v , i.e., the amount of expected deformation between the images. That is, we have only subtle residual differences and a small deformation in case of the HEART images; the estimated eigenvalues are almost constant. For the HAND and the UT images the deformations and the residual differences are larger; the changes in the estimates for the largest eigenvalue are more pronounced. Another important observation is that the most significant changes occur during the first few outer

iterations. Once we are close to the solution of our problem, the eigenvalues are almost constant.

Overall, these results suggest that we can limit the estimation of the eigenvalues to the first iteration, or—if we observe a deterioration in the performance of our preconditioner—re-estimate the eigenvalues and for the subsequent solves again keep them fixed. We can observe that we have to estimate the eigenvalues only once for a given set of images; changes in the regularization parameter can simply be accounted for by rescaling these eigenvalue estimates. This is in accordance with our theoretical understanding of how changes in the regularization parameter affect the spectrum of the Hessian operator.

Conclusions: The estimates for the largest eigenvalue do not vary significantly during the course of the inversion for the considered test problems. We can estimate the eigenvalues efficiently during the first iteration (zero initial guess) and potentially use this estimate throughout the entire inversion.

4.2.2. Convergence: KKT Solve

Purpose: To assess the rate of convergence of the KKT solve for the different schemes to precondition the reduced space Hessian.

Setup: We consider three sets of images, the test problem SMOOTH A in Fig. 2, and the BRAIN and the HAND images in Fig. 3. We solve the forward problem to setup a synthetic test problem based on the velocity field $\mathbf{v}^{h*}$ of problem SMOOTH A, i.e., we transport m_R to obtain a synthetic template image m_T . We consider a GN approximation to $\mathcal{H}^h$. We study three schemes to precondition the KKT system: (i) the regularization preconditioner $\mathbf{P}_{\text{REG}}$, (ii) the nested preconditioner $\mathbf{P}_{2L}$ the inverse action of which we compute using a PCG method, and (iii) the nested preconditioner $\mathbf{P}_{2L}$ the inverse action of which we compute based on a CHEB method. If we use a PCG method to invert the preconditioner, we have to use a higher accuracy than the one we use to solve the KKT system. We increase the accuracy by one order of magnitude; we refer to this solver as PCG(1E-1). For the CHEB method we can use a fixed number of iterations; we have tested 5, 10, and 20 iterations. We observed an overall good performance for 10 iterations. We refer to this strategy as CHEB(10).

We perform two experiments: In the first experiment we use a true solution $\tilde{\mathbf{v}}^{h*} = -0.5\mathbf{v}^{h*}$ and apply the Hessian operator to generate a synthetic right hand side $\tilde{\mathbf{b}}^h$. We solve the KKT system $\mathcal{H}^h \tilde{\mathbf{v}}^h = \tilde{\mathbf{b}}^h$ with a zero initial guess for $\tilde{\mathbf{v}}^h$ using a PCG method with a tolerance of 1E-12. We compute the (relative) ℓ^2 -norm of the difference between $\tilde{\mathbf{v}}^h$ and $\tilde{\mathbf{v}}^{h*}$ to assess if our schemes converge to the true solution with the same accuracy. We set up the KKT system based on the test problem SMOOTH A in Fig. 2.

For the second experiment, we evaluate the reduced gradient $\mathbf{g}^h$ at the true solution $\tilde{\mathbf{v}}^{h*}$ and solve the system $\mathcal{H}^h \tilde{\mathbf{v}}^h = -\mathbf{g}^h$ with a zero initial guess for $\tilde{\mathbf{v}}^h$. We solve the system using a PCG method with a tolerance of 1E-6. We consider a compressible diffeomorphism (H^2 -regularization). We use the RK2A scheme with a CFL number of 0.2 for the regularization

preconditioner and the SL scheme with a CFL number of five for the two-level preconditioner. We report results for the test problems SMOOTH A, BRAIN, and HAND. We consider different spatial resolution levels (grid convergence) and different choices for the regularization parameter β_v . An ideal preconditioner is mesh-independent and delivers the same rate of convergence irrespective of the choice of the regularization weight.

Results: We summarize the results for the first experiment in Tab. 7 and the results for the second part in Tab. 8. We illustrate the convergence of a subset of the results reported in Tab. 8 in Fig. 5.

Observations: The most important observation is that the nested preconditioner P_{2L} is very effective; it allows us to significantly reduce the number of iterations especially when turning to low regularization parameters. Our new scheme results in a speedup by—on average—one order of magnitude, with a peak performance of more than 20x.

The results in Tab. 7 demonstrate that our schemes all converge to the true solution with an error that is at least at the order of the tolerance used to invert the KKT system, i.e. $\mathcal{O}(1E - 12)$. The solver does not seem to be sensitive to the CFL number used for the SL method.

The results in Fig. 5 suggest that there are dramatic differences in the performance of our preconditioners; the number of iterations reduces significantly for P_{2L} . We can for instance reduce the number of iterations from 310 (run #102) to 7 (run #105). The differences in time to solution, however, are less pronounced. Our spectral discretization makes it in general extremely challenging to design a preconditioner that is more effective than P_{REG} given its ideal application and construction costs; inverting and applying P_{REG} is only at the cost of a spectral diagonal scaling (see §3.4.2 for details). The regularization preconditioner is effective for smooth problems and large regularization parameters β_v (see the first column and, e.g., run #29 (HAND images) or run #30 (BRAIN images) in in Tab. 8). We can observe that this preconditioner becomes less effective as we decrease β_v (see also [58, 59]). For example, the number of iterations increases from 33 to 92 to 279 if we reduce β_v from $1E-1$ to $1E-2$, and finally to $1E-3$ (run #57, run #66, and run #75 in Tab. 8, respectively). We can reduce the number of iterations by more than one order of magnitude if we use the nested preconditioner P_{2L} . If we use a PCG method with a tolerance of $1E-7$ to compute the action of the inverse of P_{2L} the preconditioner is almost ideal, i.e., the number of iterations is independent of β_v and the grid size n_x (see e.g., run #14, run #41, run #68, and run #95 in Tab. 8). The low tolerance (in our case one order of magnitude smaller than the tolerance we use to solve the reduced space KKT system) to compute the action of the inverse of P_{2L} results in significant application costs. Despite this increase in application costs we can—already for the present two-dimensional prototype implementation—reduce the time to solution for most of the test problems (see, e.g., run #67 or run #95 in Tab. 8). A significant factor is the SL scheme. We can further reduce the CPU time if we replace the PCG method for computing the action of the inverse of P_{2L} by a CHEB method with a fixed number of iterations. We can see that the effectiveness of this scheme is almost independent of the grid size n_x (compare, e.g., run #17, run #44, run #71, and run #98 in Tab. 8). Also, given that we use a fixed number of iterations the percentage of CPU time spent on applying the

preconditioner remains almost constant for β_v fixed. The nested preconditioner becomes less effective if we reduce β_v from $1\text{E}-2$ to $1\text{E}-3$; the number of iterations increases, which in turn makes the speedup less pronounced (see, e.g., run #98 vs. run #107 or run #99 vs. run #108 in Tab. 8). Although the speedup varies from case to case, we can see that P_{2L} in combination with CHEB(10) outperforms our original scheme for all experiments.

Conclusions: Our nested preconditioner allows us to reduce the number of iterations by more than one order of magnitude and the time to solution by up to a factor of more than 20. We expect these differences to be more pronounced for an optimized three-dimensional implementation, something we will investigate in a follow up paper.

4.3. Inverse Solve

Purpose—To study the rate of convergence of our scheme for the entire inverse solve.

Setup—We consider different test images to study the performance of our numerical scheme (HAND, HEART, BRAIN, UT). We terminate our solver if the gradient is reduced by two orders of magnitude or if the ℓ^2 norm of g_k^h , $k = 1, 2, \dots$, is equal or smaller than $1\text{E}-5$. We consider the regularization preconditioner with an RK2A(0.2) PDE solver and the two level preconditioner with an SL(5) PDE solver. We estimate the eigenvalues for the CHEB method only for the first iteration (zero velocity field). We report results for compressible, near incompressible, and incompressible diffeomorphisms, accounting for different regularization norms (H^1 -seminorm, H^2 -seminorm, and H^3 -seminorm). We study convergence (number of outer iterations and Hessian matvecs) as a function of the grid size, constraints, regularization parameter, and regularization norm. We choose the regularization weights empirically (based on experience from our former work [58, 59]). We report (i) the relative change of the reduced gradient, (ii) the relative change of the residual between m_R^h and m_1^h , (iii) the number of outer iterations, (iv) the number of Hessian matvecs, (v) the time to solution, and (vi) the obtained speedup compared to our original scheme.

Results—We report results for a compressible diffeomorphism (H^2 -regularization norm) in Tab. 9. We study grid convergence and convergence with respect to different regularization weights β_v . We report results for an incompressible diffeomorphism in Tab. 10 accounting for different regularization norms (H^1 -seminorm; H^2 -seminorm; and H^3 -seminorm). We report results for a near-incompressible diffeomorphism in Tab. 11.

Observations—The most important observation is that our solver remains effective for the entire inversion irrespective of the regularization weights, norms, and grid size.

The average speedup compared to the stabilized version of our original solver is about 10x (see, e.g., run #1 through run #6 in Tab. 10) with a peak performance of more than 20x (see, e.g., run #17 vs. run #18 and run #31 vs. run #33 in Tab. 9 or run #5 vs. run #6 in Tab. 11). We can, e.g., reduce the time to solution from ~ 3 hours to 10 minutes for a 256×256 image (run #11 vs. run #12 in Tab. 9). We can also infer that the reduced accuracy in time does not significantly affect the overall rate of convergence of our solver; the number of outer iterations remains almost constant.

