

Iterative Inversion of Deformation Vector Fields with Feedback Control

Abhishek Kumar Dubey,^{a)} Alexandros-Stavros Iliopoulos,^{a)} and Xiaobai Sun

Department of Computer Science, Duke University, Durham, NC 27708, USA

Fang-Fang Yin

Department of Radiation Oncology, Duke University School of Medicine, Durham, NC 27710, USA

Medical Physics Program, Duke Kunshan University, Kunshan, Jiangsu, China 215316

Lei Ren^{b)}

Department of Radiation Oncology, Duke University School of Medicine, Durham, NC 27710, USA

Purpose: Often, the inverse deformation vector field (DVF) is needed together with the corresponding forward DVF in 4D reconstruction and dose calculation, adaptive radiation therapy, and simultaneous deformable registration. This study aims at improving both accuracy and efficiency of iterative algorithms for DVF inversion, and advancing our understanding of divergence and latency conditions.

Method: We introduce a framework of fixed-point iteration algorithms with active feedback control for DVF inversion. Based on rigorous convergence analysis, we design control mechanisms for modulating the inverse consistency (IC) residual of the current iterate, to be used as feedback into the next iterate. The control is designed adaptively to the input DVF with the objective to enlarge the convergence area and expedite convergence. Three particular settings of feedback control are introduced: constant value over the domain throughout the iteration; alternating values between iteration steps; and spatially variant values. We also introduce three spectral measures of the displacement Jacobian for characterizing a DVF. These measures reveal the critical role of what we term the non-translational displacement component (NTDC) of the DVF. We carry out inversion experiments with an analytical DVF pair, and with DVFs associated with thoracic CT images of 6 patients at end of expiration and end of inspiration.

Results: NTDC-adaptive iterations are shown to attain a larger convergence region at a faster pace compared to previous non-adaptive DVF inversion iteration algorithms. By our numerical experiments, alternating control yields smaller IC residuals and inversion errors than constant control. Spatially variant control renders smaller residuals and errors by at least an order of magnitude, compared to other schemes, in no more than 10 steps. Inversion results also show remarkable quantitative agreement with analysis-based predictions.

Conclusion: Our analysis captures properties of DVF data associated with clinical CT images, and provides new understanding of iterative DVF inversion algorithms with a simple residual feedback control. Adaptive control is necessary and highly effective in the presence of non-small NTDCs. The adaptive iterations or the spectral measures, or both, may potentially be incorporated into deformable image registration methods.

Keywords: deformation vector field, inversion, deformable image registration, fixed-point iteration, spectral analysis

I. INTRODUCTION

We consider numerical inversion of a deformation vector field (DVF). Inverse DVFs are needed, together with their respective forward DVFs, to map images, structure contours, or doses back and forth in applications such as 4D image reconstruction,¹ dose accumulation calculations and multi-modality treatment planning in adaptive radiotherapy,^{2–5} and cardiac functional analysis.⁶ DVF inversion is also a fundamental operation in simultaneous and symmetric registration methods.^{7–11} An important consideration is ensuring that the forward and reverse mappings are inverse-consistent.^{10,12} Theoretical guar-

antees of convergence and computational efficiency have been a long-standing open problem with DVF inversion.

A consistent pair of forward and reverse mappings can be obtained via one-way deformable registration followed by a DVF inversion process,^{2,13} or via simultaneous, symmetric two-way registration methods.^{7–9,11} The former, asymmetric approach is often preferred in certain clinical applications with limited time window, in part because it is shown to be faster empirically, and in part because of the non-negligible asymmetry in clinical image quality. One of the images may be more adversely affected by noise or artifacts than the other. Such asymmetry

makes one-way registration better or worse depending on the mapping direction, due to the high sensitivity of registration methods to noise or variation in imaging conditions.^{2,6,14} We consider the asymmetric approach to be an effective means to counterbalancing the asymmetry in image quality.

DVF inversion is often involved also in simultaneous registration, which results in both forward and reverse mappings.^{7-9,11} Inverse consistency (IC) between forward and reverse mappings for deformable registration was addressed in the early work of Thirion¹⁵ and Christensen.¹² The IC condition has since been incorporated in various deformable registration models. It is either used as an explicit constraint attached to an optimization model,⁷ or employed implicitly and approximately in numerical iterations.⁸ The registration process may involve multiple intermediate transformations and their composition. DVF inversion is used to ensure that the transformations, intermediate as well as final, meet the IC condition. Many studies on simultaneous estimation of forward-inverse DVF pairs can be found in the survey by Sotiras et al.¹⁰ and references therein.

With a provided DVF as input, numerical inversion of the DVF can be governed by the IC condition and carried out in displacement space, rather than in image space. The inversion relationship is inherently nonlinear. One therefore resorts to iterative solution methods, except in certain cases. There is a close relationship between IC residuals (i.e., deviations from the IC condition) and inversion errors in the iterative estimates. We will leverage this relationship to improve upon the inverse DVF iterates by using the IC residuals as feedback into the iteration.

Two particular and influential iteration algorithms for DVF inversion were developed by Christensen and Johnson⁷ and Chen et al.¹³ The iteration by Christensen and Johnson⁷ is often effective, and is notable in making use of residual feedback. The iteration is closely related to the residual method by Thirion,¹⁵ which is based on a heuristic to enforce that DVF estimates are bijective (invertible). The condition under which the iteration converges or fails to converge was hitherto unknown. Chen et al.¹³ departed from heuristic design. They introduced a particular fixed-point iteration for DVF inversion, and identified a convergence condition. The iteration was not compared to the earlier algorithm by Christensen and Johnson,⁷ and there is no mention of residual feedback.

We make the following key contributions to the understanding and convergence control of iterative DVF inversion algorithms. (i) We present a framework of iteration algorithms with simple and adaptive residual feedback control. This includes the two precursor algorithms, in both format and analysis. (ii) The framework is underpinned by a unified analysis of error propagation and convergence. The analysis enables connections and comparisons among iteration algorithms for DVF inversion, and leads to the design of more effective ones. (iii) We characterize the critical role of what we introduce as the

non-translational displacement component (NTDC) in error propagation, and provide quantitative NTDC measures. When the NTDC is non-small, adaptive residual feedback control is necessary to guarantee convergence. This insight is new. We assess our findings experimentally with synthetic DVF data, and patient DVF data obtained from thoracic CT images.

The rest of the document is organized as follows. In Section II, we introduce our algorithm framework, provide formal analysis, present three practical feedback control schemes, and discuss three spectral measures for NTDC characterization. In Section III, we describe and assess the patient DVF data used in our experiments. Pre-inversion assessment of control schemes and post-inversion evaluation of results with the patient DVFs are provided in Section IV. A direct evaluation using analytical DVF data is included in Appendix A. We conclude the paper in Section V with additional discussion on the clinical utility of our algorithms and analysis.

II. METHODS

II.A. DVF inversion preliminaries

DVF inversion can be phrased as follows. A reference and a target image, denoted by I_{ref} and I_{tgt} , respectively, can be related to one another by two non-linear transformations. Denote by Ω the image domain, $\Omega \subset \mathbb{R}^3$. The forward transformation, $\mathbf{f}: \Omega \rightarrow \Omega$, maps the voxels of the reference image I_{ref} onto those of the target image I_{tgt} via $\mathbf{f}(\mathbf{x}') = \mathbf{x}' + \mathbf{u}(\mathbf{x}')$, where $\mathbf{u}(\mathbf{x}')$ is the forward 3D displacement at $\mathbf{x}' \in \Omega$. Conversely, the inverse transformation, $\mathbf{g}: \Omega \rightarrow \Omega$, maps the voxels of I_{tgt} back to I_{ref} via $\mathbf{g}(\mathbf{x}) = \mathbf{x} + \mathbf{v}(\mathbf{x})$, where $\mathbf{v}(\mathbf{x})$ is the inverse 3D displacement at $\mathbf{x} \in \Omega$. DVF inversion means obtaining \mathbf{v} given \mathbf{u} . The two transformations are inverse to each other: $(\mathbf{f} \circ \mathbf{g})(\mathbf{x}) = \mathbf{x}$, and $(\mathbf{g} \circ \mathbf{f})(\mathbf{x}') = \mathbf{x}'$ for $\mathbf{x}, \mathbf{x}' \in \Omega$. Consequently, the forward and inverse DVFs satisfy the simultaneous *inverse consistency* (IC) condition:

$$\mathbf{v}(\mathbf{x}) + \mathbf{u}(\mathbf{x} + \mathbf{v}(\mathbf{x})) = \mathbf{0}, \quad (1a)$$

$$\mathbf{u}(\mathbf{x}') + \mathbf{v}(\mathbf{x}' + \mathbf{u}(\mathbf{x}')) = \mathbf{0}. \quad (1b)$$

The IC condition governs iterative DVF inversion.

II.A.1. Inversion error & inverse consistency residual

Denote by $\hat{\mathbf{v}}$ an estimate of the inverse DVF \mathbf{v}_* . The unknown error in the estimate,

$$\mathbf{e}(\mathbf{x}) = \hat{\mathbf{v}}(\mathbf{x}) - \mathbf{v}_*(\mathbf{x}), \quad (2)$$

is manifested in the inverse consistency (IC) residual,

$$\mathbf{r}_v(\mathbf{x}) = \hat{\mathbf{v}}(\mathbf{x}) + \mathbf{u}(\mathbf{x} + \hat{\mathbf{v}}(\mathbf{x})), \quad (3)$$

which is computationally available. Qualitatively, the residual is zero if and only if the inversion error is zero. In order to use the IC residual as feedback for improving the inverse DVF estimate, we investigate the quantitative relationship between inversion error and IC residual.

Assume in analysis that the deformation transformation \mathbf{f} is differentiable. By the mean value theorem, the IC residual and the inversion error can be related by

$$\mathbf{r}_v(\mathbf{x}) = \mathbf{J}_f(\boldsymbol{\xi}) \mathbf{e}(\mathbf{x}), \quad (4)$$

where $\mathbf{J}_f(\boldsymbol{\xi})$ is the forward transformation Jacobian evaluated at $\boldsymbol{\xi}$, which lies between $\mathbf{x} + \mathbf{v}_*(\mathbf{x})$ and $\mathbf{x} + \hat{\mathbf{v}}(\mathbf{x})$. When $\mathbf{e}(\mathbf{x})$ is small, $\mathbf{J}_f(\boldsymbol{\xi})$ can be numerically approximated by $\mathbf{J}_f(\mathbf{x} + \hat{\mathbf{v}}(\mathbf{x}))$. The Jacobian \mathbf{J}_f is spatially variant over Ω , except in some special cases. Provided with a forward DVF \mathbf{u} , we will rely on the IC residual \mathbf{r}_v and relationship (4) to improve upon the inverse estimate $\hat{\mathbf{v}}$ via an iterative process.

