

Optimal Transport for Diffeomorphic Registration

Jean Feydy^{1,2}, Benjamin Charlier^{3,5}, François-Xavier Vialard^{4,6}(✉),
and Gabriel Peyré^{1,5}

¹ DMA – École Normale Supérieure, Paris, France
{jean.feydy,gabriel.peyre}@ens.fr

² CMLA – ENS Cachan, Cachan, France

³ Institut Montpellierain Alexander Grothendieck,
Univ. Montpellier, Montpellier, France
benjamin.charlier@umontpellier.fr

⁴ Univ. Paris-Dauphine - PSL Research, Paris, France
vialard@ceremade.dauphine.fr

⁵ CNRS, Paris, France

⁶ INRIA Mokaplan, Paris, France

Abstract. This paper introduces the use of unbalanced optimal transport methods as a similarity measure for diffeomorphic matching of imaging data. The similarity measure is a key object in diffeomorphic registration methods that, together with the regularization on the deformation, defines the optimal deformation. Most often, these similarity measures are local or non local but simple enough to be computationally fast. We build on recent theoretical and numerical advances in optimal transport to propose fast and global similarity measures that can be used on surfaces or volumetric imaging data. This new similarity measure is computed using a fast generalized Sinkhorn algorithm. We apply this new metric in the LDDMM framework on synthetic and real data, fibres bundles and surfaces and show that better matching results are obtained.

1 Introduction

State of the art methods in deformable registration of medical image data are often variational methods that estimate an optimal deformation by minimizing an energy which is the sum of two terms: a data fidelity term which quantifies how well the source data has been aligned with the target data; a regularization term on the deformation, which is necessary to make the registration problem well-posed, see [14] for a detailed overview. In the following, we will be interested in diffeomorphic image matching, resulting in a non-convex optimization problem and in estimation of local minima of the energy. The choice of the data fidelity is thus crucial to avoid poor local minima and to obtain meaningful registration.

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-66182-7_34](https://doi.org/10.1007/978-3-319-66182-7_34)) contains supplementary material, which is available to authorized users.

In this paper we propose a “global” fidelity making use of optimal transport theory and fast entropic approximation schemes.

Previous Works. *Data fidelity for registration.* Several similarity measures have been introduced in the literature which emphasizes its crucial role. For instance, for dense image registration, when only the intensity information is considered, the sum of squared differences (SSD) [8] or generalized versions of the cross-correlation such as normalized cross-correlation [2] are used which correct for possible intensity biases. When the images have already been segmented, a lot of interest has been devoted to matching between shape using parametrization invariant data fidelities. Let us mention currents, varifolds and more recently functional currents. All these shape metrics can be understood as a non-local SSD. Most often, the similarity measures are local or at most non-local in order to keep the computational cost low.

Unbalanced optimal transport. The basic theory of optimal transport (OT) (see [12]) defines a distance between probability distributions using the amount of effort needed to transport mass from one measure to another one. These distances are appealing for geometric problems such as shape registration because they are sensitive to spatial displacements of the shape. A recent breakthrough has been made in extending OT to the case of measures of different total mass [5, 10], which is crucial for practical applications where mass can vary because of changes in shape scales, or because of mass creation/destruction processes. The use of OT as a fidelity term is not new, see [11] for an example in machine learning. Quite surprisingly, to the best of our knowledge, it has never been used for diffeomorphic registration purpose.

Entropic regularization. A critical aspect of OT is that it was involved numerically. This situation has however radically changed in the last few years, thanks to the introduction of an efficient entropic approximation scheme [6], which is (i) efficient and easily parallelizable, (ii) and leads to a smooth differentiable data fidelity term [7]. This scheme also applies to unbalanced OT [4].

Contributions. This paper proposes a new non-local geometric similarity measure based on the recently developed theory of unbalanced OT and fast entropic solvers. The resulting hybrid pipeline is able to combine the strength of both diffeomorphic models presented in Sect. 2 and entropic unbalanced OT (convex optimization with order of magnitude faster solvers) detailed in Sect. 3. As shown by simulations on synthetic and real data in Sect. 4, OT can thus seamlessly be integrated in state-of-the-art registration methods, enabling a long-range non-local attraction term toward the target (which is crucial to avoid poor local minima) while remaining able to match intricate fine scale details. The code used to produce the figures of this article is available: <https://github.com/jeanfeydy/lddmm-ot>.