Potential options to further improve our scheme are a re-estimation of the eigenvalues during the solution process once we have made significant progress (i.e., the velocity field changed drastically) and an increased accuracy for the evaluation of the gradient and the objective. That is, we currently use the same accuracy for evaluating the Hessian and the gradient. We might be able to further improve the overall accuracy and maybe convergence if we use a more accurate SL scheme when evaluating the reduced gradient.

Conclusions—Our experiments suggest that our improved solver remains effective irrespective of the regularization norm, regularization weight, or grid size. We can achieve good performance for our compressible, incompressible and near-incompressible formulations for constrained diffeomorphic image registration. We obtain a speedup of about 10x with a peak performance of 20x compared to the stabilized version of our original solver.

5. Conclusions

With this paper we follow up on our former work on constrained diffeomorphic image registration [58, 59]. We have provided an improved numerical scheme to efficiently solve the registration problem. Our solver features a semi-Lagrangian formulation, which—combined with a two-level preconditioner for the reduced space KKT system—provides a one order of magnitude speedup compared to our original solver [58,59].

We have originally described our Newton–Krylov solver in [58]; this includes a comparison against a first order gradient descent scheme still predominantly used in many diffeomorphic registration algorithms that operate on velocity fields; see, e.g., [10,46,73].¹⁷ We have extended our original formulation [58] for constrained diffeomorphic image registration in [59]. The work in [59] features a preliminary study of registration quality as a function of regularization norms and weights. The present work focuses on numerical aspects of our solver. Our contributions are:

- The implementation of an unconditionally stable SL scheme for constrained diffeomorphic image registration.
- The implementation of a two-level preconditioner for the system in (16).
- A detailed numerical study of our new, improved solver.

We perform numerical tests on various synthetic and real-world datasets to study (i) the convergence behavior of our forward solver, (ii) the adjoint errors of our schemes, (iii) the grid convergence of our preconditioner, (iv) the sensitivity of our preconditioner to changes in terms of the regularization parameters, and (v) the overall convergence of our solver with respect to different choices for the preconditioner, forward solver, regularization norms, and constraints. We found that

¹⁷The control equation in (6) corresponds to the reduced L^2 gradient, i.e., the variation of the Lagrangian $\mathcal{L}$ in (5) with respect to $\mathbf{v}$; we use the gradient in the Sobolev space induced by the regularization operator in our gradient descent scheme in [58]; see also [10,46].

- Our original solver (spectral discretization in combination with an RK2 scheme; [58, 59]) can become unstable, even for smooth problems (see, e.g., run #36 in Tab. 2).
- Our new solver (spectral discretization in combination with an RK2A scheme and our SL scheme) remains stable for all considered test cases.
- The SL scheme results in an order of magnitude speedup due to its unconditional stability compared to the RK2A scheme, subject to a reduction in numerical accuracy (see, e.g., run #17 vs. run #20 in Tab. 2 or run #53 vs. run #66 in Tab. 3; we lose two digits accuracy by increasing the CFL number from 0.2 to 5, i.e., by switching from RK2A(0.2) to SL(5)). Our numerical study suggests that this reduction in accuracy is not critical with respect to the overall performance of our Newton–Krylov solver.
- Our new scheme delivers a reduction in the number of inner iterations (i.e., the solution of the reduced space KKT system) by more than one order of magnitude (e.g., 8 vs. 279 iterations (see run #78 vs. run #75 in Tab. 8) or 7 vs. 310 iterations (see run #105 vs. run #102 in Tab. 8); see also Fig. 5). More importantly, we observe a speedup of, on average, 10x up to more than 20x (see, e.g., run #93 vs. run #99 in Tab. 8 for an individual solve of the KKT system, and run #32 vs. run #33 in Tab. 9, or run #5 vs. run #6 in Tab. 11 for the entire inversion).

Our algorithm can be used in other applications besides medical imaging, such as weather prediction and ocean physics (for tracking Lagrangian tracers in the oceans) [52] or reconstruction of porous media flows [34]. Although our method is highly optimized for regular grids with periodic boundary conditions, many aspects of our algorithm carry over. Our current Matlab prototype implementation is not yet competitive with efficient, highly optimized implementations for diffeomorphic image registration in terms of runtime [72].¹⁸ Even with the speedup we could achieve here, we are still not competitive with the (highly optimized multi-core) implementation of the algorithm presented in [72]. We expect this to change for the implementation of our solver for the three-dimensional case (which will feature highly-optimized implementations of the computational kernels of our solver dedicated to multi-core platforms and the design of efficient grid, scale, and parameter continuation schemes to further reduce the time-to-solution). We will extend the study in [59] by comparing our solver to state-of-the-art implementations of other groups (e.g., [7, 72, 73]) in terms of time-to-solution, registration quality, and inversion accuracy in the three-dimensional setting, something we are currently actively working on. We will also investigate other formulations for large deformation diffeomorphic image registration, such as for instance the map based approach in [46] or the inversion for an initial momentum in [73].

¹⁸We provide a more detailed study in [59]. Here, we show that we can outperform existing approaches for diffeomorphic image registration in terms of registration quality with our new formulation.

Appendix A. Optimality Conditions

We can derive the optimality conditions of our problem by computing the first variation of $\mathcal{L}$ with respect to the state, adjoint, and control variables, and applying integration by parts. Our derivation will be formal only. In general, we have to specify the regularity of the underlying objects to ensure existence of an optimal solution. The choices we make for the spaces for the velocity and the images are not independent. We will discuss this in more detail below. If we assume that the objective functional $\mathcal{J}$ and the PDE constraints are continuously differentiable, and satisfy a regularity condition on the constraints (see, e.g., [51]), the following system holds true at a solution $\phi^* := (m^*, \lambda^*, p^*, w^*, v^*)$ of problem (1):

$$\partial_t m^* + \nabla m^* \cdot v^* = 0 \quad \text{in } \Omega \times (0, 1], \quad (20a)$$

$$m^* = m_T \quad \text{in } \Omega \times \{0\}, \quad (20b)$$

$$-\partial_t \lambda^* - \nabla \cdot (v^* \lambda^*) = 0 \quad \text{in } \Omega \times [0, 1), \quad (20c)$$

$$\lambda^* = m_R - m^* \quad \text{in } \Omega \times \{1\}, \quad (20d)$$

$$\nabla \cdot v^* = w^* \quad \text{in } \Omega, \quad (20e)$$

$$\beta_v \mathcal{A}[v^*] + \nabla p^* + b^* = 0 \quad \text{in } \Omega, \quad (20f)$$

$$\beta_w \mathcal{B}[w^*] + p^* = 0 \quad \text{in } \Omega, \quad (20g)$$

with periodic boundary conditions on Ω . We refer to (20a) with initial condition (20b), to (20c) with final condition (20d), and to (20f) and (20g) as *state* (variation of $\mathcal{L}$ with respect to λ), *adjoint* (variation of $\mathcal{L}$ with respect to m), and *control* (variation of $\mathcal{L}$ with respect to v and w) equations, respectively. The differential operator $\mathcal{A}$ in (20f) corresponds to the first variation of the H^k -regularization norms in (4) (see (9)). The operator $\mathcal{B}$ in (20g) corresponds to the first variation of the regularization operator for w .

We can completely eliminate the adjoint variables w and p , the control equation (20g), and the constraint (20e) from (20) by simple algebraic manipulations. This is straightforward for

$w = 0$ (see [58]) and becomes slightly more involved for a non-zero w (see [59]). This elimination introduces the pseudo-differential operator $\mathcal{K}$ in (6). We refer to our preceding work for more details [58, 59]. Overall, we will arrive at the optimality conditions presented in §3.1. Computing variations of the weak form of the optimality conditions in §3.1 (which includes the operators $\mathcal{K}$ arising from the elimination of w and p) yields the PDE operators for the Newton step.

Remark 3

The derivation of the optimality conditions (20) is formal only. We note that the presentation of existence and uniqueness proofs for an optimal solution of (20) are beyond the scope of the present paper. In order for us to ensure the well-posedness of the forward problem, the differentiability of the objective functional and the constraints, and, ultimately, the existence and uniqueness for an optimal solution of the control problem, we have to make sure that the variables in our control formulation meet certain regularity requirements; we have to specify appropriate function spaces for the input images m_I , $I \in \{R, T\}$, and the velocity field $\mathbf{v}$. Several works of other authors have addressed these theoretical requirements in the context of related (optimal control) formulations; see, e.g., [9, 10, 19, 24, 25, 27, 28, 55, 73, 74]. These, e.g., include results for formulations that model images as functions of bounded variation [25, 74] or functions of Sobolev regularity [10, 19, 73], respectively. They consider H^1 [19, 24, 27, 28], H^2 [10, 55, 73], and H^3 [25] regularization models for $\mathbf{v}$, accounting for incompressible [24, 25] or near-incompressible [19, 27] velocities. It has also been suggested to stipulate adequate regularity requirements by introducing a diffusion operator into the transport problem [9, 54].

In our formulation, we model images as compactly supported, smooth functions; we use appropriate mollification and Gaussian smoothing to ensure that we meet these requirements. Numerically, we control the smoothness of the velocity by adjusting the weights for the regularization operator for $\mathbf{v}$ to ensure that we obtain a diffeomorphic map (up to numerical accuracy). Our experimental results suggest that we stably converge to a local optimal solution using our formulation. However, we note that we are not aware of a theoretical proof that H^1 -regularity for $\mathbf{v}$ and its divergence are sufficient to guarantee the existence of an optimal solution of our control problem in the theoretical limit. We also note that we observed instabilities if we stipulate H^1 -regularity for $\mathbf{v}$ only without controlling its divergence, in our numerical experiments. Instead of directly controlling the smoothness of $\mathbf{v}$, we can also (additionally) control the curl of $\mathbf{v}$; this will add additional regularity to our solution [4, 55]. For instance, adding an additional H^1 -regularization model for the curl of $\mathbf{v}$ will ensure that $\mathbf{v}$ is an H^2 -function. A rigorous proof remains open for future work. Finally, we note that we can change the regularization operators if our formulation for near-incompressible diffeomorphisms does not meet the theoretical requirements; all the derivations and algorithmic features presented here will still apply.

