Diffeomorphic Density Registration in Thoracic Computed Tomography

Caleb Rottman[1(✉)], Ben Larson[1], Pouya Sabouri[2], Amit Sawant[2],
and Sarang Joshi[1]

[1] Scientific Computing and Imaging Institute, Department of Bioengineering,
University of Utah, Salt Lake City, USA
crottman@sci.utah.edu
[2] University of Maryland School of Medicine, Baltimore, MD, USA

Abstract. Accurate motion estimation in thoracic computed tomography (CT) plays a crucial role in the diagnosis and treatment planning of lung cancer. This paper provides two key contributions to this motion estimation. First, we show we can effectively transform a CT image of effective linear attenuation coefficients to act as a density, i.e. exhibiting conservation of mass while undergoing a deformation. Second, we propose a method for diffeomorphic density registration for thoracic CT images. This algorithm uses the appropriate density action of the diffeomorphism group while offering a weighted penalty on local tissue compressibility. This algorithm appropriately models highly compressible areas of the body (such as the lungs) and incompressible areas (such as surrounding soft tissue and bones).

Keywords: Diffeomorphisms · Thoracic motion estimation · Density action · Image registration

1 Introduction

According to the Centers for Disease Control and Prevention, lung cancer is the leading cause of cancer death, accounting for 27 % of all cancer deaths in the United States [1]. Accurate modeling of the motion and biomechanics of the lungs under respiration is essential for the diagnosis and treatment of this disease. In particular, accurate estimation of organ movement and deformations plays a crucial role in dose calculations and treatment decisions in radiation therapy of lung cancer [8,11].

Essential to our method of motion estimation is that CT images act similar to densities: i.e. they exhibit some conservation of mass properties while undergoing deformations. The relationship between CT images and densities can be clearly seen by viewing a single patient thoracic CT throughout the breathing cycle. During inhalation, lung volume increases and lung CT intensities decrease,

and during exhalation, lung volume decreases and lung CT intensities increase. Because of these changing image intensities, the L^2 image action of a diffeomorphism does not accurately reflect motion in CT imaging. Most state-of-the-art methods deal with these changing intensities by not using the L^2 metric between images but to instead use either mutual information or normalized cross correlation [2]. We incorporate these intensity changes into our deformation model by treating these images as densities. Some mass-preserving registration methods using cubic B-splines have also been introduced [9,12]. In contrast, we use the full space of diffeomorphisms equipped with an H^1 metric, and use the recently discovered link between densities and diffeomorphisms [4]. Furthermore, we show experimentally that CT images are not inherently mass preserving and must be transformed to become mass preserving.

We will first introduce the mathematical definition of a density and then describe how it relates to material density and CT images. Mathematically, a 3D density $I(x)\,dx$ is a volume form on a domain $\Omega \subseteq \mathbb{R}^3$ where $I(x)$ is a non-negative function on Ω and $dx = dx^1 \wedge dx^2 \wedge dx^3$ is the standard volume element on $\mathbb{R}^3$. The key difference between a density and a function is how a diffeomorphism $\varphi \in \text{Diff}(\Omega)$ acts on them. The left action of φ on a *function* $g(x)$ (called the L^2 action) is simply function composition:

$$\varphi_* g(x) = g \circ \varphi^{-1}(x). \tag{1}$$

The left action of φ on a *density* $I(x)\,dx$ is:

$$\varphi_*(I(x)\,dx) = I \circ \varphi^{-1}(x)|D\varphi^{-1}(x)|\,dx, \tag{2}$$

where $|D\varphi^{-1}|$ is the Jacobian determinant of the diffeomorphism.

A unique property of a density is that the total mass is conserved under the action of a diffeomorphism, where here the total mass is defined as the integral of the density over Ω:

$$\int_\Omega I \circ \varphi^{-1}(x)|D\varphi^{-1}(x)|\,dx = \int_\Omega I(y)\,dy. \tag{3}$$

This equality holds by performing a change of variables: $x = \varphi(y)$, $dx = |D\varphi(y)|dy$, and using the identity $|D\varphi^{-1}(x)| = \frac{1}{|D\varphi(y)|}$.

