Computing global minimizers to a constrained B-spline image registration problem from optimal l_1 perturbations to block match data

Edward Castillo^{a)}

Department of Radiation Physics, University of Texas MD Anderson Cancer Center, 1515 Holcombe Boulevard, Unit 56 Houston, Texas 77030 and Department of Computational and Applied Mathematics, Rice University, 6100 Main MS-134, Houston, Texas 77005

Richard Castillo^{b)}

Department of Radiation Physics, University of Texas MD Anderson Cancer Center, 1515 Holcombe Boulevard, Unit 56 Houston, Texas 77030

David Fuentes^{c)}

Department of Imaging Physics, University of Texas MD Anderson Cancer Center, 1515 Holcombe Boulevard, Unit 1902 Houston, Texas 77030

Thomas Guerrero^{d)}

Department of Radiation Oncology, University of Texas MD Anderson Cancer Center, 1515 Holcombe Boulevard, Unit 56 Houston, Texas 77030 and Department of Computational and Applied Mathematics, Rice University, 6100 Main MS-134, Houston, Texas 77005

(Received 27 June 2013; revised 14 January 2014; accepted for publication 11 February 2014; published 10 March 2014)

Purpose: Block matching is a well-known strategy for estimating corresponding voxel locations between a pair of images according to an image similarity metric. Though robust to issues such as image noise and large magnitude voxel displacements, the estimated point matches are not guaranteed to be spatially accurate. However, the underlying optimization problem solved by the block matching procedure is similar in structure to the class of optimization problem associated with B-spline based registration methods. By exploiting this relationship, the authors derive a numerical method for computing a global minimizer to a constrained B-spline registration problem that incorporates the robustness of block matching with the global smoothness properties inherent to B-spline parameterization.

Methods: The method reformulates the traditional B-spline registration problem as a basis pursuit problem describing the minimal l_1 -perturbation to block match pairs required to produce a B-spline fitting error within a given tolerance. The sparsity pattern of the optimal perturbation then defines a voxel point cloud subset on which the B-spline fit is a global minimizer to a constrained variant of the B-spline registration problem. As opposed to traditional B-spline algorithms, the optimization step involving the actual image data is addressed by block matching.

Results: The performance of the method is measured in terms of spatial accuracy using ten in-hale/exhale thoracic CT image pairs (available for download at www.dir-lab.com) obtained from the COPDgene dataset and corresponding sets of expert-determined landmark point pairs. The results of the validation procedure demonstrate that the method can achieve a high spatial accuracy on a significantly complex image set.

Conclusions: The proposed methodology is demonstrated to achieve a high spatial accuracy and is generalizable in that it can employ any displacement field parameterization described as a least squares fit to block match generated estimates. Thus, the framework allows for a wide range of image similarity block match metric and physical modeling combinations. © 2014 American Association of Physicists in Medicine. [<http://dx.doi.org/10.1118/1.4866891>]

Key words: deformable image registration, l_1 optimization, thoracic CT, COPD

1. INTRODUCTION

Deformable image registration (DIR) algorithms were first developed for computer vision applications to estimate the apparent motion depicted in an image pair.¹ DIR is essentially the process of mapping the voxels contained in a reference image onto those contained in a target image according to assumptions placed on the properties of the transformation, as well as on the relationship between the image intensity values.² DIR is now essential for a wide host of appli-

cations in medical imaging including, CT derived ventilation imaging,^{3,4} functional brain imaging,^{5,6} and image guided radiotherapy.⁷ As such, automated, robust, and efficient methods demonstrated to achieve a high spatial accuracy are required.

Block matching based registration methods determine optimal corresponding reference and target voxel locations with respect to a given image similarity metric, such as sum of squared differences or mutual information, via an exhaustive search strategy.^{8,9} Though simplistic, the block matching

strategy is known to be robust to both image noise and large voxel displacements. Additionally, the approach requires no image derivative approximations.¹⁰ Moreover, the advent of GPU (Graphics Processing Unit) computing has made the computational burden associated with exhaustive searches feasible within an acceptable time frame.¹¹ However, the main drawback to block matching strategies is the fact that an optimal point match, in terms of the similarity metric, may not be optimal in terms of spatial accuracy. For instance, the size and location of a specified search window may not contain the most spatially accurate target location, resulting in an erroneous point match. Similarly, in the presence of large magnitude image deformations, the spatially accurate point match may not be optimal in terms of the image similarity due to anatomic deformation. As a consequence, DIR methods based on block matching typically incorporate a local filtering approach^{12,13} or thresholding procedure¹⁴ to dampen the effects of erroneous point matches, often referred to as “outliers,” before computing the full DIR transformation. In the context of rigid or affine registration, robust fitting based on statistical tools, such as least squares trimming, is a popular and effective approach for recovering the transformation coefficients from block match generated data.¹⁵ In essence, this method systematically removes the contributions of point matches with high fit error. However, the approach only generalizes to parameterizations that are well defined on an unstructured spatial distribution of data points (as polynomial functions are for example). Otherwise, the data removal procedure is not guaranteed to maintain numerical stability without global regularization (as in Ref. 16).

DIR methods based on B-spline parameterizations typically recover an optimal set of coefficients by solving a nonlinear least squares optimization problem defined by a given image similarity metric.¹⁷ B-spline methods have the benefit of reducing the full DIR problem to that of a smaller coefficient recovery problem, while at the same time, incorporating the inherent smoothness properties of the B-spline basis functions into the DIR solution. The B-spline framework is also flexible enough to allow for the incorporation of additional properties, such as inverse consistency,^{18,19} sliding motion consistency,²⁰ and temporal dimension modeling,²¹ into the recovered transformation. However, the optimization problem associated with B-spline DIR is highly nonlinear and nonconvex, meaning that the objective function is likely to possess many erroneous local minima. Consequently, finding a global minimizer is in general a difficult problem and the solutions computed by gradient-based optimization methods, such as Levenberg-Marquardt, Gauss-Newton, or BFGS, are sensitive to the guess used to initialize the algorithm. Existing B-spline DIR methods often adopt a coarse-to-fine multiresolution strategy to reduce the chances of converging to an erroneous local minimizer.^{22,23} Moreover, most gradient based methods require twice differentiable objective functions.²⁴ Considering image data are typically not even continuous, the image data must first be smoothed by applying a Gaussian filter. The smoothing ensures continuous first- and second-order derivative approximations but comes at the expense of destroying fine image structures.²⁵

From a mathematical perspective, block matching is a robust strategy for computing the global minimizer to a nonlinear least squares problem defined on an image grid. The block matching solution describes the optimal displacement for a single voxel. B-spline based DIR also requires the solution of a nonlinear least squares problem defined on an image grid. However, the B-spline formulation describes a globally defined displacement field. In this study, we derive a numerical method for computing a global minimizer to a constrained B-spline registration problem by exploiting the similar structure within the block matching and B-spline formulations. The constrained formulation is defined on a point cloud of voxel locations within the reference image and can be solved by decoupling the optimization problem into a basis-pursuit optimization problem and a series of block matching problems. Using the basis pursuit formulation, it is possible to describe the minimal l_1 perturbation to block match generated displacement estimates that produces a residual B-spline fitting error within a certain tolerance. The sparsity pattern of the optimal perturbation then dictates a point cloud on which the recovered B-spline coefficients are globally optimal. The proposed approach combines the numerical robustness of block matching with the global smoothness properties inherent to B-splines without requiring *a priori* Gaussian smoothing or image derivative approximations.