By (1b), one shall also consider the other IC residual,

$$\mathbf{r}_u(\mathbf{x}') = \mathbf{u}(\mathbf{x}') + \hat{\mathbf{v}}(\mathbf{x}' + \mathbf{u}(\mathbf{x}')). \quad (5)$$

We omit a rigorous analysis of the differential relationship between \mathbf{r}_u and the estimation error \mathbf{e} . In a nutshell, the residual is spatially related to the error through the following mapping:

$$\mathbf{r}_u(\mathbf{x} + \hat{\mathbf{v}}(\mathbf{x})) = \mathbf{e}(\mathbf{x} + \mathbf{r}_v(\mathbf{x})). \quad (6)$$

The residual \mathbf{r}_v in (6) can be made sufficiently small if the iteration converges. We will use $\mathbf{r}_u(\mathbf{x} + \hat{\mathbf{v}}(\mathbf{x}))$ in addition to $\mathbf{r}_v(\mathbf{x})$ to quantitatively assess inverse DVF estimates.

II.A.2. Non-translational displacement component (NTDC)

We introduce the decomposition of a DVF into translational and non-translational components, to elucidate the relationship between inversion errors and IC residuals. By (4), the estimate error is related to the IC residual via \mathbf{J}_f . The transformation Jacobian \mathbf{J}_f is the displacement Jacobian \mathbf{J}_u shifted by the identity:

$$\mathbf{J}_f = \mathbf{I} + \mathbf{J}_u. \quad (7)$$

When $\mathbf{J}_u = \mathbf{0}$, then $\mathbf{J}_f = \mathbf{I}$ and the residual \mathbf{r}_v is equal to the estimate error \mathbf{e} . The inverse displacement \mathbf{v} can then be obtained immediately by adding the residual to the current estimate, regardless of the direction and magnitude of displacement \mathbf{u} . In this case, we consider the corresponding displacement \mathbf{u} as translational. When $\mathbf{J}_u \neq \mathbf{0}$, there is a non-translational component in the displacement. The non-translational displacement component (NTDC) is responsible for the nontrivial, non-transparent relationship between the estimate error and the IC residual.

The NTDC Jacobian is used in DVF characterization, pre-inversion convergence analysis and prediction, and adaptive feedback control design. Conceptually, we decompose the displacement into translational \mathbf{u}_t and non-translational \mathbf{u}_{nt} components: $\mathbf{u}(\mathbf{x}) = \mathbf{u}_t + \mathbf{u}_{nt}(\mathbf{x})$. These components are identified by their respective contributions to the Jacobian \mathbf{J}_u . Only non-translational components contribute to the Jacobian, i.e., $\mathbf{J}_{u_t} = \mathbf{0}$ and $\mathbf{J}_{u_{nt}} = \mathbf{J}_u$. We may therefore refer to the NTDC Jacobian as the displacement Jacobian \mathbf{J}_u , with the understanding that the translational component plays no part in it. We will introduce in Section II.C spectral measures for characterizing the NTDC, and provide an explicit criterion for considering the NTDC as non-small.

II.B. NTDC-adaptive iteration

II.B.1. Active feedback control framework

We introduce a family of fixed-point iterations for DVF inversion, using the IC residual \mathbf{r}_v as feedback. In the rest of this section, we denote the IC residual simply by \mathbf{r} . Feedback control is exercised to suppress the estimate error, based on error propagation analysis and (4). Assume a forward DVF \mathbf{u} is provided over Ω . At step $k = 0, 1, 2, \dots$, we compute the residual \mathbf{r}_k associated with the current estimate \mathbf{v}_k and get the next estimate by

$$\mathbf{v}_{k+1}(\mathbf{x}) = \mathbf{v}_k(\mathbf{x}) - \mathbf{B}_k(\mathbf{x}) \mathbf{r}_k(\mathbf{x}), \quad (8)$$

where the term $\mathbf{B}_k(\mathbf{x}) \mathbf{r}_k(\mathbf{x})$ is the modulated residual, and $\mathbf{B}_k(\mathbf{x})$ is a 3×3 feedback control matrix associated with $\mathbf{x} \in \Omega$. In this paper, we consider the control mechanism in its simplest form: $\mathbf{B}_k(\mathbf{x})$ is an isotropic scaling matrix, $\mathbf{B}_k(\mathbf{x}) = (1 - \mu_k(\mathbf{x}))\mathbf{I}$. Iteration (8) then takes the form

$$\mathbf{v}_{k+1}(\mathbf{x}) = \mathbf{v}_k(\mathbf{x}) - (1 - \mu_k(\mathbf{x})) \mathbf{r}_k(\mathbf{x}). \quad (9)$$

We refer to $\mu_k(\mathbf{x})$ as the feedback control parameter. In what follows, we will introduce three particular adaptive control schemes: constant parameter value over the domain and throughout the iteration, alternating values between iteration steps, and spatially variant values.

When the control parameter is spatially uniform, i.e., it does not vary with \mathbf{x} , iteration (9) becomes

$$\mathbf{v}_{k+1}(\mathbf{x}) = \mathbf{v}_k(\mathbf{x}) - (1 - \mu_k) \mathbf{r}_k(\mathbf{x}). \quad (10)$$

Control is stationary if $\mu_k = \mu$ for some constant μ , and non-stationary otherwise. The iterations with constant-value control can be further divided into non-adaptive (pre-fixed constants) and adaptive ones. The two precursor algorithms for DVF inversion^{7,13} mentioned in Section I both adhere to the form of iteration (10) with constant non-adaptive (pre-fixed) control values. Specifically, $\mu_k = 0.5$ yields the algorithm of Christensen and Johnson,⁷ and $\mu_k = 0$ yields the algorithm of Chen et al.¹³

II.B.2. Spectral analysis

We provide analytical apparatus for designing residual feedback control in the simple form of (9) in order to guarantee convergence and improve convergence speed. By (4) and (9), inverse estimate errors (2) propagate throughout the iteration by the equation

$$\mathbf{e}_{k+1}(\mathbf{x}) = \mathbf{P}_k(\mathbf{x}; \mu) \mathbf{e}_k(\mathbf{x}), \quad (11a)$$

$$\mathbf{P}_k(\mathbf{x}; \mu) = \mathbf{I} - (1 - \mu) \mathbf{J}_f(\boldsymbol{\xi}_k), \quad (11b)$$

where μ is short for $\mu_k(\mathbf{x})$, and $\mathbf{P}_k(\mathbf{x}; \mu)$ is the one-step error propagation matrix at $\boldsymbol{\xi}_k$, which lies between $\mathbf{x} + \mathbf{v}_*(\mathbf{x})$ and $\mathbf{x} + \mathbf{v}_k(\mathbf{x})$. The propagation matrix depends on the value of μ and varies during the iteration. If

$$\rho(\mathbf{P}_k(\mathbf{x}; \mu)) \leq \rho_{\text{sup}}, \quad k = 0, 1, 2, \dots, \quad (12)$$

for some $\rho_{\text{sup}} \in [0, 1)$, then the iteration converges. Here, $\rho(\mathbf{P})$ denotes the spectral radius of the propagation matrix \mathbf{P} , i.e., the ratio between successive errors in magnitude.

Consider the special case where the displacement is translational, i.e., $\mathbf{J}_{\mathbf{u}}(\mathbf{x}) = \mathbf{0}$, $\mathbf{J}_{\mathbf{f}}(\mathbf{x}) = \mathbf{I}$, and $\mathbf{P}_k(\mathbf{x}; \mu) = \mu \mathbf{I}$. The iteration converges with any $\mu \in (-1, 1)$, and converges faster when μ is closer to 0. At the mid-range value, $\mu = 0$, the iteration renders the inverse DVF in one step. We focus our study on deformations with non-translational components, without excluding the case of translational displacement.

Feedback control design based on convergence analysis is challenging. The point-wise error sequences (11a), each associated with a voxel location \mathbf{x} , cover the whole domain Ω . Although the sequences depend on the initial guess, a sequence $\mathbf{e}_k(\mathbf{x})$, $k = 0, 1, 2, \dots$, converges to zero if it can be guaranteed that the associated non-negative scalar sequence $\rho(\mathbf{P}_k(\mathbf{x}; \mu))$ at $\xi_k(\mathbf{x})$, $k = 0, 1, 2, \dots$, is bounded from above below 1; see (11b) and (12). Even with a fixed value of μ , tracking every scalar sequence of $\rho(\mathbf{P}_k(\mathbf{x}; \mu))$ at unknown and spatially varying mean-value locations $\xi_k(\mathbf{x})$ is implausible. We surmount this challenge by the following novel approach.

We take a covering-and-partitioning approach. Specifically, with any value of the control parameter μ , we consider the deformation Jacobian everywhere over Ω . We define the infinitesimal contraction matrix at all $\mathbf{x} \in \Omega$:

$$\mathbf{Q}(\mathbf{x}; \mu) = \mathbf{I} - (1 - \mu)\mathbf{J}_{\mathbf{f}}(\mathbf{x}) = \mu \mathbf{I} - (1 - \mu)\mathbf{J}_{\mathbf{u}}(\mathbf{x}). \quad (13)$$

Any particular value of the control parameter μ partitions Ω into a contraction region,

$$\Omega_c(\mu) = \{\mathbf{x} \mid \rho(\mathbf{Q}(\mathbf{x}; \mu)) < 1, \mathbf{x} \in \Omega\}, \quad (14)$$

and its complement $\Omega - \Omega_c(\mu)$, the non-contraction region. The error (11a) converges to zero if all the mean-value locations ξ_k fall within the contraction region. If $\Omega_c(\mu) = \Omega$, i.e., the contraction region covers the entire domain (or, more strictly speaking, $\sup_{\mathbf{x} \in \Omega} \rho(\mathbf{Q}(\mathbf{x}; \mu)) < 1$), then the iteration converges. By the covering-and-partitioning approach, we reduce feedback control design to finding the control parameter values that yield the largest contraction region over Ω .

Consider the special value $\mu = 0$, with which $\rho(\mathbf{Q}(\mathbf{x}; 0)) = \rho(\mathbf{J}_{\mathbf{u}}(\mathbf{x}))$. If $\rho(\mathbf{J}_{\mathbf{u}}(\mathbf{x})) < 1$, then \mathbf{x} lies in the contraction region. Otherwise, \mathbf{x} is in the non-contraction region,

$$\rho(\mathbf{J}_{\mathbf{u}}(\mathbf{x})) \geq 1, \quad (15)$$

and a positive feedback control value ($\mu > 0$) is warranted. Formally, we define by (15) the concept of a non-small NTDC at \mathbf{x} : the NTDCs are considered non-small where the spectral radius of the displacement Jacobian is equal to or greater than 1. Whenever the mean-value location ξ_k , at which $\mathbf{J}_{\mathbf{u}}(\xi_k) = \mathbf{P}_k(\mathbf{x}; 0)$, falls in a non-small-NTDC region, the error in the next iterate is magnified if active feedback control is not applied.