2 Diffeomorphic Registration

Data Representation with Measures. After a pre-processing step, it is often possible to efficiently represent medical image data as measures over a conveniently chosen space \mathbb{X} . This representation of data is essentially motivated by methodological and numerical purposes to help the design of fidelity terms.

For registration of segmented shapes (curves, surfaces), we advocate the use of a lifted features space $\mathbb{X} = \mathbb{R}^d \times \mathbb{S}^{d-1}$: each segment (*respectively* triangle) of the curve (*resp.* surface) can be represented as a Dirac of mass p_i equal to its length (*resp.* area), placed at location $(a_i, u_i) \in \mathbb{X}$, where a_i is the center and u_i the unit-length normal of the shape element.

Eventually, source and target are thus represented as

$$\mu = \sum_{i \in I} p_i \delta_{x_i} \quad \text{and} \quad \nu = \sum_{j \in J} q_j \delta_{y_j} \quad (1)$$

where $x_i, y_j \in \mathbb{X}$ are the sampled features, $p_i, q_j \geq 0$ are the associated masses, and δ_x denotes the Dirac mass at some point $x \in \mathbb{X}$.

Diffeomorphic Registration and Data Fidelity. Variational diffeomorphic registration of shape μ onto the shape ν consists in the minimization of

$$\mathcal{E}(f) \stackrel{\text{def.}}{=} \mathcal{R}(f) + \mathcal{L}(f_*\mu, \nu) \quad (2)$$

where \mathcal{R} is the regularization term on the diffeomorphism $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and \mathcal{L} is the data fidelity. The notation $f_*\mu$ stands for the data μ deformed by f .

The most simple data fidelity terms are derived from Euclidean norms using a smoothing operation against a kernel G_σ of width σ

$$\mathcal{L}(\mu, \nu) = \int_{\mathbb{X}} \left(\int_{\mathbb{X}} G_\sigma(x, y) (d\mu(x) - d\nu(x)) \right)^2 dy. \quad (3)$$

This class of losses has been used extensively for shape matching (see for instance [3, 15]) and is also popular in machine learning under the name Kernel mean embedding [13]. Its limitations for diffeomorphic registration are discussed in Sect. 4.

Optimization Scheme. In all the cases of interest, the action $f_*\mu$ of f on a finite discrete measure of the form (1) can be written using a finite dimensional vector θ of parameters. This formulation includes non-parametric methods since the parametrization may depend on the input shape, as in LDDMM methods. We write down this action as $f_*\mu = \mu_\theta \stackrel{\text{def.}}{=} \sum_i p(\theta)_i \delta_{x(\theta)_i}$.

Registration is achieved by computing a local minimizer of

$$\min_{\theta} \mathcal{E}(\theta) = \mathcal{R}(\theta) + \mathcal{L}(\mu_\theta, \nu) \quad (4)$$

using a descent-based method, typically initialized for $\theta = 0$ using $\mu_0 = \mu$, i.e. $p(0)_i = p_i, x(0)_i = x_i$. From a computational perspective, we simply need to provide the gradient of the functional which reads, thanks to the chain rule

$$\nabla \mathcal{E}(\theta) = \nabla \mathcal{R}(\theta) + [\partial x(\theta)]^* (\nabla_x \mathcal{L}(\mu_\theta, \nu)) + [\partial p(\theta)]^* (\nabla_p \mathcal{L}(\mu_\theta, \nu)), \quad (5)$$

where $([\partial x(\theta)]^*, [\partial p(\theta)]^*)$ are the adjoints of the Jacobians of the maps $(\theta \mapsto x(\theta), \theta \mapsto p(\theta))$ and $(\nabla_x \mathcal{L}(\mu, \nu), \nabla_p \mathcal{L}(\mu, \nu))$ are the gradients of the maps $x \mapsto \mathcal{L}(\sum_i p_i \delta_{x_i}, \nu)$ and $p \mapsto \mathcal{L}(\sum_i p_i \delta_{x_i}, \nu)$.

3 Optimal Transport Data Fidelity

This section proposes a new class of data fidelity $\mathcal{L}(\mu, \nu)$ between measures, using the recently proposed framework of unbalanced optimal transport between positive measures.

Unbalanced Regularized Optimal Transport. We consider two input discrete measures (1). In OT, the transportation can be described by a joint distribution defined on $\mathbb{X} \times \mathbb{X}$ coupling the two measures. For discrete inputs, this coupling is an array $\gamma = (\gamma_{i,j})_{i,j} = \text{“displaced mass from } x_i \text{ to } y_j\text{”}$ of positive numbers, whose marginals $\gamma \mathbf{1}_J \stackrel{\text{def.}}{=} (\sum_j \gamma_{i,j})_i$ and $\gamma^\top \mathbf{1}_I \stackrel{\text{def.}}{=} (\sum_i \gamma_{i,j})_j$ should be equal (for classical balanced) or close (for unbalanced) to the input measures μ (source) and ν (target).