Appendix B. Stabilized RK2 Scheme

Here, we present the derivation of the stabilized RK2 scheme introduced in §3.3 for the transport equations that appear in our optimality system. We refer to [35, 53] for a general

discussion of this scheme. We start by deriving the antisymmetric form of the forward problem in (7a). Inserting and subtracting the term $\frac{1}{2}\nabla \cdot m\boldsymbol{\nu}$ and by using the identity $\nabla \cdot m\boldsymbol{\nu} = \nabla m \cdot \boldsymbol{\nu} + m\nabla \cdot \boldsymbol{\nu}$ we obtain

$$\partial_t m + \frac{1}{2}(\nabla m \cdot \boldsymbol{\nu} + \nabla \cdot m\boldsymbol{\nu} - m\nabla \cdot \boldsymbol{\nu}) = 0 \quad \text{in } \Omega \times (0, 1],$$

with periodic boundary conditions on Ω . We use this antisymmetric form as the forward operator. Computing first and second variations using this model results in

$$-\partial_t \lambda - \frac{1}{2}(\nabla \cdot \lambda\boldsymbol{\nu} + \nabla \lambda \cdot \boldsymbol{\nu} + \lambda\nabla \cdot \boldsymbol{\nu}) = 0 \quad \text{in } \Omega \times [0, 1),$$

$$\partial_t \tilde{m} + \frac{1}{2}(\nabla \tilde{m} \cdot \boldsymbol{\nu} - \tilde{m}\nabla \cdot \boldsymbol{\nu} + \nabla \cdot (\tilde{m}\boldsymbol{\nu} + m\tilde{\boldsymbol{\nu}}) + \nabla m \cdot \tilde{\boldsymbol{\nu}} - m\nabla \cdot \tilde{\boldsymbol{\nu}}) = 0 \quad \text{in } \Omega \times (0, 1],$$

$$-\partial_t \tilde{\lambda} - \frac{1}{2}(\nabla \tilde{\lambda} \cdot \boldsymbol{\nu} + \tilde{\lambda}\nabla \cdot \boldsymbol{\nu} + \nabla \cdot (\tilde{\lambda}\boldsymbol{\nu} + \lambda\tilde{\boldsymbol{\nu}}) + \nabla \lambda \cdot \tilde{\boldsymbol{\nu}} + \lambda\nabla \cdot \tilde{\boldsymbol{\nu}}) = 0 \quad \text{in } \Omega \times [0, 1),$$

for the adjoint, incremental state and incremental adjoint equation (notice, that some terms in the above equations will drop for the incompressible case, i.e., for $\nabla \cdot \boldsymbol{\nu} = 0$). Similarly, we obtain the integro-differential operators

$$\mathbf{b} = \int_0^1 \frac{1}{2}(\lambda\nabla m - m\nabla \lambda + \nabla(\lambda m))dt$$

and

$$\tilde{\mathbf{b}} = \int_0^1 \frac{1}{2}(\lambda\nabla \tilde{m} - \tilde{m}\nabla \lambda + \nabla(\lambda\tilde{m}) + \tilde{\lambda}\nabla m - m\nabla \tilde{\lambda} + \nabla(\tilde{\lambda}m))dt$$

for the reduced gradient in (6) and the Hessian matvec in (10), respectively.

Appendix C. Computational Complexity

We report the computational complexity as a function of the number of FFTs we have to compute in Tab. 12. A comparison of the timings for applying a single FFT and for one cubic spline interpolation step with respect to different grid sizes can be found in Tab. 13. We compute the characteristic $\mathbf{X}$ for the forward and the adjoint problems only once per iteration. When we evaluate the objective and the gradient we have to solve the state and the adjoint equation. This is when we compute the characteristic; we do not recompute it during the incremental solves. We assign these costs to the solution of the state and adjoint equation.

References

1. Adavani SS, Biros G. Fast algorithms for source identification problems with elliptic PDE constraints. *SIAM Journal on Imaging Sciences*. 2008; 3:791–808.
2. Adavani SS, Biros G. Multigrid algorithms for inverse problems with linear parabolic PDE constraints. *SIAM Journal on Scientific Computing*. 2008; 31:369–397.
3. Amit Y. A nonlinear variational problem for image matching. *SIAM Journal on Scientific Computing*. 1994; 15:207–224.
4. Amrouche C, Bernardi C, Dauge M, Girault V. Vector potentials in three-dimensional nonsmooth domains. *Mathematical Methods in the Applied Sciences*. 1998; 21:823–864.
5. Ashburner J. A fast diffeomorphic image registration algorithm. *NeuroImage*. 2007; 38:95–113. [PubMed: 17761438]
6. Ashburner J, Friston KJ. Diffeomorphic registration using geodesic shooting and Gauss-Newton optimisation. *NeuroImage*. 2011; 55:954–967. [PubMed: 21216294]
7. Avants BB, Epstein CL, Brossman M, Gee JC. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*. 2008; 12:26–41. [PubMed: 17659998]
8. Axelsson, O. Iterative solution methods. Cambridge University Press; 1996.
9. Barbu V, Marinocchi G. An optimal control approach to the optical flow problem. *Systems & Control Letters*. 2016; 87:1–9.
10. Beg MF, Miller MI, Trounev A, Younes L. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*. 2005; 61:139–157.
11. Benzi M. Preconditioning techniques for large linear systems: A survey. *Journal of Computational Physics*. 2002; 182:418–477.
12. Benzi M, Golub GH, Liesen J. Numerical solution of saddle point problems. *Acta Numerica*. 2005; 14:1–137.
13. Benzi M, Haber E, Taralli L. A preconditioning technique for a class of PDE-constrained optimization problems. *Advances in Computational Mathematics*. 2011; 35:149–173.
14. Biegler LT, Ghattas O, Heinkenschloss M, van Bloemen Waanders B. Large-scale PDE-constrained optimization, Springer. 2003; 3
15. Biros G, Doan G. A multilevel algorithm for inverse problems with elliptic PDE constraints. *Inverse Problems*. 2008; 24
16. Biros G, Ghattas O. Parallel Lagrange-Newton-Krylov-Schur methods for PDE-constrained optimization—Part I: The Krylov-Schur solver. *SIAM Journal on Scientific Computing*. 2005; 27:687–713.
17. 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 Journal on Scientific Computing*. 2005; 27:714–739.
18. Borzi A, Ito K, Kunisch K. An optimal control approach to optical flow computation. *International Journal for Numerical Methods in Fluids*. 2002; 40:231–240.
19. Borzi A, Ito K, Kunisch K. Optimal control formulation for determining optical flow. *SIAM Journal on Scientific Computing*. 2002; 24:818–847.
20. Borzi A, Schulz V. Multigrid methods for PDE optimization, *SIAM Review*. 2009; 51:361–395.
21. Borzi, A., Schulz, V. Computational optimization of systems governed by partial differential equations. SIAM; Philadelphia, Pennsylvania, US: 2012.
22. Burger M, Modersitzki J, Ruthotto L. A hyperelastic regularization energy for image registration. *SIAM Journal on Scientific Computing*. 2013; 35:B132–B148.
23. Cao Y, Miller MI, Winslow RL, Younes L. Large deformation diffeomorphic metric mapping of vector fields, *Medical Imaging. IEEE Transactions on*. 2005; 24:1216–1230.
24. Chen, K. PhD thesis. University of Bremen; 2011. Optimal control based image sequence interpolation.
25. Chen K, Lorenz DA. Image sequence interpolation using optimal control. *Journal of Mathematical Imaging and Vision*. 2011; 41:222–238.

26. Christensen GE, Geng X, Kuhl JG, Bruss J, Grabowski TJ, Pirwani IA, Vannier MW, Allen JS, Damasio H. Introduction to the non-rigid image registration evaluation project. *Proc Biomedical Image Registration*. 2006; LNCS 4057:128–135.
27. Crippa, G. PhD thesis. University of Zürich; 2007. The flow associated to weakly differentiable vector fields.
28. DiPerna RJ, Lions PL. Ordinary differential equations, transport theory and Sobolev spaces. *Inventiones Mathematicae*. 1989; 98:511–547.
29. Dontchev AL, Hager WW, Veliov VM. Second-order Runge–kutta approximations in control constrained optimal control. *SIAM Journal on Numerical Analysis*. 2000; 38:202–226.
30. Droske M, Rumpf M. A variational approach to non-rigid morphological registration. *SIAM Journal on Applied Mathematics*. 2003; 64:668–687.
31. Dupuis P, Gernander U, Miller MI. Variational problems on flows of diffeomorphisms for image matching. *Quarterly of Applied Mathematics*. 1998; 56:587–600.
32. Ewing RE, Wong H. A summary of numerical methods for time-dependent advection-dominated partial differential equations. *Journal of Computational and Applied Mathematics*. 2001; 128:423–445.
33. Fischer B, Modersitzki J. Ill-posed medicine – an introduction to image registration. *Inverse Problems*. 2008; 24:1–16.
34. Fohring J, Haber E, Ruthotto L. Geophysical imaging for fluid flow in porous media. *SIAM Journal on Scientific Computing*. 2014; 36:S218–S236.
35. Fornberg B. On a Fourier method for the integration of hyperbolic equations. *SIAM Journal on Numerical Analysis*. 1975; 12:509–527.
36. Frigo M, Johnson SG. The design and implementation of FFTW3. *Proc of the IEEE*. 2005; 93:216–231.
37. Gholami A, Mang A, Biros G. An inverse problem formulation for parameter estimation of a reaction-diffusion model of low grade gliomas. *Journal of Mathematical Biology*. 2016; 72:409–433. [PubMed: 25963601]
38. Gill PE, Murray W, Wright MH. *Practical optimization*, Academic Press, Waltham, Massachusetts, US. 1981; 9:12.
39. Gunzburger, MD. *Perspectives in flow control and optimization*. SIAM; Philadelphia, Pennsylvania, US: 2003.
40. Gurtin, ME. *An introduction to continuum mechanics*. Academic Press; 1981. vol 158 of *Mathematics in Science and Engineering*
41. Haber E, Ascher UM. Preconditioned all-at-once methods for large, sparse parameter estimation problems. *Inverse Problems*. 2001; 17:1847–1864.
42. Haber E, Modersitzki J. Numerical methods for volume preserving image registration. *Inverse Problems*. 2004; 20:1621–1638.
43. Haber E, Modersitzki J. Image registration with guaranteed displacement regularity. *International Journal of Computer Vision*. 2007; 71:361–372.
44. Hager WW. Runge-Kutta methods in optimal control and the transformed adjoint system. *Numerische Mathematik*. 2000; 87:247–282.
45. Hajnal, JV., Hill, DLG., Hawkes, DJ. *Medical Image Registration*. CRC Press; Boca Raton, Florida, US: 2001.
46. Hart GL, Zach C, Niethammer M. An optimal control approach for deformable registration. *Proc IEEE Conference on Computer Vision and Pattern Recognition*. 2009:9–16.
47. Hernandez M. Gauss-Newton inspired preconditioned optimization in large deformation diffeomorphic metric mapping. *Physics in Medicine and Biology*. 2014; 59:6085–6115. [PubMed: 25254606]
48. Hernandez M, Bossa MN, Olmos S. Registration of anatomical images using paths of diffeomorphisms parameterized with stationary vector field flows. *International Journal of Computer Vision*. 2009; 85:291–306.
49. Herzog R, Kunisch K. Algorithms for PDE-constrained optimization. *GAMM Mitteilungen*. 2010; 33:163–176.

