

Published in final edited form as:

Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit. 2010 June ; 2010: 2933–2939. doi:10.1109/CVPR.2010.5540035.

Image Atlas Construction via Intrinsic Averaging on the Manifold of Images

Yuchen Xie, Jeffrey Ho, and Baba C. Vemuri

Department of Computer and Information Sciences and Engineering University of Florida

Yuchen Xie: yxie@cise.ufl.edu; Jeffrey Ho: jho@cise.ufl.edu; Baba C. Vemuri: vemuri@cise.ufl.edu

Abstract

In this paper, we propose a novel algorithm for computing an atlas from a collection of images. In the literature, atlases have almost always been computed as some types of means such as the straightforward Euclidean means or the more general Karcher means on Riemannian manifolds. In the context of images, the paper's main contribution is a geometric framework for computing image atlases through a two-step process: the localization of mean and the realization of it as an image. In the localization step, a few nearest neighbors of the mean among the input images are determined, and the realization step then proceeds to reconstruct the atlas image using these neighbors. Decoupling the localization step from the realization step provides the flexibility that allows us to formulate a general algorithm for computing image atlas. More specifically, we assume the input images belong to some smooth manifold \mathbf{M} modulo image rotations. We use a graph structure to represent the manifold, and for the localization step, we formulate a convex optimization problem in \mathbb{R}^k (k the number of input images) to determine the crucial neighbors that are used in the realization step to form the atlas image. The algorithm is both unbiased and rotation-invariant. We have evaluated the algorithm using synthetic and real images. In particular, experimental results demonstrate that the atlases computed using the proposed algorithm preserve important image features and generally enjoy better image quality in comparison with atlases computed using existing methods.

1. Introduction

Computing an atlas from a collection of images is a fundamental problem in medical imaging and computer vision. (e.g., [10,7,6,21]). Since the atlas is supposed to be an informative representative of the given set of images, in most existing approaches, the atlas construction problem is often formulated as some kind of mean estimation problem. For example, one of the popular existing methods for computing an atlas is the unbiased diffeomorphic atlas construction algorithm proposed in [7]. If $\{\mathbf{I}_1, \dots, \mathbf{I}_k\}$ denote the input images, [7] solves the atlas construction problem using a variational approach based on the large deformation diffeomorphism framework described in [12,3,13]. The image atlas \mathbf{I} is given as:

$$\mathbf{I} \equiv \operatorname{argmin}_{\tilde{\mathbf{I}} \in \mathcal{I}} \sum_{i=1}^k \mathbf{d}_T^2(\mathbf{I}_i, \tilde{\mathbf{I}}), \quad (1)$$

where $d_{\mathcal{I}}$ is a metric on the space \mathcal{I} of images. Geometrically, Equation 1 can be interpreted as computing the mean of the collection of points (images) $\{\mathbf{I}_1, \dots, \mathbf{I}_k\}$ in the metric space $(\mathcal{I}, d_{\mathcal{I}})$.

Figure 1 shows the result of applying the above method to a collection of 2D images. The result is clearly unsatisfactory in that the atlas image is blurry and many important details are missing or are corrupted. A geometric explanation for this result is offered in Figure 2: The smooth variation of the images allows us to assume that these images belong to some smooth submanifold \mathbf{M} in \mathcal{I} . [7] computes the mean of the images in the ambient space \mathcal{I} , and as the figure shows, there is no guarantee that the result will be on \mathbf{M} . Clearly, the correct mean in this case should be the mean of the images considered as the points on \mathbf{M} . Based on the seminal paper of Karcher [9], numerical computation of means on Riemannian manifolds has received considerable attention in the past few years (e.g., [20,14]), and there is already a body of recent works that applies these techniques to vision problems (e.g., [5,22]). Unfortunately, all of these methods require the precise knowledge of the Riemannian metric, and this limits the current techniques to mostly symmetric Riemannian manifolds, examples of which include the Grassmanians and the symmetric positive-definite matrices, etc.. In our case, the manifold \mathbf{M} and its metric are both unknown. Therefore, we need a method that can estimate the mean on \mathbf{M} without using the metric on \mathbf{M} , a requirement that seems contradictory at first. However, Figure 2 also reveals one thing that is often valid in that the mean on \mathbf{M} can be approximated well by, for example, averaging a few of its neighbors. The challenge here is, of course, to identify these neighbors in \mathbf{M} without knowing explicitly the Riemannian structure of \mathbf{M} . In particular, this also suggests that in the context of working with images, computing image atlases should be a two-step process: localize the mean first by determining these neighbors and realize this *abstract* mean as the atlas image using these neighbors. The two steps are qualitatively different in that the localization step utilizes the entire set of input images. It uses the global geometry to determine these crucial neighbors among the input images, and the realization step that follows uses only these neighbors for constructing the image.