We describe in the rest of this section how to materialize the sufficient condition $\Omega_c(\mu) = \Omega$ for guaranteed convergence of all iterative sequences. We translate the condition into expressions that relate the parameter μ to the eigenvalues of the transformation Jacobian, $\lambda_j(\mathbf{x}) = \lambda_j(\mathbf{J}_{\mathbf{f}}(\mathbf{x}))$, $j = 1, 2, 3$.

Consider first the case where the eigenvalues are all real and positive over Ω . The condition $\Omega_c(\mu) = \Omega$ is achieved by any specific value of μ in the range

$$\max \left\{ -1, 1 - 2 \min_{\mathbf{x} \in \Omega, j=1,2,3} \frac{1}{\lambda_j(\mathbf{x})} \right\} < \mu < 1. \quad (16)$$

When μ is negative, $(1 - \mu) > 1$ and the impact of $\mathbf{J}_{\mathbf{u}}$ on error contraction is over-relaxed; see (13). When $\mu < -1$, a small spectral radius of $\mathbf{Q}(\mathbf{x}; \mu)$ is the result of algebraic cancellation. To avoid severe cancellation and subsequent instability in numerical computations, we bound the control parameter value from below by -1 .

Consider next the presence of complex eigenvalues, which are prevalent in DVFs associated with patient CT images (see Table I in Section III.C). This should not be a surprise; the eigenvalues of a plane rotation matrix are complex, for instance. Complex eigenvalues of a real-valued matrix exist in conjugate pairs. For any fixed value $\mu < 1$, the error contraction condition $\rho(\mathbf{Q}(\mathbf{x}; \mu)) < 1$ becomes

$$2 \operatorname{Re}(\lambda_j(\mathbf{x})) > (1 - \mu)|\lambda_j(\mathbf{x})|^2, \quad j = 1, 2, 3, \quad (17)$$

for $\mathbf{x} \in \Omega$. This condition immediately rejects singular Jacobians, and gives rise to the local feasible parameter range

$$\max \{-1, 1 - 2\gamma(\mathbf{x})\} < \mu < 1, \quad (18)$$

where

$$\gamma(\mathbf{x}) = \min_j \frac{\operatorname{Re}(\lambda_j(\mathbf{x}))}{|\lambda_j(\mathbf{x})|^2} = \min_j \operatorname{Re}(\lambda_j^{-1}(\mathbf{x})) > 0. \quad (19)$$

We refer to (19) as the controllability condition. It has two equivalent expressions: one in terms of the eigenvalues, and one in terms of the reciprocal eigenvalues.

A few remarks about condition (19) are in order. The condition rejects any Jacobian with eigenvalues on the imaginary axis or in the left half of the complex plane, in which case the determinant of the Jacobian may be negative, zero, or even positive. The condition is necessary and sufficient for the parameter range (18) to be non-empty, and hence for feasible values of μ to exist. One can locate the few cases where the condition is violated. Violations of the controllability condition anywhere in the image domain most likely manifest as artifacts introduced by the forward DVF generation process. In other words, this condition shall be recognized as a rule for local regularization, with respect to DVF inversion, not a limitation imposed by the control mechanism (9).

II.B.3. Maximal error suppression

We consider now how to determine, over an infinitesimal neighborhood of \mathbf{x} , the control parameter value




FIG. 1 Geometric outline of (22a) and (22b). See the explanation in Section II.B.3. The notation $\lambda(\mu)$ is short for $\lambda(\mathbf{Q}(\mu))$.

$\mu_*(\mathbf{x})$ such that estimate errors are suppressed as much as possible. Specifically, we minimize the spectral radius of the local infinitesimal contraction matrix, i.e., $\mu_*(\mathbf{x}) = \arg \min_{\mu} \rho(\mathbf{Q}(\mathbf{x}; \mu))$. By (13), the spectrum of $\mathbf{Q}(\mathbf{x}; \mu)$ is that of the displacement Jacobian $\mathbf{J}_u(\mathbf{x})$ scaled by $(\mu - 1)$ and shifted by 1. Since the analysis concerns the variation of $\rho(\mathbf{Q}(\mathbf{x}; \mu))$ with parameter μ at any fixed point \mathbf{x} , we simplify the notation $\mathbf{Q}(\mathbf{x}; \mu)$ to $\mathbf{Q}(\mu)$, and similarly with notation for other related quantities, omitting \mathbf{x} for the rest of this sub-section. To further simplify the expressions below, we use the eigenvalues of \mathbf{J}_f , which are the NTDC eigenvalues shifted by 1; see (7).

Assume the controllability condition (19) holds, i.e., $\gamma > 0$; otherwise, there exists no feasible parameter value. When the three eigenvalues of \mathbf{J}_f are all real and positive, it is straightforward to verify that

$$\mu_* = 1 - \frac{2}{\lambda_{\max} + \lambda_{\min}}, \quad (\text{case } R) \quad (20a)$$

and

$$\rho_* = \min_{\mu} \rho(\mathbf{Q}(\mu)) = \frac{\lambda_{\max} - \lambda_{\min}}{\lambda_{\max} + \lambda_{\min}} < 1. \quad (20b)$$

In contrast, in the presence of complex eigenvalues, which appear in a conjugate pair, λ_c and $\bar{\lambda}_c$, the geometric positions of the conjugate pair relative to the real eigenvalue λ_r are not maintained after scaling and shifting. Define

$$\gamma_c = \frac{\text{Re}(\lambda_c)}{|\lambda_c|^2} = \text{Re}(\lambda_c^{-1}); \quad (21)$$

then, $\gamma = \min\{\gamma_c, \lambda_r^{-1}\}$. By our analysis, the geometric relationship between the complex and real eigenvalues can be put into two mutually exclusive cases: case C1 if $|1 - \gamma_c \lambda_r| \leq |1 - \gamma_c \lambda_c|$, and case C2 otherwise. The optimal control parameter value for each case is

$$\mu_* = \begin{cases} 1 - \gamma_c, & (\text{case } C1) \\ 1 - \frac{\text{Re}(\lambda_c - \lambda_r)}{|\lambda_c| - \lambda_r} \frac{2}{|\lambda_c| + \lambda_r}. & (\text{case } C2) \end{cases} \quad (22a)$$

The optimal spectral radius is, for both cases,

$$\rho_* = \min_{\mu} \rho(\mathbf{Q}(\mu)) = |1 - (1 - \mu_*)\lambda_c| < 1. \quad (22b)$$

The proof of (22) could be lengthy and tedious in words. We provide instead a geometric explanation with two drawings in Fig. 1; one for each case. The drawings illustrate how the eigenvalues of \mathbf{J}_f are scaled and shifted onto those of $\mathbf{Q}(\mu)$ in the complex plane, and how to locate the optimal scaling value. In each drawing, we locate first the complex eigenvalues of $\mathbf{Q}(0) = -\mathbf{J}_u$. We connect them by straight lines to $(1, 0)$, where the multiple eigenvalues of $\mathbf{Q}(1)$ reside. For any μ , the conjugate pair $\lambda_c(\mathbf{Q}(\mu))$ lie on the two conjugate lines. In case C1, the optimal value μ_* maps λ_c onto the very point at which these lines meet the tangential circle centered at the origin; this gives the minimal spectral radius ρ_* , since the real eigenvalue by the same scaling and shifting falls inside the circle. In case C2, the real eigenvalue is outside the circle (dashed blue in Fig. 1b). We rescale the eigenvalues until all three eigenvalues are on the same circle, which is of the minimal spectral radius.

II.B.4. Control parameter schemes

II.B.4.a. Mid-range parameter value. We present an adaptive scheme for determining a constant value, including its existence condition, for the control parameter such that uniform convergence is guaranteed over a sub-region or neighborhood $\mathcal{N} \subseteq \Omega$. First, we extend point-wise quantities to region-wise quantities. Specifically, let $\gamma(\mathcal{N}) = \min_{\mathbf{x} \in \mathcal{N}} \gamma(\mathbf{x})$. Computationally, $\gamma(\mathcal{N})$ can be obtained easily with a minimum filter. If $\gamma(\mathcal{N}) > 0$, then the control parameter range, $(\max\{-1, 1 - 2\gamma(\mathcal{N})\}, 1)$, for uniform convergence over \mathcal{N} is non-empty. Any value in this range can serve as a constant value for the control parameter, guaranteeing convergence. We may use

in particular the mid-range value. When $1 - 2\gamma > -1$, the mid-range value over \mathcal{N} is simply

$$\mu_m(\mathcal{N}) = 1 - \gamma(\mathcal{N}), \quad (23)$$

and can be easily determined. If $\gamma(\mathcal{N}) = 1$, then $\mu_m(\mathcal{N}) = 0$; if $\gamma(\mathcal{N}) = 0.5$, then $\mu_m(\mathcal{N}) = 0.5$. The mid-range scheme (23) is adaptive to any DVF with $\gamma > 0$ over \mathcal{N} . It also leads to the next control scheme.

II.B.4.b. Alternating parameter values. Convergence can be made faster, within the parameter range for uniform convergence, by a simple modification: allowing the control parameter to take two alternating values. The idea is to exploit the non-uniform spectral structure of the DVF while keeping the control spatially uniform. As we shall show in Section III with DVFs from patient images, the local mid-range values, or locally optimal values, over the entire image domain Ω may be grouped into two sub-ranges in the convergence parameter range; one at the lower end, and one at the higher end. A simple alternating scheme is to use a value μ_e at even steps and another value μ_o at odd steps. Convergence can be analyzed via two-step error propagation,

$$\mathbf{e}_{2(k+1)}(\mathbf{x}) = \underbrace{\mathbf{P}_{2k+1}(\mathbf{x}; \mu_o) \mathbf{P}_{2k}(\mathbf{x}; \mu_e)}_{\mathbf{P}_{oe}(\mathbf{x})} \mathbf{e}_{2k}(\mathbf{x}), \quad (24)$$

where $\mathbf{P}_{k'}(\mathbf{x}; \mu) = \mathbf{I} - (1 - \mu_{k'})\mathbf{J}_f(\xi_{k'})$ and $\xi_{k'}$ lies between $\mathbf{x} + \mathbf{v}_*(\mathbf{x})$ and $\mathbf{x} + \mathbf{v}_{k'}(\mathbf{x})$, at iteration step k' . The spectral radius ρ_{oe} of the two-step propagation matrix can be bounded from above by $\rho_o\rho_e$, where ρ_o and ρ_e are the spectral radii of the odd- and even-step propagation matrices ($\mathbf{P}_{2k+1}(\mathbf{x}; \mu_o)$ and $\mathbf{P}_{2k}(\mathbf{x}; \mu_e)$), respectively. Thus, the contraction condition is maintained. Improvement in convergence speed is due to the suppression of local errors, which are aggressively suppressed at odd (even) steps without being enlarged at even (odd) steps.