In this study, we demonstrate the spatial accuracy of the proposed algorithm using the ten thoracic computed tomography inhale/exhale image pairs and corresponding sets of expert determined landmark point pairs previously described in Ref. 26 (available at www.dir-lab.com). These ten challenging validation cases were acquired from the COPDgene study (www.copdgene.org).²⁷ Registration of the COPDgene image pairs is challenging for several reasons (Fig. 1). First, the inhale image is acquired at maximum effort, while the exhale image is acquired at end normal expiration. This results in large (on the order of 100 voxel) displacements, changes in density and consequently CT value, and changes in anatomic shape of the vasculature due to the large volume change. Second, in order to reduce the radiation dose to the study subjects, the exhale images are acquired at one-fourth the CT tube current as the inhale image. As a consequence, the exhale image noise is doubled. We compare the proposed algorithm on

FIG. 1. Coronal slices for the inhale (left) and exhale (middle) image pair represented by COPD Case 5 on the www.dir-lab.com website. The exhale image was acquired at one-quarter the CT tube current used to acquire the inhale image. Consequently, the exhale image contains more image noise. A fusion of the two coronal slices (right) illustrates the magnitude of the deformation between the two images. The scale is in xy-plane voxels.

the same data set with the top-performing algorithm from the EMPIRE10 study.²⁸

2. B-SPLINE AND BLOCK MATCHING DIR AS NONLINEAR LEAST SQUARES PROBLEMS

Pairwise DIR algorithms determine a spatial mapping

$$\psi(\mathbf{x}) = \mathbf{x} + \phi(\mathbf{x}), \quad \psi : \mathbb{R}^3 \rightarrow \mathbb{R}^3, \quad (1)$$

that relates the voxels in the reference image $R(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}$, to those within a target image $T(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}$. Recovery of the unknown displacement field $\phi(\mathbf{x})$ is then the goal of a DIR algorithm. A common strategy is to assume the displacement field can be described as a linear combination of N B-spline basis functions

$$\phi(\mathbf{x}; \mathbf{a}) = \sum_{i=1}^N \mathbf{a}_i \beta_i(\mathbf{x}), \quad \mathbf{a}_i \in \mathbb{R}^3, \quad \beta_i : \mathbb{R}^3 \rightarrow \mathbb{R}, \quad (2)$$

where each β_i is centered on a control point (also referred to as a knot) $\mathbf{c}_i$ is located on a uniformly spaced voxel grid dictated by the support radius of the basis function. Formulation (2) reduces the registration problem to identifying the optimal set of basis function coefficients that best relates the intensity information in the reference image to that of the target image. For instance, employing the popular “sum of squared differences” metric yields the following nonlinear least squares problem:

$$\min_{\mathbf{a}} \sum_{\mathbf{x}_j \in X} [T(\mathbf{x}_j + \phi(\mathbf{x}_j; \mathbf{a})) - R(\mathbf{x}_j)]^2, \quad (3)$$

where the voxels contained in the set X are taken from a region of interest within the reference image. Typically, X is the set of control points corresponding to the spline parameterization, though this need not be the case. The solution to Problem (3) defines a displacement field describing the motion of every voxel within the support of the B-spline parameterization.

Block matching employs a simplistic modeling approach describing the displacement vector of a single voxel $\mathbf{x}_i$ as the solution to the following nonlinear least squares problem

$$\min_{\mathbf{d}} \sum_{\mathbf{x}_j \in \Omega_{\sigma}(\mathbf{x}_i)} [T(\mathbf{x}_j + \mathbf{d}) - R(\mathbf{x}_j)]^2, \quad (4)$$

where $\Omega_{\sigma}(\mathbf{x}_i)$, often referred to as the “matching window,” is a voxel neighborhood with radius σ centered on $\mathbf{x}_i$. Assuming the unknown displacement vector is integer valued, the global minimizer can be computed robustly for any similarity or dissimilarity metric F using an exhaustive search strategy

$$\hat{\mathbf{d}}_i = \arg \min_{\mathbf{d} \in B_s} \sum_{\mathbf{x}_j \in \Omega_{\sigma}(\mathbf{x}_i)} F^2(\delta; \mathbf{x}_j, R, T), \quad (5)$$

where $B_s = \{\mathbf{d} \in \mathbb{Z}^3 : \|\mathbf{d}\|_{\infty} \leq s\}$ is the “search window.”

Since Problem (5) only describes the displacement of a single voxel, block matching DIR strategies are based on computing a data set of estimated point correspondences by solving Problem (5) for a subset of voxel locations. However, block matching is not guaranteed to always produce spatially accurate data.^{12,14,15} Consequently, the process of generating a full displacement field from the estimated point matches

typically requires the dampening or removal of erroneous data “outliers.”

Problems (3) and (5) both describe unknown motion as the solution to a nonlinear least squares problem defined by image data. The simplicity of the block matching formulation allows for robust exhaustive search, but only describes the motion of a single voxel. Though the B-spline formulation describes the entire displacement field, a more sophisticated optimization algorithm is required to solve the problem and care must be taken to address the difficulties associated with nonlinear, nonconvex optimization. In order to combine the global smoothness properties of B-splines with the numerical robustness of block matching, we propose reformulating the B-spline DIR Problem (3) as a constrained nonlinear least squares problem based on the block matching formulation (5).

3. A CONSTRAINED NONLINEAR LEAST SQUARES FORMULATION FOR B-SPLINE DIR

A B-spline parameterization is defined in part by the resolution of the control point grid. Given a displacement vector estimate for each voxel control point, the spline coefficients yielding an exact interpolation are described by square linear systems of equations. This type of interpolation strategy assumes the estimates are spatially accurate. However, this property is not guaranteed for block match generated data.

In addition to control point estimates, incorporating the displacement estimates for a point cloud of voxel positions equally spaced throughout a region of interest (as illustrated in Fig. 2), results in an overdetermined linear system for each spatial dimension component of the parameterization. Specifically, given a set of voxel locations $X = \{\mathbf{x}_i\}_{i=1}^N$ and a B-spline parameterization defined by the control points $C = \{\mathbf{c}_i\}_{i=1}^M$ with $C \subset X$, the least squares fit to the displacement estimates $D = \{\mathbf{d}_i\}_{i=1}^N$ is given as

$$\mathbf{a}_{lsq}^{(k)} = \arg \min_{\mathbf{a}} \frac{1}{2} \|\mathbf{A}\mathbf{a} - \mathbf{D}^{(k)}\|_2^2, \quad (6)$$

FIG. 2. The B-spline control point grid (green crosses) and a point cloud of voxel locations (red dots) distributed throughout a region of interest (lung volume). The full 3D control point grid is displayed (bottom right) within a volume rendering of the full 3D ROI.

where $k = 1, 2, 3$ indicates the spatial dimension and $A_{ij} = \beta_j(\mathbf{x}_i)$. The entries in the right hand side vector $D^{(k)} = (\mathbf{d}_1^{(k)}, \mathbf{d}_2^{(k)}, \dots, \mathbf{d}_N^{(k)})^T$ are the displacement estimates (in the k th spatial dimension) for the voxels in X .

Though robust to Gaussian noise, standard least squares fitting is still susceptible to the effects of data outliers.²⁹ However, the magnitude of the residual (fitting error) for Problem (6) provides an overall estimate for the quality of the fit. For each spatial dimension component, the magnitude of the least square residual can be derived from the “QR” factors of A , in terms of the right hand side vector $D^{(k)}$,

$$\|Aa_{lsq}^{(k)} - D^{(k)}\|_2^2 = \|QD^{(k)}\|_2^2, \quad (7)$$

where $A = \bar{Q}r$ with $r \in \mathbb{R}^{M \times N}$, $\bar{Q} = [Q_1, Q_2] \in \mathbb{R}^{N \times N}$ is orthonormal and $Q = Q_2^T \in \mathbb{R}^{N-M \times N}$ (see Ref. 24 for full derivation). The benefit of expression (7) is that it provides a straightforward manner for computing the magnitude of the least squares residual for any set of displacement estimates. Thus, Eq. (7) provides an expression for relating block match generated estimates to the quality of the corresponding least squares B-spline fit.