50. Hestenes MR, Stiefel E. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*. 1952; 49:409–436.
51. Hinze, M., Pinnau, R., Ulbrich, M., Ulbrich, S. *Optimization with PDE constraints*. Springer; Berlin, DE: 2009.
52. Kalany, E. *Atmospheric modeling, data assimilation and predictability*. Oxford University Press; 2002.
53. Kreiss HO, Olinger J. Comparison of accurate methods for the integration of hyperbolic equations. *Tellus*. 1972; 24:199–215.
54. Kunisch K, Lu X. Optimal control for multi-phase fluid Stokes problems. *Nonlinear Analysis*. 2011; 74:585–599.
55. Lee E, Gunzburger M. An optimal control formulation of an image registration problem. *Journal of Mathematical Imaging and Vision*. 2010; 36:69–80.
56. Lions, JL. *Optimal control of systems governed by partial differential equations*. Springer; 1971.
57. Lorenzi M, Pennec X. Geodesics, parallel transport and one-parameter subgroups for diffeomorphic image registration. *International Journal of Computer Vision*. 2013; 105:111–127.
58. Mang A, Biros G. An inexact Newton–Krylov algorithm for constrained diffeomorphic image registration. *SIAM Journal on Imaging Sciences*. 2015; 8:1030–1069. [PubMed: 27617052]
59. Mang A, Biros G. Constrained H^1 -regularization schemes for diffeomorphic image registration. *SIAM Journal on Imaging Sciences*. 2016; 9:1154–1194. [PubMed: 29075361]
60. 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. *Medical Physics*. 2012; 39:4444–4459. [PubMed: 22830777]
61. Modersitzki, J. *Numerical methods for image registration*. Oxford University Press; New York: 2004.
62. Modersitzki, J. *FAIR: Flexible algorithms for image registration*. SIAM; Philadelphia, Pennsylvania, US: 2009.
63. Nocedal, J., Wright, SJ. *Numerical Optimization*. Springer; New York, New York, US: 2006.
64. Ruhnau P, Schnörr C. Optical Stokes flow estimation: An imaging-based control approach. *Experiments in Fluids*. 2007; 42:61–78.
65. Sdika M. A fast nonrigid image registration with constraints on the Jacobian using large scale constrained optimization. *Medical Imaging, IEEE Transactions on*. 2008; 27:271–281.
66. Simoncini V. Reduced order solution of structured linear systems arising in certain PDE-constrained optimization problems. *Computational Optimization and Applications*. 2012; 53:591–617.
67. Sotiras A, Davatzikos C, Paragios N. Deformable medical image registration: A survey. *Medical Imaging, IEEE Transactions on*. 2013; 32:1153–1190.
68. Staniforth A, Côté J. Semi-Lagrangian integration schemes for atmospheric models—A review. *Monthly Weather Review*. 1991; 119:2206–2223.
69. Stoll M, Breiten T. A low-rank in time approach to PDE-constrained optimization. *SIAM Journal on Scientific Computing*. 2015; 37:B1–B29.
70. Trounev A. Diffeomorphism groups and pattern matching in image analysis. *International Journal of Computer Vision*. 1998; 28:213–221.
71. Vercauteren T, Pennec X, Perchant A, Ayache N. Symmetric log-domain diffeomorphic registration: A demons-based approach. *Proc Medical Image Computing and Computer-Assisted Intervention*. 2008; LNCS 5241:754–761.
72. Vercauteren T, Pennec X, Perchant A, Ayache N. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*. 2009; 45:S61–S72. [PubMed: 19041946]
73. Vialard FX, Risser L, Rueckert D, Cotter CJ. Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation. *International Journal of Computer Vision*. 2012; 97:229–241.
74. Vialard FX, Santambrogio F. Extension to BV functions of the large deformation diffeomorphism matching approach. *Comptes Rendus de l'Académie des Sciences*. 2009; 347:27–32.

Fig. 1.

Tracing the characteristic X in a semi-Lagrangian scheme. We start with a regular grid Ω^h (dark orange points on the right) consisting of coordinates x at time point t^j . We assume we have already computed the intermediate solution u^h of a given transport equation at time point t^{j-1} ; we know the input data on the regular grid at time point t^{j-1} (the regular grid nodes are illustrated in light gray). In a first step, we trace back the characteristic X by solving (14) backward in time subject to the initial condition $X(t = t^j) = x$. Once we have found the characteristic (black line in the figure on the left) we can—in a second step—assign the value of u^h at t^j given at the departure point X_D (dark orange point on the left) to x at t^j based on some interpolation model. We illustrate the grid of departure points in light orange and the original grid in gray (left figure).

Fig. 2.

Synthetic test problems. From left to right: reference image m_R ; template image m_T ; v^1 component of velocity field; and v^2 component of velocity field. The intensity values of the images are in $[0, 1]$. The magnitude of the velocity field is in $[-0.5, 0.5]$ (top row) and $[-1, 1]$ (bottom row), respectively.

Fig. 3.

Registration problems. Top left: *UT* images (synthetic problem); top right: *HAND* images [3, 62]; bottom left: *HEART* images; bottom right: *BRAIN* images [26]. The intensity values for these images are normalized to $[0, 1]$. We provide (from left to right for each set of images) the reference image m_R , the template image m_T , and the residual differences between these images prior to registration.

Fig. 4.

Self-convergence for the forward solver. We illustrate solutions of the forward problem (state equation; see (7a)) for the synthetic test problems in Fig. 2 (top row: *SMOOTH A*;

bottom rows: *SMOOTH B*). We report results for different grid sizes $\mathbf{n}_x = (n_x^1, n_x^2)^\top$. We use the same number of time steps (*CFL* number of 0.2) for all *PDE* solvers.

Fig. 5.

Convergence results for different strategies to precondition the reduced space KKT system.

We report exemplary trends of the relative residual $\|r_k\|_{rel} := \|r_k\|_2 / \|r_0\|_2$ with respect to the iteration number k . We report results for different images (top row: *BRAIN*; bottom row: *HAND*; grid size $(256, 256)^T$) with respect to varying regularization weights β_v (left column: $\beta_v = 1E-1$; middle column: $\beta_v = 1E-2$; right column: $\beta_v = 1E-3$). We solve the system at the true solution available for the considered synthetic test problems. We use an H^2 -regularization model (compressible diffeomorphism). We use a PCG method with a tolerance of $1E-6$ to solve this system. We report results for the regularization preconditioner P_{REG} (red curve) and the nested preconditioner P_{2L} . We use two solvers to invert the preconditioner: $PCG(1E-1)$ (blue curve) and $CHEB(10)$ (green curve). The results correspond to those reported in Tab. 8.

Table 1

Commonly used notation and symbols.