II.B.4.c. Spatially variant parameter values. The non-uniform spectral structure of the transformation Jacobian can be better exploited by letting the control parameter vary spatially with \mathbf{x} over Ω . This can be achieved by determining the parameter value at $\mathbf{x} \in \Omega$ by a local neighborhood $\mathcal{N}(\mathbf{x})$, such that the entire image domain is covered, $\bigcup_{\mathbf{x}} \mathcal{N}(\mathbf{x}) = \Omega$. The neighborhood size need not be greater than the maximal displacement length, which is known in advance. The parameter value at \mathbf{x} can be the mid-range value over $\mathcal{N}(\mathbf{x})$, or the locally optimal value when the neighborhood is small enough. The iteration with spatially variant control is essentially non-stationary, because the value of $\mu_k(\mathbf{x})$ depends on the location of $\mathbf{x} + \mathbf{v}_k(\mathbf{x})$.

Locally optimal control, as described in Section II.B.3, assumes implicitly that the remaining displacement refinement is small. This assumption is not a restriction, as it can be met easily in practice by simply preceding the scheme by a few iterations with another control scheme. In particular, we may use spatially variant mid-range values over larger neighborhoods, which are not necessarily uniform in size nor shape.

Spatially variant schemes incur extra but modest cost in two parts. First, we create a parameter map once for all iteration steps, by calculating the Jacobians, eigenvalues, and control parameter values over Ω . This preprocessing step takes about the same time in execution as a single iteration step. Then, the look-up of parameter values over the domain at each iteration step takes no more than 5% of the cost for 3D vector-field interpolation. The overall cost is outweighed by the gain in practice; see Section IV.

II.C. Spectral NTDC characterization

We discuss in this section how we characterize and evaluate non-translational displacement components in a given DVF, with respect to DVF inversion. The NTDCs over Ω can be fully described by the eigenvalues of the displacement Jacobians $\mathbf{J}_u(\mathbf{x})$ in the complex plane. Rather, we employ the following three real-valued scalar functions for their informative properties. (i) The determinant of $\mathbf{J}_f(\mathbf{x})$, which is commonly used for deformation characterization.^{7,13} The transformation \mathbf{f} is invertible if and only if $|\mathbf{J}_f(\mathbf{x})| \neq 0$. (ii) The spectral radius of $\mathbf{J}_u(\mathbf{x})$. Where $\rho(\mathbf{J}_u(\mathbf{x})) = 0$, the DVF is locally translational; where $\rho(\mathbf{J}_u(\mathbf{x})) \geq 1$, the NTDC is non-small. (iii) The algebraic control index

$$1 - 2\gamma(\mathbf{x}), \quad (25)$$

where $\gamma(\mathbf{x})$ is defined in (19).

The algebraic control index is informative in several ways. First, it offers an equivalent criterion to (15) on whether the NTDC at \mathbf{x} is non-small:

$$1 - 2\gamma(\mathbf{x}) > 0 \iff \rho(\mathbf{J}_u(\mathbf{x})) \geq 1. \quad (26)$$

Second, the necessary and sufficient condition for the existence of feasible control parameter values is

$$1 - 2\gamma(\mathbf{x}) < 1 \iff \gamma(\mathbf{x}) > 0. \quad (27)$$

That is, the algebraic control index distinguishes non-small NTDCs from small ones and furthermore tells whether or not a non-small NTDC can be put under control by a single-parameter control mechanism (9). Third, the index is a lower bound to all feasible control values; see (18). It can be employed directly for locating the mid-range parameter value (Section II.B.4.a), as well as for selecting alternating values (Section II.B.4.b).

The algebraic control index (25) falls a little short of replacing entirely the roles of the other two measures, in two particular circumstances. When $|\mathbf{J}_f(\mathbf{x})| = 0$, the transformation is locally singular; the controllability condition is violated in this case and the index is not well-defined. When $\rho(\mathbf{J}_u(\mathbf{x})) = 0$, the displacement at \mathbf{x} is locally translational; the algebraic control index is equal to -1 in this case, but the converse is not necessarily true. We use all three measures for data assessment in Section III.C.

III. EXPERIMENTS

We report experimental results on numerical DVF inversion, using clinical and analytical data. In this section, we describe the experimental set-up and present characterization measures of the clinical DVF data. Results with different inversion iteration schemes are presented in Section IV with the clinical DVF data, and in Appendix A with the analytical DVF data.

III.A. Dataset description

The experiments presented here are carried out with 6 pairs of thoracic CT images at end of expiration (EE) and end of inspiration (EI). The image data are from the CT image collection available through the website of the Deformable Image Registration Laboratory (DIR-Lab) at the University of Texas Medical Branch.^{16,17} The data collection includes two sets of patient images, with 10 patients in each set. One dataset contains 4DCT images acquired at the University of Texas MD Anderson Cancer Center as part of radiotherapy planning for thoracic malignancy treatment.¹⁶ In-slice spatial resolution is $(0.96 \text{ mm})^2$ and slice thickness is 2.5 mm. The other set contains EE and EI breath-hold CT images, taken from the COPDGene study archive of the National Heart, Lung, and Blood Institute.¹⁷ In-slice spatial resolution ranges from $(0.586 \text{ mm})^2$ to $(0.647 \text{ mm})^2$ and slice thickness is 2.5 mm. We refer to the data associated with each patient by the assigned label in the DIR-Lab website collections. We have selected the images of 1 patient from the 4DCT set and 5 patients from the COPD set.

Our selection was based on variations in displacement and spectral measures (Section II.C) of associated DVFs, such that we present the cases that pose a bigger challenge to DVF inversion. The forward DVF for each patient was obtained by one-way deformable image registration with the Velocity software (Varian Medical Systems, Palo Alto, CA, USA). The EE image is used as reference (primary) and the EI image as target (secondary), since the EE image is less susceptible to respiratory phase binning or motion artifacts. The 4DCT DVFs exhibit only small NTDCs, posing little challenge to effective feedback control design, whereas the COPD DVFs exhibit more diverse spatial and spectral measures. We select the five COPD DVFs with the largest amplitudes and variation in displacements to test the inversion iterations with the more challenging cases. We also select one DVF as representative of the 4DCT DVFs, using the same criterion. In what follows, we will provide measures and results in summaries for each of the 6 DVFs. We will also provide in detail results for COPD4 DVF, which is the most challenging case among the six DVFs with respect to inversion, as indicated by the highest control index values in Table I. The COPD4 DVF is displayed in Fig. 2 via displacement-vector quivers in axial, coronal, and sagittal slices, over a magenta-green overlay of the reference and target images.




FIG. 2 The forward DVF with patient COPD4, visualized by quiver-plots in axial, coronal, and sagittal slices, against an overlay of the reference CT image (EE, magenta) and target image (EI, green). The displayed quiver is spatially down-sampled by a factor of 12 along the LR and AP axes and a factor of 6 along the SI axis.

III.B. Evaluation measures

We assess feedback control parameter settings for the DVF inversion iteration with pre- and post-inversion evaluation measures. For pre-inversion evaluation, we use the spectral radii of the infinitesimal contraction matrix (13) over Ω to assess the error contraction area and ratio by each control setting. For post-inversion evaluation, we use the two IC residuals, \mathbf{r}_v and \mathbf{r}_u (Section II.A.1), measured in point-wise magnitudes. In addition to inversion evaluation measures, we report characterization measures of the forward DVFs. All measures are scalar fields over the displacement domain.

We define the valid displacement domain Ω as follows. Given the image domain Ω_0 for an input DVF, we get

$$\Omega = \Omega_0 - (\Omega_1 \cup \Omega_2), \quad (28)$$

where $\Omega_1 = \{\mathbf{x} \mid \mathbf{x} \in \Omega_0, \mathbf{f}(\mathbf{x}) \notin \Omega_0\}$ and $\Omega_2 = \Omega_0 - \{\mathbf{f}(\mathbf{x}) \mid \mathbf{x} \in \Omega_0\}$. That is, we exclude regions that are either mapped outside the original domain or not overlapping with the transformed domain.

Regarding summaries of scalar-field measures, we address two issues at once. First, a summary shall take into consideration the uncertainty in numerically provided DVF data due to regional delineation, noise, artifacts, and outliers. Second, it shall reflect spatial variation in the measure field and not obscure non-negligible changes in relatively small regions. Taking into account these two concerns, we summarize voxel-wise scalar measures over Ω via multiple percentiles (upper-bound values). Specifically, let ϕ be a scalar field over Ω , bounded from below by ϕ_{\min} . Let $\Omega(\phi = \tau) = \{\mathbf{x} \mid \phi(\mathbf{x}) = \tau, \mathbf{x} \in \Omega\}$ be the level set or iso-contour set of ϕ at value τ , and $p(\phi = \tau) = |\Omega(\phi = \tau)|/|\Omega|$ be the density of the level set. The β -th percentile value of ϕ is defined as

$$\phi[\beta\%] = \inf_{\tau} \left\{ \tau \mid \int_{\phi_{\min}}^{\tau} p(\phi = \tau') d\tau' > \beta\% \right\}. \quad (29)$$

In practical computation, we approximate $\phi[\beta\%]$ via a discrete histogram of ϕ . For all six patient DVFs, we report evaluation summaries with box-and-whisker plots showing 2nd, 10th, 50th, 90th, and 98th percentiles; and DVF characterization summaries in a table with 50th, 90th, and 98th percentiles. For the COPD4 DVF, we also display image slices of each volumetric measure field.

TABLE I Characterization summary of 6 patient DVFs by displacement lengths (along LR, AP, and SI axes) and spectral measures (determinant of the deformation Jacobian, spectral radius of the displacement Jacobian, and algebraic control index (25)). The rightmost column shows the fraction of voxels associated with complex eigenvalues. Percentiles in the other columns are defined as per (29).

Patient	β	$\mathbf{u}(\mathbf{x})$ (mm)			spectral measures			$\text{Im}(\lambda) \neq 0$
		LR	AP	SI	$ \mathbf{J}_f ^\S$	$\rho(\mathbf{J}_u)$	$1-2\gamma$	
4DCT7	50	0.7	0.9	1.3	1.0	0.1	-0.9	74%
	90	2.3	3.1	7.8	0.8	0.3	-0.7	
	98	4.3	5.3	13.3	0.6	0.4	-0.5	
COPD1	50	3.0	12.1	7.7	0.9	0.3	-0.8	51%
	90	8.3	32.1	25.9	0.4	0.7	-0.3	
	98	14.0	37.4	33.6	0.1	1.0	0.0	
COPD4	50	3.6	8.0	11.9	0.9	0.4	-0.7	44%
	90	10.2	23.9	34.6	0.3	0.9	-0.1	
	98	14.9	29.1	49.5	0.0	1.4	0.3	
COPD5	50	2.3	8.7	10.1	1.0	0.3	-0.7	58%
	90	7.7	30.9	31.2	0.4	0.8	-0.2	
	98	12.5	39.2	43.0	0.1	1.2	0.1	
COPD6	50	2.7	21.6	8.9	0.9	0.3	-0.7	69%
	90	7.9	26.0	21.6	0.4	0.7	-0.3	
	98	11.9	32.7	33.1	0.1	1.0	0.0	
COPD8	50	1.8	6.2	6.7	0.9	0.3	-0.7	57%
	90	6.3	15.2	23.5	0.4	0.8	-0.3	
	98	10.0	20.3	38.2	0.0	1.3	0.3	

\S (100 - β)-th percentiles (50th, 10th, 2nd) shown in the case of $|\mathbf{J}_f|$

III.C. DVF characterization

Table I lists spatial and spectral measures of the six forward DVFs. The spatial measures are displacement lengths along the LR, AP, and SI axes. The spectral measures were described in Section II.C.