In many applications, it is reasonable to demand rotation invariance for the atlas, and Figure 3 shows one example. More specifically, suppose \mathbf{I} is the atlas for $\{\mathbf{I}_1, \dots, \mathbf{I}_k\}$ and if k different rotations g_i are applied to these images to form a new collection $\{g_1(\mathbf{I}_1), \dots, g_k(\mathbf{I}_k)\}$, then the atlas $\tilde{\mathbf{I}}$ for the latter collection should be related to \mathbf{I} by some rotation. This requirement is desirable because the image contents of \mathbf{I}_i and $g_i(\mathbf{I}_i)$ are essentially the same, and modulo the rotations, these two image collections are completely equivalent. Therefore, it is reasonable to require that their atlases are also equivalent modulo rotations, i.e., $\tilde{\mathbf{I}} = g(\mathbf{I})$ for some rotation g . However, this requirement means that we can no longer assume that the input images belong to some submanifold \mathbf{M} in \mathcal{I} . Instead, the correct assumption should be that modulo rotations, the images belong to a submanifold \mathbf{M} in the quotient space \mathbf{Q} of \mathcal{I} [17,16].

In this paper, we propose a novel approach to the atlas construction problem that incorporates the ideas discussed above. Starting with a rotation-invariant metric $d_{\mathcal{I}}$ on \mathcal{I} , it induces a metric $d_{\mathbf{Q}}$ on the quotient space \mathbf{Q} . The proposed algorithm computes the atlas as an approximation of the mean m of the images considered as points on \mathbf{M} . In general, the metric on \mathbf{M} is unknown; however, the pairwise geodesic distances $m_{ij} = d_{\mathbf{M}}(\mathbf{I}_i, \mathbf{I}_j)$ between images can be estimated using the metric in \mathbf{Q} [8]. An important element in the proposed algorithm is to estimate the geodesic distances $d_{\mathbf{M}}(m, \mathbf{I}_i)$ between the mean m and images \mathbf{I}_i using the pairwise geodesic distances m_{ij} as the inputs. We formulate this as a convex optimization problem in \mathbb{R}^k , and the solution to this optimization problem will provide an approximations to $d_{\mathbf{M}}(m, \mathbf{I}_i)$. This will allow us to locate the mean m on the manifold \mathbf{M} and

identify images that are neighbors of m (belonging to a neighborhood of m). These neighbors (images) are used, together with their geodesic distances $\mathbf{d}_{\mathbf{M}}(m, \mathbf{I}_i)$ to the mean, to form the image for m , which, before this point, is simply an abstract point in \mathbf{M} . It is straightforward to show that the proposed algorithm is both unbiased and rotation-invariant. Preliminary experimental results have shown that, with greater preservation of important image features and details, the image atlases produced by the proposed algorithm enjoy better image quality when compared with atlases obtained using exiting methods.

The rest of this paper is organized as follows: in Section 2, we describe the theoretical framework for image atlas construction. Experimental results for synthetic and real image data are presented in Section 3. Finally, we conclude in section 4.