Symbol/Notation	Description
CFL	Courant–Friedrichs–Lewy (condition)
FFT	Fast Fourier Transform
GN	Gauss–Newton
KKT	Karush–Kuhn–Tucker (system)
matvec	(Hessian) matrix–vector product
PCG	Preconditioned Conjugate Gradient (method)
PCG(ϵ)	PCG, where $\epsilon > 0$ indicates the used tolerance
PDE	partial differential equation
PDE solve	solution of a hyperbolic transport equation
RK2	2nd order Runge–Kutta (method)
RK2(c)	RK2 method, where c indicates the employed CFL number
RK2A	RK2 scheme based on an antisymmetric form
RK2A(c)	RK2A method, where c indicates the employed CFL number
SL	semi-Lagrangian (method)
SL(c)	SL method, where c indicates the employed CFL number
d	spatial dimensionality; typically $d \in \{2, 3\}$
Ω	spatial domain; $\Omega := (-\pi, \pi)^d \subset \mathbf{R}^d$ with boundary $\partial\Omega$ and closure $\overline{\Omega} := \Omega \cup \partial\Omega$
$\mathbf{x}$	spatial coordinate; $\mathbf{x} := (x^1, \dots, x^d)^\top \in \mathbf{R}^d$
m_R	reference image; $m_R: \overline{\Omega} \rightarrow \mathbf{R}$
m_T	template image; $m_T: \overline{\Omega} \rightarrow \mathbf{R}$
m	state variable (transported intensities); $m: \overline{\Omega} \times [0, 1] \rightarrow \mathbf{R}$
m_1	deformed template image (state variable at $t = 1$); $m_1: \overline{\Omega} \rightarrow \mathbf{R}$
λ	adjoint variable (transport equation); $\lambda: \overline{\Omega} \times [0, 1] \rightarrow \mathbf{R}$
p	adjoint variable (incompressibility constraint); $p: \overline{\Omega} \rightarrow \mathbf{R}$
$\mathbf{v}$	control variable (stationary velocity field); $\mathbf{v}: \overline{\Omega} \rightarrow \mathbf{R}^d$
w	control variable (mass source); $w: \overline{\Omega} \rightarrow \mathbf{R}$
$\mathbf{b}$	body force; $\mathbf{b}: \overline{\Omega} \rightarrow \mathbf{R}^d$
$\mathcal{H}$	(reduced) Hessian
$\mathbf{g}$	(reduced) gradient
$\mathbf{y}$	Eulerian (pullback) deformation map
$\mathbf{F}$	deformation gradient at $t = 1$ (computed from $\mathbf{v}$); $\mathbf{F}: \overline{\Omega} \rightarrow \mathbf{R}^{d \times d}$; $\mathbf{F} := (\nabla \mathbf{y})^{-1}$

Symbol/Notation	Description
β_v	regularization parameter for the control v
β_w	regularization parameter for the control w
$\mathcal{A}$	regularization operator (variation of regularization model acting on v)
∂_i	partial derivative with respect to $x^i, i = 1, \dots, d$
∂_t	partial derivative with respect to time
D_t	Lagrangian derivative
∇	gradient operator (acts on scalar and vector fields)
Δ	Laplacian operator (acts on scalar and vector fields)
$\nabla \cdot$	divergence operator (acts on vector and 2nd order tensor fields)
$\langle \cdot, \cdot \rangle_{L^2(\mathcal{X})}$	L^2 inner product on $\mathcal{X}$

Table 2

Self-convergence for the *RK2*, the *RK2A*, and the *SL* method for the numerical integration of the state (see (7a); results reported in top block) and the adjoint (see (8a); results reported in bottom block) equation. We report the relative $\mathcal{L}$ -error $\|\delta u^h\|_{\text{rel}}$ between solutions for the state ($u^h = m_1^h$) and the adjoint ($u^h = \lambda_0^h$) equation computed on a grid of size n_x and a grid of size $2n_x$. We use a *CFL* number of 0.2 for the *RK2* and the *RK2A* method, and a *CFL* number of 0.2, 1, and 5, for the *SL* method; we provide the associated number of time points n_t . We report errors for different grid sizes and test problems (top block: *SMOOTH A*; bottom block: *SMOOTH B*; see Fig. 2);

		RK2(0.2)				RK2A(0.2)				SL(0.2)				SL(1)				SL(5)				
		n_x	run	n_t	$\ \delta u^h\ _{\text{rel}}$	time	run	n_t	$\ \delta u^h\ _{\text{rel}}$	time	run	n_t	$\ \delta u^h\ _{\text{rel}}$	time	run	n_t	$\ \delta u^h\ _{\text{rel}}$	time	run	n_t	$\ \delta u^h\ _{\text{rel}}$	time
STATE EQ	A	64		26	1.25E-5	3.25E-1		26	1.25E-5	6.64E-1		26	3.79E-6	8.53E-1		6	4.93E-5	2.70E-1		2	3.19E-4	5.88E-2
		128		51	3.28E-6	1.63		51	3.28E-6	3.20		51	8.63E-7	3.85		11	1.52E-5	1.32		3	1.69E-4	2.41E-1
		256		102	8.19E-7	1.07E+1		102	8.19E-7	2.04E+1		102	2.00E-7	3.04E+1		21	4.26E-6	6.81		5	6.75E-5	1.56
		512		204	2.05E-7	8.44E+1		204	2.05E-7	1.61E+2		204	4.81E-8	2.53E+2		41	1.14E-6	5.09E+1		9	2.22E-5	1.14E+1
	B	64		102	2.36E-1	9.20E-1		102	7.67E-2	2.38		102	9.93E-2	3.33		21	8.64E-2	6.78E-1		5	7.35E-2	2.24E-1
		128		204	2.81E-2	6.44		204	9.83E-3	1.26E+1		204	1.06E-2	1.77E+1		41	9.67E-3	4.60		9	9.81E-3	1.15
		256		408	***	4.23E+1		408	1.20E-4	7.86E+1		408	4.03E-4	1.16E+2		82	3.67E-4	2.51E+1		17	1.19E-3	6.94
		512		815	***	3.43E+2		815	3.88E-6	6.52E+2		815	1.88E-5	1.06E+3		163	2.74E-5	1.96E+2		33	3.10E-4	4.23E+1
	A	64		26	1.28E-4	2.14E-1		26	1.28E-4	6.60E-1		26	8.33E-5	7.33E-1		6	1.39E-3	3.05E-1		2	8.36E-3	1.64E-1
		128		51	3.40E-5	1.36		51	3.40E-5	3.42		51	2.25E-5	6.04		11	4.47E-4	1.30		3	4.61E-3	5.30E-1
		256		102	8.52E-6	9.29		102	8.52E-6	2.36E+1		102	5.66E-6	3.58E+1		21	1.28E-4	8.24		5	1.93E-3	2.08
		512		204	2.13E-6	7.09E+1		204	2.13E-6	1.63E+2		204	1.42E-6	2.92E+2		41	3.46E-5	6.37E+1		9	6.54E-4	1.56E+1
	B	64		102	4.04E-1	8.21E-1		102	4.19E-1	2.23		102	6.07E-1	3.27		21	5.30E-1	7.79E-1		5	4.34E-1	2.82E-1
		128		204	5.72E-2	6.31		204	5.94E-2	1.14E+1		204	6.56E-2	1.80E+1		41	6.06E-2	3.39		9	5.69E-2	1.12
		256		408	***	4.20E+1		408	7.31E-4	7.78E+1		408	2.37E-3	1.46E+2		82	2.04E-3	2.31E+1		17	3.92E-3	4.81
		512		815	***	2.81E+2		815	1.17E-5	6.42E+2		815	1.08E-4	1.19E+3		163	1.13E-4	2.29E+2		33	1.01E-3	5.03E+1

ADJOINT EQ

SLAM J Sci Comput. Author manuscript; available in PMC 2018 November 21.

*** indicates that the solver became unstable (not due to a violation of the CFL condition; see text for details). We also report the time to solution (in seconds).

Table 3

Convergence of the *SL* method to a reference solution computed via the *RK2A* scheme. We compute the reference solution on a grid of size $n_x = (512, 512)^T$ using the *RK2A* scheme with a *CFL* number of 0.2. We report results for varying discretization sizes. We report the *CFL* number c , the associated number of time steps n_t , and the relative ℓ_2 -error between the solution for the *SL* scheme and the reference solution computed via the *RK2A* scheme. We also report errors for the *RK2A* method as a reference (self convergence). We report results for the state equation (two blocks on the left; see (7a)) and the adjoint equation (two blocks on the right; see (8a)). We consider the test problems *SMOOTH A* and *SMOOTH B* in Fig. 2 for the velocity field and to set up the initial and terminal conditions, respectively.

STATE EQ															ADJOINT EQ								
SMOOTH A															SMOOTH B			SMOOTH A			SMOOTH B		
n_x^i	c	n_t	run	SL	run	RK2A	n_t	run	SL	run	RK2A	n_t	run	SL	run	RK2A	n_t	run	SL	run	RK2A		
64	10	2		5.84E-4		-	3		8.11E-2		-	2		1.57E-2		-	3		4.18E-1		-		
	5	2		5.84E-4		-	5		7.58E-2		-	2		1.57E-2		-	5		4.41E-1		-		
	2	3		2.67E-4		-	11		7.99E-2		-	3		7.44E-3		-	11		4.90E-1		-		
	1	6		7.04E-5		-	21		8.57E-2		-	6		2.01E-3		-	21		5.31E-1		-		
	0.2	26		5.10E-6		1.66E-5	102		9.81E-2		7.66E-2	26		1.14E-4		1.71E-4	102		6.06E-1		4.19E-1		
128	10	2		5.84E-4		-	5		1.78E-2		-	2		1.57E-2		-	5		7.10E-2		-		
	5	3		2.67E-4		-	9		1.08E-2		-	3		7.43E-3		-	9		5.91E-2		-		
	2	6		6.90E-5		-	21		9.55E-3		-	6		2.01E-3		-	21		5.90E-2		-		
	1	11		2.11E-5		-	41		9.84E-3		-	11		6.22E-4		-	41		6.14E-2		-		
	0.2	51		1.33E-6		4.10E-6	204		1.07E-2		9.83E-3	51		3.08E-5		4.25E-5	204		6.65E-2		5.94E-2		
256	10	3		2.67E-4		-	9		5.06E-3		-	3		7.43E-3		-	9		1.55E-2		-		
	5	5		9.86E-5		-	17		1.60E-3		-	5		2.84E-3		-	17		5.21E-3		-		
	2	11		2.10E-5		-	41		5.02E-4		-	11		6.21E-4		-	41		2.09E-3		-		
	1	21		5.96E-6		-	82		3.93E-4		-	21		1.75E-4		-	82		2.12E-3		-		
	0.2	102		4.92E-7		8.19E-7	408		4.21E-4		1.20E-4	102		8.61E-6		8.52E-6	408		2.47E-3		7.31E-4		
512	10	5		9.86E-5		-	17		1.47E-3		-	5		2.84E-3		-	17		4.60E-3		-		
	5	9		3.11E-5		-	33		4.15E-4		-	9		9.18E-4		-	33		1.35E-3		-		
	2	21		5.94E-6		-	82		8.61E-5		-	21		1.75E-4		-	82		2.85E-4		-		

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

ADJOINT EQ											
STATE EQ											
SMOOTH A						SMOOTH B					
SMOOTH A						SMOOTH B					
η_x^i	c	n_t	run	SL	run	SL	run	RK2A	n_t	run	SL
1	41	163	-	RK2A	run	SL	run	RK2A	run	SL	run
0.2	204	815	0	RK2A	run	SL	run	RK2A	run	SL	run

Table 4

Convergence of the reduced gradient computed via the SL method to the gradient computed via the $RK2A$ method. We evaluate the reference gradient on a grid of size $\mathbf{n}_x = (512, 512)^T$ via the $RK2A$ method with a CFL number of 0.2. For the SL method the reduced gradient is computed on a grid of size $\mathbf{n}_x = (256, 256)^T$ and $\mathbf{n}_x = (512, 512)^T$ with a varying number of time steps n_t . We report the CFL number c , the associated number of time steps n_t , the relative $\mathcal{L}$ -error between numerical approximations to the reduced gradient $\mathbf{g}^h$, and the wall-clock time for the evaluation of $\mathbf{g}^h$. We consider the test problems $SMOOTH A$ and $SMOOTH B$ in Fig. 2 as input data. As a reference, we also provide relative errors for the $RK2A$ scheme.

