Quantifying Tensor Field Similarity with Global Distributions and Optimal Transport

Arnold D. Gomez[1(✉)], Maureen L. Stone[2], Philip V. Bayly[3], and Jerry L. Prince[1]

[1] Electrical and Computer Engineering Department, Jonhs Hopkins University, Baltimore, USA
adgomez@jhu.edu

[2] Department of Neural and Pain Sciences, University of Maryland, Baltimore, USA

[3] Mechanical Engineering Department, Washington University in St. Louis, St. Louis, USA

Abstract. Strain tensor fields quantify tissue deformation and are important for functional analysis of moving organs such as the heart and the tongue. Strain data can be readily obtained using medical imaging. However, quantification of similarity between different data sets is difficult. Strain patterns vary in space and time, and are inherently multidimensional. Also, the same type of mechanical deformation can be applied to different shapes; hence, automatic quantification of similarity should be unaffected by the geometry of the objects being deformed. In the pattern recognition literature, shapes and vector fields have been classified via global distributions. This study uses a distribution of mechanical properties (a 3D histogram), and the Wasserstein distance from optimal transport theory is used to measure histogram similarity. To evaluate the method's consistency in matching deformations across different objects, the proposed approach was used to sort strain fields according to their similarity. Performance was compared to sorting via maximum shear distribution (a 1D histogram) and tensor residual magnitude in perfectly registered objects. The technique was also applied to correlate muscle activation to muscular contraction observed via tagged MRI. The results show that the proposed approach accurately matches deformation regardless of the shape of the object being deformed. Sorting accuracy surpassed 1D shear distribution and was on par with residual magnitude, but without the need for registration between objects.

Keywords: Strain · Tensor fields · Tagged MRI · Organ deformation

1 Introduction

Tissue deformation is necessary for vital bodily functions such as cardiac blood pressurization [1], locomotion via the musculoskeletal system [2], and motion in

A. F. Frangi et al. (Eds.): MICCAI 2018, LNCS 11071, pp. 428–436, 2018.
https://doi.org/10.1007/978-3-030-00934-2_48

the tongue (involved in breathing, swallowing, and speech) [3]. The strain tensor quantifies deformation, and can produce biomarkers related to organ function [1,3]. For this reason, many acquisition and processing methods have been developed for using medical imaging to observe motion (as in Fig. 1) and to measure strain. Strain can be reduced to global metrics to simplify clinical or research applications, e.g., radial or circumferential strains in cardiac deformation [1], or fiber strains for analysis of tongue motion [3]. However, the state of deformation across an organ can only be truly captured by a volumetric strain field.

Statistical analysis of mechanical deformation is relevant when studying disease progression and functional mechanisms, particularly with computational models aimed at predicting deformation, such as in traumatic brain injury research [4], speech analysis [5], the study of muscular contraction [2], and for assessing of cardiac contractility [6]. Ideally, these models must agree with experimental data; hence, similarity quantification between tensor fields is useful in: (i) determining the uncertainty of experimental measurements, (ii) measuring the descriptive or predictive power of a simulation, and (iii) assessing the sensitivity of strain results to processing or modeling parameters [7,8]. Strain comparisons are also useful to align data for atlas generation [9].

Similarity measurement between strain fields is challenging [8]. As with scalar fields (e.g., images), the concept of similarity is diverse, encompassing not only correspondence and alignment, but also directionality, magnitude, and orientation [8]. Mechanical deformation also changes in time and across patients, and is defined in many different discretization schemes. For instance, image measurements generally result in pixel wise values [1] while computational models often use finite elements (FE) [10] or (more recently) point clouds [4]. This variety of configurations and formats can make it difficult to compare strain fields by means of evaluating field properties in matching comparison points. (Fields can be defined across different geometric domains, which can have topological or discretization incompatibilities, limiting or even preventing registration.)

This paper describes a novel approach to measure similarity between strain tensor fields. The methodology was developed to overcome the need to register domains, while capturing the spatial distribution of strain within the organ.

