Fast Multigrid Optimal Mass Transport for Image Registration and Morphing

Tauseef ur Rehman, Gallagher Pryor and Allen Tannenbaum
Georgia Institute of Technology, Atlanta GA, USA

Abstract

In this paper we present a novel, computationally efficient algorithm for non-rigid 2D image registration based on the work of Haker *et al.*[1, 2]. We formulate the registration task as an Optimal Mass Transport (OMT) problem based on the Monge-Kantorovich theory. This approach gives a number of advantages over other conventional registration methods: **(1)** It is parameter free and no landmarks need to be specified, **(2)** it is symmetrical and the energy functional has a unique minimiser, and **(3)** it can register images where brightness constancy is an invalid assumption. Our algorithm solves the Optimal Mass Transport program via multi-resolution, multi-grid, and parallel methodologies on a consumer graphics processing unit (GPU). Although solving the OMT problem has been shown to be computationally expensive in the past, we show that our approach is almost two orders magnitude faster than previous work and is capable of finding transport maps with optimality measures (mean curl) previously unattainable by other works (which directly influences the quality of registration). We give results where the algorithm was used to register 2D short axis cardiac MRI images and to morph two image sets from a SOHO solar flare image sequence.

1 Introduction

1.1 Image Registration and Morphing

Image registration and morphing are amongst the most common image processing problems. Registration is necessary in order to compare or integrate image data obtained from different measurements while image morphing, on the other hand, is a class of techniques that deals with the metamorphosis of one image into another (also known as image interpolation). Given two related images, these two techniques can be used together to generate a sequence of intermediate images in which an image gradually changes into the other over time. Critical to the success of this process is the quality of the warping function generated by the registration procedure. Morphing is achieved by shifting all pixels in the image according to the vector field defined by this warping function and will not be reasonable if the underlying warping function is not correct.

1.2 Optimal Mass Transport

The optimal mass transport problem was first formulated by a French mathematician Gasper Monge in 1781, and was given a modern formulation in the work of Kantorovich [3] and, therefore, is now known as the *Monge-Kantorovich problem*. The original problem concerned finding the optimal way to move a pile of soil from one site to another in the sense of minimal transportation cost. Hence, the Kantorovich-Wasserstein distance is also commonly referred to as the *Earth Mover's Distance (EMD)*.

Recently, Haker et al. have applied the optimal mass transport approach to certain medical image registration problems [1, 2]. Rigorous mathematical details for their algorithm are given by Angenent et al. [4]. Although there have been a number of algorithms in the literature for computing an optimal mass transport, the method by Haker et al. computes the optimal warp from a first order partial differential equation, which is a computational improvement over earlier proposed higher order methods [5, 6] and computationally complex discrete methods based on linear programming. The OMT method has a number of distinguishing characteristics: **(1)** It is parameter free and no landmarks need to be specified, **(2)** it is symmetrical and the energy functional has a unique minimiser, and **(3)** it can register images where brightness constancy is an invalid assumption. Conducting registration via OMT also utilizes all of the greyscale data in both images and places the two images on equal footing. It is thus symmetrical; the optimal mapping from image A to image B is the inverse of the optimal mapping from B to A. Furthermore, OMT does not require that landmarks be specified and the minimizer of the distance functional involved is unique; there are no other local minimisers. Finally, OMT is specifically designed to take into account changes in densities that result from changes in area or volume.

However, solving the Optimal Mass Transport problem is computationally expensive and a typical 2D image registration using the method of [1, 2] has been reported to require a matter of minutes. It has also been shown that this method may not necessarily converge to an optimal solution within a feasible amount of time. In our previous work [16] we showed considerable improvement in computation time when implemented using multi-resolution strategy.

Contribution. In this paper we present a new algorithm based on [1, 2] for solving the Optimal Mass Transport Problem using a *coarse-to-fine* strategy and a *full multigrid solver*, which we implement on a consumer *graphics processing unit*. We demonstrate two orders of magnitude improvement in computation time and that the algorithm is capable of computing transport maps with optimality measures (mean curl) previously unattainable by other works. The latter development has a direct influence on the quality of registration and image warps obtained by the algorithm while the speed-up achieved by our algorithm makes the use of OMT on 3D grids and other large bodies of data feasible and furthers the applicability of OMT for image processing tasks.

