Volume-Preserving Mapping and Registration for Collective Data Visualization

Jiaxi Hu, Guangyu Jeff Zou and Jing Hua

Abstract—In order to visualize and analyze complex collective data, complicated geometric structure of each data is desired to be mapped onto a canonical domain to enable map-based visual exploration. This paper proposes a novel volume-preserving mapping and registration method which facilitates effective collective data visualization. Given two 3-manifolds with the same topology, there exists a mapping between them to preserve each local volume element. Starting from an initial mapping, a volume restoring diffeomorphic flow is constructed as a compressible flow based on the volume forms at the manifold. Such a flow yields equality of each local volume element between the original manifold and the target at its final state. Furthermore, the salient features can be used to register the manifold to a reference template by an incompressible flow guided by a divergence-free vector field within the manifold. The process can retain the equality of local volume elements while registering the manifold to a template at the same time. An efficient and practical algorithm is also presented to generate a volume-preserving mapping and a salient feature registration on discrete 3D volumes which are represented with tetrahedral meshes embedded in 3D space. This method can be applied to comparative analysis and visualization of volumetric medical imaging data across subjects. We demonstrate an example application in multimodal neuroimaging data analysis and collective data visualization.

Index Terms—Volume-preserving mapping, data regularization, data transformation

1 Introduction

Three-dimensional imaging data is accumulated at a faster than ever speed in science, engineering, biomedicine, and real-world sensing. The complexity of the acquired digital data usually overwhelms the useful information and makes it extremely difficult to derive true understanding from it. Especially for scientific imaging data, it is typically collected from individual subjects, and many consist of sophisticated 3D spatial structures and other high-dimensional, heterogeneous features. In many scientific areas, such as biomedical imaging, material science and engineering, geology, physical simulations, etc., assessment of similarity and disparity from the heterogeneous 3D imaging data across subjects plays a central role in modern scientific discovery. This kind of tasks poses increasing challenges in data modeling and visualization. Comparative analysis and visualization on a large number of cross-subject volume data is even more complicated, which we refer as *collective data visualization*. For collective data visualization, effectively and efficiently identifying the important interplay between data samples, data types and features in the collective data is of utter most importance. Apparently, intrinsic geometric structures and shapes extracted from 3D imaging of real-world objects are very important in linking and indexing individual objects for further interpretation and collective data visualization. In order to visualize and analyze such complex collective data, complicated geometric structure of each data is desired to be mapped onto a canonical domain to enable a map-based visual exploration. Instead of working on the complicated geometric structure directly, all the operations are on simpler domains, such as planes and structured grid volumes. The mapping will help improving the efficiency of the data indexing and processing.

However, a manifold can not be mapped to another domain without any distortion. Thus, different mapping methods have been proposed to preserve certain local geometries [7, 16]. In many visualization applications, such as medical data visualization, it is oftentimes important to maintain local volume as well as their correspondence across subjects during the mapping procedure in order to achieve accurate sampling and computation of volumetric measurements.

In this work, we focus on a mapping method which preserves the local volume element. Thus, similar geometric structures across a group of subjects can be mapped onto the same canonical domain for collective data visualization. Note that, the collective data visualization also requires the salient points or regions to be registered and aligned as well. These demands post two challenges for an effective computational and visualization method: **(1)** mapping between two given manifolds with local volume elements preserved and **(2)** registering salient features to achieve an alignment without breaking the shape of the manifold and the volume-preserving constraint.

We present a general method to solve global rigorous volume-preserving mappings among closed connected manifolds with salient feature registration. The proposed techniques are essential for spatially regularizing and indexing cross-subject data for efficient data access, query and collective data visualization. The main contributions of this paper can be summarized as follows:

- **Volume-preserving mapping for general manifolds supported by a rigorous continuous theory**. By introducing fluid dynamics to manifold mapping, the evolution of volume elements can be precisely quantified by the compressible flow. Computing a volume-preserving mapping is therefore equivalent to deriving a time-dependent vector field that corrects volume distortions, i.e., dilation/shrinkage induced by an initial mapping, through a diffeomorphic flow.