An approximate OT discrepancy is obtained by looking for an optimal coupling as the *regularized* optimal transport cost is given by

$$W_{\varepsilon, \rho}(\mu, \nu) \stackrel{\text{def.}}{=} \min_{\gamma \in \mathbb{R}_+^{I \times J}} \sum_{i,j} c(x_i, y_j) \gamma_{i,j} - \varepsilon H(\gamma) + \rho \text{KL}(\gamma \mathbf{1}_J | p) + \rho \text{KL}(\gamma^\top \mathbf{1}_I | q). \quad (6)$$

Here the Kullback-Leibler divergence and entropy read

$$\text{KL}(h|p) \stackrel{\text{def.}}{=} \sum_i h_i \log \left(\frac{h_i}{p_i} \right) - h_i + p_i, \quad H(\gamma) \stackrel{\text{def.}}{=} - \sum_{i,j} \gamma_{i,j} (\log(\gamma_{i,j}) - 1),$$

and $c(x_i, y_j)$ is the cost of displacing a *unit* amount of mass between positions x_i and y_j . In the problem (6), $\varepsilon \geq 0$ controls the degree of regularization, and setting $\varepsilon = 0, \rho = +\infty$ recovers the usual OT. The influence of both parameters ε and ρ is discussed in Sect. 4.

Generalized Sinkhorn Algorithm. Following [4], one can check that the optimal γ can be written in the form

$$\gamma = \exp(K(u, v)) \text{ where } \forall (i, j) \in I \times J, \quad K(u, v)_{i,j} \stackrel{\text{def.}}{=} \frac{u_i + v_j - c(x_i, y_j)}{\varepsilon}, \quad (7)$$

which thus only requires the computation of two “dual” vectors $(u, v) \in \mathbb{R}^I \times \mathbb{R}^J$. In addition, these two vectors can be computed using an adaptation of the classical Sinkhorn algorithm. Introducing the log-sum-exp operator $\text{LSE}_I(K) = \log \sum_{i \in I} \exp(K_{i,j}) \in \mathbb{R}^J$ (and similarly for LSE_J , doing summation over $j \in J$), and starting from $(u, v) = (0_I, 0_J)$, Sinkhorn’s iterations read

$$\begin{aligned} u &\leftarrow \lambda u + \varepsilon \lambda \log(p) - \varepsilon \lambda \text{LSE}_J(K(u, v)) \\ v &\leftarrow \lambda v + \varepsilon \lambda \log(q) - \varepsilon \lambda \text{LSE}_I(K(u, v)) \end{aligned} \quad (8)$$

where we defined $\lambda \stackrel{\text{def.}}{=} \frac{\rho}{\rho+\varepsilon}$. The output of the algorithm is then $\gamma = \exp(K(u, v))$. Note that when $\rho = +\infty$ (balanced case), $\lambda = 1$ and these iterations correspond to a stable implementation over the log-domain of the well-known Sinkhorn algorithm (which is usually written using multiplicative updates, which is unstable for small ε). This algorithm is known to converge linearly to the optimal coupling.

Derivatives of the OT Fidelity. A proof similar to the balanced case (see [7]) shows that for $\varepsilon > 0$, $(p, x) \mapsto W_{\varepsilon, \rho}(\mu, \nu)$ is smooth and its gradient reads

$$\nabla_p W_{\varepsilon, \rho}(\mu, \nu) = \rho(1 - e^{-\frac{u}{\rho}}) \quad \text{and} \quad \nabla_x W_{\varepsilon, \rho}(\mu, \nu) = (\sum_j \gamma_{i,j} \partial_1 c(x_i, y_j))_i \quad (9)$$

(with the convention $(1 - e^{-\frac{u}{\rho}}) = u$ for $\rho = +\infty$) where γ is the solution of (6) and u is the limit of Sinkhorn iterations (8). Here $\partial_1 c$ is the derivative of the cost c with respect to the first variable.

4 Numerical Results

In this section, we showcase the use of unbalanced OT as a versatile fidelity term for registration. Our code is available: github.com/jeanfeydy/lddmm-ot.