The process of applying the block matching strategy to each $\mathbf{x}_i \in X$ can be expressed mathematically as a nonlinear least squares optimization problem

$$\min_D \sum_{i=1}^N \sum_{\mathbf{x}_j \in \Omega_\sigma(\mathbf{x}_i)} [T(\mathbf{x}_j + \mathbf{d}_i) - R(\mathbf{x}_j)]^2. \quad (8)$$

Problem (8) is fully decoupled, meaning that the i th nonlinear equation is not influenced by any other optimization variables aside from $\mathbf{d}_i$. Combining Formulation (8) with the QR representation of the least squares residual in Eq. (7) results in the following constrained nonlinear least squares problem:

$$\begin{aligned} \min_D \sum_{i=1}^N \sum_{\mathbf{x}_j \in \Omega_\sigma(\mathbf{x}_i)} [T(\mathbf{x}_j + \mathbf{d}_i) - R(\mathbf{x}_j)]^2 \\ \text{s.t. } \|QD^{(k)}\|_2^2 \leq \rho, \text{ for } k = 1, 2, 3. \end{aligned} \quad (9)$$

Problem (9) represents a B-spline DIR problem defined solely by the displacements of the voxels in X . The inequality constraints require that the displacement estimates yield a least squares fitting error within a given tolerance $\rho \geq 0$. As a result, the B-spline coefficients are essentially absorbed into the inequality constraints. At first glance, the increased number of unknowns as well as the introduction of linear inequality constraints make Problem (9) appear more challenging than the standard B-spline Formulation (3). However, a global minimizer to a relaxed variant of Problem (9) can be computed by exploiting the structure of the problem and incorporating robust l_1 optimization.

4. MINIMAL l_1 PERTURBATION TO BLOCK MATCH DATA

The motivation for this work comes from sparse signal recovery methods where the goal is to recover a large and sparse signal using a relatively few number of linear measurements.^{30,31} Previous registration methodologies

based on robust statistical regression operate under the assumption that at least 50% of block match generated data is spatially accurate.¹⁵ Under the additional assumption that the true displacement field is accurately described by a B-spline parameterization, the error vector denoting the distance between the block match estimates and the true B-spline fit is sparse. Thus, a strategy for computing a global minimizer to Problem (9) is to find the B-spline coefficients satisfying the inequality constraints that yield the sparsest residual error vector for a given set of block match estimates. In general, this process is NP-hard and requires a combinatorial search of the block match data. Instead, we propose solving an l_1 optimization problem describing the minimal block match estimate perturbations required to satisfy each of the B-spline inequality constraints. Specifically, given a set of block match computed displacement estimates $\hat{D} = \{\hat{\mathbf{d}}_i\}_{i=1}^N$ (as defined in Eq. (5)) for the voxel set X , the optimal l_1 perturbation for the k th spatial dimension is defined as

$$\begin{aligned} P^{(k)} = \arg \min_Y \|Y\|_1 \\ \text{s.t. } \|Q(\hat{D}^{(k)} + Y)\|_2^2 \leq \rho. \end{aligned} \quad (10)$$

For $\rho = 0$, Problem (10) is known as a “basis-pursuit” problem and can be solved by several numerical methods (see Ref. 32, for example).

The resulting perturbed displacements $\tilde{\mathbf{d}}_i = \hat{\mathbf{d}}_i + (P_i^{(1)}, P_i^{(2)}, P_i^{(3)})^T$ satisfy the B-spline constraints (as illustrated in Fig. 3). However, only estimates paired with a zero magnitude perturbation are optimal in terms of the image similarity metric, since block match displacement estimates are each globally optimal. Considering that block match displacements are defined on an integer grid, a sufficiently small magnitude perturbation preserves the global optimality of the estimate. The source locations corresponding to the estimates satisfying this property

$$X^* = \{\mathbf{x}_i \in X : |P_i^{(k)}| \leq \epsilon, k = 1, 2, 3\} \quad (11)$$

FIG. 3. Optimal l_1 perturbation to block match data. (Left) The original block match estimates are illustrated within a sagittal projection (test case 1 in Table I). The darker (red) estimates denote spatially inaccurate outliers. (Right) The optimal l_1 perturbation obtained by solving Problem (10) is applied to the original data. The darker (blue) estimates denote the perturbed positions of the erroneous estimates. The remaining (yellow) data are the estimates corresponding to a zero magnitude perturbation. The perturbed data satisfy the B-spline fitting constraint and is therefore smoother than the original block match data.

generate a B-spline coefficient matrix $A^* \in \mathbb{R}^{N^* \times M}$ with

$$A_{ij}^* = \beta_j(\mathbf{x}_i), \quad \text{for all } \mathbf{x}_i \in X^*, \quad (12)$$

and $N^* = |X^*|$. The matrix A^* can also be derived by removing the rows corresponding to the voxels $\mathbf{x}_i \notin X^*$ from A [as defined in Eq. (6)]. If A^* is full rank, then the block match solutions corresponding to the voxels in X^* ,

$$\hat{D}^* = \{\hat{\mathbf{d}}_i \in \hat{D} : |P_i^{(k)}| \leq \epsilon, \quad k = 1, 2, 3\} \quad (13)$$

represent a global minimizer to Problem (9) restricted to the voxel subset X^* ,

$$\begin{aligned} \min_D \sum_{\mathbf{x}_i \in X^*} \left(\sum_{\mathbf{x}_j \in \Omega_\epsilon(\mathbf{x}_i)} [T(\mathbf{x}_j + \mathbf{d}_i) - R(\mathbf{x}_j)]^2 \right) \\ \text{s.t. } \|Q^* D^{(k)}\|_2^2 \leq \rho, \quad \text{for } k = 1, 2, 3, \end{aligned} \quad (14)$$

where Q^* is the least squares residual matrix corresponding to the QR factorization of A^* . Therefore, the block matching solutions corresponding to the sparsity pattern of the optimal l_1 perturbation [as defined by Eq. (11)] represent a global minimizer to the constrained B-spline nonlinear least squares problem defined on X^* . This result follows from the fact that each block match is optimal in terms of the image similarity metric and the perturbed estimates satisfy the inequality fitting constraint. The matrix A^* is full rank, which is required to ensure the resulting B-spline fit is well-defined, when the support of each B-spline basis function contains at least one $\mathbf{x}_i \in X^*$.

5. BLOCK MATCH FILTERING FROM MINIMAL l_1 PERTURBATION

Though the optimal spline coefficients corresponding to the solution of Problem (14) can easily be computed from the perturbed positions, the optimal B-spline fit is not guaranteed to exactly interpolate the displacements in $\hat{D}^*$. In fact, this is only the case when $\rho = 0$ and Problem (14) reduces to a standard l_1 fitting problem. However, the minimal l_1 perturbation approach can be viewed as a point match filtering strategy where the solution sets X^* and D^* represents trusted landmark data. As suggested in Ref. 12, the trusted data can then be interpolated confidently to compute the full DIR using a method such as moving least squares (MLS).³³ With this type of approach, the optimal B-spline fitting parameters are ultimately inconsequential.