- **Salient feature registration without breaking the shape of manifold or the volume-preserving constraint**. With a computed volume-preserving mapping, a registration process can be derived by constructing incompressible flow. Vertices are moving along such a field, which will not break the volume-preserving constraint.

- **Supporting for map-based collective data visualization and extensibility to 2D or higher dimensions**. The volume-preserving mapping and registration do not focus on 3D cases only. The theories and algorithms are general to all dimensional cases. It can facilitate effective map-based collective data visualization.

The remainder of the paper is organized as follows: Section 2 reviews the related work. The volume-preserving mapping is presented

• *Jiaxi Hu, Guangyu Jeff Zou and Jing Hua are with the Department of Computer Science, Wayne State University, Detroit, MI, 48202. E-mail: {jiaxihu,gzou,jinghua}@wayne.edu*

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014. Date of publication 11 Aug. 2014; date of current version 9 Nov. 2014.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.
Digital Object Identifier 10.1109/TVCG.2014.2346457

in Section 3. The discrete algorithm to map a tetrahedral mesh is introduced in Section 4. After that, a volume-preserving registration follows in Section 5. The performance assessment as well as a demonstrative application to the manifold-based multimodality analytics of brain imaging data across subjects is given in Section 6. Finally, in Section 7, we summarize our work and conclude with a discussion about the limitation, possible improvements as well as a few potential future directions of this method.

2 Related Work

Manifold mapping is a powerful tool to reduce the complexity of arbitrary manifolds onto canonical domains such as unit cube or sphere [9, 11]. Mapping general manifolds between each other will introduce different kinds of distortions. In geometry processing, reducing volume distortion is a major interest. There are many approaches for reducing the volume distortions of the mappings. One mapping approach is called conformal method, which always provides a unique mapping preserving local angle geometries. While preserving the angles, the conformal mapping introduces large volume distortions. In order to reduce the volume distortions, cuts are employed to improve the final results. The local volume elements are still not exactly preserved due to the nature of the conformal mapping. Additionally, cuts also break the continuous structures of the original manifold. Jin et al. [12] proposed to search for the optimal global conformal parameterization in the space of Möbius transformations. As conformality is invariant through Möbius transformations, the resulting parameterization is still conformal, which cannot completely eliminate area distortions. Based on the fact that surface parameterization is closely related to the topological constraints, delicate topological modifications can sometimes improve the uniformity of the parameterization. Towards this end, Gu and Yau punctured small holes at the tip of long appendages [8]. Cone singularities were introduced with non-vanishing Gaussian curvature in [14, 3]. Surface cuts were repeatedly augmented according to the geometric stretches generated through the course of tentative parameterizations [7]. Although such schemes provide excellent remedies to some applications, they are generally not acceptable to the rest. In terms of volume preservation, these methods can be considered as heuristic approaches. Rigorous volume-preserving mapping is in general not achievable along this direction.

Another major branch is categorized as functional methods. These methods derive a mapping by optimizing some predefined penalty functions. The penalties can be defined with the volume distortions so that the optimization leads to a mapping minimizing such volume distortions. This kind of approaches rely on different parameters and halting criteria in the optimization procedures and oftentimes suffer from local optimum problem. In general, no quantities are exactly preserved by the methods of this type. The functional mapping methods offer the volume distortion minimization instead of exact volume-preserving.

Similarly, to prevent severe geometric stretches, the original model can be decomposed into many charts which are topologically equal to an open disk to make interior Gaussian curvature close to zero [16, 25, 27, 31]. Even distribution of intrinsic distortions [24, 27, 12, 31] are studied to achieve uniformity metrics mapping. Based on the singular values of the Jacobian matrix, Sander et al. [24] optimized the parametric location of each vertex within its 1-ring neighborhood to reduce local stretches. Sorkine et al.'s bounded-distortion mapping [27] made heavy use of mesh cuts to keep distortions below some preset threshold. Zhang et al. [31] identified an anisotropic stretch term from a 2×2 tensor metric closely related to the one derived in [24] and applied it to guide the vertex optimization. They considered 2D area distortion with an energy term. The purpose is to minimize the 2D area distortion. No conclusion is drawn whether an absolute area-preserving patch mapping can be practically achieved. Kraevoy and Sheffer also presented a "low distortion" optimization for cross-parameterization and remeshing of 3D models [15]. Desbrun et al. [4] also studied minimizing area distortion for the intrinsic mapping of triangle meshes. In their work, an intuitive area-preserving functional was devised. Since the functional tried to preserve the area structure of the original 1-ring, the optimal mapping derived was termed "Discrete Authalic Parameterization (DAP)". Due to the fact that the parameterization is only a critical point to the functional, it cannot exactly preserve areas across the mesh.