Practical Implementation of OT Fidelity. To use OT on real data, one simply needs: an appropriate mapping from the dataset to the space of measures on a features space \mathbb{X} ; a cost function $c(x, y)$ on $\mathbb{X} \times \mathbb{X}$, with values in \mathbb{R}^+ .

In the case of curves/surfaces, following the construction of Sect. 2, one needs to choose a cost function on the (positions, normals) product $\mathbb{X} = \mathbb{R}^d \times \mathbb{S}^{d-1}$. One can use the canonical distance between $x = (a, u)$ and $y = (b, v)$

$$c(x, y) = \|a - b\|^2 + \alpha d_{\mathbb{S}^{d-1}}^2(u, v), \quad (10)$$

$$\text{or, for instance, } c(x, y) = \|a - b\|^2 \cdot \left(1 + \alpha (1 - \langle u, v \rangle)^k\right) \quad (11)$$

where $\alpha \geq 0$ and $d_{\mathbb{S}^{d-1}}$, k parametrizes the angular selectivity of the registration. Doing so, choosing $\alpha = 0$ allows one to retrieve the standard Wasserstein distance between shapes, whereas using $d_{\mathbb{S}^{d-1}}(u, v) = (1 - \langle u, v \rangle)$, $k = 1$ (resp. $d_{\mathbb{S}^{d-1}}(u, v) = (2 - 2\langle u, v \rangle^2)$, k even) can be seen as using globalized variants of the classical currents [15] (resp. varifold [3]) costs.

The registration is then obtained using the fidelity $\mathcal{L} = W_{\varepsilon, \rho}$ in the registration problem (4). The impact of this change (with respect to using more classical fidelity terms such as (3)) simply requires to input the expressions (9) of the gradients in the chain rule (5), which can be evaluated after running Sinkhorn algorithm (8) to compute the optimal γ and u needed in formula (9). In order to get non-negative fidelities, one can also discard the entropy and KL divergence from the final evaluation of the cost \mathcal{E} , and compute its derivatives using an autodiff library such as Theano [1]: this is what was used for Fig. 1.

Synthetic Dataset. Figure 1 showcases the results of our methods on a difficult 2-D curve registration problem. The first curve (rainbow colors, represented as a

measure μ) is deformed into the purple one (measure ν). Both curves are rescaled to fit into the unit square, the background grid of (a) is graduated every .05, and the cost function used is that of equation (11) with $\alpha = 1, k = 4$. The diffeomorphism is computed with an LDDMM sum of Gaussian kernels, with $k(x, y) = 1. \exp(-\|x - y\|^2 / (2 \cdot .025^2)) + .75 \exp(-\|x - y\|^2 / (2 \cdot .15^2))$.

RKHS fidelity: first row (b), (c). (b) and (c) have been computed using a *kernel-varifold* fidelity, with a spatial Gaussian kernel of deviation σ and an acute angular selectivity in $\cos^4(\theta)$ – with θ the angle between two normal directions. As shown in (b), such an RKHS fidelity method performs well with a large bandwidth σ . Unfortunately, trying to increase the precision by lowering the value of σ leads to the creation of undesirable local minima. In the eventual registration (c) the arms are not transported, but shrunk/expanded, as indicated by the *color code*. Classical workarounds involve decreasing σ during the shape optimization in a coarse-to-fine scheme, which requires a delicate parameters tuning. The main contribution of this paper is that it provides a principled solution to this engineering problem, which is *independent* of the underlying optimization/gradient descent toolbox, and can be adapted to any non-local fidelity term.

OT fidelity: second row. In sharp contrast with this observed behavior of RKHS fidelity terms, the OT data attachment term overcomes local minima through the