2 Optimal Mass Transport Theory

2.1 Formulation of the Problem

We will briefly provide an introduction to modern formulation of the Monge-Kantorovich problem. We assume we are given, *a priori,* two sub-domains Ω_0 and Ω_1 of $\boldsymbol{R}^d$ with smooth boundaries, and a pair of positive density functions, μ_0 and μ_1 defined on Ω_0 and Ω_1 respectively. We assume that,

$$\int_{\Omega_0} \mu_0 = \int_{\Omega_1} \mu_1 \tag{1}$$

This ensures that we have same total mass in both the domains. We now consider *diffeomorphisms* $\tilde{u}$ from Ω_0 to Ω_1 which map one density to other in the sense that,

$$\mu_0 = \left| D\tilde{u} \right| \mu_1 \circ \tilde{u} \tag{2}$$

which we call the *mass preservation* (MP) property, and write $\tilde{u} \in MP$. Equation (2211) is called the *Jacobian equation.* Here, $|D \tilde{u}|$ denotes the determinant of the Jacobian map $D \tilde{u}$, and $\circ$ denotes composition of functions. It basically implies that if a small region in Ω_0 is mapped to a larger region in Ω_1, then there must be a corresponding decrease in density in order for the mass to be preserved. There may be many such mappings, and we want to pick an optimal one in some sense. Accordingly, we define the squared L^2 Monge-Kantorovich distance as following:

$$d_2^2(\mu_0, \mu_1) = \inf_{\tilde{u} \in MP} \int \left\| \tilde{u}(x) - x \right\|^2 \mu_0(x) dx \tag{3}$$

The *optimal MP map* is a map which minimizes this integral while satisfying the constraint (2). The Monge-Kantorovich functional (3) is seen to place a penalty on the distance the map $\tilde{u}$ moves each bit of material, weighted by the material's mass. A fundamental theoretical result[7-9], is that there is a unique optimal $\tilde{u} \in MP$ transporting μ_0 to μ_1, and that $\tilde{u}$ is characterized as the gradient of a convex function ω, i.e., $\tilde{u} = \nabla \omega$. This theory translates into a practical advantage, since it means that there are no non-global minima to stall our solution process.

2.2 Computing the Transport Map

We will describe here only the algorithm for finding the optimal mapping $\tilde{u}$. The details of this method can be found in [2]. The basic idea for finding the optimal warping function is first to find an initial MP mapping u^0 and update it iteratively to decrease an energy functional. When the pseudo time t goes to ∞, the optimal u will be found, which is $\tilde{u}$. Basically, there are two steps:

2.2.1 Finding an Initial Mapping

The first step in this algorithm is to find an initial mass preserving mapping. This can be done for general domains using the method of Moser [10] or the algorithm proposed in [2]. The later method can simply be interpreted as the solution of a one-dimensional Monge-Kantorovich problem in the *x*-direction followed by the solution

of a family of one-dimensional Monge-Kantorovich problems in *y*-direction. We used this later method for our implementation and experiments.

2.2.2 Finding the Minimizer

The second step is to adjust the initial mapping found above iteratively using gradient descent in order to minimize the functional defined in Equation (3), while constraining *u* so that it continues to satisfy Equation (2). This process iteratively removes the curl from the initial mapping *u* and, thereby, finds the polar factorization of *u*. For details on this technique, please refer to [2]. The overall algorithm is summarized graphically in Figure (1).

2.3 Defining the Warping Map

In elastic registration applications, one usually wants to visualize the explicit warping between the two images where one image smoothly deforms to the other. This has been shown to be easily done using the solution from the Monge-Kantorovich problem described above using the following relationship [2]:

$$X(x,t) = x + t(\widetilde{u}^*(x) - x) \qquad (4)$$

where, *X(x,t)* defines our continuous warping map between densities μ_0 and μ_1. Note that when *t = 0, X* is the identity map and when *t=1*, it is the solution $\widetilde{u}^*$ to the Monge-Kantorovich problem. All the morphs in the results section have been created using this equation.
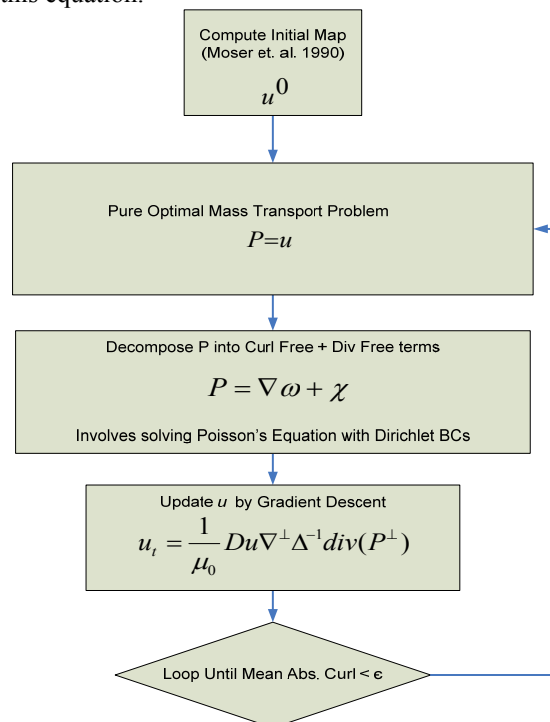