The COPD DVFs have non-small NTDCs over 2% to 10% of the domain, by the criterion $\rho(\mathbf{J}_u) \geq 1$ or the equivalent criterion on the control index, $1 - 2\gamma \geq 0$. They also have regions with zero or negative determinants. Regions with non-small NTDCs seem to concentrate primarily, but not exclusively, around the diaphragm and chest wall where there is substantial motion due to inspiration. In order to guarantee convergence of the inversion iteration over no less than 98% of Ω , it is necessary to exercise adequate residual feedback control.

We present in Fig. 3 a detailed view of the spectral measures of the COPD4 DVF, via contoured heat-maps over image slices. As discussed in Section II.C, the control index maps (bottom row) are indeed most informative. Regions with non-small NTDCs ($1 - 2\gamma(\mathbf{x}) \geq 0$) are in the gray-to-orange color range, and the salient orange spots indicate where the controllability condition (19) is violated ($1 - 2\gamma(\mathbf{x}) \geq 1$). The other two maps provide complementary or mutually confirming information. The determinant maps (top row) show negative and zero values in blue and black, respectively; these are within the orange regions in the bottom maps. The maps in the

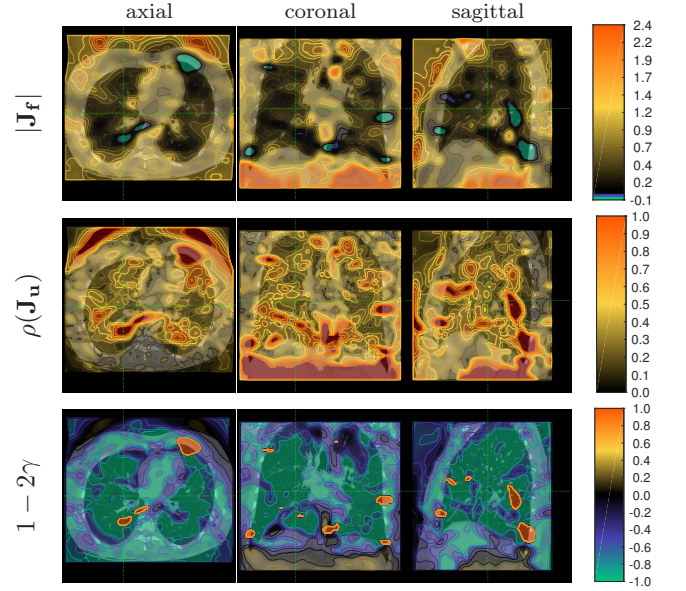


FIG. 3 Spatial variation of three spectral measures (Section II.C) with the COPD4 DVF, in volumetric, contoured heat-maps over the reference domain: $|\mathbf{J}_f(\mathbf{x})|$ (top); $\rho(\mathbf{J}_u(\mathbf{x}))$ (middle); and $1 - 2\gamma(\mathbf{x})$ (bottom). The display range for the determinant map is determined by its 90th percentile value as per (29). Zero-valued (black) regions in the determinant show where the deformation transformation is non-invertible. The maps of $\rho(\mathbf{J}_u)$ show that NTDCs are observed almost everywhere. The orange spots in the bottom maps show where the control index is greater than one; they coincide with the black and blue spots in the top maps, which indicate zero or negative determinant values.

middle row show the spectral radii of NTDCs. Regions with non-small NTDCs ($\rho(\mathbf{J}_u(\mathbf{x})) \geq 1$) are highlighted in red. The controllability condition holds over more than 98% of Ω . We suspect that the problematic spots where it is violated were artifacts of the registration process.

The 4DCT7 DVF has small NTDCs: $\rho(\mathbf{J}_u)$ is small, and $|\mathbf{J}_f|$ is not far from 1. The DVF satisfies the controllability condition virtually everywhere.

IV. RESULTS

We present results for iterative DVF inversion by (9) with three types of feedback control schemes: (i) uniform, constant parameter μ , with each of 4 values in $\{0, 0.3, 0.5, 0.7\}$, which includes the two precursor algorithms^{7,13}; (ii) alternating parameter values $\mu_{oe} = (\mu_o, \mu_e)$, where μ_o and μ_e are adaptively set as the mid-range values at the 50th and 98th percentiles, respectively (see Sections II.B.4.a and II.B.4.b, and $1 - 2\gamma$ in Table I); and (iii) spatially variant control $\mu_*(\mathbf{x})$ with locally optimal values by (20a) or (22a). All six control parameter settings are applied to each of the six DVFs.

Adaptive feedback control with the mid-range parameter value (Section II.B.4.a), although not reported as another control scheme along the 6 schemes listed above, is actually used in more than one way in the experiments. First, the 50th and 98th percentile mid-range values are employed in alternating-values control scheme. Second,




FIG. 4 Error contraction ratio percentiles over Ω for each of the 6 DVFs, with 4 feedback control parameter settings: constant $\mu = 0$ and $\mu = 0.5$, alternating μ_{oe} , and spatially variant $\mu_*(\mathbf{x})$. Box whiskers indicate the 2nd and 98th percentiles; the 10th and 90th percentiles are at the low and high ends of each box; and the median (50th percentile) is marked by a horizontal bar through each box.

the initial guess for each DVF inversion iteration is set to $\mathbf{v}_0(\mathbf{x}) = (\mu_m[98\%] - 1)\mathbf{u}(\mathbf{x})$ with all 6 control schemes, where $\mu_m[98\%]$ is the 98th percentile mid-range value. This is equivalent to taking a single step with mid-range value control and zero initial guess.

This data-adaptive initialization yields a better initial estimate, and mitigates an out-of-boundary issue with small-magnitude control parameter values. Consider, for instance, the iteration with $\mu = 0$ and zero initial guess. Even within the contraction region, it is likely that $\mathbf{x} - \mathbf{u}(\mathbf{x})$ falls outside the image boundary, and the iteration fails at the second step, $\mathbf{v}_2(\mathbf{x}) = -\mathbf{u}(\mathbf{x} - \mathbf{u}(\mathbf{x}))$, over a large number of voxels, such as those on or below the diaphragm, close to the inferior boundary.

We present in Section IV.A pre-inversion evaluation of contraction area and ratio with each control scheme, and in Section IV.B post-inversion evaluation of IC residuals with the inverse DVF estimates.

IV.A. Contraction area & ratio

We provide in Fig. 4 a summary of error contraction ratios, i.e., the spectral radii of the infinitesimal contraction matrices, for each of the 6 DVFs and with each of the following 4 control settings: $\mu = 0$, $\mu = 0.5$, alternating values μ_{oe} , and spatially variant values $\mu_*(\mathbf{x})$. The latter two are adaptive. Lower contraction ratios indicate faster error suppression. With the 4DCT7 DVF, the contraction ratios indicate convergence by all schemes, with the iteration with $\mu = 0.5$ at a much slower pace. With the COPD DVFs, except COPD4, the two constant, non-adaptive schemes are comparable to each other at the 90th percentile. At the 95th (not shown) and 98th percentiles, the scheme with $\mu = 0$ fails to contract over non-small NTDC regions, whereas the scheme with $\mu = 0.5$ maintains contraction ratios under 1. The value 0.5 happens to be in the control parameter range of each COPD DVF. The adaptive control scheme with alternating values for each DVF is better than the non-adaptive schemes up to the 90th percentile and in fact up to the 95th percentile (not shown). If it is desirable that the contraction region cover 98% or more of the domain, one shall use the




FIG. 5 Volumetric heap-maps of the error contraction ratio $\rho(\mathbf{Q}(\mathbf{x}; \mu))$ over the reference domain, with the COPD4 DVF. Comparison between 4 different feedback control parameter settings: $\mu = 0$; $\mu = 0.5$; alternating values, μ_{oe} , with $\mu_o = 0.15$ at odd steps and $\mu_e = 0.65$ at even steps; and spatially variant values, $\mu_*(\mathbf{x})$. Errors are suppressed more aggressively over the darker regions where $\rho(\mathbf{Q}(\mathbf{x}; \mu))$ is small, but are enlarged over the red regions, where $\rho(\mathbf{Q}(\mathbf{x}; \mu)) \geq 1$.

high-percentile mid-range value, which happens to be between 0.5 and 0.65 for each COPD DVF. The spatially variant scheme yields the lowest contraction ratios at all percentiles and with each DVF.

We display in Fig. 5 a comparison between the 4 control parameter settings in the spatial variation of error contraction ratios with the COPD4 DVF. The regions in red are non-contraction regions. Failure to contract is either due to violation of the controllability condition or due to inadequate feedback control (cf. Figs. 3 and 5). The violation regions are common to all heat-maps, and correspond to the orange spots in the control index maps in Fig. 3. The scheme with $\mu = 0$ has the smallest contraction area, indicating the failure of this non-adaptive control setting over the controllable region. Spatially variant control $\mu_*(\mathbf{x})$ yields the largest contraction area and lowest contraction ratios.