2. A Theoretical Framework for Atlas Construction

In the following discussion, \mathcal{I} will denote the space of images. Images in \mathcal{I} are defined formally as L^2 -functions on a finite domain $\Omega \subset \mathbb{R}^n$, where Ω is the underlying coordinate system for the images in \mathcal{I} ($n = 2$ for 2D images and $n = 3$ for 3D images). Let $\mathcal{C} \equiv \{\mathbf{I}_1, \dots, \mathbf{I}_k\}$ denote a collection of k images and $g(\mathbf{I})(x) = \mathbf{I}(g^{-1}(x))$ for $g \in \mathbf{SO}(n)$ and $x \in \Omega$. In this section, we will define the image atlas \mathbf{I} for the image collection \mathcal{C} and outline an algorithm for computing \mathbf{I} . We require the algorithm to satisfy the following two properties:

1. **Unbiasedness:** The algorithm output should not depend on the particular ordering of the images in the collection \mathcal{C} .
2. **Rotation Invariance:** If $\mathcal{C}' \equiv \{\mathbf{I}'_1, \dots, \mathbf{I}'_k\}$ are images obtained by applying k rotations to images in \mathcal{C} : $\mathbf{I}'_i = g_i(\mathbf{I}_i)$ for $g_i \in \mathbf{SO}(n)$, $1 \leq i \leq k$, then the atlas \mathbf{I}' for the image collection \mathcal{C}' is related to \mathbf{I} by a rotation $g \in \mathbf{SO}(n)$: $\mathbf{I}' = g(\mathbf{I})$.

2.1. Geometry of the Quotient Space \mathbf{Q}

The demand for rotation invariance requires us to work not in \mathcal{I} but in \mathbf{Q} , the quotient space of \mathcal{I} by the rotation $\mathbf{SO}(n)$ (see [17,16] for details). Basically, \mathbf{Q} is a space that parameterizes the $\mathbf{SO}(n)$ -orbits in \mathcal{I} and in this paper, we will assume that \mathbf{Q} has a manifold structure induced from the manifold structure of \mathcal{I} . Let $\pi: \mathcal{I} \rightarrow \mathbf{Q}$ denote the canonical projection map that sends each image $\mathbf{I} \in \mathcal{I}$ to the unique $\mathbf{SO}(n)$ -orbit $[\mathbf{I}] \in \mathbf{Q}$ containing \mathbf{I} : $\pi(\mathbf{I}) = [\mathbf{I}]$. The points $\{[\mathbf{I}_1], \dots, [\mathbf{I}_k]\}$ are assumed to belong to a smooth submanifold \mathbf{M} in \mathbf{Q} . As a submanifold of \mathbf{Q} , \mathbf{M} is equipped with the natural induced Riemannian metric and we let $\mathbf{d}_{\mathbf{M}}([\mathbf{I}_i], [\mathbf{I}_j])$ denote the geodesic distance function on \mathbf{M} . Let m denote the mean of $\{[\mathbf{I}_1], \dots, [\mathbf{I}_k]\}$ with respect to $\mathbf{d}_{\mathbf{M}}$, and the atlas $[\mathbf{I}]$ computed by our algorithm will be a suitable approximation of m .

$\mathbf{SO}(n)$ -invariant metric is defined as

$$\mathbf{d}_{\mathcal{I}}(\mathbf{I}_1, \mathbf{I}_2) = \mathbf{d}_{\mathcal{I}}(g(\mathbf{I}_1), g(\mathbf{I}_2)),$$

for any two images $\mathbf{I}_1, \mathbf{I}_2 \in \mathcal{I}$ and $g \in \mathbf{SO}(n)$. Many metrics can be easily shown to be $\mathbf{SO}(n)$ -invariant, such as L^2 -metric and the following metric defined in [2]:

$$\mathbf{d}_{\mathcal{I}}(\mathbf{I}_1, \mathbf{I}_2) = \min_{h_i} \left(\int_{\Omega} |\mathbf{I}_1(h_i(x)) - \mathbf{I}_2|^2 dx + \int_0^1 \int_{\Omega} \|L_{V_i}\|^2 dx ds \right)^{1/2} \quad (2)$$