Previous block matching methods based on outlier detection operate under the assumption that at least half the computed block match data are spatially accurate.^{12,15} Thus, the value of the ρ used to define the B-spline fitting constraint should be chosen to reflect this property. In general, a meaningful value for ρ is difficult to set *a priori*. An alternative approach is to solve an unconstrained penalty function problem

$$\min_Y \|Y\|_1 + \frac{1}{\lambda} Q(\hat{D}^{(k)} + Y)^2. \quad (15)$$

For every value λ , it can easily be shown that there exists a value of ρ such that Problems (15) and (14) are equivalent.³⁴ For instance, setting $\lambda = \infty$ reduces Problem (15) to the basic

block matching Problem (8), or equivalently, to Problem (9) with $\rho = \infty$. However, too high a value of λ could potentially relax the B-spline constraint to the point of allowing erroneous block matches into the final solution set.

A filtered block match dataset representing a solution to the B-spline DIR Problem (14) can be generated by the following algorithm:

Algorithm: Minimal l_1 Perturbation (MILO) to Block Match Data

1. Set $\lambda = \lambda_0$, $d\lambda > 0$, and the voxel locations $X = \{\mathbf{x}_i\}$, $i = 1, 2, \dots, N$.
2. Compute the block match estimates $\hat{D} = \{\hat{\mathbf{d}}_i\}$.
3. Solve Problem (15)
4. If $|X^*| > \gamma|X|$ stop, otherwise, set $\lambda = \lambda + d\lambda$ and go to step 3.
5. Generate a full displacement field from the final solutions sets X^* and $\hat{D}^*$.

The value of $\gamma \in (0, 1)$ represents the percentage of block matches assumed to be spatially accurate.

6. NUMERICAL IMPLEMENTATION

The computational workload required by the MILO algorithm is dominated by the block matching exhaustive search component. Considering the embarrassingly parallel nature of block matching, our current implementation is written in the Compute Unified Device Architecture (CUDA) C++ language for use with massively parallel graphics processing units (GPU).³⁵ The l_1 optimization Problem (15) is solved in MATLAB using the open source YALL1 (Your ALgorithms for L1) software package available at www.yall1.blogs.rice.edu and described in Ref. 34. The B-spline component of the algorithm utilizes cubic splines and is also implemented in MATLAB.

A full DIR is computed from the filtered block matches returned by the MILO algorithm using MLS. A modified version of the MLS method described in Ref. 36 is used to generate a full displacement field from trusted block match landmarks. The modification is a preprocessing heuristic designed to guard against folding in the MLS transformation. Specifically, an estimate of the Jacobian determinant is computed from the MLS “diffuse derivative” (as defined in Ref. 37) at each voxel. Data points with a negative estimate are replaced by the MLS approximate computed solely from data points with positive estimates. The full MLS method described in Ref. 36 is then applied to the modified data set. Though this approach does not fully guarantee that the resulting MLS transformation will be free of folding, the approach yields transformations with no detectable folding for all the test cases detailed in Table III. The MLS software is implemented for first degree polynomials in parallel using C++ and the “OpenMP” library.

The cost of computing the QR factorization of the B-spline matrix A becomes increasingly expensive for larger values of N . In order to reduce the computational burden and increase the spatial accuracy of the block match generated

point matches, we employ a “mutual minimizer” strategy where the solution $\hat{\mathbf{d}}_i$ to the block match Problem (5) is considered acceptable if and only if

$$-\hat{\mathbf{d}}_i = \arg \min_{\delta \in B_i} \sum_{\mathbf{x}_j \in \Omega_\sigma(\mathbf{x}_i + \hat{\mathbf{d}}_i)} F^2(\delta; \mathbf{x}_j, T, R). \quad (16)$$

The enforcement of this constraint provides robustness to the basic grid search procedure by only selecting point matches that correspond to image intensity blocks that are consistent in both the reference and target image for a given search window size. Intuitively, the constraint is more difficult to satisfy for larger search windows. The mutual minimizer constraint is only applied to noncontrol point voxel locations in order to maintain a full rank parameterization matrix. Instead of the sum of squares difference formulation, normalized cross-correlation is used as the block matching metric in order to account for the substantial changes in voxel intensities depicted in the validation image pairs.²⁶

7. NUMERICAL EXPERIMENTS AND RESULTS

The spatial accuracy of the proposed DIR strategy was measured using ten challenging thoracic computed tomography inhale/exhale image pairs from the COPDgene database and corresponding sets of expert-determined landmark point pairs. The validation data are publicly available at the www.dir-lab.com website and described fully in Ref. 26. Table I details the image and landmark characteristics for each of the ten cases.

The MILO algorithm was initialized using a two-step multiresolution scheme. First, an affine function is fit to the block match estimates acquired on low resolution (one-quarter the full resolution in the xy-plane and full resolution in the z-dimension) images. The low resolution block match windows were set to $7 \times 7 \times 3$ voxels and the search window was set to $50 \times 50 \times 50$ voxels, considering the large average magnitude displacements detailed in Table I. A midresolution registration was then obtained by applying the MILO algorithm to images with half resolution in the xy-plane and full resolution in the z-dimension using a $7 \times 7 \times 3$ voxel match window with a search window of $31 \times 31 \times 11$ voxels.

At full resolution, the matching window was set to $11 \times 11 \times 3$ voxels, which is approximately an isotropic neighborhood of $7.0 \times 7.0 \times 7.0$ mm, with a search window of $31 \times 31 \times 11$ voxels (a larger window was used in the direction of largest displacement). Previous block matching studies indicate that the results are not overly influenced by the match window size.¹² Similar to previous B-spline DIR studies, the B-spline control points were uniformly spaced 30mm apart²⁰ within the inhale image. In addition to the control points, a point cloud of uniformly spaced voxel locations (spaced two voxels apart in all directions) was then determined by a “dart throwing” procedure³⁸ applied to simple histogram segmentations of the inhale and exhale lung volumes (see Ref. 39 for details on segmentation procedure), as illustrated in Fig. 2. The MILO algorithm registration was applied to the

segmented lung volumes with a target of 60% sparsity for the optimal perturbation vector. This percentage overall yielded the best results based on numerical experiments.

All experiments were performed on a Silicon Mechanics Hyperform HPCg R2504.v2 workstation with two Intel Xeon X5675 Six-Core 3.06GHz processors and a NVIDIA Tesla 2070 GPU. The run times for the MILO algorithm ranged from 10 to 20 min depending on the number of block matches computed and the corresponding l_1 optimization problem size. The precise timings for each component of the algorithm are detailed in Table II.

For each of the ten cases, a full DIR was computed from the trusted block match data generated by the MILO algorithm using the stable MLS method described in Ref. 36. The MLS displacement field for Case 1 is given in Fig. 4. In addition to the MLS generated displacement field, the cubic B-spline field corresponding to the optimal block match data was also generated. Table I reports the spatial accuracies achieved by the MILO algorithm for each of the ten cases using the MLS field and the B-spline field. As expected, the MLS approach achieves a lower spatial error than the B-spline generated field for all cases. Though the errors achieved by the B-spline field are low considering the large magnitude of expert-determined landmarks, the MLS method maintains a higher data fidelity to the trusted block match estimates and consequently achieves a lower average error. This result suggests that the filtered block match data generated by the MILO algorithm is spatially accurate.