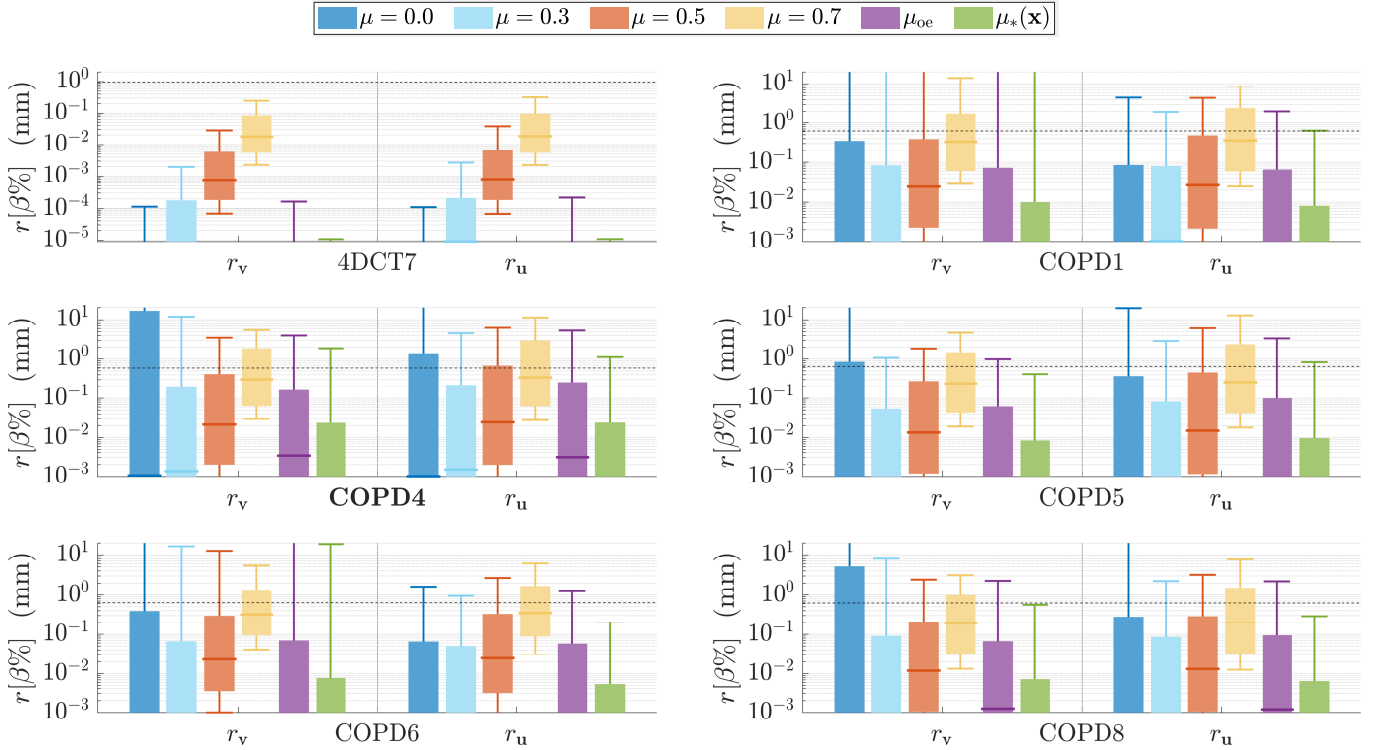


FIG. 6 Inverse consistency residuals $r_v(\mathbf{x})$ and $r_u(\mathbf{x} + \mathbf{v}(\mathbf{x}))$, reported in magnitude percentiles (29) in log-scale over the target domain, at the 10th iteration step with each patient and feedback control setting. Box whiskers indicate the 2nd and 98th percentiles; the 10th and 90th percentiles are at the low and high ends of each box; and the median (50th percentile) is marked by a horizontal bar through each box. The y -axis range is truncated to $[10^{-5}, 2]$ mm for the 4DCT7 plot, and to $[0.001, 20]$ mm for the COPD plots; percentile values outside these ranges are not shown. The horizontal dashed line in each plot indicates the in-slice resolution of the corresponding CT image.

IV.B. Inverse consistency residuals

We summarize in Fig. 6 post-inversion evaluation of the iteration (9) by 6 different feedback control schemes, for each of the 6 DVFs, in terms of the two IC residuals in Section II.A.1. The residuals are calculated at the 10th iteration step, measured by magnitude $r(\mathbf{x}) = \sqrt{r_{LR}^2(\mathbf{x}) + r_{AP}^2(\mathbf{x}) + r_{SI}^2(\mathbf{x})}$, and reported as percentile values in log-scale.

With the 4DCT7 DVF, IC residuals are below the in-slice resolution with every scheme. They are larger with $\mu = 0.5$ due to slower convergence, as expected by the pre-inversion evaluation, and even larger with $\mu = 0.7$.

With the COPD DVFs, we note first that substantial spatial variation of the residuals is observed for each DVF. The residuals at the 50th and lower percentiles are well below 1 mm with all control settings. How far below 1 mm is related to the convergence pace, since the contraction region for each control setting covers much more than half of the domain. The iteration with constant $\mu = 0.7$ is the slowest. We observe the following among the 4 constant, non-adaptive control schemes. With the same initialization, residual reduction does not correlate linearly with the values of μ : specifically for the 5 COPD DVFs, the $\mu = 0$ and $\mu = 0.7$ settings are inferior to $\mu = 0.3$ and $\mu = 0.5$ at high percentiles. The latter two settings are comparable at the 98th percentile, and the scheme with $\mu = 0.3$ yields smaller residuals up to

the 90th. No single constant-value scheme is as good as the adaptive scheme with alternating values. The residuals by the scheme with spatially variant values are much smaller than by other schemes, by roughly an order of magnitude or more at the 90th percentile. More remarkably, the 98th percentile residuals are reduced to below 2 mm by only 10 iteration steps with spatially variant control.

We provide in Figs. 7 and 8 a detailed comparison between 4 feedback control settings with the COPD4 DVF. Fig. 7 shows the spatial variation of the IC residuals at the 10th iteration step, visualized as volumetric quiver-plot and magnitude heat-map slices, overlaid on the target CT image. Fig. 8 shows the progression of residual magnitudes in 90th, 95th, and 98th percentile values for the first 15 iteration steps. Both figures highlight the advantage of NTDC-adaptive iterations over non-adaptive ones. In particular, the spatially variant control scheme reduces the residuals more rapidly and over the largest area. At the 10th step, by Fig. 7, it has removed or substantially reduced the problematic spots visible in the other heat maps. In fact, it renders residuals below the in-slice resolution (0.59 mm) by 7 steps at the 95th percentile, and by 15 steps at the 98th; the non-adaptive schemes take twice as many or more steps to reach the same range, or fail to reach it.

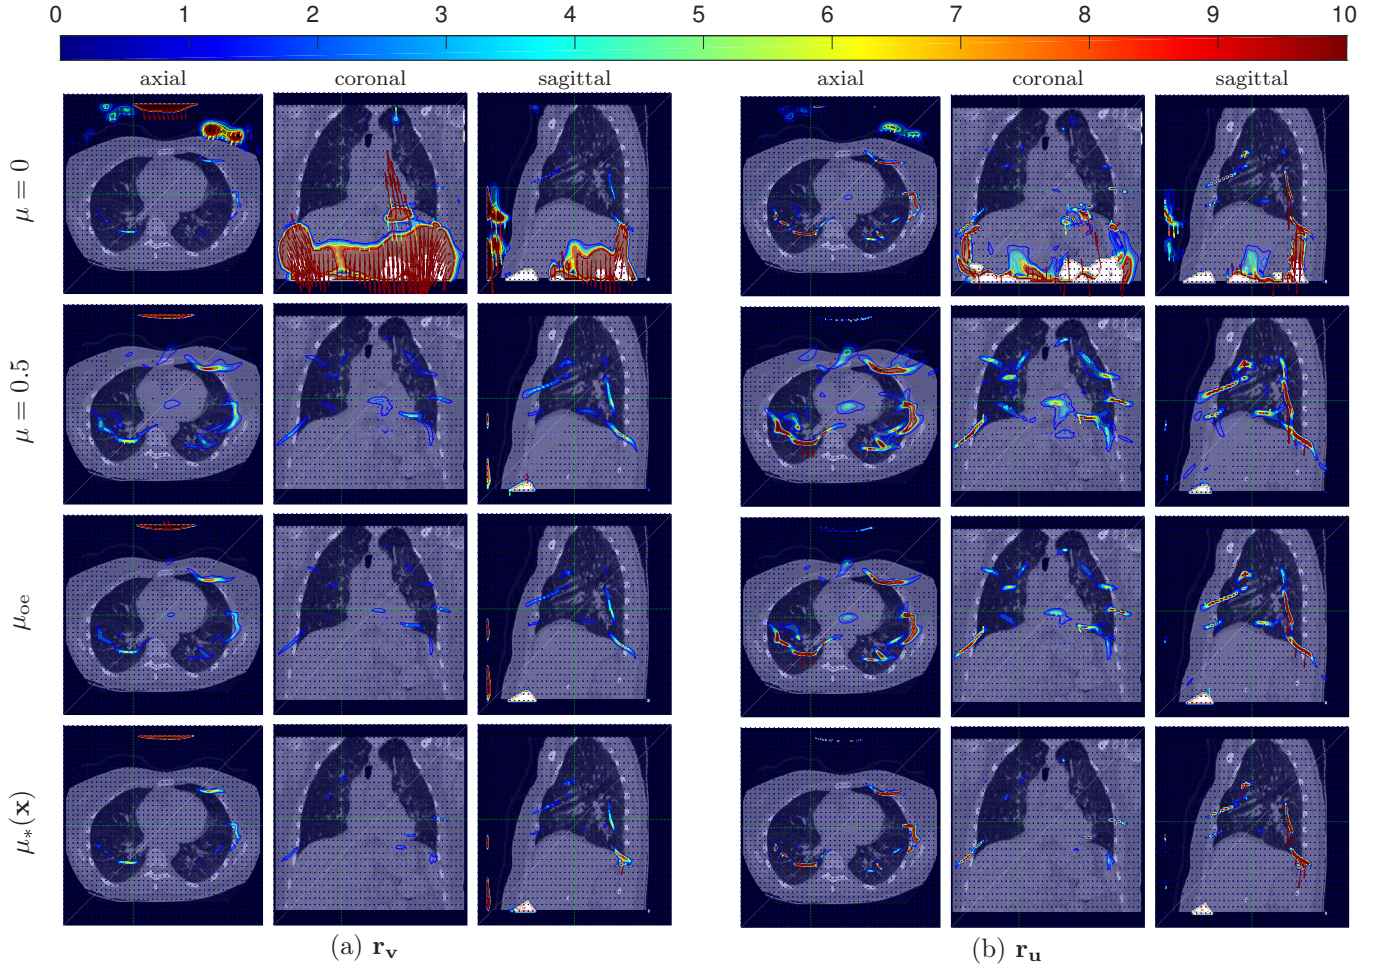


FIG. 7 Volumetric quiver-plots of $\mathbf{r}_v(\mathbf{x})$ and $\mathbf{r}_u(\mathbf{x} + \mathbf{v}(\mathbf{x}))$, with contoured heat-maps of magnitude (in mm), at the 10th iteration step with the COPD4 DVF, overlaid on the target image. Comparison between four feedback control schemes: constant $\mu = 0$, constant $\mu = 0.5$, alternating $\mu = \mu_{oe}$ (with $\mu_o = 0.15$ and $\mu_e = 0.65$), and spatially variant $\mu_*(\mathbf{x})$. The heat-map display range is truncated at 10 mm for visual inspection. White spots indicate regions where residual feedback entailed out-of-bounds values during the iteration.