where v_i is the time-dependent vector field that defines the diffeomorphic flow from the identity to the diffeomorphism h_i , and L is a second-order elliptic operator. Any $\mathbf{SO}(n)$ -invariant metric $\mathbf{d}_{\mathcal{I}}$ on \mathcal{I} induces a metric $\mathbf{d}_{\mathbf{Q}}$ in \mathbf{Q} using the formula [16]

$$\mathbf{d}_{\mathbf{Q}}([x], [y]) = \min_{s \in [x], t \in [y]} \mathbf{d}_{\mathcal{I}}(s, t) \quad (3)$$

for any two points $[x], [y] \in \mathbf{Q}^1$. The two related metrics $\mathbf{d}_{\mathcal{I}}$, $\mathbf{d}_{\mathbf{Q}}$ allows us to go between \mathcal{I} and \mathbf{Q} . For example, computing the mean in \mathbf{Q} minimizes the following function,

$$[\mathbf{I}] = \operatorname{argmin}_{[\mathbf{I}'] \in \mathbf{Q}} \sum_{i=1}^k \mathbf{d}_{\mathbf{Q}}^2([\mathbf{I}_i], [\mathbf{I}']).$$

The corresponding function to minimize in \mathcal{I} would be (using Equation 3)

$$\mathbf{I} = \operatorname{argmin}_{\mathbf{I}' \in \mathcal{I}, s_1, \dots, s_k \in \mathbf{SO}(n)} \sum_{i=1}^k \mathbf{d}_{\mathcal{I}}^2(s_i(\mathbf{I}_i), \mathbf{I}').$$

In particular, we have $\rho(\mathbf{I}) = [\mathbf{I}]$ and $\mathbf{I} \in \pi^{-1}([\mathbf{I}])$. Following the similar approach in Isomap [8], we could compute pair-wise geodesic distance $\mathbf{d}_{\mathbf{M}}([\mathbf{I}_i], [\mathbf{I}_j])$ using the induced metric $\mathbf{d}_{\mathbf{Q}}$.

2.2. Estimating Mean from Pairwise Geodesic Distances

Suppose $\{x_1, \dots, x_k\}$ are points on a manifold \mathbf{M} with $\mathbf{d}_{\mathbf{M}}$ denoting the Riemannian geodesic distance. One possible way to determine the mean m of x_i on \mathbf{M} is the following. Let $a_i = \mathbf{d}_{\mathbf{M}}(x_i, m)$, $1 \leq i \leq k$. Determining a_i is of course equivalent to locating m on \mathbf{M} , and a_i can be determined as the solution to an optimization problem given by:

$$\begin{aligned} \min \quad & a_1^2 + a_2^2 + \dots + a_k^2 \\ \text{constraints} \quad & a_i \geq 0, i=1, \dots, k \\ & a_i + a_j \geq m_{ij}, \\ & a_i + m_{ij} \geq a_j, \\ & a_j + m_{ij} \geq a_i, i, j=1, \dots, k \\ & \text{HDC}_{\mathbf{M}}(a_i, m_{ij}), \end{aligned} \quad (4)$$

where $m_{ij} = \mathbf{d}_{\mathbf{M}}(x_i, x_j)$ and $\text{HDC}_{\mathbf{M}}(a_i, m_{ij})$ denote the remaining higher-degree constraints among a_i, m_{ij} . The collection of constraints in $\text{HDC}_{\mathbf{M}}$ depends on \mathbf{M} as a metric space and in general, for two different manifolds, their corresponding HDC will be different. The linear inequality constraints between a_i and m_{ij} are imposed because $\mathbf{d}_{\mathbf{M}}$ satisfies the triangle inequality.