For comparison, the top performing algorithm obtained from the EMPIRE10 study,²⁸ the greedy diffeomorphic registration (GSYN),⁴⁰ was applied to all COPD cases using the open source ANTs⁴¹ library and the parameters reported in Ref. 42. Briefly, the lung masks (same masks employed by the MILO algorithm) were used to zero intensity values outside the lungs. A linear adjustment was applied to the values within the lungs, I , to normalize intensity values,

$$\frac{\max(I) - I}{\max(I) - \min(I)} \in [0, 1].$$

The GSYN registration consisted of (1) an affine registration step and (2) a diffeomorphic step. A multiresolution approach was applied in both steps. Intuitively, the lowest resolutions are computationally efficient and capture bulk movement. Each increase resolution recovers more fine scale features. A four-level and five-level image pyramid was applied to the affine registration and diffeomorphic registration steps, respectively. A gradient descent was used in the optimization of the affine and diffeomorphic registration. The neighborhood radius in computing cross correlation was 2 voxels. The Gaussian regularization kernel width was 6 voxels. The observed run time of the ANTs solver were similar to those previously reported.⁴² The GSYN and MILO registrations were comparable, as reported in Table I, with MILO achieving a higher spatial accuracy on eight of the ten cases and a lower average error overall. The error vectors produced by both methods for case 8 (the case with the largest difference in spatial accuracy between the two methods) are illustrated in Fig. 5. Additionally, the amount of folding present

TABLE I. The image and expert determined landmark characteristics, as well as the average millimeter error (standard deviation) computed by the MILO Algorithm for ten test cases are given. The voxel dimensions for each case denote the voxel spacing in terms of *mm*. The observer errors refer to the reproducibility of the expert-determined landmarks as described in Ref. 26. The average millimeter error achieved by the MILO algorithm when using Moving Least Squares and the optimal B-spline fit to generate the full displacement field are given in the RL (Right/Left), AP (Anterior/Posterior), and SI (Superior/Inferior) component directions, as well as in three-dimensional Euclidean space (3D). For comparison, the results produced by the greedy diffeomorphic registration (GSYN) described in Ref. 40 are also given.

Case	Voxel dimensions	Observer error 3D		No DIR	MLS	Spline	GSYN
1	$0.625 \times 0.625 \times 2.50$	0.65 (0.73)	RL	4.38 (3.27)	0.38 (0.47)	0.48 (0.58)	0.39 (0.78)
			AP	24.05 (11.06)	0.41 (0.52)	0.55 (0.60)	0.82 (2.10)
			SI	7.54 (6.08)	0.37 (0.88)	0.55 (1.06)	0.51 (1.30)
			3D	26.33 (11.44)	0.93 (0.92)	1.22 (1.07)	1.33 (2.45)
2	$0.645 \times 0.645 \times 2.50$	1.06 (1.51)	RL	3.18 (2.33)	0.58 (0.74)	0.69 (0.76)	0.71 (0.99)
			AP	19.64 (6.70)	0.98 (1.36)	1.10 (1.27)	0.93 (1.35)
			SI	6.38 (5.47)	0.83 (1.58)	1.29 (1.99)	1.41 (3.26)
			3D	21.77 (6.47)	1.77 (1.92)	2.28 (2.09)	2.23 (3.44)
3	$0.652 \times 0.652 \times 2.50$	0.58 (0.87)	RL	2.05 (1.55)	0.40 (0.48)	0.52 (0.53)	0.41 (0.60)
			AP	6.34 (3.63)	0.42 (0.51)	0.50 (0.55)	0.48 (0.95)
			SI	9.97 (6.44)	0.39 (0.91)	0.61 (1.09)	0.48 (1.14)
			3D	12.64 (6.40)	0.99 (0.91)	1.27 (1.02)	1.05 (1.45)
4	$0.590 \times 0.590 \times 2.50$	0.71 (0.96)	RL	5.46 (4.28)	0.40 (0.50)	0.63 (0.63)	0.41 (0.73)
			AP	16.96 (10.20)	0.50 (0.54)	0.66 (0.67)	0.57 (1.45)
			SI	20.31 (13.82)	0.53 (1.06)	1.23 (1.73)	0.34 (0.91)
			3D	29.58 (12.95)	1.14 (1.04)	1.92 (1.57)	1.02 (1.74)
5	$0.647 \times 0.647 \times 2.50$	0.65 (0.87)	RL	4.24 (3.29)	0.34 (0.44)	0.57 (0.64)	0.38 (0.77)
			AP	23.79 (10.66)	0.55 (0.93)	0.86 (1.28)	0.49 (1.29)
			SI	15.28 (11.91)	0.40 (0.96)	0.65 (1.21)	0.33 (1.00)
			3D	30.08 (13.36)	1.02 (1.23)	1.56 (1.60)	0.92 (1.70)
6	$0.633 \times 0.633 \times 2.50$	1.06 (2.38)	RL	3.35 (2.53)	0.39 (0.46)	0.52 (0.67)	0.39 (0.63)
			AP	24.77 (6.24)	0.52 (0.77)	0.69 (0.93)	0.74 (2.00)
			SI	11.64 (9.39)	0.33 (0.90)	0.67 (1.53)	0.34 (1.17)
			3D	28.46 (9.17)	0.99 (1.08)	1.44 (1.67)	1.18 (2.28)
7	$0.625 \times 0.625 \times 2.50$	0.65 (0.78)	RL	3.29 (2.73)	0.40 (0.47)	0.54 (0.54)	0.41 (0.71)
			AP	9.47 (5.90)	0.52 (0.69)	0.72 (0.79)	0.53 (1.18)
			SI	18.11 (7.50)	0.39 (0.97)	0.83 (1.26)	0.43 (1.13)
			3D	21.60 (7.74)	1.03 (1.08)	1.59 (1.22)	1.06 (1.64)
8	$0.586 \times 0.586 \times 2.50$	0.96 (3.07)	RL	3.79 (2.87)	0.53 (0.58)	0.68 (0.77)	0.71 (1.30)
			AP	14.03 (6.02)	0.68 (1.21)	0.97 (1.48)	1.06 (2.36)
			SI	19.74 (15.18)	0.53 (1.40)	0.99 (1.72)	1.03 (2.85)
			3D	26.46 (13.24)	1.31 (1.76)	1.89 (2.13)	1.92 (3.80)
9	$0.664 \times 0.664 \times 2.50$	1.01 (2.54)	RL	3.87 (2.75)	0.35 (0.46)	0.47 (0.53)	0.33 (0.50)
			AP	7.60 (4.44)	0.47 (0.77)	0.62 (0.86)	0.53 (1.26)
			SI	10.05 (10.79)	0.24 (0.79)	0.51 (1.04)	0.36 (1.27)
			3D	14.86 (9.82)	0.86 (1.06)	1.24 (1.20)	0.97 (1.74)
10	$0.742 \times 0.742 \times 2.50$	0.87 (1.65)	RL	6.05 (4.20)	0.54 (0.74)	0.66 (0.78)	0.64 (1.13)
			AP	7.06 (4.86)	0.59 (0.80)	0.67 (0.77)	0.70 (1.18)
			SI	18.18 (11.31)	0.45 (1.04)	0.93 (1.37)	0.74 (1.98)
			3D	21.81 (10.51)	1.23 (1.27)	1.70 (1.40)	1.53 (2.39)
Average 3D:				23.38 (11.86)	1.13 (1.29)	1.61 (1.57)	1.32 (2.42)

in the transformations produced by each method was quantified by computing the number of voxel locations within the image mask with a negative Jacobian determinant. The partial derivatives of each transformation were estimated using first order centered differences applied to the values on the voxel grid. The results, as detailed in Table III, indicate that while the B-spline and GSYN transformations contain only a fraction of voxel locations with detectable folding, the MLS field contains no areas of detectable folding.