SMOOTH A										SMOOTH B			
n_x^i	c	run	n_t	SL	time	RK2A	time	run	n_t	SL	time	RK2A	time
256	10		3	2.54E-3	4.51E-1	-	-		9	2.28E-2	1.29E0	-	-
	5		5	9.08E-4	1.02E0	-	-		17	2.22E-2	2.19E0	-	-
	2		11	2.19E-4	1.77E0	-	-		41	2.21E-2	3.87E0	-	-
	1		21	1.31E-4	2.59E0	-	-		82	2.20E-2	9.55E0	-	-
	0.2		102	1.21E-4	1.17E+1	1.21E-4	8.57E0		408	2.20E-2	3.00E+1	2.19E-2	2.74E+1
512	10		5	9.00E-4	3.66E0	-	-		17	1.42E-3	7.59E0	-	-
	5		9	2.74E-4	3.74E0	-	-		33	3.94E-4	1.19E+1	-	-
	2		21	4.93E-5	8.72E0	-	-		82	7.89E-5	2.86E+1	-	-
	1		41	1.23E-5	1.49E+1	-	-		163	3.60E-5	6.13E+1	-	-
	0.2		204	4.87E-7	7.32E+1	0	4.96E+1		815	2.90E-5	2.65E+2	0	2.04E+2

Table 5

Relative adjoint error δ_{ADJ} (see text for details) for a grid size of $\mathbf{n}_x = (256, 256)^\top$ and a varying number of time steps n_t . We consider the test problem *SMOOTH A* in Fig. 2. We report the *CFL* number c , the associated number of time steps n_b and the relative errors for the *SL* and the *RK2A* scheme.

c	n_t	SL	RK2A
10	3	3.28E-3	–
5	5	1.26E-3	–
2	11	2.75E-4	–
1	21	7.72E-5	–
0.2	102	3.30E-6	1.24E-16

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Estimates for the largest eigenvalue $e_{\max}$ of $\mathbf{P}_{2L}$ during the course of the inversion. We limit this experiment to a compressible diffeomorphism (H^2 -regularization). We report results for the $\mathcal{U}T(256 \times 256)$, the $HAND(128 \times 128)$, the $HEART(192 \times 192)$, and the $BRAIN(256 \times 300)$ images (see Fig. 3). We terminate the inversion if the relative change of the gradient is equal or smaller than three orders of magnitude or if the ℓ^2 -norm of the reduced gradient is smaller or equal to $1E-5$. We consider different regularization weights β_v . We estimate $e_{\max}$ every time we apply the preconditioner $\mathbf{P}_{2L}$. We report the initial estimate for $e_{\max}$ (zero velocity field), and the min, mean, and max values of the estimates computed during the course of the entire inversion.

Table 6

	run	β_v	$e_{\max,0}$	min	max	mean
UT		1E-1	1.07E+2	9.26E+1	1.07E+2	9.34E+1
		1E-2	1.06E+3	8.19E+2	1.06E+3	8.59E+2
		1E-3	1.06E+4	7.88E+3	1.06E+4	8.36E+3
HAND		1E-1	2.73E+1	2.50E+1	2.73E+1	2.54E+1
		1E-2	2.64E+2	2.20E+2	2.64E+2	2.24E+2
		1E-3	2.63E+3	2.13E+3	2.63E+3	2.16E+3
HEART		1E-1	4.22E+1	4.22E+1	4.23E+1	4.23E+1
		1E-2	4.13E+2	4.13E+2	4.14E+2	4.14E+2
		1E-3	4.13E+3	4.12E+3	4.13E+3	4.13E+3
BRAIN		1E-1	2.97E+1	2.92E+1	2.97E+1	2.94E+1
		1E-2	2.88E+2	2.82E+2	2.90E+2	2.89E+2
		1E-3	2.87E+3	2.81E+3	2.87E+3	2.83E+3

Error between the true solution $\hat{q}^{h*}$ and the numerical solution $\hat{q}^h$ of the KKT system for different schemes to precondition the reduced space Hessian. We report the absolute and the relative $\mathcal{L}$ -error between $\hat{q}^{h*}$ and $\hat{q}^h$. We report results for different preconditioners ($\mathbf{P}_{REG}, \mathbf{P}_{2L}$), different choices for the PDE solver ($SL(c)$ for different CFL numbers c and $RK2A(0.2)$), and different choices for the method to solve for the action of the inverse of the preconditioner ($CHEB(10)$ and $PCG(1E-1)$). We solve for $\hat{q}^h$ using a PCG method with a tolerance of $1E-12$. We consider the test problem $SMOOTH A$ as in Fig. 2 to set up the problem. We solve the system on a grid of size 256×256 .

Table 7

run	PC	PDE solver	PC solver	$\ \delta\ _2$	$\ \delta\ _{2,rel}$
$\mathbf{P}_{REG}$	$\mathbf{P}_{2L}$	RK2A(0.2)	–	7.49E-13	4.14E-14
		SL(0.2)	CHEB(10)	8.17E-14	4.51E-15
$\mathbf{P}_{2L}$		SL(1)	PCG(1E-1)	1.63E-12	9.03E-14
		SL(1)	CHEB(10)	1.23E-13	6.77E-15
		SL(2)	PCG(1E-1)	1.14E-12	6.30E-14
		SL(2)	CHEB(10)	1.89E-13	1.05E-14
		SL(5)	PCG(1E-1)	4.49E-12	2.48E-13
		SL(5)	CHEB(10)	5.22E-13	2.88E-14
		SL(10)	PCG(1E-1)	2.06E-12	1.14E-13
		SL(10)	CHEB(10)	2.06E-12	1.14E-13

Table 8

Convergence results for different strategies to precondition the reduced space KKT system. We report results for two preconditioners—our original preconditioner based on the regularization operator (P_{REG}) and the proposed, nested preconditioner (P_{2L}). We solve the reduced space KKT system via a PCG method with a tolerance of $1E-6$. We use two different solvers for the latter to invert the preconditioner—a PCG method with a tolerance that is $1E-1$ times the tolerance of the PCG method used to solve the reduced space KKT system (i.e., a tolerance of $1E-7$) and a $CHEB$ method with a fixed number of 10 iterations. We consider a compressible diffeomorphism with an H^2 -regularization model. We report results for different images ($SMOOTH$ A, BRAIN, and HAND), for different regularization weights β_v , and a varying grid sizes n_x (grid convergence; number of unknowns $n=2n_x^1 n_x^2$). We solve the reduced space KKT system at the true solution v^{h*} ; the velocity field v^{h*} corresponds to the test problem $SMOOTH$ A. We consider the $RK2A$ method with a CFL number of 0.2 for the regularization preconditioner and the SL scheme for the nested preconditioner with a CFL number of 5. We report (i) the number of PCG iterations until convergence, (ii) the time spent on the Hessian matvecs (in seconds), (iii) the percentage of that time spent on inverting the preconditioner (if applicable), and (iv) the speedup compared to our original preconditioner (regularization preconditioner in combination with the $RK2A$ scheme).

n	β_v	P	PDE solver	PC solver	SMOOTH A					HAND					BRAIN				
					run	iter	time	% PC	speedup	run	iter	time	% PC	speedup	run	iter	time	% PC	speedup
8192	1E-1	P_{REG}	RK2A(0.2)	—	—	4	2.67	—	—	—	19	7.12	—	—	—	21	8.39	—	—
		P_{2L}	SL(5)	PCG(1E-1)	—	2	3.08	64.04%	8.66E-1	—	7	1.67E+1	90.42%	4.27E-1	—	7	1.66E+1	91.15%	5.05E-1
		P_{2L}	SL(5)	CHEB(10)	—	2	2.36	55.81%	1.13	—	6	6.19	75.56%	1.15	—	7	6.44	75.69%	1.30
	1E-2	P_{REG}	RK2A(0.2)	—	—	6	3.14	—	—	—	47	1.59E+1	—	—	—	53	1.76E+1	—	—
		P_{2L}	SL(5)	PCG(1E-1)	—	2	4.31	73.56%	7.29E-1	—	8	4.60E+1	95.95%	3.47E-1	—	8	4.39E+1	96.10%	3.99E-1
		P_{2L}	SL(5)	CHEB(10)	—	3	3.19	62.15%	9.83E-1	—	8	8.26	79.30%	1.93	—	7	6.98	76.15%	2.52
	1E-3	P_{REG}	RK2A(0.2)	—	—	16	6.11	—	—	—	138	4.26E+1	—	—	—	161	4.96E+1	—	—
		P_{2L}	SL(5)	PCG(1E-1)	—	2	7.77	84.55%	7.86E-1	—	10	1.66E+2	98.69%	2.56E-1	—	11	2.24E+2	98.98%	2.22E-1
		P_{2L}	SL(5)	CHEB(10)	—	6	5.93	75.31%	1.03	—	24	2.16E+1	83.93%	1.97	—	22	2.04E+1	82.94%	2.43
32768	1E-1	P_{REG}	RK2A(0.2)	—	—	4	1.89E+1	—	—	—	22	4.31E+1	—	—	—	26	5.93E+1	—	—
		P_{2L}	SL(5)	PCG(1E-1)	—	2	3.53	47.19%	5.37	—	7	2.73E+1	85.64%	1.58	—	7	2.95E+1	88.16%	2.01
		P_{2L}	SL(5)	CHEB(10)	—	2	4.16	50.82%	4.56	—	6	9.70	64.92%	4.44	—	6	9.04	71.28%	6.56
	1E-2	P_{REG}	RK2A(0.2)	—	—	6	2.35E+1	—	—	—	54	1.07E+2	—	—	—	74	1.56E+2	—	—
		P_{2L}	SL(5)	PCG(1E-1)	—	3	7.14	64.44%	3.29	—	7	6.50E+1	94.88%	1.65	—	7	7.24E+1	96.04%	2.16
		P_{2L}	SL(5)	CHEB(10)	—	3	5.16	55.89%	4.55	—	9	1.26E+1	69.26%	8.51	—	10	1.44E+1	69.92%	1.09E+1