For our problem, the manifold \mathbf{M} and its metric $\mathbf{d}_{\mathbf{M}}$ are unknown, and consequently, the constraints in $\text{HDC}_{\mathbf{M}}$ cannot be known. However, as an approximation, we can try to solve

¹Note that $[x], [y]$ are realized in \mathcal{I} as $\mathbf{SO}(n)$ -orbits. $\mathbf{d}_{\mathbf{Q}}$ computes the distance between the two orbits in \mathcal{I} as measured by $\mathbf{d}_{\mathcal{I}}$

the following optimization problem in \mathbb{R}^k and use the solution a_1, \dots, a_k as an approximation to the solution of the above problem:

$$\begin{aligned} \min \quad & a_1^2 + a_2^2 + \dots + a_k^2 \\ \text{constraints} \quad & a_i \geq 0, i=1, \dots, k \\ & a_i + a_j \geq m_{ij}, \\ & a_i + m_{ij} \geq a_j, \\ & a_j + m_{ij} \geq a_i, i, j=1, \dots, k \end{aligned} \quad (5)$$

The optimization problem is clearly convex since the objective function is strictly convex (the Hessian is positive-definite everywhere) and the domain, which is the intersection of half-spaces, is also convex. This particular type of optimization problem (with quadratic cost function and linear inequality constraints) can be solved efficiently even with a large number of variables and constraints. And because the object function is strictly convex, the solution is unique [1]. Furthermore, the solution is stable with respect to the input parameters m_{ij} in the sense that small perturbations of m_{ij} will not significantly change the solution a_1, \dots, a_k [1]. For computing robust L^1 -median [4] instead of mean, we simply need to change the cost function to $a_1 + a_2 + \dots + a_k$.

2.3. Computing the Image Atlas I

The non-negative number a_i provides an estimate on the geodesic distance $\mathbf{d}_M(m, [\mathbf{I}_i])$. The last step is to use a_i to determine a small number of points in $\{[\mathbf{I}_1], \dots, [\mathbf{I}_k]\}$ that are close to m as measured by \mathbf{d}_M , and m is then approximated from these points with respect to the metric \mathbf{d}_Q in \mathbf{Q} . In practice, a positive integer K is specified and K points $[\mathbf{I}_1], \dots, [\mathbf{I}_K]$ with shortest distances to m are chosen. We require that $K \geq 1 + \text{dimension of the manifold}$.

We could approximate m with a weighted mean of the K points $[\mathbf{I}_j]$, with weights w_j , $1 \leq j \leq K$:

$$[\mathbf{I}] = \operatorname{argmin}_{[\mathbf{I}'] \in \mathbf{Q}} \sum_{j=1}^K w_j \mathbf{d}_Q^2([\mathbf{I}_j], [\mathbf{I}'])$$

Since a point closer to m contributes more in atlas realization, the weights w_j can be constructed from the estimated geodesic distances a_{ij} :

$$w_j = b_j / \left(\sum_{k=1}^K b_k \right)$$

where $b_j = \exp(-a_{ij}^2 / \sigma^2)$, $1 \leq j \leq K$, for some $\sigma > 0$. Using the metric defined by Equation 2, the corresponding atlas \mathbf{I} is computed by solving the variational problem

$$\min_{\tilde{\mathbf{I}}, h_j, s_j} \sum_{j=1}^K w_j \int_{\Omega} \|\mathbf{I}_j(h_j s_j(x)) - \tilde{\mathbf{I}}\|^2 dx + \int_0^1 \int_{\Omega} \|L v_j\|^2 dx ds \quad (6)$$

3. Experiments

In this section, we present three sets of experimental results that evaluate the proposed atlas construction algorithm. In these experiments, we will use both synthetic and real image data, and comparisons will be made with atlases computed using existing methods [4,7,21] that are based on non-rigid group-wise registrations.

3.1. Experimental Results Using COIL-20

In the first set of experiments, we compute the image atlas using images from the Columbia Object Image Library dataset (COIL-20) [18]. COIL-20 contains twenty objects, and images for each object are obtained by placing the object on a turntable and rotating it in five-degree increments in front of a fixed camera. This gives 72 images for each object. As the viewpoint varies on a circle, the appearance variation of the images can be modelled using an one-dimensional submanifold \mathbf{M} in the image space. In particular, the dataset has been used frequently to evaluate recognition and manifold learning algorithms (e.g., [19]). In this experiment, the number of nearest neighbors K is set to four, and the image atlases for five objects in COIL-20 are shown in Figure 4.