Approaches in [1] backed by a general result from [23] guarantee the existence of an area-preserving diffeomorphism between two surfaces with the same total surface area. However, the discussion was restricted to a spherical domain. Based on the Monge-Kantorovich theory of *optimal mass transport*, Haker et al. [10] developed an image registration and warping technique. Lying at the core is a decomposition of the deformation into a divergence-free vector field plus a curl-free one, called the *polar factorization*. A mass (e.g., area/volume) preserving mapping was explicitly sought in $\mathbf{R}^n$ by a gradient descent method to the Monge-Kantorovich function. Their method has been used for area-preserving corrections on top of conformally flattened vessel surfaces in $\mathbf{R}^2$, implemented on a regular Cartesian grid [32]. Recently, the technique of optimal mass transport has also been applied to texture mapping of closed genus zero surfaces [5]. Along this direction, Zou et al. [33] presented a practical method to compute a group of global 2D area-preserving mapping mathematically with *Lie advection*, a frequently appearing concept in classical mechanics. However, it is not clear that how it can be applied to higher dimensions in theory and in discrete implementation. There also exist harmonic volumetric mappings among objects [18]. Li et al. [19] aligned those harmonic mappings with salient features. Wan et al. [30] registered objects with polycube mapping.

Note that, the previous volume-preserving mapping techniques focused only on preserving volume/area elements on the manifold. The registration capability is not provided in any of the work mentioned above. Funck et al. [29] discussed vector-field based deformation algorithm for volume-preserving shape editing though. In theory, their algorithm has its root in 2D domain. The vector field kernel has only two dimensions. The 3D applications are simple extensions by either rotating or translating such a 2D kernel. On the contrary, HAMMER algorithm [26], volume flow registration [13], the Large Deformation Diffeomorphic Metric Mapping (LDDMM) [2], and LogDemons [28] methods can register the feature constraints without breaking local structures, but volume distortions are inevitable. Mansi et al. proposed iLogDemons [21] to register volumes with divergence-free whirls, which preserve the interior local volume, but the final boundary shape (i.e., the surface) is not considered. Achieving registration while still preserving volume/area and the target boundary shape is extremely important. It will facilitate volume/area preserving data analysis across aligned subjects.

Compared to exiting literature, particularly [33], this paper discusses a more general manifold mapping in R^n, in both theory and discrete algorithms for different manifold representations, especially for 3-manifolds. It can map a shape to any desired target domain with exact local volume-preserving constraint for both interior and the boundary of the volume. More importantly, volume-preserving registration is also introduced, based on incompressible flow theory. If a manifold flows under a divergence-free velocity field, the local volume element will not change. This paper proves the existence of the solution by constructing such a field that the divergence is guaranteed to be zero and help salient features register to their target positions.

3 Volume-Preserving Mapping

This section presents a rigorous volume-preserving manifold mapping, which preserves the local volume element among manifolds. Given two connected manifolds, there are diffeomorphisms between them. Different mappings have different purposes. They either preserve local geometry or minimize some deforming energy. This section shows a solution to find a mapping to preserve the local volume.