2 Background

Assuming that tissue behaves as a continuum for macroscopic deformation, a particle in the reference configuration $\boldsymbol{X} = \boldsymbol{X}(t_o) \in \mathbb{R}^3$ at time t_0 is associated to a single point in the deformed configuration $\boldsymbol{x}(t)$ via a Lagrangian displacement field $\boldsymbol{u}(\boldsymbol{X}, t) = \boldsymbol{X} - \boldsymbol{x}(\boldsymbol{X}, t)$. Displacement information is obtainable via medical imaging e.g., speckle tracking in ultrasound, registration of CINE MRI, or tagged MRI [1,11], and biomechanics simulations [10]. Relative displacement between neighboring particles is related to the deformation gradient tensor

$$\boldsymbol{F} = \boldsymbol{F}(\boldsymbol{X}, \boldsymbol{x}, t) = \frac{\partial \boldsymbol{x}(t)}{\partial \boldsymbol{X}(t_0)} = \mathbf{1} + \frac{\partial \boldsymbol{u}(t)}{\partial \boldsymbol{X}(t_0)}, \tag{1}$$

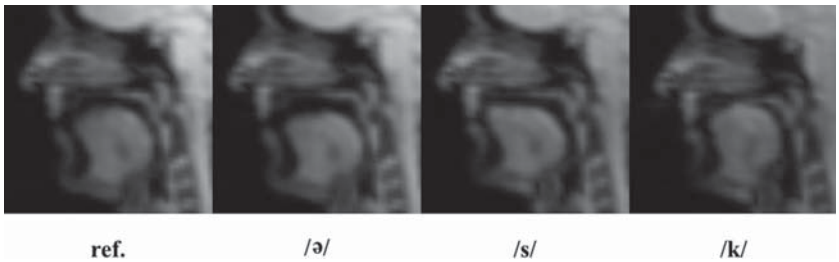

Fig. 1. Sequence from magnetic resonance imaging (MRI) during speech generation (sagittal). Deformation can be seen by comparing the shape of the tongue at the beginning of the sound /ə/ (/ə/ means "a" as in "car"), which is the reference (ref.) versus the shapes during the sounds /ə/, /s/ (as in "sound"), and /k/ (as in "kiosk").

where $\mathbf{1}$ is the identity tensor [12]. $\boldsymbol{F}$ differs from $\mathbf{1}$ in rigid motion and rotation; thus, the direct application of $\boldsymbol{F}$ in quantifying deformation is limited. Instead, we can define strain tensors insensitive to rigid transformations. These include the (left) Cauchy-Green strain tensor

$$\boldsymbol{C} = \boldsymbol{C}(\boldsymbol{X}, \boldsymbol{x}, t) = \boldsymbol{F}^T \boldsymbol{F}, \tag{2}$$

and the Green-Lagrange strain tensor $\boldsymbol{E} = \frac{1}{2}(\boldsymbol{C} - \mathbf{1})$ [12]. Eigenvalue decomposition of these quantities yields principal components along key directions such as the radial and circumferential strains used in cardiac mechanics [1].

Methods for quantification of similarity between strain fields are varied. One common approach is to apply point-wise norms or patch correlations [8]. While this is an effective similarity measure, it requires matching points or regions between two or more geometries. If registration is needed, it can introduce computational expense and (or) registration error. Another common approach is to compare the distributions of scalar metrics [4,8], which does not require registration and is effective in cases where similarity is determined largely by magnitude. However, traditional global distributions are spatially underdetermined resulting in loss of structural information.

In pattern recognition, structural distributions have emerged as a fast and effective means of identifying objects [13,14], and these have also been extended to vector fields [15]. The main idea behind these techniques is the construction of a *structurally sensitive* distribution (a shape or vector field signature), which is populated by sampling properties between points inside each domain. Examples include 1D histograms constructed from the distance between two random points within an object [13], and 2D histograms populated by sampling magnitude and angle differences along streamlines in a vector field [15].

3 Quantifying Tensor Field Similarity

Based on previous work in pattern recognition [13–15], we propose the construction of a structurally sensitive global distribution for each strain field. This

distribution acts as a *strain signature* (or a set of features), enabling similarity quantification. In the context of pattern identification, a metric of similarity is derived from a normalized measure of distance between distributions.