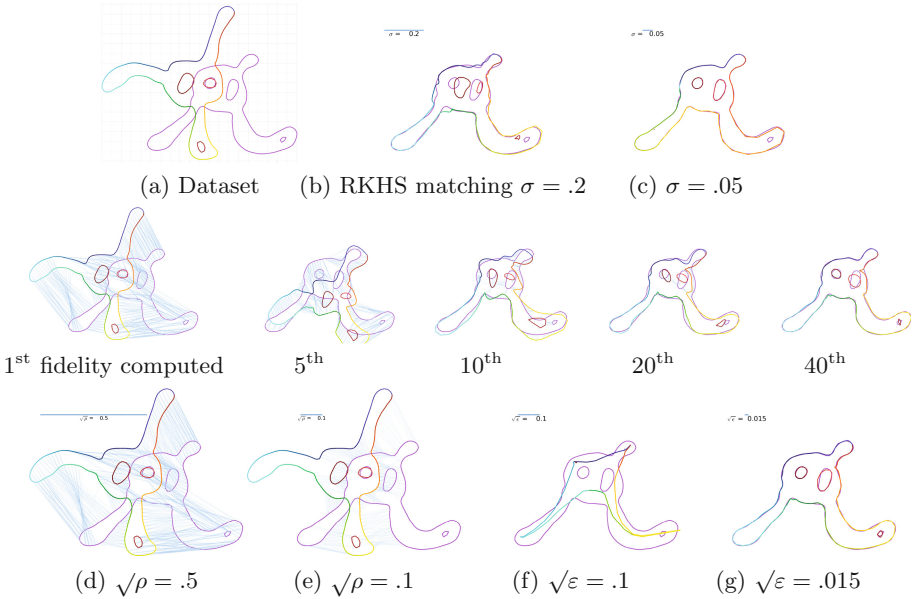


Fig. 1. First row: presentation of a difficult registration problem. Even though it looks precise, (c) completely mismatches the shapes’ arms as evidenced by the *color code*. **Second row:** evolution of the registration algorithm (minimizing \mathcal{E}). **Third row:** influence of (ρ, ϵ) .

computation of a *global* transport plan, displayed in light blue. Note that since the cost function used in this section is *quadratic*, ρ and ε should be interpreted as squared distances, and we used here $\sqrt{\varepsilon} = .015$, $\sqrt{\rho} = .5$. The final matching is displayed (g).

Influence of ρ , third row (d), (e). Here, we used $\sqrt{\varepsilon} = .03$. The value of $\sqrt{\rho}$ acts as a “cutoff scale”, above which OT fidelity tends to favour mass destruction/creation over transport. This result in a “partial” and localized transport plan, which is useful when dealing with outliers, large mass discrepancies which should not be explained through transport.

Influence of ε , third row (f), (g). Here, $\sqrt{\rho} = .5$. $\sqrt{\varepsilon}$ should be understood as a diffusion, blurring scale on the optimal transport. The resulting matching can therefore only capture structure up to a scale $\sqrt{\varepsilon}$: in (f), the “skeleton” mean axis of the shape.

Computational cost. The number of steps needed to compute a transport plan roughly scales like $O(\rho/\varepsilon)$. In this experiment, an evaluation of the fidelity term and its gradient was 100 (resp. 1000) times as long to compute as a RKHS loss of the form (3), for $\sqrt{\varepsilon} = .1$ (resp. $.015$). It thus has roughly the same complexity (resp. one order of magnitude slower) than evaluating the LDDMM diffeomorphism itself. As shown in the second row of Fig. 1, an efficient optimization routine may only require to evaluate the OT fidelity a handful of times to be driven to an appropriate rough deformation. Although not used here, a heuristic to drastically reduce the computational workload is a two-step scheme: first, use OT with a large ε to find the good basin of attraction; then, use a fast non-local fidelity (e.g. (3) with small σ) to increase precision.

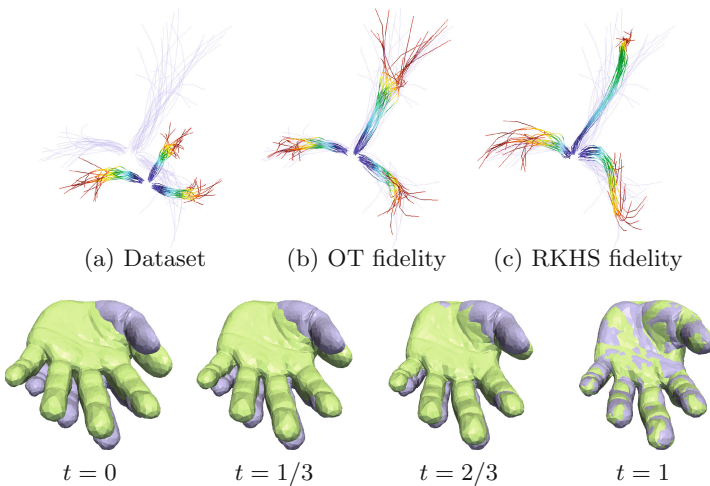


Fig. 2. First row: Matching of fibres bundles. **Second row:** Matching of two hand surfaces using a balanced OT fidelity. Target is in purple.

Fibres Bundles Dataset. The numerical experiment presented Fig. 2 illustrates the problem of registration of fibres bundles in 3d. It is often difficult to compute convincing registration of fibers bundles as the ends of the fibres are in practice difficult to align. This toy example may be seen as a very simple prototype of white matter bundles common in brain imaging. Currents-based distance together with a LDDMM framework were already used to analyze this kind of data, see e.g.; [9].