SMOOTH A																			HAND				BRAIN			
n	β_c	P	PDE solver	PC solver	run	iter	time	% PC	speedup	run	iter	time	% PC	speedup	run	iter	time	% PC	speedup							
131072	1E-3	P_{REG}	RK2A(0.2)	—	16	3.85E+1	—	—	—	160	3.45E+2	—	—	—	224	5.47E+2	—	—	—							
		$P_{2\text{L}}$	SL(5)	PCG(1E-1)	3	1.27E+1	77.88%	3.04	—	9	2.53E+2	98.27%	1.36	—	10	3.49E+2	98.82%	1.57	—							
		$P_{2\text{L}}$	SL(5)	CHEB(10)	6	8.51	67.56%	4.52	—	27	3.92E+1	73.91%	8.80	—	31	4.10E+1	76.74%	1.34E+1	—							
131072	1E-1	P_{REG}	RK2A(0.2)	—	4	6.62E+1	—	—	—	25	2.71E+2	—	—	—	33	3.52E+2	—	—	—							
		$P_{2\text{L}}$	SL(5)	PCG(1E-1)	2	1.24E+1	31.60%	5.33	—	6	7.17E+1	78.58%	3.77	—	6	1.02E+2	82.62%	3.46	—							
		$P_{2\text{L}}$	SL(5)	CHEB(10)	2	1.33E+1	46.49%	4.97	—	5	2.58E+1	53.36%	1.05E+1	—	5	3.07E+1	56.96%	1.15E+1	—							
	1E-2	P_{REG}	RK2A(0.2)	—	6	8.38E+1	—	—	—	63	6.61E+2	—	—	—	92	9.42E+2	—	—	—							
		$P_{2\text{L}}$	SL(5)	PCG(1E-1)	2	1.56E+1	39.61%	5.37	—	7	2.17E+2	92.20%	3.05	—	7	2.51E+2	94.07%	3.76	—							
		$P_{2\text{L}}$	SL(5)	CHEB(10)	3	2.06E+1	47.60%	4.06	—	11	5.75E+1	57.39%	1.15E+1	—	12	5.86E+1	59.85%	1.61E+1	—							
	1E-3	P_{REG}	RK2A(0.2)	—	16	1.66E+2	—	—	—	188	1.74E+3	—	—	—	279	2.29E+3	—	—	—							
		$P_{2\text{L}}$	SL(5)	PCG(1E-1)	3	3.08E+1	75.16%	5.39	—	9	6.73E+2	97.64%	2.58	—	8	9.54E+2	98.19%	2.40	—							
		$P_{2\text{L}}$	SL(5)	CHEB(10)	6	3.04E+1	57.77%	5.46	—	33	1.37E+2	65.15%	1.27E+1	—	38	1.74E+2	63.03%	1.32E+1	—							
524288	1E-1	P_{REG}	RK2A(0.2)	—	4	5.40E+2	—	—	—	25	1.97E+3	—	—	—	37	2.68E+3	—	—	—							
		$P_{2\text{L}}$	SL(5)	PCG(1E-1)	2	6.20E+1	29.01%	8.71	—	5	3.54E+2	77.07%	5.58	—	5	4.25E+2	83.31%	6.31	—							
		$P_{2\text{L}}$	SL(5)	CHEB(10)	2	7.15E+1	35.60%	7.55	—	4	1.21E+2	46.50%	1.63E+1	—	5	1.60E+2	50.11%	1.68E+1	—							
	1E-2	P_{REG}	RK2A(0.2)	—	6	6.49E+2	—	—	—	67	4.61E+3	—	—	—	103	7.12E+3	—	—	—							
		$P_{2\text{L}}$	SL(5)	PCG(1E-1)	2	7.12E+1	36.97%	9.11	—	6	8.72E+2	90.68%	5.29	—	6	1.20E+3	93.45%	5.92	—							
		$P_{2\text{L}}$	SL(5)	CHEB(10)	3	8.31E+1	38.18%	7.81	—	11	2.96E+2	55.59%	1.56E+1	—	14	3.51E+2	56.00%	2.03E+1	—							
	1E-3	P_{REG}	RK2A(0.2)	—	16	1.31E+3	—	—	—	196	1.31E+4	—	—	—	310	2.10E+4	—	—	—							
		$P_{2\text{L}}$	SL(5)	PCG(1E-1)	2	1.27E+2	59.84%	1.03E+1	—	7	2.93E+3	96.85%	4.47	—	7	4.58E+3	97.99%	4.58	—							
		$P_{2\text{L}}$	SL(5)	CHEB(10)	6	1.68E+2	49.40%	7.77	—	35	8.89E+2	60.45%	1.47E+1	—	46	1.17E+3	61.08%	1.80E+1	—							

Table 9

Convergence results for the inversion using our formulation for a compressible diffeomorphism (H^2 -regularization). We report results for registering different sets of images using our original preconditioner (P_{REG} ; $RK2A$ scheme with a CFL number of 0.2) and the proposed preconditioner (P_{2L} ; SL scheme with a CFL number of 5 and 10; $CHEB$ method with a fixed number of 10 iterations). We report results for different registration problems: $HAND$ (grid sizes: 128×128 ; 256×256 ; and 512×512), $HEART$ (grid size 192×192), and $BRAIN$ (grid size: 256×300); see Fig. 3. We study convergence as a function of the grid size ($HAND$ images; number of unknowns $n=2n_x n_y^2$) and as a function of the regularization parameter β_v ($HAND$, $HEART$, and $BRAIN$ images). We terminate the inversion if the relative change of the $\mathcal{P}$ -norm of the reduced gradient is at least two orders of magnitude or if the $\mathcal{P}$ -norm of the gradient is smaller or equal to $1E-5$. We report (i) the relative change of the reduced gradient $\|g^{\star}\|_{\mathcal{P}}$, (ii) the relative change of the residual $\|r\|_{rel}$ (L^2 -distance between m_R and m_I), (iii) the number of outer iterations, (iv) the number of Hessian matvecs, (v) the time to solution, and (vi) the speedup compared to our original scheme.

n	β	run	P	PDE solver	PC solver	$\lg^{\star}\ r\ _{rel}$	iter	matvecs	time	speedup	
HAND											
32768	1.00E-1		P_{REG}	RK2A(0.2)	—	4.85E-3	2.42E-1	8	58	9.45E+1	—
			P_{2L}	SL(5)	CHEB(10)	6.82E-3	2.42E-1	8	21	1.72E+1	5.50
			P_{REG}	RK2A(0.2)	—	8.39E-3	1.00E-1	8	97	1.55E+2	—
			P_{2L}	SL(5)	CHEB(10)	5.42E-3	9.99E-2	9	30	2.23E+1	6.95
			P_{REG}	RK2A(0.2)	—	8.59E-3	6.48E-2	11	401	9.89E+2	—
			P_{2L}	SL(5)	CHEB(10)	8.61E-3	6.50E-2	11	67	6.63E+1	1.49E+1
131072	1.00E-1		P_{REG}	RK2A(0.2)	—	9.04E-3	3.32E-1	12	113	7.71E+2	—
			P_{2L}	SL(5)	CHEB(10)	7.27E-3	3.32E-1	13	39	9.31E+1	8.28
			P_{REG}	RK2A(0.2)	—	9.67E-3	2.00E-1	11	159	1.17E+3	—
			P_{2L}	SL(5)	CHEB(10)	3.52E-3	1.99E-1	14	60	1.72E+2	6.81
			P_{REG}	RK2A(0.2)	—	9.59E-3	1.57E-1	17	758	1.06E+4	—
			P_{2L}	SL(5)	CHEB(10)	8.44E-3	1.57E-1	18	150	6.00E+2	1.76E+1
524288	1.00E-1		P_{REG}	RK2A(0.2)	—	1.06E-2	3.40E-1	14	134	6.38E+3	—
			P_{2L}	SL(5)	CHEB(10)	9.83E-3	3.40E-1	15	46	4.81E+2	1.33E+1
			P_{REG}	RK2A(0.2)	—	1.01E-2	2.11E-1	13	208	1.22E+4	—
			P_{2L}	SL(5)	CHEB(10)	1.09E-2	2.11E-1	16	65	9.72E+2	1.25E+1