8. DISCUSSION

B-spline DIR methods are typically based on a nonlinear least squares optimization problem. These problems are in general difficult to solve due to both the standard challenges that arise in nonlinear nonconvex optimization problems, such as local minima, and the fact that raw image data are discontinuous. As such, the application of a Newton's method type optimization algorithm requires *a priori* smoothing of the data

TABLE II. The MILO algorithm consists of three computational components: block matching, l_1 optimization, and MLS. For each case, the run times corresponding to spatial accuracies reported in Table I are given, as well as the run times for each algorithm component. The algorithm is initialized by a two-step multiresolution scheme. The run time for initialization step is also given. All timings are reported in seconds. The number of voxels for which a mutual minimizer block match search was performed (as described in the Numerical Implementation) is also provided.

Case	Initialization time	Block matching (number of voxels)	l_1 Solver	MLS	Total run time
1	2.77	4.33 (237, 843)	9.45	3.28	19.85
2	2.43	4.13 (228, 218)	3.58	3.18	13.33
3	2.50	3.78 (207, 824)	12.12	1.77	20.17
4	3.13	3.50 (193, 280)	5.97	4.15	16.75
5	2.63	3.65 (200, 762)	6.27	2.72	15.27
6	2.92	2.68 (147, 862)	4.95	1.40	11.95
7	2.48	2.57 (140, 855)	4.92	1.05	11.07
8	4.32	3.22 (177, 738)	6.08	1.73	15.35
9	2.62	2.47 (136, 702)	4.97	0.87	10.92
10	3.68	3.42 (187, 372)	3.63	1.63	12.37

in order to generate meaningful image derivative approximations needed by the algorithm. In addition, Newton methods are only guaranteed to converge to a local minimizer. The proposed MILO algorithm framework is based on addressing this shortcoming by finding global minimizers of a reformulated B-spline problem defined on a subset of voxel locations. Incorporating the simplistic block match formulation into the B-spline problem removes the dependence on image smoothing and gradient approximations all together and allows for

FIG. 4. Representative slices are shown in trans-axial (top), coronal (middle), and sagittal (bottom) section for test case 1. Each image shows color fusion of the inhale and exhale breath-hold datasets, with displacement vectors within the lung sampled on a uniform grid shown overlain. The sampled vectors are shown to scale, and depict the 2D inplane displacement components corresponding to the displayed image sections.

FIG. 5. The differences between the 300 expert-determined landmark positions and the DIR estimated positions are displayed as error vectors within volume renderings of test case 8 for the GSYN algorithm (left) and the MILO algorithm (right). The scale depicts the magnitude (in millimeters) of the error vector. The largest difference in average error computed by the two algorithms occurred on test case 8 with the MILO algorithm achieving an average error of 1.31 mm and GYSN achieving an average error of 1.92 mm.

the straightforward derivation of simple conditions for finding a global minimizer. In contrast, the greedy diffeomorphic algorithm⁴⁰ depends on a gradient-descent optimization scheme. The method performed comparably in terms of spatial accuracy, but overall produced a higher average error than the MILO algorithm.

Unlike traditional B-spline registration methods, the MILO solution is defined in terms of individual voxel displacement vectors and not the basis function coefficients. The minimal l_1 perturbation required to satisfy the B-spline fitting constraint essentially acts as a data filter with trusted block matches corresponding to the sparsity pattern. As a corollary, data corresponding to large magnitude perturbations represent erroneous outliers. Previous methodologies for removing spatially inaccurate block matches employing outlier detection or removal schemes borrowed from robust statistical regression have been shown to be effective for low dimensional parametrized data fittings.^{12,15} However, these approaches break down numerically for parameterizations defined by basis functions with compact support. The proposed MILO algorithm avoids this problem by operating solely on the right hand side vectors defined by the block match estimates, and is therefore stable for any parameterization that can be described as a least squares fit to estimated voxel displacements.

The idea of incorporating outlier removal into a regularized DIR formulation is known to be an effective DIR approach. The method proposed in Ref. 16 essentially fits a biomechanical model to a set of outlier-free block matches determined by applying the least trimmed squares algorithm. The proposed MILO strategy is similar in that a DIR model (in this case B-splines) is used to

TABLE III. The percentage of voxels with a negative Jacobian determinant produced by the MLS, B-spline, and GSYN transformations is given for each of the ten test cases. First order centered differences were used to approximate the partial derivatives. The MILO MLS field does not produce any negative Jacobian determinant values.

	1	2	3	4	5	6	7	8	9	10
MLS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Spline	3.88×10^{-2}	3.72×10^{-3}	0.00	4.12×10^{-1}	4.62×10^{-1}	5.67×10^{-2}	0.00	1.01×10^{-1}	0.00	6.56×10^{-3}
GSYN	1.19×10^{-2}	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

remove erroneous block matches. However, the MILO approach employs a sparsity promoting l_1 -optimization formulation that explicitly defines the relationship between the outlier-free data and the corresponding image registration optimization problem. Moreover, the derivation of the MILO algorithm is only dependent on a general least squares relationship between the block matches and the DIR model. Therefore, it is possible to substitute the biomechanical model for the B-spline model within the MILO framework to compute the DIR.

The method described in Ref. 43 is based on interpolating block match estimates with thin plate splines. However, the issue of erroneous data is not explicitly addressed. Instead, the method relies on first determining a set of voxels corresponding to easily identifiable landmarks. This idea could also be incorporated into the MILO framework since it does not require any structure on the spatial distribution of data. Reducing the number of data points has the benefit of reducing the l_1 problem size, and consequently, the overall computation workload of the algorithm. The current MILO implementation adopts a mutual minimizer block matching approach (as described in the Numerical Implementation) to achieve an effect similar to the one produced by the easily identifiable point search used in Ref. 43. A more sophisticated block match constraint is introduced in Ref. 44. The idea is to categorize block match estimates according to criteria defined by the characteristics of the image similarity objective function. More specifically, a block match is considered acceptable if the associated global minimum value is below a predefined threshold and the estimated objective function is convex at the minimizer. Unacceptable match estimates are adjusted according to an elastic energy regularization model. The full displacement field is then interpolated from the accepted and adjusted block match data using a B-spline parameterization. In contrast, trusted block match data determined from the sparsity pattern of the optimal l_1 perturbation within the MILO algorithm are used to determine the full displacement field via MLS. Though the optimal l_1 perturbation provides adjusted position estimates for the untrusted block matches as well, the perturbed locations are not guaranteed to be optimal in terms of the image similarity metric. As such, the MILO algorithm ignores the adjusted data points. In this way, the resulting MLS field is determined solely from the block match data corresponding to a global minimizer of the constrained B-spline optimization problem. As evidenced by the lower spatial accuracy achieved by the B-spline field, the MLS approach is superior.

The MILO sparse residual formulation for block match filtering does not place any constraints on the distribution of the filtered output data. MLS is in general an approximation method design for scattered “point cloud” data.⁴⁵ At any given location $\mathbf{x}$ within the image domain, the MLS field is evaluated by locally fitting (in the least squares sense) a polynomial to the scattered data and then evaluating the optimal polynomial at $\mathbf{x}$. The polynomial fit is typically weighted to ensure a smoothly varying approximation. Error bounds derived for the MLS approximate depend locally on the density of the scattered data.^{33,46} Given that the filtered data computed by the MILO algorithm is not guaranteed to possess a uniform spatial distribution, standard MLS is inappropriate. The MLS variant, referred to as “stable MLS,” introduced in Ref. 36 defines a Voronoi weighting scheme designed to maintain approximation accuracy and smoothness for inhomogeneous data distributions. However, in order to guard against folding in the MLS transformation, our implementation of the stable MLS procedure applies a heuristic correction based on the sign of the estimated Jacobian determinant computed from the diffuse derivative of the MLS approximation. Assuming the MILO generated data are spatially accurate, as the numerical results suggest, the inherent smoothness typically assumed of thoracic deformations is maintained by the modified stable MLS procedure, despite an unknown data distribution. The fact that the corresponding optimal B-spline solutions are less spatially accurate than the MLS fields, despite their smoothness, implies the 30 mm spacing of the spline control points used for this study does not provide enough degrees of freedom to fully describe the underlying motion associated with the COPDgene data set.