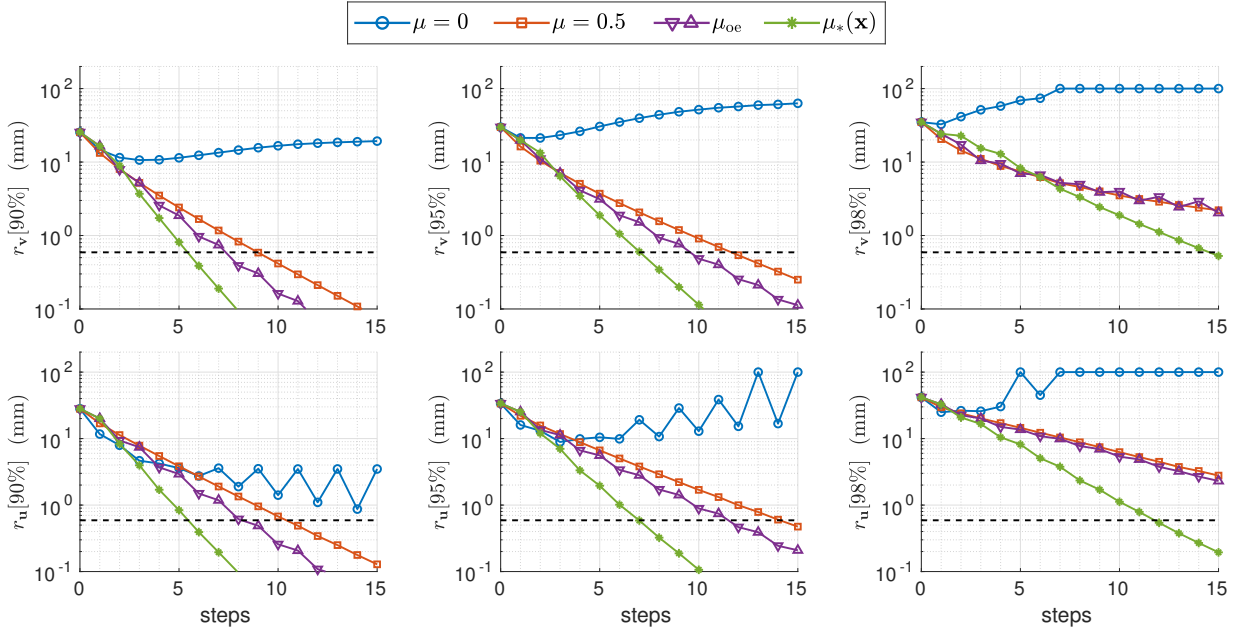


FIG. 8 Pace of inverse consistency residual suppression during the first 15 iteration steps with the COPD4 DVF, reported in percentiles of $r_v(\mathbf{x})$ and $r_u(\mathbf{x} + \mathbf{v}(\mathbf{x}))$ in log-scale. Comparison between four feedback control schemes: constant $\mu = 0$, constant $\mu = 0.5$, alternating $\mu = \mu_{oe}$ (with $\mu_o = 0.15$ and $\mu_e = 0.65$), and spatially variant $\mu_*(\mathbf{x})$. The black horizontal dashed lines mark the in-slice image resolution.

V. DISCUSSION AND CONCLUSION

We have elucidated and characterized the central role of non-translational displacement components (NTDC) in iterative DVF inversion. We have developed a framework of NTDC-adaptive algorithms for DVF inversion with a simple residual feedback control mechanism, and completed the framework with rigorous convergence analysis. Experimental results demonstrate the superior performance of our adaptive control methodology, in both convergence area and speed. We have also found remarkable agreement between pre-inversion assessment of control schemes, as arising from our analysis, and post-evaluation of inversion results.

The clinical utility of our algorithms and analysis can be reflected in multiple ways. (i) NTDC-adaptive iteration algorithms enable quick and accurate estimation of an inverse DVF, which is valuable in a number of clinical applications, such as 4D image reconstruction and adaptive radiotherapy. (ii) The spectral measures in Section II.C can be used, independently of any inversion task, for evaluating DVFs generated with existing software under clinical application conditions. (iii) The inversion algorithms, presented here as an asymmetric approach for generating the inverse of a forward mapping, can be incorporated with ease into existing software for simultaneous deformable image registration, which may consist of multiple forward and backward transformation stages.¹⁰ It is also plausible to employ the spectral measures as local regularization terms at each transformation stage, complementary to global ones. Additional, systematic studies are needed for such extensions. Potential benefits include rapid refinement in inverse consistency, increased robustness to the asymmetry in image quality between registered images, and reduced registration artifacts.

Lung deformations are of great clinical concern.^{1,2} DVFs obtained from thoracic CT images are used in turn to test our theory and algorithms. Further testing, with similar as well as different types of deformations, will better underscore the scope and impact of our theory and algorithms.

ACKNOWLEDGEMENTS

This work was supported by the National Institutes of Health Grant No. R01-CA184173. We thank the anonymous reviewers, Associate Editors, and Editor for their valuable comments on our previous manuscripts. We thank Dr. Richard Castillo for provision of the CT data used in this study, and Dr. Animesh Srivastava for technical assistance.

DISCLOSURE OF CONFLICTS OF INTEREST

The authors have no conflicts of interest to disclose.

REFERENCES

- a) A.K. Dubey and A.S. Iliopoulos have contributed equally to this study, and are considered as co-first authors.
- b) Author to whom correspondence should be addressed. Electronic mail: lei.ren@duke.edu
- ¹ L. Ren, Y. Zhang, and F.-F. Yin, "A limited-angle intrafraction verification (LIVE) system for radiation therapy," *Med. Phys.* **41**, 020701 (2014).
- ² C. Yan, H. Zhong, M. Murphy, E. Weiss, and J. V. Siebers, "A pseudoinverse deformation vector field generator and its applications," *Med. Phys.* **37**, 1117–1128 (2010).
- ³ T. Vercauteren, W. De Gerssem, L. A. Olteanu, I. Madani, F. Duprez, D. Berwouts, B. Speleers, and W. De Neve, "Deformation field validation and inversion applied to adaptive radiation therapy," *Phys. Med. Biol.* **58**, 5269–5286 (2013).
- ⁴ D. Rivest-Hénault, N. Dowson, P. B. Greer, J. Fripp, and J. A. Dowling, "Robust inverse-consistent affine CT-MR registration in MRI-assisted and MRI-alone prostate radiation therapy," *Med. Image Anal.* **23**, 56–69 (2015).
- ⁵ C. Gustafsson, F. Nordström, E. Persson, J. Brynolfsson, and L. Olsson, "Assessment of dosimetric impact of system specific geometric distortion in an MRI only based radiotherapy workflow for prostate," *Phys. Med. Biol.* **62**, 2976–2989 (2017).
- ⁶ G. S. K. Fung, L. Ciuffo, H. Ashikaga, and K. Taguchi, "Motion estimation for cardiac functional analysis using two x-ray computed tomography scans," *Med. Phys.* **44**, 4677–4686 (2017).
- ⁷ G. E. Christensen and H. J. Johnson, "Consistent image registration," *IEEE Trans. Med. Imag.* **20**, 568–582 (2001).
- ⁸ A. Leow, S.-C. Huang, A. Geng, J. Becker, S. Davis, A. Toga, and P. Thompson, "Inverse consistent mapping in 3D deformable image registration: its construction and statistical properties," in *IPMI 2005* (Glenwood Springs, CO, USA) pp. 493–503.
- ⁹ B. B. Avants, C. L. Epstein, M. Grossman, and J. C. Gee, "Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain," *Med. Image Anal.* **12**, 26–41 (2008).
- ¹⁰ A. Sotiras, C. Davatzikos, and N. Paragios, "Deformable medical image registration: a survey," *IEEE Trans. Med. Imag.* **32**, 1153–1190 (2013).
- ¹¹ M. P. Heinrich, I. J. Simpson, B. W. Papież, S. M. Brady, and J. A. Schnabel, "Deformable image registration by combining uncertainty estimates from supervoxel belief propagation," *Med. Image Anal.* **27**, 57–71 (2016).
- ¹² G. E. Christensen, "Consistent linear-elastic transformations for image matching," in *IPMI 1999* (Visegrád, Hungary) pp. 224–237.
- ¹³ M. Chen, W. Lu, Q. Chen, K. J. Ruchala, and G. H. Olivera, "A simple fixed-point approach to invert a deformation field," *Med. Phys.* **35**, 81–88 (2008).
- ¹⁴ M. R. Sabuncu, B. T. T. Yeo, K. Van Leemput, T. Vercauteren, and P. Golland, "Asymmetric image-template registration," in *MICCAI 2009* (London, UK) pp. 565–573.
- ¹⁵ J.-P. Thirion, "Image matching as a diffusion process: an analogy with Maxwell's demons," *Med. Image Anal.* **2**, 243–260 (1998).
- ¹⁶ R. Castillo, E. Castillo, D. Fuentes, M. Ahmad, A. M. Wood, M. S. Ludwig, and T. Guerrero, "A reference dataset for deformable image registration spatial accuracy evaluation using the COPDgene study archive," *Phys. Med. Biol.* **58**, 2861–2877 (2013).
- ¹⁷ E. Castillo, R. Castillo, J. Martinez, M. Shenoy, and T. Guerrero, "Four-dimensional deformable image registration using trajectory modeling," *Phys. Med. Biol.* **55**, 305 (2010).




FIG. 9 Image-space visualization of the analytical DVFs of (A1). (a) Reference image of concentric rings over $[-34, 34]^2$. (b) Target image obtained by deforming the reference image with the forward DVF of (A1a) with parameter values $m = 8$ and $b = 0.5$. (c) Target image by the forward DVF with parameter values $m = 8$ and $b = 0.8$.

Appendix A: Numerical inversion errors with an analytical DVF pair

We provide comparisons in numerical inversion errors among several DVF inversion algorithms in the framework of (9). The study of inversion errors, unlike IC residuals, requires the ground-truth inverse DVF. We use the 2D analytical DVFs introduced by Chen et al.¹³:

$$\mathbf{u}(\mathbf{x}') = \left(\frac{1}{1 + b \cos(m\theta(\mathbf{x}'))} - 1 \right) \mathbf{x}', \quad (\text{A1a})$$

$$\mathbf{v}(\mathbf{x}) = b \cos(m\theta(\mathbf{x})) \mathbf{x}, \quad (\text{A1b})$$

where \mathbf{x}, \mathbf{x}' lie in a 2D domain Ω , $b \in (0, 1)$ is a radial (stretch) parameter, $m \in \mathbb{N}$ is an angular (oscillation) parameter, and $\theta(\mathbf{x}) \in [0, 2\pi)$ is the angular coordinate of \mathbf{x} in the polar representation $\mathbf{x}^T = \|\mathbf{x}\|(\cos \theta(\mathbf{x}), \sin \theta(\mathbf{x}))$. The angular coordinate is well-defined everywhere except at $\mathbf{x} = \mathbf{0}$. It is straightforward to verify that the analytical DVFs are inverse to each other. The DVFs were visualized in image space by Chen et al.¹³ via deforming a specific reference image of concentric rings; we re-create such images in Fig. 9.

In numerical experiments, we discretize the analytical DVFs on a grid of finite resolution, apply each inversion algorithm to the discretized forward DVF, and compare the numerical inverse estimate to the inverse DVF. Specifically, we set $m = 8$ and $b = 0.8$, and discretize \mathbf{u} on a 2D grid over $[-34, 34]^2$ with spatial resolution 0.05. We let Ω be the sub-grid over $[-17, 17]^2$, where both DVFs are valid in the sense of (28). We assess the spectral properties of the numerical DVF and present in Fig. 10 its spectral maps by the three characterization measures introduced in Section II.C.