The source and target shape have 3 bundles containing around 20 fibers each. The diameter of the dataset is normalized to fit in a box of size 1. The cost function used for the OT fidelity is (10) with the orientation dependant distance between normals. We use the unbalanced framework with $\sqrt{\rho} = 1$ and $\sqrt{\varepsilon} = 0.07$. Using this OT fidelity with LDDMM allows to recover the shape of the target bundles (see Fig. 2 first row) whereas the RKHS fidelity based registration (we use a Gaussian kernel width $\sigma = 0.8$) converges toward a poor local minimum.

Hands Dataset. OT fidelity may be used with large datasets thanks to an efficient implementation of Sinkhorn iterations. The two hand shape surfaces of Fig. 2 contain more than 5000 triangles. The registration takes less than 1 h on a GPU. This numerical experiment shows that OT fidelity may be used to register surfaces with features at different scales.

5 Conclusion

In this article, we have shown that optimal transport fidelity leads to more robust and simpler diffeomorphic registration, avoiding poor local minima. Thanks to the fast Sinkhorn algorithm, this versatile tool has proven to be usable and scalable on real data and we illustrated its efficiency on curves, surfaces and fibres bundles. We plan to extend it to segmented volumetric image data.

References

1. Al-Rfou, R., Alain, G., Almahairi, A., Angermüller, C., Bahdanau, D., Ballas, N., Zhang, Y., et al.: Theano: A python framework for fast computation of mathematical expressions. CoRR abs/1605.02688 (2016)
2. Avants, B., Epstein, C., Grossman, M., Gee, J.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* **12**(1), 26–41 (2008)
3. Charon, N., Trounev, A.: The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM J. Imaging Sci.* **6**(4), 2547–2580 (2013)
4. Chizat, L., Peyré, G., Schmitzer, B., Vialard, F.X.: Scaling algorithms for unbalanced transport problems. arXiv preprint [arXiv:1607.05816](https://arxiv.org/abs/1607.05816) (2016)
5. Chizat, L., Schmitzer, B., Peyré, G., Vialard, F.X.: An interpolating distance between optimal transport and Fisher-Rao metrics. *Found. Comput. Math.* 1–44 (2016). doi:[10.1007/s10208-016-9331-y](https://doi.org/10.1007/s10208-016-9331-y)
6. Cuturi, M.: Sinkhorn distances: lightspeed computation of optimal transportation. *Proc. NIPS* **26**, 2292–2300 (2013)

7. Cuturi, M., Doucet, A.: Fast computation of Wasserstein barycenters. In: Proceedings of the 31st International Conference on Machine Learning (ICML), JMLR W&CP, vol. 32 (2014)
8. Gee, J.C., Reivich, M., Bajcsy, R.: Elastically deforming a three-dimensional atlas to match anatomical brain images. *J. Comput. Assist. Tomogr.* **17**, 225–236 (1993)
9. Gori, P., et al.: A prototype representation to approximate white matter bundles with weighted currents. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) MICCAI 2014. LNCS, vol. 8675, pp. 289–296. Springer, Cham (2014). doi:[10.1007/978-3-319-10443-0_37](https://doi.org/10.1007/978-3-319-10443-0_37)
10. Liero, M., Mielke, A., Savaré, G.: Optimal entropy-transport problems and a new Hellinger-Kantorovich distance between positive measures. arXiv e-prints (2015)
11. Montavon, G., Müller, K.R., Cuturi, M.: Wasserstein training of restricted Boltzmann machines. In: *Advance in Neural Information Processing Systems* (2016)
12. Santambrogio, F.: Optimal transport for applied mathematicians. *Progress in Non-linear Differential Equations and Their Applications*, vol. 87. Springer, Cham (2015)
13. Sejdinovic, D., Sriperumbudur, B., Gretton, A., Fukumizu, K.: Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *Ann. Statist.* **41**(5), 2263–2291 (2013)
14. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: a survey. *IEEE Trans. Med. Imaging* **32**(7), 1153–1190 (2013)
15. Vaillant, M., Glaunès, J.: Surface matching via currents. In: Christensen, G.E., Sonka, M. (eds.) IPMI 2005. LNCS, vol. 3565, pp. 381–392. Springer, Heidelberg (2005). doi:[10.1007/11505730_32](https://doi.org/10.1007/11505730_32)