Consider an n dimension manifold P, a local volume element is defined on the local coordinates $x_1, x_2, \ldots, x_n$ as

$$\tau_P = \rho_P(\mathbf{x}) d\mathbf{x}, \tag{1}$$

where $d\mathbf{x} = dx_1 \wedge dx_2 \wedge \ldots \wedge dx_n$ is the wedge product along all the axes in the n dimensional space, representing the local volume element, and

$\rho_P(\mathbf{x}) > 0$ is a volume density function. There exists a diffeomorphism f which maps P to another manifold Q, $f : P \to Q$. Without loss of generality, the global volume of Q can be scaled to that of P, which means $\int_P \tau_P = \int_Q \tau_Q$. After the mapping, the local volume element τ_P deforms to τ_Q. Suppose there is another diffeomorphism g which maps Q to itself, $g : Q \to Q$. g also maps the local volume element τ_Q to τ'_Q. If $\tau_P \equiv \tau'_Q$ everywhere, $g \circ f$ is such a diffeomorphism, which maps P to Q preserving the local volume.

According to Eq. (1), $\tau_P \equiv \tau'_Q$ can be converted to $\rho'_Q \equiv \rho_P$. g not only maps Q to itself, but also maps the volume density back to the original one in P. This density change in Q is described as fluid dynamics. Suppose Q follows a velocity field V defined at itself. The material derivative can be derived from V

$$\frac{D\rho}{Dt} = -\rho \nabla \cdot V. \tag{2}$$

The material derivative describes the density changes along the stream line in V. Note that, each element moves along the stream line in V, so that the local density ρ_Q changes to ρ_P. This change can be interpolated linearly as

$$\rho(t) = (1-t)\rho_Q + t\rho_P, \quad t \in [0,1]. \tag{3}$$

Thus, V must satisfy

$$-\rho_Q + \rho_P = -\rho \nabla \cdot V. \tag{4}$$

In order to solve this equation of V, we introduce a scalar field $\varphi(x,t)$ such that

$$V = \nabla \varphi. \tag{5}$$

Then, Eq. (4) is converted into a Poisson equation of $\varphi(x,t)$

$$-\rho_Q + \rho_P = -\rho \Delta \varphi. \tag{6}$$

This Poisson equation has a unique solution. Once the scalar field is solved, the path of each element on Q is an integral curve $\mathbf{x}(t)$ such that

$$\begin{cases} \frac{d\mathbf{x}(t)}{dt} = V(\mathbf{x}(t),t) \\ \mathbf{x}(0) = \mathbf{x}_0, \end{cases} \tag{7}$$

where $\mathbf{x}_0$ is the initial position of the curve. If each element follows the curve $\mathbf{x}(t)$ from $t = 0$ to $t = 1$, the density ρ on it changes linearly from ρ_Q to ρ_P. Therefore, the curve $\mathbf{x}(t)$ yields a desired diffeomorphism g, and $g \circ f : P \to Q$ is a diffeomorphism which maps manifold P to Q with local volume preserved.

4 Discrete Volume-Preserving Mapping Algorithm

The previous section showed that with a given diffeomorphism $f : P \to Q$ there also exits a diffeomorphism from P to Q preserving each local volume element. This section will show how to apply it to the discrete data. We will develop an iterative algorithm handling those manifolds with discrete representations. Manifolds in 3D can be represented with tetrahedral volume meshes. Note that, this algorithm can also be extended to 2D or higher dimensions. The volume-preserving mapping algorithm can be described as following. First, an initial mapping is constructed, such as a discrete harmonic mapping. Then, the time interval is divided into the discrete time steps for the iterations. At the beginning of each iteration, the current local density is calculated on each vertex. The Poisson equation, Eq. (6), is solved within current time step. All the vertices move a time step, along the vector field V guided by the solved scalar field. The iteration continues for each time step.

Assume a 3D volume is described with a tetrahedral mesh $T = \{V,E,F,C\}$. $V = \{v_i\}$ denotes the vertex set, $E = \{e_{ij}\}$ the edge set, $F = \{f_{ijk}\}$ the face set, and $C = \{c_{ijkl}\}$ the tetrahedral cell set with $1 \le i,j,k,l \le m = |V|$. Given a scalar function φ on the volume, its discrete version is a vector $\vec{\varphi} = [\varphi_1, \varphi_2, ..., \varphi_m]^T$, defined at the vertices. The volume element is represented with the voronoi volume $|\Omega_i|$ of that vertex.