Four iterative inversion algorithms are examined. The same algorithms are compared in IC residuals in Section IV.B. The first two are existing work^{7,13}; they are non-adaptive and globally constant, with the control parameter values $\mu = 0$ and $\mu = 0.5$. The next two are adaptive to the DVF. One is non-stationary, using the alternating parameter values μ_o and μ_e (Section II.B.4.b), adaptively set to 50th and 98th percentile mid-range values over Ω . The other is spatially variant $\mu_*(\mathbf{x})$ (Section II.B.4.c), using locally optimal values (per (20a)




FIG. 10 Spectral maps of three characterization measures (Section II.C) of the discretized DVF (A1a) over $\Omega = [-17, 17]^2$. (a) Determinant map. Small (black) and large (orange) values show volume shrinkage and expansion, respectively. (b) Spectral NTDC radius map. Non-small NTDC regions, per (15) and (26), are highlighted in red. (c) Algebraic control index map. The index range shows that global convergence is feasible with adaptive feedback control; see (18) and (27).

and (22a), customized to 2D DVFs), except in a small neighborhood around $\mathbf{x} = \mathbf{0}$ where the mid-range value is used.

Inversion errors are calculated pixel-wise by $\mathbf{e}_k(\mathbf{x}) = \mathbf{v}_k(\mathbf{x}) - \mathbf{v}(\mathbf{x})$, with the estimate at the k -th iteration step and $\mathbf{x} \in \Omega$; see (2). They are presented in two complementary views. In the summary view (Fig. 11), we provide the sequence of inversion error percentiles (29) over each iteration process up to the 15th step. We show in Fig. 12 the spatial distribution of errors via snapshots of inversion error maps at three iteration steps, $k = 1, 8, 15$. In order to take into consideration also the response of each algorithm to the initial guess, we present the error-map snapshots with two different initial guesses: $\mathbf{v}_0^{[a]}(\mathbf{x}) = \mathbf{0}$ and $\mathbf{v}_0^{[b]}(\mathbf{x}) = 0.08 \cos(8\theta)[-x_2 \ x_1]^T$.

The two non-adaptive iterations fail in reducing the errors over the 8 radial ridge regions with $b = 0.8$, as shown in Fig. 12. The divergence regions correspond to, and can be predicted by, those with non-small NTDCs and high control index values in Fig. 10. The two adaptive iterations successfully suppress and annihilate the inverse errors. With the spatially variant adaptive control, inversion errors are reduced to sub-pixel length in no more than 3 steps. If the analytical values of $\mu_*(\mathbf{x})$ are used, the algorithm renders the inverse in a single step.

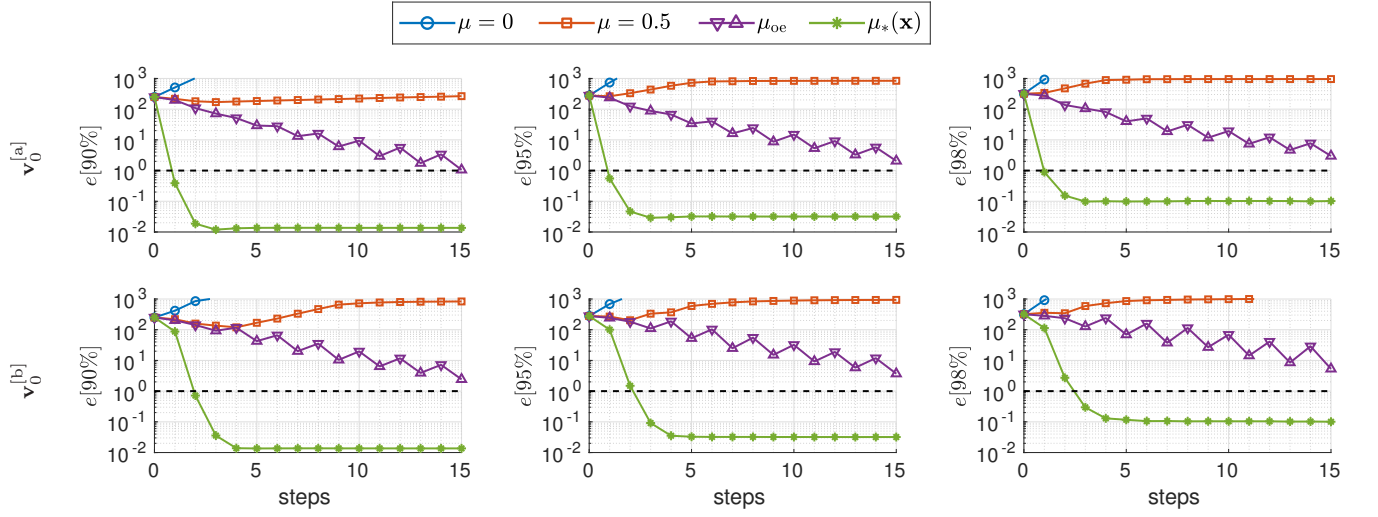


FIG. 11 Inversion error magnitudes in three percentiles (90th left, 95th middle, and 98th right) during the first 15 steps of each iteration with the discretized DVF, with two different initial guesses (top and bottom). Each plot shows the error sequence with each of the following four control schemes: constant $\mu = 0$, constant $\mu = 0.5$, alternating μ_{oe} with $\mu_o = 0$ and $\mu_e = 0.77$, and spatially variant $\mu_*(\mathbf{x})$. Error magnitudes, $e(\mathbf{x}) = \sqrt{e_{LR}^2 + e_{AP}^2 + e_{SI}^2}$, are measured in pixel-length unit and plotted in log-scale.

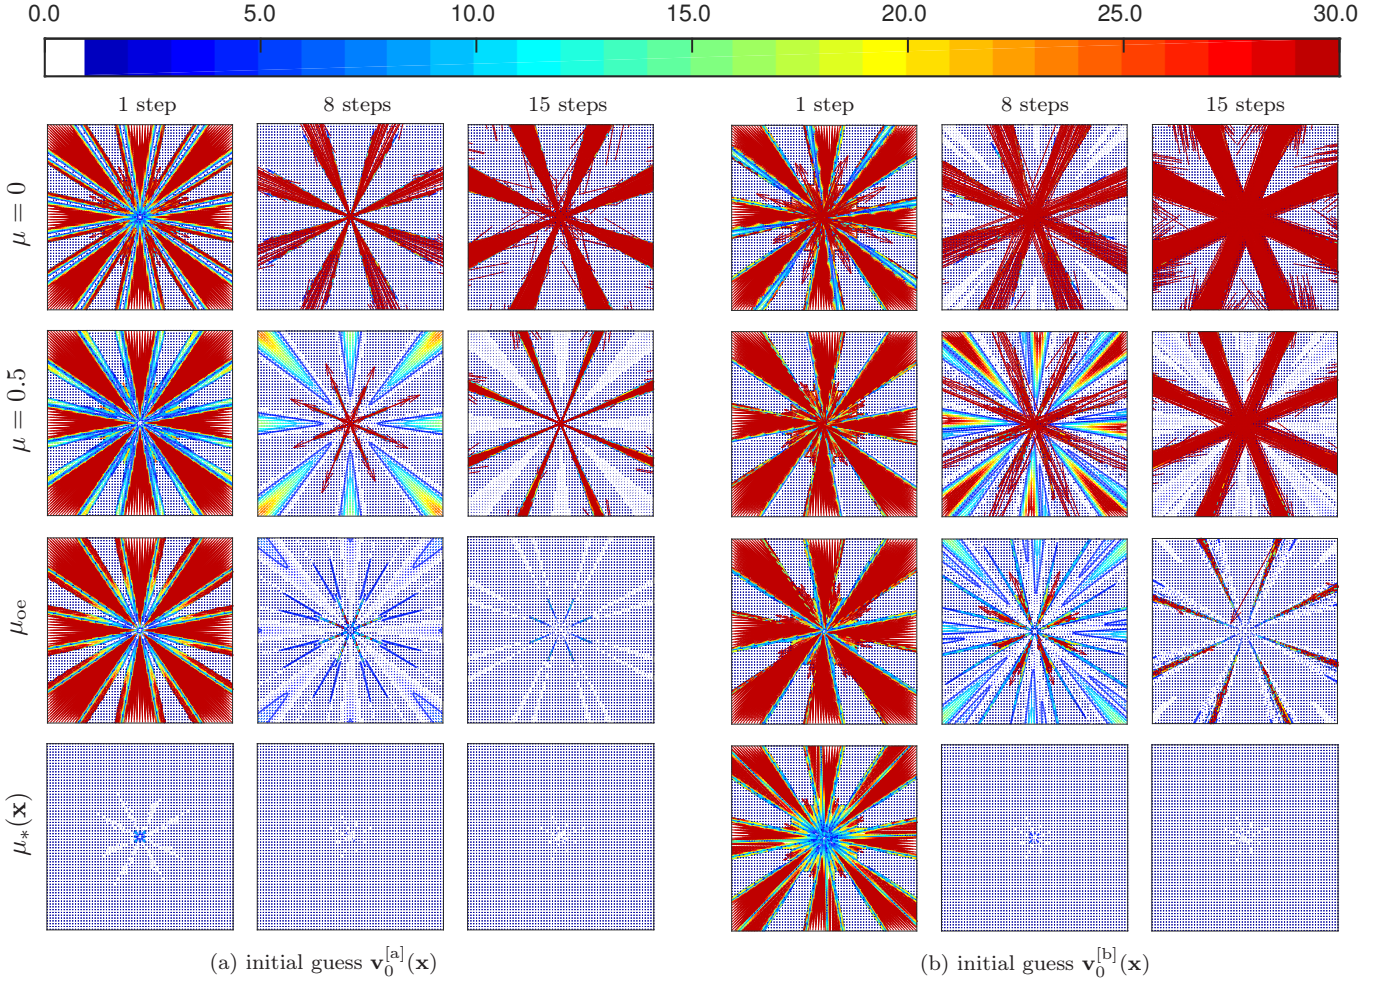


FIG. 12 Inversion error map snapshots at three steps ($k = 1, 8, 15$) during iterative inversion of the discretized DVF, with initial guess $\mathbf{v}_0^{[a]}$ (left) and $\mathbf{v}_0^{[b]}$ (right). Each row contains error maps with one of the following four control schemes (from top to bottom): constant $\mu = 0$, constant $\mu = 0.5$, alternating μ_{oe} with $\mu_o = 0$ and $\mu_e = 0.77$, and spatially variant $\mu_*(\mathbf{x})$. Error magnitudes are measured in pixel-length unit. Errors within a pixel are shown in white, and errors beyond 30 pixels are shown in red. The first two schemes fail in error suppression around the 8 radial ridge lines; the next two successfully suppress the errors over the entire domain. The scheme with $\mu_*(\mathbf{x})$ is robust to the change in the initial guess.