Figure 1: Optimal Mass Transport Algorithm

3 Multi-resolution OMT

Performing image registration using a multi-resolution approach is widely used to improve speed, accuracy and robustness. The basic idea is that registration is first performed at a coarse scale. The spatial mapping determined at the coarse level is then used to initialize registration at the next finer scale. This process is repeated until it reaches the finest scale. This *coarse-to-fine strategy* greatly improves the registration success rate and also increases robustness by eliminating local optima at coarser scales [11]. Our coarse-to-fine hierarchy comprised of three levels and uses bi-cubic interpolation to interpolate the solution from the coarse to fine grids.

Figure 2: Coarse-to-fine Approach

This multi-resolution technique diminishes the mean curl of u remarkably fast as compared to the normal case as is shown in Figure 3 below. These results correspond to the heart registration example given in Figure 8. Figure 3 shows the mean curl of u achieved in the same number of iterations with and without multi-resolution. It should be noted that same number of iterations of multi-resolution complete in less than 1/3 the time.


Figure 3: Convergence with and without multi-resolution approach

3.1 Multigrid

We used the *Full Multigrid* (FMG) algorithm in our implementation of the Poisson solver also known as the *nested iteration*. In the case of FMG, the initial approximation is obtained by interpolating from a coarse-grid solution. For discretization we used the standard five-point discrete approximation of the Laplacian given by:

$$\nabla^2 u_{ij} \approx u_{i-1,j} + u_{i+1,j} + u_{i,j+1} + u_{i,j-1} - 4u_{i,j} \qquad (5)$$

where, i and j are row and column indices into the grid. We chose the most frequently used Standard Coarsening, i.e. doubling the mesh size in each direction. The same discretization of the Laplace operator is used on all the grid sizes. This is known as *Discrete Coarse Grid Approximation* (DCA). We *restrict* the residual from the fine-to-coarse transfer using *Half Weighting* technique and for *interpolation* from coarse-to-fine transfers we use *Bilinear Interpolation.* For a comprehensive overview of multigrid methods, we refer to [12, 13].

The FMG algorithm is graphically illustrated in the following figure. The thick arrows indicate use of a higher order interpolation (bi-cubic) scheme instead of the Bilinear Interpolation. This is due to the fact that when we increment to the next to coarsest level, we do not have to use an interpolation complementing the Full weighting scheme to maintain symmetry in a cycle.


Figure 4: Structure of one Full Multigrid cycle

3.2 GPU Implementation

An advantage of our solution to the OMT problem is that it is particularly suited for implementation on parallel computing architectures making optimization with multiple processing cores highly advantageous. Over the past few years, it has been shown that GPUs are particularly suited for these types of parallelizable problems [14, 15].

Implementing the PDE-based OMT solver on the GPU comes quite naturally as it can be abstracted down to a series of convolutions and simple point-wise arithmetic between data grids. From this standpoint, the GPU has much to gain over the CPU implementation: while the CPU computes updates on data grids one element at a time, the GPU is capable of updating entire grids in one pass due to their massively parallel architecture. This is graphically illustrated in Figure 5 below. Such simplicity allowed us to implement the entire algorithm via an OpenGL/fragment shader approach using the same methodology as [14, 15] for GPU computation. In contrast to [14,15], however, the OMT solver is much more

complex in that a multigrid solver is only a component of the algorithm in addition to additional computation of temporary quantities (curl, divergence, coarse-to-fine interpolation) and an upwinding scheme for state updates.


Figure 5: CPU vs. GPU solutions of PDEs.

The GPU implementation of the algorithm gave considerable performance gains over the CPU. For instance, on a 512^2 grid, the time required to perform one iteration of the solver by Haker et. al. was 0.42 seconds while the GPU implementation required 0.0159 seconds, representing an improvement of 2681 percent (Figure 6). This improvement increases with grid size to 7060 percent on a 2048^2 grid. All results were obtained on a Dual Xeon 1.6GHz and nVidia GeForce 7950 GX2 graphics card with a 3DMark score of 6747.