The properties $\boldsymbol{f}$ for the construction of features are obtained from points with randomly assigned coordinates $\boldsymbol{R} \in \mathcal{D} \in \mathbb{R}^3$, where $\mathcal{D}$ is the domain in which the strain field is defined. At each point,

$$\boldsymbol{f}(\boldsymbol{R}) = \begin{bmatrix} \gamma_{\max} \\ \theta \\ \bar{\gamma} \end{bmatrix}. \tag{3}$$

The quantity $\gamma_{\max} = \gamma_{\max}(\boldsymbol{R}) = E_1 - E_3$ (where E_1 and E_3 represent the first and third eigenvalues of $\boldsymbol{E}$) is the maximum shear strain. Since $\gamma_{\max}$ is directly evaluated at $\boldsymbol{R}$, it provides a *global* component to the distribution. The other two properties are associated with $\boldsymbol{R}$ and $\boldsymbol{P} \in \mathcal{D}$—another point chosen at random. The local state of deformation at these two locations is coupled via $\boldsymbol{F}_{RP} = \boldsymbol{F}(\boldsymbol{R})^{-1}\boldsymbol{F}(\boldsymbol{P})$, which is analogous to the distance between random internal points used to construct shape distributions [13]. $\boldsymbol{F}_{RP}$ represents a mechanical deformation; thus, θ is the net rotation from $\boldsymbol{F}_{RP}$ obtained via singular value decomposition. Finally, $\bar{\gamma} = \frac{1}{2}\Sigma(\boldsymbol{C}_{RP})_{ij}$ $(i \neq j)$ is the mean shear strain, which is extracted substituting $\boldsymbol{F}_{RP}$ into (2). Both θ and $\bar{\gamma}$ are relative properties, providing the *structural* component of the distribution. Shear and rotation values are used due to the incompressible nature of tissue, which favors isochoric deformation—this intrinsic feature affects similarity comparison indirectly.

To approximate the distribution of $\boldsymbol{f}(\boldsymbol{R}) \forall \boldsymbol{R} \in \mathcal{D}$, N samples $\boldsymbol{f}_i$ $(i = 1, 2, ..., N)$ are binned into a 3D histogram $\boldsymbol{p}_f$ with d bins. Sample histograms appear in Fig. 2, and show differences based on the deformation pattern, and not necessarily on the orientation and geometry of the object being deformed (Fig. 2(a) is different to Fig. 2(b) given the same object). There are also similarities that can be used to define distance (Fig. 2(a) is closer to Fig. 2(c)).

By edge normalization, any two histograms to be compared, say $\boldsymbol{p}_f$ and $\boldsymbol{p}_g$, are defined in a common simplex $\Sigma_d = \{x \in \mathbb{R}_+^d : x^T\mathbf{1}_d = 1\}$ [16]. ($\mathbf{1}_d$ denotes an array of ones.) We can thus write the transportation, or coupling, polytope $U(\boldsymbol{p}_f, \boldsymbol{p}_g) = \{Z \in \mathbb{R}_+^{d\times d} | Z\mathbf{1}_d = \boldsymbol{p}_f, Z^T\mathbf{1}_d = \boldsymbol{p}_g\}$, which allows posing an optimal transport (OT) problem as seen in some registration and shape similarity approaches [16–18]. Letting C be a cost matrix reflecting the l^2 distance between the centers of the histogram bins,

$$W_c(\boldsymbol{p}_f, \boldsymbol{p}_g) = \min_{Z\in U} \langle Z, C \rangle \tag{4}$$

is a discrete transport to measure the distance—or dissimilarity—between $\boldsymbol{p}_f$ and $\boldsymbol{p}_g$. (Angle brackets denote the Frobenius product.) Specifically, (4) defines the Wasserstein distance after obtaining an optimal transport $Z^\star$.

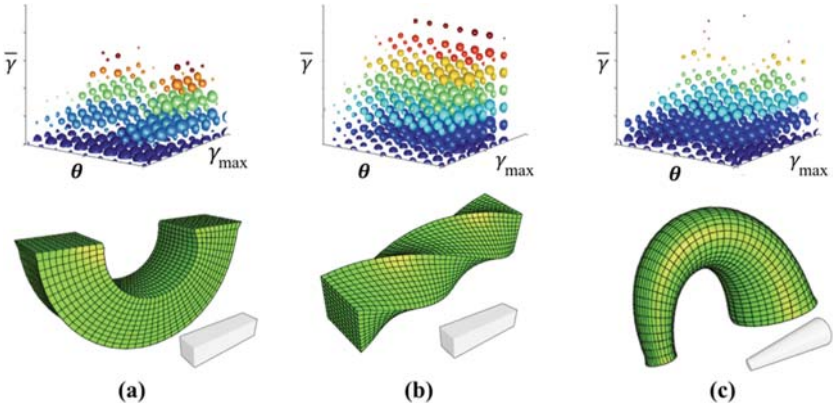

Fig. 2. Field property distributions. The top row shows 3D histograms, where the count scales sphere radii logarithmically, and color indicates height. The corresponding deformed shapes include: (a) a rod with square cross-section being bent upwards, (b) the same rod being twisted, and (c) a cylindrical rod being bent downwards. The undeformed shape appears below.