This conservation of mass property extends to its traditional meaning in a physical mass density $\rho(x)$ (units g/cm^3). Physical mass density integrated over a domain becomes physical mass (units g). Similarly, the narrow beam X-ray linear attenuation coefficient (LAC) for a single material (units cm^{-1}) is defined as $\mu(x) = m\rho(x)$, where m is a material-specific property called the mass attenuation coefficient (units cm^2/g) that depends on the energy of the X-ray beam. In a mixture of materials, the total linear attenuation coefficient is $\mu(x) = \sum_i m_i \rho_i(x)$. Integrating $\mu(x)$ over a domain gives us the total LAC, which we will call the LAC mass (units cm^2). Therefore, conservation of physical mass implies conservation of LAC mass in a closed system.

During respiration, we assume that the change in physical lung mass due to air in the lungs is negligible. We then expect conservation of LAC mass in the lungs.

2   CT Images as Densities

We have shown that LAC mass is theoretically conserved. Unfortunately, CT image intensities do not represent true narrow beam linear attenuation coefficients. Instead, modern CT scanners use wide beams that yield secondary photon effects at the detector. CT image intensities reflect *effective* linear attenuation coefficients as opposed to the true narrow beam linear attenuation coefficient.

To see the relationship between effective LAC and true narrow beam LAC, we ran a Monte Carlo simulation using an X-ray spectrum and geometry from a Philips CT scanner at various densities of water (since lung tissue is very similar to a mixture between water and air) [6]. The nonlinear relationship between effective LAC and narrow beam LAC relationship is clear (see Fig. 1).

If we have conservation of mass within a single subject in a closed system, we expect an inverse relationship between average density in a region Ω and volume of that region: $D_t = \frac{M}{V_t}$. Here $V_t = \int_{\Omega_t} 1dx$, $D_t = \int_{\Omega_t} I_t(x)dx/V_t$, Ω_t is the domain of the closed system (that moves over time), and t is a phase of the breathing cycle. This relationship becomes linear in log space with a slope of -1:

$$\ln(D_t) = \ln(M) - \ln(V_t) \quad (4)$$

Fig. 1. Effective LAC from Monte Carlo simulation (solid line) and NIST reference narrow beam LAC (dashed line). The true relationship between effective LAC and narrow beam LAC is nonlinear.

Our experimental results confirm the Monte Carlo simulation in that lungs imaged under CT do not follow this inverse relationship. Rather, the slope found in these datasets in log space is consistently greater than -1 (see Fig. 3). Because of this, we seek a nonlinear intensity transformation for CT images such that CT mass is preserved under deformation.

In this paper, we model this intensity transformation as an exponential function, i.e. $I(x) \mapsto I(x)^\alpha$, and we solve for the α that yields the best mass preservation. We chose an exponential model because it is a single parameter monotonic function that preserves the density of air at zero. Furthermore, using an exponential function makes our analysis invariant to image scaling. That is, if $I_t(x)^\alpha$ exhibits conservation of mass, so does $(cI_t(x))^\alpha$, for $c \in \mathbb{R}^+$.