We validate the algorithm using a known ground truth. Taking the ‘duck’ object in COIL-20, we choose 36 images with pose (viewing angle) ranging from 5° to 90° and -5° to -90° . Several sample images are shown in Figure 5. As the viewpoint varies from -90° to 90° , the manifold \mathbf{M} underlies the 36 images has the topology of a semi-circle, and the mean in \mathbf{M} can be taken to be the image with 0-degree pose.

Figure 6 displays the image atlases (as \mathbf{L}^2 -mean and \mathbf{L}^1 -median) computed using our method and the algorithms proposed in [7,4] that use non-rigid group-wise registration. The algorithm in [7] computes the atlases as the \mathbf{L}^2 -mean while [4] constructs the atlases using \mathbf{L}^1 -median. In both cases, it is clear that our results are much closer to the ground truth than the atlases produced by the other methods. It is not surprising since the two existing methods does not respect the manifold structure and they in fact compute the atlases as the mean and median of the ambient space. Therefore, their images are typically blurry with substantial loss of image details.

Finally, we test the rotational invariance of the proposed algorithm. Given the 36 images used to compute the duck atlas above, we randomly rotate these images with angle between 20° and 30° . Since the algorithm factored out the rotations by working in the quotient space, the new and the old atlases should be the same point in the quotient space. That is, they should differ only by a planar rotation. As shown in Figure 6 (Second Row), this is indeed the case.

3.2. Experimental Results Using Synthetic Images

In the second experiment, we synthetically generate 56 images from the Stanford bunny mesh model. The model is centered at the origin and we sample 56 camera locations on the viewing sphere surrounding the model. These 56 locations are uniformly sampled with respect to a central (apex) location $([0\ 0\ 1])$, with maximum angular deviation from the apex of 30 degrees. These points are plotted in Figure 8. In all the renderings, we assume simple Lambertian reflectance model with the distant lighting direction parallel to the viewpoint (camera location). The 56 images are considered as a collection of sampled points from a two-dimensional submanifold in the image space, and we choose a value of 5 for the parameter K in this experiment. Figure 7 shows the atlas constructed by our method and the two comparisons that use straightforward averaging of intensity values and the unbiased

diffeomorphic atlas [7]. Note that outputs from the two competing methods are not satisfactory as many details and features are lost or smoothed out. The camera locations corresponding to the five neighbors ($K = 5$) closest to the mean determined by our algorithm are shown by the five red points in Figure 8. For this example, it is evident that these five points are near the expected mean for this collection of points on the sphere. Correspondingly, our algorithm produces the image atlas using only images taken from these five viewpoints.

3.3. Experimental Results Using OASIS

In the third experiment, we apply the proposed algorithm to construct atlases from collections of MRI images. Specifically, the image data used in the experiment are the MRI images from the freely available Open Access Series of Imaging Studies (OASIS) [11]. OASIS contains MRI brain images from a cross-sectional population of 416 subjects. Each brain image has a resolution of $176 \times 208 \times 176$ voxels. The ages of the subjects range from 18 to 96, and it is known that the structural difference in the brain across different age groups can be significant [15]. Therefore, we divide the 416 subjects from OASIS dataset into three groups: youth group (subjects younger than 40), middle-aged group (between 40 and 60) and elderly group (older than 60), and one atlas is computed for each age group. Figure 10 displays sample images from the three age groups, and the computed atlases (using $K = 12$ nearest neighbors) are shown in Figure 9. For comparison, we use the EM-based iCluster algorithm proposed in [21]. iCluster can automatically cluster a set of images and compute the atlas (template) for each cluster. Applied to the OASIS dataset [21], iCluster clusters the images into three clusters, which differ slightly from the three age groups defined above. However, the two atlases visually look similar and comparable, especially for the youth group since the composition of the youth group in [21] includes 201 subjects (aged 31.2 ± 14.5 years), which is very similar to ours. As shown in Figure 9, the atlases computed by our method are substantially sharper with more structural details. Again, this is not surprising because our atlases are computed as means on the manifolds, instead of the means in the ambient space. Thus, the atlases computed by our method retain important brain structures and they genuinely look like real brains when compared with the real samples in Figure 10. In particular, our method keeps these structures mostly intact and clear while other existing methods (e.g., [21,7]) blur them to various degrees. For applications such as atlas-based segmentation and registration, prominent structures such as the ventricles are potentially very useful, and therefore, our atlases with more contrasting details should provide better templates for these applications.