Computing $Z^\star$ for large multidimensional histograms is computationally prohibitive; thus, we implemented an approximation via the Sinkhorn-Knopp algorithm [16]. This approach uses an entropy constraint enforced by a Lagrangian multiplier λ, i.e., we find $W_c^\lambda(\boldsymbol{p}_f, \boldsymbol{p}_g) = \langle Z^\lambda, C \rangle$, where

$$Z^\lambda = \arg\min_{Z \in U} \langle Z, C \rangle - \frac{1}{\lambda} h(Z), \tag{5}$$

with $h(Z) = -\sum_{ij} Z_{ij}(\log(Z_{ij}) - 1)$. For the experiments in this study, we find that Z^λ is a reasonable approximation of $Z^\star$ for values of λ near 200, while being faster by nearly two orders of magnitude.

4 Experiments and Results

Sorting Test Strain Fields: Similarity-based sorting was used to determine accuracy and consistency against other methods. We created a dictionary of 27 simulated deformation fields containing bending, torsion, and extension (nine cases each). The fields were defined in geometries with different cross-section and discretization primitives using FE software [10]. A reference (or query) deformation pattern was matched to dictionary entries. Ideally, similar deformations will be the closest to the query, not only in top match, but the top $k-1$ matches—as there are $k = 9$ cases with similar deformations. Every field was used as query against all others to quantify the percentage of successes. For comparison, sorting was performed via: (i) 3D distributions using OT distance (3DOT), (ii) 3D distributions using the Euclidean distance (3DEU) [15], and (iii) 1D distributions of $\gamma_{\max}$ using Euclidean distance (1DEU) [4]. Given the random nature of

the sampling, the comparisons were performed 10 times using different seeds, and compared using unpaired t-test with significance level at $p = 0.01$.

The 3DOT approach was able to identify strain patterns regardless of geometry and discretization (see Fig. 3). Note that registration between the shapes—e.g., a square and a quarter-circular cross-section—would be difficult to perform. The proposed method was capable of identifying the top five matches with perfect accuracy, which reduced to 55% for the top eight matches. This is in stark contrast compared to 3DEU (accuracy dropped below 80% on the top three, and was 22% for the top eight), which demonstrates the utility of the OT approach. 1DEU was the least accurate likely due to structural information loss. Statistically significant performance differences between 3DOT and the alternatives were observed for $k = 3$ through $k = 8$. Sampling 10^4 points and calculation of W_c^λ took 90 s on a standard 12-core server.

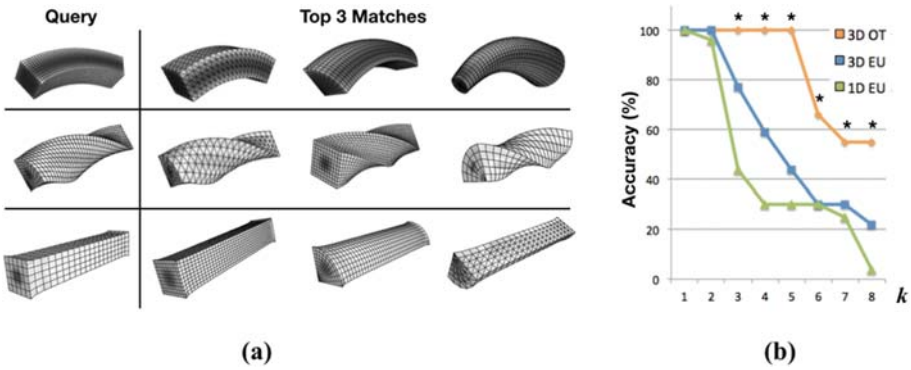

Fig. 3. Strain field sorting results. (a) Typical sorting results using the proposed algorithm. (b) Accuracy of the top k results across all simulations. A start denotes statistically significant differences between 3DOT and both of the alternatives.