Note that the field of view of the CT scanner is not a closed system, as portions of the body leave and enter the field of view during respiration. We therefore evaluate the accuracy of this intensity transformation inside the lungs, which is essentially a closed system. We evaluate our methods using the Deformable
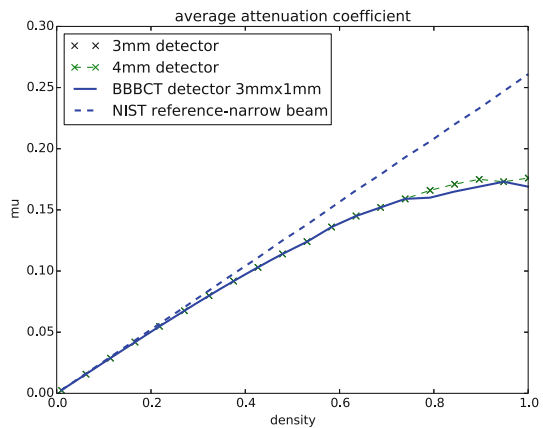
Image Registration (DIR) Laboratory dataset (http://www.dir-lab.com/) [7], which consists of ten subjects with ten 4DCT timepoints each. Second, we evaluate on a set of 30 subjects with ten 4DCT timepoints each procured at UT Southwestern Medical Center. The lungs for each patient at each timepoint are segmented with active contours using ITK-SNAP [13] (http://www.itksnap.org/) combined with an intensity based segmentation to remove high-density regions in the lungs and around the lung border due to imperfect initial segmentations.

For each subject, we perform a linear regression of the measured LAC density and calculated volume in log space. Let $\mathbf{d}(\alpha) = \log\left(\int_{\Omega_t} I_t(x)^\alpha dx / \int_{\Omega_t} 1 dx\right)$ (the log density) and $\boldsymbol{v} = \log(\int_{\Omega_t} 1 dx)$ (the log volume), where again t is a breathing cycle timepoint. The linear regression then models the relationship in log space as $\mathbf{d}(\alpha) \approx a\boldsymbol{v} + b$. Let $a_j(\alpha)$ be the slope solved for in this linear regression for the j^{th} subject. To find the optimal α for the entire dataset, we solve

$$\alpha = \arg\min_{\alpha'} \sum_j (a_j(\alpha') + 1)^2, \tag{5}$$

which finds the value of α that gives us an average slope closest to -1. We solve for α using a brute force search.

Applying this exponential function to the CT data allows us to perform our density matching algorithm described in the next section.

3   Weighted Diffeomorphic Density Matching

Mathematically, our problem is to find a diffeomorphic (bijective and smooth) transformation between two densities I_0 and I_1, using our exponentially transformed CT images defined in the previous section as our densities. We use the Fisher-Rao metric on densities which has the unique property that it is the only metric between densities that is invariant to the action of a diffeomorphism [3]. When vol(Ω) is infinite, the Fisher-Rao metric between two densities becomes the Hellinger distance:

$$d_F^2(I_0 dx, I_1 dx) = \int_\Omega (\sqrt{I_0} - \sqrt{I_1})^2 dx. \tag{6}$$

The Riemannian geometry of the diffeomorphism group with a Sobolev H^1 metric is intimately linked to the geometry of the space of densities with the Fisher-Rao metric. In particular, there are Sobolev H^1 metrics on the diffeomorphism group that descend to the Fisher-Rao metric on the space of densities [4]. This descending property from Diff(Ω) to Dens(Ω) allows us to compute the distance on Diff(Ω) by using the Fisher-Rao metric. Since the space of densities is flat, we can solve for the distance in Diff(Ω) in closed form: we do not need to time-integrate velocity fields or solve for adjoint equations as is necessary in LDDMM [5].

We therefore seek to minimize the following energy functional:

$$E(\varphi) = d_F^2(\varphi_*(f\,dx),(f\circ\varphi^{-1})dx) + d_F^2(\varphi_*(I_0\,dx),I_1\,dx)) \qquad (7)$$

$$= \underbrace{\int_\Omega (\sqrt{|D\varphi^{-1}|}-1)^2\,f\circ\varphi^{-1}\,dx}_{E_1(\varphi)} + \underbrace{\int_\Omega \left(\sqrt{|D\varphi^{-1}|I_0\circ\varphi^{-1}}-\sqrt{I_1}\right)^2 dx}_{E_2(\varphi)}\,.$$

$$(8)$$

To better understand this energy functional, we describe its two terms. The first term $E_1(\varphi)$ is the metric on the regularity of the deformation, descended from the Sobolev H^1 metric. This penalizes φ as it becomes non volume preserving: a unitary Jacobian determinant at a location indicates that the transformation is volume preserving. Furthermore, the density $f(x)\,dx$ is a positive weighting on the domain Ω: regions where $f(x)$ is high have a higher penalty on non-volume preserving deformations and regions where $f(x)$ is low have a lower penalty on non-volume preserving deformations.