To model the variation within each age group with respect to the atlas, we can use a field of spherical functions that encodes the point-wise deformation. More specifically, we use pairwise nonrigid registration to align each subject in a given age group to its atlas. For each voxel (point) x , these registrations provide a set of displacement vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ that characterizes the local variation relative to the atlas at x . These displacement vectors can then be interpolated using a non-negative spherical function $f: \mathbf{S}^2 \rightarrow \mathbb{R}$. This results in a field of non-negative spherical functions, and together with the atlas, they provide an effective characterization of the variation in MRI brain images within a given population group. Figure 11 displays a visualization of the field of non-negative spherical functions using a method similar to [23]. Each spherical function is computed from 195 local displacement vectors, and its graph on the sphere is plotted at each point.

4. Conclusion

We have proposed a novel approach to the atlas construction problem, an important problem in medical imaging as well as computer vision. We assume that modulo image rotations, the

input images $\mathbf{I}_1, \dots, \mathbf{I}_k$ belong to a smooth manifold \mathbf{M} , and the atlas computed by the proposed algorithm in principle can be considered as an approximation of the mean of $\mathbf{I}_1, \dots, \mathbf{I}_k$ with respect to the metric on \mathbf{M} . We proposed a two-step algorithm that first locates the mean among the input images and then realizes the atlas image using a few of its neighbors. We have provided several experimental results that validate the proposed algorithm, and in particular, comparisons with existing methods have shown that the atlases computed using our algorithm enjoy better image quality with preservation of important image features and details.

Acknowledgments

This research is partially supported by NSF grant IIS 0916001 and NIH grant NS046812 to BCV.

References

1. Boyd, A.; Vandenberghe, L. *Convex Optimization*. Cambridge University Press; 2004.
2. Davis B, Fletcher P, Bullitt E, Joshi S. Population shape regression from random design data. *IEEE International Conference on Computer Vision (ICCV)*. 2007
3. Dupuis P, Grenander U. Variational problems on flows of diffeomorphisms for image matching. *Quarterly of applied Mathematics*. 1998; 2:587–600.
4. Fletcher P, Venkatasubramanian S, Joshi S. Robust statistics on riemannian manifold via the geometric median. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2008
5. Goh A, Vidal R. Clustering and dimensionality reduction on riemannian manifolds. *Proc IEEE Conf on Comp Vision and Patt Recog*. 2008
6. Guimond, A.; Meunier, F.; Thirion, J. Technical Report, 3731, INRIA. 1999. Average brain models: A convergence study.
7. Joshi S, Davis B, Jomier M, Gerig G. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*. 2004; 23:S151–S160. [PubMed: 15501084]
8. Tenenbaum J, de Silva V, Langford J. A global geometric framework for nonlinear dimensionality reduction. *Science*. 2000; 290(5500):2319–2323. [PubMed: 11125149]
9. Karcher H. Riemannian centers of mass and mollifier smoothing. *Comm on Pure and Appl Math*. 1977; 30:509–541.
10. Lorenzen P, Prastawa M, Davis B, Gerig G, Bullitt E, Joshi S. Multi-modal image set registration and atlas formation. *Medical Image Analysis MEDIA*. 2006; 10(3):440–451.
11. Marcus D, Wang T, Parker J, Gsernansky J, Morris J, Buckner R. Open access series of imaging studies (oasis): Cross-section mri data in young, middle aged, nondemented, and demented older adults. *Journal of Cognitive Neuroscience*. 2007; 19:1498–1507. [PubMed: 17714011]
12. Miller M, Younes L. Group actions, homeomorphisms and matching: a general framework. *International Journal of Computer Vision*. 2001; 41:61–84.
13. Miller M, Younes L. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int J Computer Vision*. 2005; 51(2)
14. Moakher M. A differential geometry approach to the geometric mean of symmetric positive-definite matrices. *SIAM J MATRIX ANAL APPL*. 2005; 26:735–747.
15. Mortamet B, Zeng D, Gerig G, Prastawa M, Bullitt E. Effects of healthy aging measured by intracranial compartment volumes using a designed mr brain database. *MICCAI*. 2005
16. Mumford, D.; Fogarty, J.; Kirwan, F. *Geometric Invariant Theory*. Springer; 1994.
17. Munkres, J. *Topology*. Prentice Hall; 2000.
18. Nene, S.; Nayar, S.; Murase, H. Technical Report 005-96, CUCS. February. 1996 Columbia object image library (coil-20).
19. Murase H, Nayar S. Visual learning and recognition of 3-d objects from appearance. *IJCV*. 1995; 14(1):5–24.