Relating Image Measurements to Simulations: We used the proposed method to associate strain to muscular activation patterns in the tongue. Activation cannot be reliably measured but can be simulated; thus, it is possible to infer which muscles are active by matching experimental data to simulations. Tongue deformation during the sequence in Fig. 1 was measured via tagged MRI in healthy adult volunteers ($n = 7$, 4 male, 3 female, with informed consent) [5]. Strain fields were sampled and 3DOT was used to compare the experimental results to a dictionary composed of 74 FE simulations. The simulations included muscular activation and deformation as linear combinations of four categories: dorsiflexion (D), which moves the tip of the tongue upwards; protrusion (P), which moves the tip forward; retraction (R), which moves the tip backward; and ventriflexion (V), which lowers it. The experimental data was decomposed into percentages of each category (D, P, R, and V) and overall activation level (AL) based on the top three matches—AL = 0 indicates no activation. Also, the

dictionary entries (which are perfectly registered with one another) were sorted within themselves to compare the top matches via 3DOT, 1DEU, and the deformation gradient residual magnitude $||\boldsymbol{F}_i - \boldsymbol{F}_j||$ (sampling was also repeated 10 times).

The decomposition results appear in Fig. 4. Deformation from the reference configuration (early /ə/) to /ə/ was associated largely to retraction ($75 \pm 5\%$) and to a lesser extent to protrusion ($25 \pm 5\%$) with AL = 3.0. This low value of AL indicates a small, retraction-like deformation, as observed in Figs. 1 and 4. Deformation to /s/ was associated largely with protrusion ($64 \pm 10\%$), and less so with retraction and retroflexion ($15 \pm 21\%$ and $17 \pm 11\%$, respectively) with AL = 5.3. This is consistent with that observed forward motion of the tip of the tongue during this sound. During /k/ the observed deformation was associated mostly with retraction ($66 \pm 10\%$), and to a lesser extent with protrusion ($34 \pm 10\%$) with AL = 5. Compared to the reference, the tongue appears to retract during /k/ in the images. Both of the latter sounds result in larger AL values, as more activation is needed to achieve larger deformations. When comparing the top results of dictionary sorting, 3DOT agreed with residual magnitude in $56 \pm 2\%$ of the top results, compared to $22 \pm 3\%$ for 1DEU (statistically significant). Sampling 2×10^4 points and distance measurement took roughly 150 s.

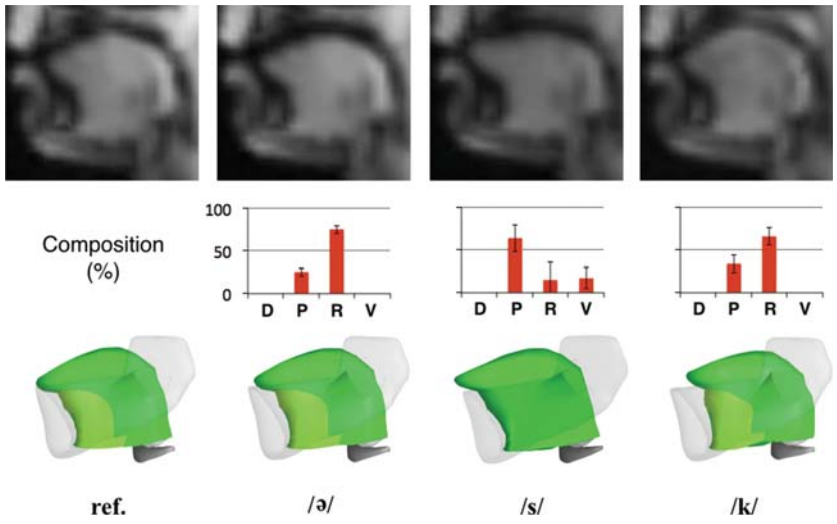

Fig. 4. Decomposition results. The close-up of a typical volunteer (top) is aligned with the decomposition shown as bar charts, and the models with the closest simulation.

5 Discussion

The proposed method is able to discriminate deformation patterns across different geometries and discretizations, its accuracy surpasses the alternatives stud-

ied, and it is also fast. Distance measurement via OT also shows clear benefits. The results from motion measurement via tagged MRI are consistent with visual inspection. The tongue shape varies considerably from subject to subject; thus, avoiding registration will simplify analysis of motion as larger datasets are produced. Sorting error was still observed likely indicating that the selection of features has a level of incompleteness via interdependence or data loss [8]. Also the distribution approach has no inverse mapping, which would be useful for machine learning. However, the methodology can be framed in different ways to address these limitations, e.g., via constructing distributions by sampling along hyperstreamlines, or applying OT methodology to tensor entries directly, as done in shape matching [17]. The proposed method remains a viable alternative given the seeding and computational requirements of these potential approaches.