Physiologically, we know the lungs are quite compressible as air enters and leaves. Surrounding tissue including bones and soft tissue, on the other hand, is essentially incompressible. Therefore, our penalty function $f(x)$ is low inside the lungs and outside the body and high elsewhere. For our penalty function, we simply implement a sigmoid function of the original CT image: $f(x) = \text{sig}(I_0(x))$.

The second term $E_2(\varphi)$ is the Fisher-Rao distance between the deformed density and the target density.

We take the Sobolev gradient with respect to the energy functional which is given by

$$\delta E = -\Delta^{-1}\Big(-\nabla\big(f\circ\varphi^{-1}(1-\sqrt{|D\varphi^{-1}|})\big)$$
$$-\sqrt{|D\varphi^{-1}|\,I_0\circ\varphi^{-1}}\nabla(\sqrt{I_1}) + \nabla\big(\sqrt{|D\varphi^{-1}|\,I_0\circ\varphi^{-1}}\big)\sqrt{I_1}\Big). \quad (9)$$

Then, the current estimate of φ^{-1} is updated directly via a Euler integration of the gradient flow [10]:

$$\varphi_{j+1}^{-1}(x) = \varphi_j^{-1}(x+\epsilon\delta E) \qquad (10)$$

for some step size ϵ. Since we take the Sobolev gradient the resulting deformation is guaranteed to be invertable with a sufficiently small ϵ.

4   Results

For the DIR dataset, we used the method from Sect. 2 to solve for the exponent that yields conservation of mass. We solved for $\alpha = 1.64$ that gives us the best fit. Without using the exponential fit, the average slope of log density log volume plot was -0.66 (SD 0.048). After applying the exponential to the CT intensities, the average slope is -1.0 (SD 0.054). The log-log plots of all ten patients in the DIR dataset as well as box plots of the slope is shown in Fig. 2.
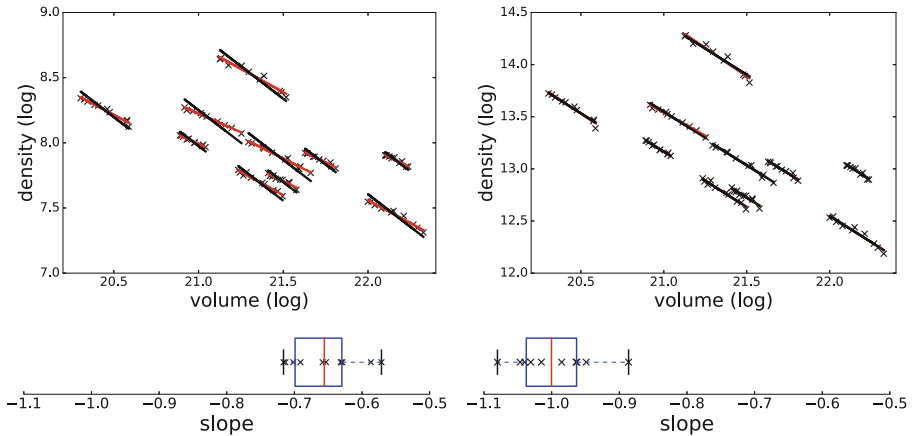
Fig. 2. Density and volume log-log plots. Upper left: log-log plots without applying the exponential correction for all ten DIR subjects. The best fit line to each dataset is in red and the mass-preserving line (slope $= -1$) is in black. Upper right: log-log plots after applying the exponential correction $I(x)^{\alpha}$ to the CT images. In this plot, the best fit line matches very closely to the mass-preserving line. Bottom row: corresponding box plots of the slopes found in the regression. (Color figure online)

For the 30 subject dataset, we solved for $\alpha = 1.90$ that gives us conservation of mass. Without using the exponential fit, the average slope of the log-log plot was -0.59 (SD 0.11).