An underlying assumption of the MILO algorithm is that the model parameterization can sufficiently describe the true deformation field. Though derived within the specific context of B-spline parameterizations, the generality of the MILO framework easily allows for the incorporation of more sophisticated biomechanical modeling approaches, such as linear elasticity,⁴⁷ the Laplacian heat transfer equation,¹ and more complex biomechanics. Employing these types of sophisticated models could better describe the motion of diseased lung tissue, thereby improving the overall spatial accuracy of the method. For instance, within the context of thoracic image registration, sliding motion along the thoracic cage boundary represents a discontinuity in the deformation field that is not accurately represented by straightforward B-spline parameterizations. The spatial accuracies reported in this study were obtained by registering segmented lung volumes. Because the

MILO algorithm assumes the employed DIR model can accurately describe the true deformation field, erroneous block matches can be represented as data outliers with respect to the optimal l_1 perturbation. However, a discontinuous deformation field generates a structured and consistent signal within the block match dataset. Consequently, the filtered data set will depend on the proportion of block matches corresponding to each competing signal. Employing a more physically realistic DIR model, such as the direction dependent B-spline parameterization suggested in Ref. 20, would alleviate this problem. Moreover, considering the generality of the block matching strategy, employing different combinations of metrics and modeling can broaden the method's applicability to other anatomic sites and imaging modalities.

Though the bulk of the computation workload is represented by the block matches, the majority of the computation time is spent solving the l_1 optimization problem (as described in Table II). Block matching easily lends itself to GPU implementations and has been studied extensively (see Ref. 48, for example). Though the run time for the block matching component of the algorithm is comparable to previous DIR block matching studies,¹² employing a more powerful GPU or multiple GPUs could greatly improve the computational performance of the method. However, the current implementation of the MILO algorithm utilizes a MATLAB based sequential code (described in Ref. 34) for the l_1 optimization problem. The software is based on the Alternating Direction Method of Multipliers⁴⁹ (ADMM), which is known to parallelize for certain classes of problems.⁵⁰ Thus, an immediate area of future research is the implementation of a parallelized ADMM for use within the MILO algorithm.

Recent work in the field of COPD requires accurate registration of image pairs similar to the ones used for validation of the MILO algorithm.^{51,52} However, as detailed in Ref. 26, such cases pose significant challenges for DIR methods, in particular, large magnitude deformations, voxel intensity variations, and image noise. Despite these challenges, the MILO algorithm demonstrates a high spatial accuracy on a significantly complex image set. Given this high accuracy, CT-derived ventilation imaging^{3,4} can be reliably computed from MILO solutions as a means for quantifying COPD and potentially aid in the identification of imaging biomarkers, in addition to improving the quality of existing COPD analysis methods.

9. CONCLUSION

In this study, we introduce the *Minimal l_1 Perturbation* method for DIR. Using cubic B-splines as the parameterization model, we demonstrate the method's high spatial accuracy on ten cases from the COPDgene image database using large set of expert determined landmark point pairs. The method was demonstrated a high spatial accuracy despite the large magnitude voxel deformations, significant variations in voxel intensities, and the high amount of image noise common to COPDgene image pairs. The proposed methodology is generalizable in that it can employ any displacement field parameterization described as a least squares fit to block match

generated estimates. As such, an immediate area of future research is to investigate the effectiveness of different parameterizations and image similarity metrics on both thoracic CT images, as well as images from other anatomical sites and different imaging modalities.

ACKNOWLEDGEMENTS

We wish to thank the COPDgene Ancillary Studies Committee and Executive Committee who approved this study. This work was partially funded by the National Institutes of Health through an NIH Director's New Innovator Award DP2OD007044 and an NIH Research Supplement to Promote Diversity in Health Related Research 3DP2OD007044-01S1. Richard Castillo was partially supported by an NIH Research Scientist Development Award K01CA181292.

- a)Electronic mail: ecastillo3@mdanderson.org. URL: www.caam.rice.edu/~ec.
- b)Electronic mail: ricastillo@mdanderson.org. URL: www.dir-lab.com.
- c)Electronic mail: dtfuentes@mdanderson.org. URL: www.dir-lab.com.
- d)Electronic mail: tguerrero@mdanderson.org. URL: www.dir-lab.com.
- ¹B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artif. Intell.* **17**, 185–203 (1981).
- ²J. Modersitzki, *Numerical Methods for Image Registration* (Oxford University Press, New York, 2004).
- ³T. Guerrero, K. Sanders, E. Castillo, Y. Zhang, L. Bidaut, T. Pan, and R. Komaki, "Dynamic ventilation imaging from four-dimensional computed tomography," *Phys. Med. Biol.* **51**, 1–15 (2006).
- ⁴R. Castillo, E. Castillo, J. Martinez, and Th. Guerrero, "Ventilation from four-dimensional computed tomography: Density versus jacobian methods," *Phys. Med. Biol.* **55**, 4661–4685 (2010).
- ⁵A. Toga and P. Thompson, "The role of image registration in brain mapping," *Image Vis. Comput.* **19**, 3–24 (2001).
- ⁶A. Gholipour, N. Kehtarnavaz, R. Briggs, M. Devous, and K. Gopinath, "Brain functional localization: A survey of image registration techniques," *IEEE Trans. Med. Imaging* **26**, 427–451 (2007).
- ⁷M. Foskey, B. Davis, L. Goyal, S. Chang, E. Chaney, N. Strehl, S. Tomei, J. Rosenman, and S. Joshi, "Large deformation three-dimensional image registration in image-guided radiation therapy," *Phys. Med. Biol.* **50**, 5869–5892 (2005).
- ⁸P. Anandan, "A computational framework and an algorithm for the measurement of visual motion," *Int. J. Comput. Vis.* **2**, 283–310 (1989).
- ⁹J. Barron, D. Fleet, and S. Beauchemin, "Performance of optical flow techniques," *Int. J. Comput. Vis.* **12**, 43–77 (1994).
- ¹⁰S. S. Beauchemin and J. L. Barron, "The computation of optical flow," *ACM Computing Surveys (CSUR)* **27**(3), 433–466 (1995).
- ¹¹G. Pratz and L. Xing, "GPU computing in medical physics: A review," *Med. Phys.* **38**, 2685–2697 (2011).
- ¹²E. Castillo, R. Castillo, B. White, J. Rojo, and T. Guerrero, "Least median of squares filtering of locally optimal point matches for compressible flow image registration," *Phys. Med. Biol.* **57**, 4827–4833 (2012).
- ¹³O. Commowick, V. Arsigny, A. Isambert, J. Costa, F. Dhermain, F. Bidault, P.-Y. Bondiau, N. Ayache, and G. Malandain, "An efficient locally affine framework for the smooth registration of anatomical structures," *Med. Image Anal.* **12**, 427–441 (2008).
- ¹⁴V. Garcia, O. Commowick, and G. Malandain, "A robust and efficient block-matching framework for non linear registration of thoracic CT images," in *Proceedings of Grand Challenges on Medical Image Analysis Clinic* (CreateSpace, 2010), pp. 137–146.
- ¹⁵S. Ourselin, A. Roche, S. Prima, and N. Ayache, "Block matching: A general framework to improve robustness of rigid registration of medical images," *Medical Image Computing and Computer-Assisted Intervention MICCAI 2000* (Springer, Berlin, 2000).
- ¹⁶O. Clatz, H. Delingette, I.-F. Talos, A. Golby, R. Kikinis, F. Jolesz, N. Ayache, and S. Warfield, "Robust nonrigid registration to capture brain shift from intraoperative MRI," *IEEE Trans. Med. Imaging* **24**, 1417–1427 (2005).