References

1. Ibrahim, E.S.H.: Myocardial tagging by cardiovascular magnetic resonance: evolution of techniques-pulse sequences, analysis algorithms, and applications. J. Cardiovasc. Magn. Reson. **13**(1), 36–42 (2011)
2. Moerman, K.M., Sprengers, A.M.J., Simms, C.K., Lamerichs, R.M., Stoker, J., Nederveen, A.J.: Validation of tagged MRI for the measurement of dynamic 3D skeletal muscle tissue deformation. Med. Phys. **39**(4), 1793–1810 (2012)
3. Parthasarathy, V., Prince, J.L., Stone, M., Murano, E.Z., NessAiver, M.: Measuring tongue motion from tagged cine-MRI using harmonic phase (HARP) processing. J. Acoust. Soc. Am. **121**(1), 491–504 (2007)
4. Ganpule, S., et al.: A 3D computational human head model that captures live human brain dynamics. J. Neurotrauma **34**(13), 2154–2166 (2017)
5. Gomez, A.D., Xing, F., Chan, D., Pham, D.L., Bayly, P., Prince, J.L.: Motion estimation with finite-element biomechanical models and tracking constraints from tagged MRI. In: Wittek, A., Joldes, G., Nielsen, P.M.F., Doyle, B.J., Miller, K. (eds.) Computational Biomechanics for Medicine: From Algorithms to Models and Applications, pp. 81–90. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-54481-6_7
6. Wenk, J.F., et al.: Regional left ventricular myocardial contractility and stress in a finite element model of posterobasal myocardial infarction. J. Biomech. Eng. **133**(4), 1–14 (2011)
7. Henninger, H.B., Reese, S.P., Anderson, A.E., Weiss, J.A.: Validation of computational models in biomechanics. Proc. Inst. Mech. Eng. Part H **224**(7), 801–812 (2010)
8. Tian, Y., Nearing, G.S., Peters-Lidard, C.D., Harrison, K.W., Tang, L.: Performance metrics, error modeling, and uncertainty quantification. Mon. Weather Rev. **144**(2), 607–613 (2016)
9. Woo, J., Xing, F., Lee, J., Stone, M., Prince, J.L.: Construction of an unbiased spatio-temporal atlas of the tongue during speech. In: Ourselin, S., Alexander, D.C., Westin, C.-F., Cardoso, M.J. (eds.) IPMI 2015. LNCS, vol. 9123, pp. 723–732. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-19992-4_57
10. Maas, S.A., Ellis, B.J., Ateshian, G.A., Weiss, J.A.: FEBio: finite elements for biomechanics. J. Biomech. Eng. **134**(1), 011005 (2012)

11. Keszei, A.P., Berkels, B., Deserno, T.M.: Survey of non-rigid registration tools in medicine. J. Digit. Imaging **30**(1), 102–116 (2017)
12. Spencer, A.J.M.: Continuum Mechanics, 1995th edn. Dover Books, Essex (1985)
13. Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D.: Shape distributions. ACM Trans. Graph. **21**(4), 807–832 (2002)
14. Ohbuchi, R., Minamitani, T., Takei, T.: Shape-similarity search of 3D models by using enhanced shape functions. Int. J. Comput. Appl. Technol. **23**(2/3/4), 70–78 (2005)
15. Dinh, H.Q., Xu, L.: Measuring the similarity of vector fields using global distributions. In: da Vitoria, L.N. (ed.) SSPR/SPR 2008. LNCS, vol. 5342, pp. 187–196. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-89689-0_23
16. Cuturi, M.: Sinkhorn distances: lightspeed computation of optimal transportation distances. Adv. Neural Inf. Process. Syst. **26**, 2292–2299 (2013)
17. Su, Z., et al.: Optimal mass transport for shape matching and comparison. IEEE Trans. Pattern Anal. Mach. Intell. **37**(11), 2246–2259 (2015)
18. ur Rehman, T., Haber, E., Pryor, G., Melonakos, J., Tannenbaum, A.: 3D nonrigid registration via optimal mass transport on the GPU. Med. Image Anal. **13**(6), 931–940 (2009)