We applied our proposed weighted density matching algorithm to the first subject from the DIR dataset. This subject has images at 10 timepoints and has a set of 300 corresponding landmarks between the full inhale image and the full exhale image. These landmarks were manually chosen by three independent observers. Without any deformation, the landmark error is 4.01 mm (SD 2.91 mm). Using our method, the landmark error is reduced to 0.88 mm (SD 0.94 mm), which is only slightly higher than the observer repeat registration error of 0.85 mm (SD 1.24 mm).

We implement our algorithm on the GPU and plot the energy as well as the Fisher-Rao metric with and without applying the deformation. These results are shown in Fig. 3. In this figure, we show that we have excellent data match, while the deformation remains physiologically realistic: inside the lungs there is substantial volume change due to respiration, but the deformation outside the lungs is volume preserving. With a $256 \times 256 \times 94$ voxel dataset, our algorithm takes approximately nine minutes running for four thousand iterations on a single nVidia Titan Z GPU.
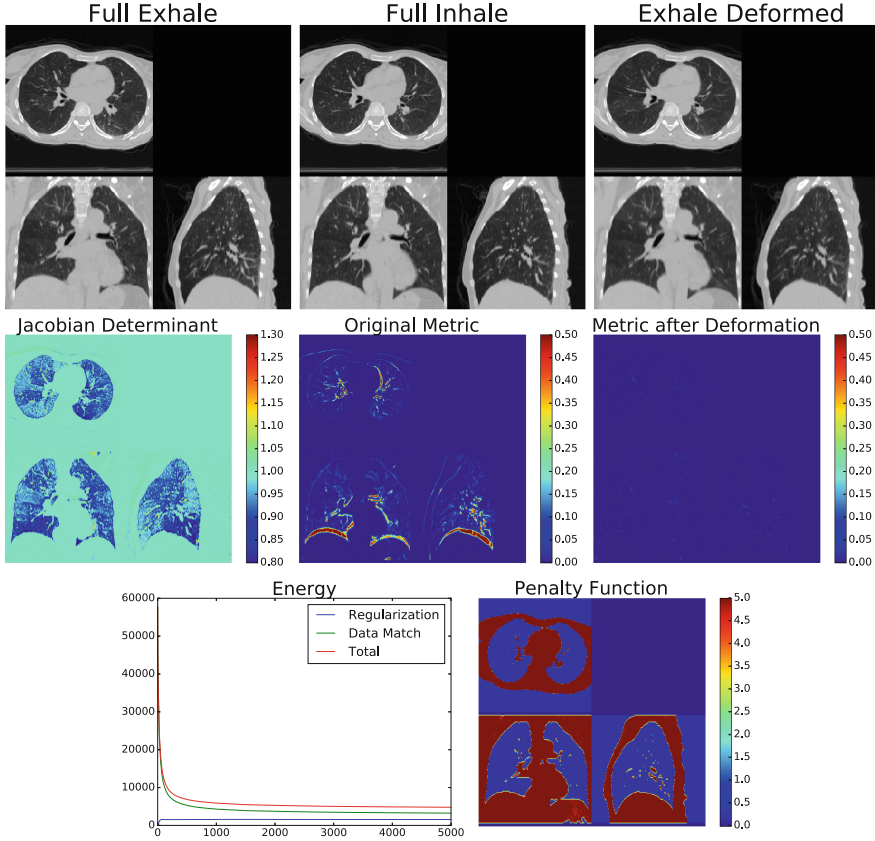
Fig. 3. Registration results. Top row: full inhale, full exhale, and the deformed exhale density estimated using our method. Middle row: Jacobian determinant of the transformation, initial Fisher-Rao metric, and Fisher-Rao metric after applying the density action. Notice that outside the lungs the estimated deformation is volume preserving. Bottom row: Energy as a function of iterations, and penalty function.