- ¹⁷R. Szeliski and J. Coughlan, "Spline-based image registration," *Int. J. Comput. Vis.* **22**, 199–218 (1997).
- ¹⁸I. Arganda-Carreras, C. S. Sorzano, R. Marabini, J. Carazo, C. Ortiz-de Solorzano, and J. Kybic, "Consistent and elastic registration of histological sections using vector-spline regularization," in *Computer Vision Approaches to Medical Image Analysis* (Springer, Berlin, 2006).
- ¹⁹Y. Yin, E. Hoffman, and C.-L. Lin, "Mass preserving nonrigid registration of CT lung images using cubic B-spline," *Med. Phys.* **36**, 4213–4222 (2009).
- ²⁰V. Delmon, S. Rit, R. Pinho, and D. Sarrut, "Registration of sliding objects using direction dependent B-splines decomposition," *Phys. Med. Biol.* **58**, 1303–1314 (2013).
- ²¹C. T. Metz, S. Klein, M. Schaap, T. van Walsum, and W. J. Niessen, "Non-rigid registration of dynamic medical imaging data using nD+t B-splines and a groupwise optimization approach," *Med. Image Anal.* **15**, 238–249 (2011).
- ²²Z. Xie and G. Farin, "Image registration using hierarchical B-splines," *IEEE Trans. Visual. Comput. Graphics* **10**, 1–10 (2004).
- ²³S. Kabus, T. Netsch, B. Fischer, and J. Modersitzki, "B-spline registration of 3D images with Levenberg-Marquardt optimization," *Proc. SPIE* **5370**, 304–313 (2004).
- ²⁴J. Nocedal and S. Wright, *Numerical Optimization* (Springer, Berlin, 1999).
- ²⁵A. Bruhn, J. Weickert, and C. Schnorr, "Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods," *Int. J. Comput. Vis.* **61**, 211–231 (2005).
- ²⁶R. Castillo, E. Castillo, D. Fuentes, M. Ahmad, A. Wood, M. 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).
- ²⁷E. A. Regan, J. E. Hokanson, J. R. Murphy, B. Make, D. A. Lynch, T. H. Beaty, D. Curran-Everett, E. K. Silverman, and J. D. Crapo, "Genetic epidemiology of COPD (COPDGene) study design," *Int. J. Chronic Obstruct. Pulm. Dis.* **7**, 32–43 (2010).
- ²⁸K. Murphy, B. Van Ginneken, J. Reinhardt, S. Kabus, K. Ding, X. Deng, K. Cao, K. Du, G. Christensen, and V. Garcia, "Evaluation of registration methods on thoracic CT: The Empire10 challenge," *IEEE Trans. Med. Imaging* **30**, 1901–1920 (2011).
- ²⁹P. J. Rousseeuw, "Least median of squares regression," *J. Am. Stat. Assoc.* **79**, 871–880 (1984).
- ³⁰E. J. Cands, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.* **59**, 1207–1223 (2006).
- ³¹D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory* **52**, 1289–1306 (2006).
- ³²W. Yin, S. Osher, D. Goldfarb, and J. Darbon, "Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing," *SIAM J. Imaging Sci.* **1**, 143–168 (2008).
- ³³D. Levin, "The approximation power of moving least squares," *Math. Comput.* **67**, 1517–1531 (1998).
- ³⁴J. Yang and Y. Zhang, "Alternating direction algorithms for ℓ_1 -problems in compressive sensing," *SIAM J. Sci. Comput.* **33**, 250–278 (2011).
- ³⁵J. Sanders and E. Kandrot, *CUDA By Example: An Introduction to General-Purpose GPU Programming* (Addison-Wesley, Boston, MA, 2011).
- ³⁶Y. Lipman, "Stable moving least-squares," *J. Approx. Theory* **161**, 371–384 (2009).
- ³⁷P. Breitkopf, A. Rassineux, G. Touzot, and P. Villon, "Explicit form and efficient computation of MLS shape functions and their derivatives," *Int. J. Num. Methods Eng.* **48**, 451–466 (2000).
- ³⁸D. Cline, S. Jeschke, K. White, A. Razdan, and P. Wonka, "Dart throwing on surfaces," *Comput. Graphics Forum* **28**, 1217–1226 (2009).
- ³⁹R. Castillo, E. Castillo, R. Guerra, V. Johnson, T. McPhail, A. Garg, and T. Guerrero, "A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets," *Phys. Med. Biol.* **54**, 1849–1870 (2009).
- ⁴⁰B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, and J. C. Gee, "A reproducible evaluation of ANTs similarity metric performance in brain image registration," *Neuroimage* **54**, 2033–2044 (2011).
- ⁴¹B. B. Avants, N. Tustison, and G. Song, "Advanced normalization tools (ANTs)," *Insight J.* (2009), see <http://hdl.handle.net/10380/3113>.
- ⁴²G. Song, N. Tustison, B. Avants, and J. C. Gee, "Lung CT image registration using diffeomorphic transformation models," in *Proceedings of Grand Challenge on Medical Image Analysis for the Clinic* (CreateSpace, 2010), pp. 23–32.
- ⁴³U. Malsch, C. Thieke, P. E. Huber, and R. Bendl, "An enhanced block matching algorithm for fast elastic registration in adaptive radiotherapy," *Phys. Med. Biol.* **51**, 4789–4806 (2006).
- ⁴⁴M. Sohn, M. Birkner, Y. Chi, J. Wang, D. Yan, B. Berger, and M. Alber, "Model-independent, multimodality deformable image registration by local matching of anatomical features and minimization of elastic energy," *Med. Phys.* **35**, 866–878 (2008).
- ⁴⁵R. Franke and G. Nielson, "Smooth interpolation of large sets of scattered data," *Int. J. Num. Methods Eng.* **15**, 1691–1704 (1980).
- ⁴⁶H. Wendland, "Local polynomial reproduction and moving least squares approximation," *IMA J. Num. Anal.* **21**, 285–300 (2001).
- ⁴⁷G. Christensen and H. Johnson, "Consistent image registration," *IEEE Trans. Med. Imaging* **20**, 568–582 (2001).
- ⁴⁸F. Massanes, J. Brankov, and M. Cadennes, "Compute-unified device architecture implementation of a block-matching algorithm for multiple graphical processing unit cards," *J. Electron. Imaging* **20**, 033004–1–033004–10 (2011).
- ⁴⁹S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.* **3**, 1–122 (2010).
- ⁵⁰J. F. C. Mota, J. Xavier, P. M. Q. Aguiar, and M. Puschel, "Distributed basis pursuit," *IEEE Trans. Signal Process.* **60**, 1942–1956 (2012).
- ⁵¹V. Gorbunova, P. Lo, H. Ashraf, A. Dirksen, M. Nielsen, and M. Bruijine, "Weight preserving image registration for monitoring disease progression in lung CT," in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2008* (Springer, Berlin, 2008).
- ⁵²C. Galbn, M. Han, J. Boes, K. Chughtai, C. Meyer, T. Johnson, S. Galbn, A. Rehmulla, E. Kazerooni, F. Martinez, and B. Ross, "Computed tomography based biomarker provides unique signature for diagnosis of COPD phenotypes and disease progression," *Nat. Med.* **18**, 1711–1715 (2012).