20. Pennec X, Fillard P, Ayache N. A riemannian framework for tensor computing. *Int J Computer Vision*. 2006; 66(1):41–46.
21. Sabuncu M, Balci S, Golland P. Discovering modes of an image population through mixture modeling. *MICCAI*. 2008
22. Tyagi A, Davis JW. A recursive filter for linear systems on riemannian manifolds. *Proc IEEE Conf on Comp Vision and Patt Recog*. 2008
23. Barmoutis A, Vemuri BC, Howland D, Forder JR. Extracting Tractosemas from a displacement probability field for tractography in DW-MRI. *MICCAI*. 2008




Figure 1.
Left: Images of an object. **Center:** The atlas computed using [7]. **Right:** The atlas computed using the proposed algorithm.




Figure 2. Ten points on an one-dimensional manifold. **Left** Mean for the ten points computed using the metric of the ambient space. **Right** Mean computed using the metric of the submanifold \mathbf{M} . Note that the mean can be approximated by using just two points on \mathbf{M} , the two points that are closest to the mean on \mathbf{M} .




Figure 3.

Top Row Images of a car. The smooth variation of the images allows the assumption that these images belong to a smooth submanifold \mathcal{M} in \mathcal{I} . **Bottom Row** Rotations are applied to the images in the row above. The appearance variation in this collection is erratic, and it is difficult to assume that they belong to a smooth submanifold in \mathcal{I} .




Figure 4. Image atlases for five objects from COIL-20. The images in the **first row** are the atlases computed using [7]. Our results are shown in the **second row**.




Figure 5.
Samples of input duck images.




Figure 6. **First Row:** L^2 -mean images obtained using [7] (left) and our algorithm (middle). The image on the right is the ground truth. **Second Row:** L^1 -median images obtained using [4] (left) and our algorithm (middle). The right image is the atlas obtained by our algorithm using the same set of images that are randomly rotated between 20° and 30° .




Figure 7.
Left to right straightforward averaging, unbiased diffeomorphic atlas and our atlas.




Figure 8. 56 Camera locations sampled on the viewing sphere. The camera locations of the five nearest neighbors to the mean determined by our method are shown by the red points.




Figure 9. Axial views of atlases. **Top Row:** Atlases generated using our method for youth, middle-aged and elderly groups (from left to right). **Bottom Row:** Youth, older middle-aged and elderly templates generated by iCluster [21]. The images shown here are taken directly from [21].




Figure 10. Sample images from from the three age groups. From left to right: youth, middle-aged and elderly.




Figure 11.

Left: Visualization of the field of non-negative spherical functions. The atlas is from the elderly group with 195 subjects. The spherical functions are represented by their graphs on the sphere: the value $f(\mathbf{u})$ at $\mathbf{u} \in \mathbf{S}^2$ gives the height of f at the point \mathbf{u} , and the resulting height surface over the sphere is shown. **Right:** A close-up view of the field.