n	β_i	run	P	PDE solver	PC solver	$\lg^* \text{frel}$	$\ r\ _{\text{rel}}$	iter	matvecs	time	speedup
HEART											
73728	1.00E-3		P_{REG}	RK2A(0.2)	—	1.11E-2	1.65E-1	19	853	8.43E+4	—
			P_{2L}	SL(5)	CHEB(10)	1.13E-2	1.65E-1	23	171	3.78E+3	2.23E+1
	1.00E-2		P_{REG}	RK2A(0.2)	—	9.10E-3	8.01E-1	20	473	5.46E+2	—
			P_{2L}	SL(5)	CHEB(10)	9.49E-3	8.00E-1	20	105	1.63E+2	3.36
			P_{REG}	RK2A(0.2)	—	9.30E-3	5.09E-1	31	1659	3.77E+3	—
			P_{2L}	SL(5)	CHEB(10)	9.92E-3	5.09E-1	31	410	7.52E+2	5.02
1.00E-4		P_{REG}	RK2A(0.2)	—	9.81E-3	2.98E-1	81	14455	6.17E+4	—	
		P_{2L}	SL(5)	CHEB(10)	8.86E-3	2.96E-1	76	2865	6.45E+3	9.57	
BRAIN											
153600	1.00E-1		P_{REG}	RK2A(0.2)	—	9.05E-3	4.82E-1	21	269	1.75E+3	—
			P_{2L}	SL(5)	CHEB(10)	9.16E-3	4.82E-1	21	70	1.69E+2	1.04E+1
	1.00E-2		P_{2L}	SL(10)	CHEB(10)	9.22E-3	4.82E-1	21	70	1.50E+2	1.17E+1
			P_{REG}	RK2A(0.2)	—	8.75E-3	3.21E-1	74	2645	2.60E+4	—
			P_{2L}	SL(5)	CHEB(10)	9.29E-3	3.21E-1	79	619	2.09E+3	1.25E+1
			P_{2L}	SL(10)	CHEB(10)	8.99E-3	3.22E-1	80	624	1.91E+3	1.37E+1
1.00E-3		P_{REG}	RK2A(0.2)	—	9.00E-3	2.05E-1	110	10306	1.21E+5	—	
		P_{2L}	SL(5)	CHEB(10)	9.54E-3	2.05E-1	74	1156	5.15E+3	2.35E+1	
		P_{2L}	SL(10)	CHEB(10)	9.21E-3	2.05E-1	76	1239	4.53E+3	2.67E+1	

Table 10

Convergence results for the inversion using our formulation for a fully incompressible diffeomorphism (linear Stokes regularization). We report results for registering the UT images (see Fig. 3; grid size: 256×256) using our original preconditioner (P_{REG} ; $RK2A$ scheme with a CFL number of 0.2) and the proposed preconditioner (P_{2L} ; SL scheme with a CFL number of 5; $CHEB$ method with a fixed number of 10 iterations). We consider different regularization norms: an H^1 -seminorm; an H^2 -seminorm; and an H^3 -seminorm (from top to bottom). We terminate the inversion if the change in the $\mathcal{L}^0$ -norm of the reduced gradient g_k^h , $k = 1, 2, \dots$, is at least two orders of magnitude or if the $\mathcal{L}^0$ -norm of g_k^h is smaller or equal to $1E-5$. We report (i) the relative change of the reduced gradient $\|g^{\star}\|_{red}$, (ii) the relative change of the residual $\|r\|$ rel (L^2 -distance between m_R and m_1), (iii) the number of outer iterations, (iv) the number of Hessian matvecs, (v) the time to solution, and (vi) the speedup compared to our original scheme.

run	norm	P	PDE solver	PC solver	$\lg \star \ r\ _{red}$	$\ r\ _{red}$	iter	matvecs	time	speedup
H^1	P_{REG}	P_{REG}	RK2A(0.2)	—	8.04E-3	1.40E-2	12	137	1.20E+3	—
	P_{2L}	P_{2L}	SL(5)	CHEB(10)	6.40E-3	1.38E-2	13	43	8.88E+1	1.35E+1
	P_{REG}	P_{REG}	RK2A(0.2)	—	9.05E-3	1.90E-1	17	177	1.38E+3	—
H^2	P_{2L}	P_{2L}	SL(5)	CHEB(10)	8.81E-3	1.90E-1	16	53	1.06E+2	1.30E+1
	P_{REG}	P_{REG}	RK2A(0.2)	—	9.14E-3	6.24E-1	34	402	2.60E+3	—
H^3	P_{2L}	P_{2L}	SL(5)	CHEB(10)	9.30E-3	6.24E-1	36	111	2.22E+2	1.17E+1

Table 11

Convergence results for the inversion using our formulation for a near-incompressible diffeomorphism (linear Stokes regularization). We consider the *HAND* images in Fig. 3. We report results for our original preconditioner (P_{REG} ; *RK2A* scheme with a *CFL* number of 0.2) and the proposed preconditioner (P_{2L} ; *SL* scheme with a *CFL* number of 5; *CHEB* method with a fixed number of 10 iterations). We terminate the inversion if the change in the ℓ^∞ -norm of the reduced gradient g_k^h , $k = 1, 2, \dots$, is at least two orders of magnitude or if the ℓ^∞ -norm of g_k^h is smaller or equal to $1E-5$. We report (i) the relative change of the reduced gradient $\|g^{\star}\|_{rel}$, (ii) the relative change of the residual $\|r\|_{rel}/(L^2$ -distance between m_R and m_I), (iii) the number of outer iterations, (iv) the number of *Hessian* matvecs, (v) the time to solution, and (vi) the speedup compared to our original scheme.

β_i	β_w	run	P	PDE solver	PC solver	$\ g^{\star}\ _{rel}$	$\ r\ _{rel}$	iter	matvecs	time	speedup
1.00E-1	1.00E-3		P_{REG}	RK2A(0.2)	—	9.20E-3	1.55E-1	10	119	1.76E+2	—
			P_{2L}	SL(5)	CHEB(10)	5.55E-3	1.57E-1	10	29	2.40E+1	7.32
			P_{REG}	RK2A(0.2)	—	9.49E-3	1.48E-1	9	99	1.63E+2	—
			P_{2L}	SL(5)	CHEB(10)	9.02E-3	1.50E-1	9	25	1.82E+1	8.98
1.00E-2	1.00E-3		P_{REG}	RK2A(0.2)	—	9.14E-3	6.56E-2	14	731	1.50E+3	—
			P_{2L}	SL(5)	CHEB(10)	9.57E-3	6.60E-2	13	60	7.05E+1	2.13E+1
			P_{REG}	RK2A(0.2)	—	8.52E-3	5.31E-2	13	513	1.09E+3	—
			P_{2L}	SL(5)	CHEB(10)	9.58E-3	5.37E-2	13	60	6.28E+1	1.74E+1

Table 12

Computational complexity of our solver for the compressible case. We report this complexity as a function of the number of *FFTs* and interpolation steps within the key building blocks of our solver. We provide these counts for (i) the hyperbolic transport equations that appear in the optimality system (state equation (7a): *SE*; adjoint equation (8a): *AE*; incremental state equation (11a): *incSE*; incremental adjoint equation (12a): *incAE*), the evaluation of the objective $\mathcal{J}^h$ in (2a), the evaluation of the gradient $\mathbf{g}^h$ in (6), and the Hessian matvec in (10). We report numbers for the full Newton case (*FN*) and the Gauss–Newton approximation (*GN*). The costs for evaluating the objective include the costs for the forward solve. We assign the costs of the adjoint solve to the evaluation of the gradient. The costs for the Hessian matvec include the solution of the *incSE* and *incAE*.

	RK2		RK2A		SL	
	FFTs	IPs	FFTs	IPs	FFTs	IPs
SE	$2(d+1)n_t$	–	$2(d+1)+4(d+1)n_t$	–	–	$d+n_t$
AE	$2(d+1)n_t$	–	$2(d+1)+4(d+1)n_t$	–	$d+1$	$d+n_t+1$
incSE	$4(d+1)n_t$	–	$4(d+1)+6(d+1)n_t$	–	$(d+1)n_t$	$d+(d+1)n_t$
incAE (FN)	$2(d+1)n_t$	–	$4(d+1)+6(d+1)n_t$	–	$(d+1)n_t$	$2n_t+1$
incAE (GN)	$2(d+1)n_t$	–	$2(d+1)+4(d+1)n_t$	–	$d+1$	n_t+1
compressible	$2d+2(d+1)n_t$	–	$2d+2(d+1)+4(d+1)n_t$	–	$2d$	$d+n_t$
	$\mathcal{J}^h$					
$\mathbf{g}^h$	$2d+3(d+1)n_t$	–	$2d+2(d+1)+7(d+1)n_t$	–	$2d+(d+1)(n_t+1)$	$d+n_t+1$
matvec (FN)	$2d+8(d+1)n_t$	–	$2d+8(d+1)+18(d+1)n_t$	–	$2d+4(d+1)n_t$	$d+(d+3)n_t+1$
matvec (GN)	$2d+7(d+1)n_t$	–	$2d+9(d+1)+13(d+1)n_t$	–	$2d+(d+1)(2n_t+1)$	$d+(d+2)n_t+1$

Table 13

Wall clock times for applying one *FFT* or one cubic spline interpolation step with respect to different grid sizes. The timings are obtained for the `fftn` and `interp2` functions in *Matlab R2013a* on a *Linux cluster* with *Intel Xeon X5650 Westmere* EP 6-core processors at 2.67GHz with 24GB DDR3-1333 memory.

n_x^i	FFTs	IPs	factor
16	3.03E-5	4.98E-3	1.64E+2
32	3.79E-5	2.43E-3	6.40E+1
64	1.21E-4	2.90E-3	2.40E+1
128	3.22E-4	5.14E-3	1.59E+1
256	8.86E-4	2.24E-2	2.53E+1
512	3.83E-3	1.42E-1	3.70E+1
1024	1.30E-2	6.07E-1	4.69E+1