5   Discussion

In this paper, we have shown that although the narrow beam linear attenuation acts as a density, the effective linear attenuation coefficient found in CT does not act as a density. However, applying a simple exponential function transforms the CT dataset into a set of images that exhibit conservation of mass. This simple non-linear approximation yields excellent results even when using the same exponential function for multiple subjects in a single dataset.

We suspect that the biggest cause of the nonlinearity between true linear attenuation and effective attenuation is the presence of X-ray scatter and secondary photons, which are dependent on the scanner geometry and the energy spectrum. Therefore, we do not necessarily expect that the same α parameter of

the exponential functions works across different CT scanners; however we found that the same α parameter accurately corrects the nonlinearity across multiple subjects on the same scanner.

We have also shown that we can use these corrected images as densities with a great deal of accurately. Our method uses the appropriate density action when dealing with CT images, and our weighting function on the domain constrains the deformation to be physiologically realistic.

References

1. CDC - basic information about lung cancer. http://www.cdc.gov/cancer/lung/basic_info/index.htm. Accessed 06 Mar 2016
2. Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C.: A reproducible evaluation of ANTs similarity metric performance in brain image registration. NeuroImage **54**(3), 2033–2044 (2011)
3. Bauer, M., Bruveris, M., Michor, P.W.: Uniqueness of the Fisher-Rao metric on the space of smooth densities. submitted (2015)
4. Bauer, M., Joshi, S., Modin, K.: Diffeomorphic density matching by optimal information transport, pp. 1–35 (2015)
5. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. Int. J. Comput. Vis. **61**(2), 139–157 (2005)
6. Boone, J.M., Seibert, J.A.: An accurate method for computer-generating tungsten anode x-ray spectra from 30 to 140 kv. Med. phys. **24**(11), 1661–1670 (1997)
7. Castillo, R., Castillo, E., Guerra, R., Johnson, V.E., McPhail, T., Garg, A.K., Guerrero, T.: A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. Phys. Med. Biol. **54**(7), 1849–1870 (2009)
8. Geneser, S.E., Hinkle, J.D., Kirby, R.M., Wang, B., Salter, B., Joshi, S.: Quantifying variability in radiation dose due to respiratory-induced tumor motion. Med. Image Anal. **15**(4), 640–649 (2011)
9. Gorbunova, V., Sporring, J., Lo, P., Loeve, M., Tiddens, H.A., Nielsen, M., Dirksen, A., de Bruijne, M.: Mass preserving image registration for lung CT. Med. Image Anal. **16**(4), 786–795 (2012)
10. Rottman, C., Bauer, M., Modin, K., Joshi, S.C.: Weighted diffeomorphic density matching with applications to thoracic image registration. In: 5th MICCAI Workshop on Mathematical Foundations of Computational Anatomy (MFCA 2015), pp. 1–12 (2015)
11. Sawant, A., Keall, P., Pauly, K.B., Alley, M., Vasanawala, S., Loo, B.W., Hinkle, J., Joshi, S.: Investigating the feasibility of rapid MRI for image-guided motion management in lung cancer radiotherapy. BioMed. Res. Int. 2014 (2014)
12. Yin, Y., Hoffman, E.A., Lin, C.L.: Mass preserving nonrigid registration of CT lung images using cubic B-spline. Med. Phys. **36**(9), 4213–4222 (2009)
13. Yushkevich, P.A., Piven, J., Hazlett, H.C., Smith, R.G., Ho, S., Gee, J.C., Gerig, G.: User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. Neuroimage **31**(3), 1116–1128 (2006)