Figure 6: Comparison of execution speeds at each grid size.

4 Examples and Results

We illustrate our method using two examples. In the first case, we interpolate between two greyscale SOHO solar flare images shown in colour here to highlight the details. We take Figure 7(a) and Figure 7(b) to be the starting and ending images, respectively. Figure 8 shows the intermediate images generated by our morphing algorithm at times t=0.25, 0.5 and 0.75. In Figure 9 we show that a classical non-rigid registration based on optical flow fails to warp the flare images. The procedure cannot account for the new information present in the second image versus the first image and thus the warp fails.

In the second case, we applied our method to two different 2D short axis cardiac MRI images (256^2 pixels) taken from a sequence of images (Figure 10). The images were inverted for improved performance. Note that the deformed grid

(Figure 11) shows areas of compression of the ventricle. Note that in each case, total computation time was only a matter of seconds.


(a) Source Image   (b) Target Image

Figure 7: Source and target solar flare images (512^2 grid)

(a) $t = 0.25$ (b) $t = 0.50$

(c) $t = 0.75$ (d) $t = 1.00$

Figure 8: OMT-interpolated solar flare images (12 sec, μ curl $10^{-3.2}$)

Figure 9: Failure of classical non-rigid registration based on optical flow warp.

Figure 10: Two frames from 2D cardiac MRI image sequence.

Figure 11: Deformed grid for cardiac example (64 sec, μ curl $10^{-3.6}$).

5 Conclusions

In this paper, we presented a computationally efficient method for image registration and morphing based on the classical problem of optimal mass transportation. The speed-up achieved by our algorithm makes the use of OMT on 3D grids and other large bodies of data feasible and furthers the applicability of OMT for image processing tasks including compression and coding.

References

[1]   S. Haker and A. Tannenbaum, "Optimal Mass Transport and Image Registration," in *IEEE Conference on Variational and Level Set methods in Computer Vision*, 2001.

[2]   S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent, "Optimal Mass Transport for Registration and Warping," *International Journal of Computer Vision,* vol. 60, pp. 225-240, 2004.

[3]   L. V. Kantorovich, "On a problem of Monge," *Uspekhi Mat. Nauk.,* vol. 3, pp. 225-226, 1948.

[4]   S. Angenent, S. Haker, and A. Tannenbaum, "Minimizing flows for the Monge-Kantorovich problem," *SIAM Journal of Mathematical Analysis,* vol. 35, pp. 61-97, 2003.

[5]   G. E. Christensen, "A deformable neuroanatomy handbook based on viscous fluid mechanics," in *27th Annual Conference on Information Sciences and Systems*, 1993, pp. 211-216.

[6]   G. E. Christensen, R. D. Rabbit, and M. Miller, "Deformable templates using large deformation kinetics," *IEEE Transactions on Image Processing,* vol. 5, pp. 1435-1447, 1996.

[7]   Y. Brenier, "Polar factorization and monotone rearrangement of vector-valued functions," *Comm. Pure Appl. Math.,* vol. 64, pp. 375-417, 1991.

[8]   W. Gangbo and R. McCann, "The geometry of optimal transportation," *Acta Math.,* vol. 177, pp. 113-161, 1996.

[9]   M. Knott and C. Smith, "On the optimal mapping of distributions," *J. Optim. Theory,* vol. 43, pp. 39-49, 1984.

[10]   J. Moser, "On the volume elements on a manifold," *Transactions of American Mathematics Society,* vol. 120, pp. 286-294, 1965.

[11]   T. S. Yoo, *Insight into images, Principles and Practice for Segmentation, Registration and Image Analysis*: A. K. Peters Ltd., 2004.

[12]   W. L. Briggs, V. E. Hensen, and S. F. McCormick, *A Multigrid Tutorial*: SIAM, 2000.

[13]   U. Trottenberg, C. Oosterelee, and A. Shuller, *Multigrid*: Academic Press, 2001.

[14]   J. Bolz and e. al., "Sparse matrix solvers on the GPU: Conjugate Gradients and Multigrid," in *SIGGRAPH*, 2003, pp. 917-924.

[15]   G. Nolan and e. al.. "A multigrid solver for boundary value problems using programmable graphics hardware," in *SIGGRAPH*, 2003, pp. 102-111.

[16]   T. Rehman and A. Tannenbaum, "Multigrid Optimal Mass Transport for Image Registration and Morphing", in *SPIE Conference on Computational Imaging V*, 2007