Based on this representation, some basic operations have to be defined in order to construct such a volume-preserving mapping, which include Poisson equation, gradient operator, curve integration. These key discrete operations are described in the following subsections.

4.1 Initial Mapping

We assume the target canonical domain is a unit sphere. There exists a mapping between any two manifolds if both of them can be mapped onto a unit sphere with registration. We use an initial mapping as a starting point of the volume-preserving mapping. In practice, we have chosen to use harmonic mapping as the initialization of our system because of its robustness. The initialization has two steps: First, mapping the boundary of the volume to the unit sphere surface, which can be achieved by applying an area-preserving mapping from a close surface to the unit sphere [33]; Second, applying volumetric harmonic mapping on the interior [17].

4.2 Discrete Poisson Equation and Gradient Operation

The discrete Laplace-Beltrami operator is linearly approximated at each vertex. Suppose edge e_{ij} is shared by n tetrahedrons, the volumetric $\Delta\varphi$ is estimated at v_i as [22]

$$\Delta\varphi(v_i) = \sum_{j \in N(i)} w_{ij}(\varphi(v_i) - \varphi(v_j)), \tag{8}$$

where w_{ij} is the volumetric edge weight defined on edge e_{ij} and $N(i)$ is the one-ring neighbour of vertex i. As the scalar function φ is represented with a vector, the Laplace operator Δ can be represented with a matrix. The discrete Laplace operator is defined with a weight matrix as

$$L_{ij} = \begin{cases} \sum_{k \in N_1(i)} w_{ij} & \text{if } i = j \\ -w_{ij} & \text{if } e_{ij} \in E \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

It is easy to verify that the matrix L satisfies Eq. 8. Suppose an edge e_{ij} is shared by n tetrahedrons, the edge weight w_{ij} is defined as [6]

$$w_{ij} = \frac{1}{|\Omega_i|}\left(\frac{1}{6}\sum_{q=1}^{n} l_q \cot(\theta_q)\right), \tag{10}$$

where $|\Omega_i|$ is the vonoroi volume of v_i, $l_q = l_{kl}$ is the length of the edge e_{kl} to the opposite edge e_{ij}, and $\theta_q = \theta_{kl}$ is the dihedral angle on the edge e_{kl}. When considering all vertices of a mesh, Eq. (8) can be written as a linear system:

$$Lx = b, \tag{11}$$

where $x = \vec{\varphi}$ and $b = (\rho_Q - \rho_P)/\rho$.

The matrix L is sparse. For this reason, Eq. (11) can be solved efficiently in linear time, e.g., using the preconditioned bi-conjugate gradient method. In practice, we have observed that solving the normal equations $L^T Lx = L^T b$ yields more robust numerical behavior than directly solving Eq. (11).

By solving the discrete Poisson equation, we obtain the scalar field φ on the unit sphere volume. The gradient of φ is computed on the tetrahedral mesh. We consider a tetrahedral cell c_{ijkl} with its four corners lying at $\mathbf{v}_i$, $\mathbf{v}_j$, $\mathbf{v}_k$, and $\mathbf{v}_l$ in $\mathbf{R}^3$. The gradient is assumed constant inside the tetrahedral cell. The gradient vector can be easily computed by solving a 3×3 linear system:

$$\begin{bmatrix} \mathbf{v}_j - \mathbf{v}_i \\ \mathbf{v}_k - \mathbf{v}_i \\ \mathbf{v}_l - \mathbf{v}_i \end{bmatrix} \nabla\varphi = \begin{bmatrix} \varphi_j - \varphi_i \\ \varphi_k - \varphi_i \\ \varphi_l - \varphi_i \end{bmatrix}, \tag{12}$$

for which an analytic solution exists. To obtain a unique vector at each vertex, $\nabla\varphi$ at vertex v_i is defined as

$$\nabla\varphi_i = \frac{1}{\sum_{c_{ijkl} \in N_1(i)} \alpha^i_{jkl}} \sum_{c_{ijkl} \in N_1(i)} \alpha^i_{jkl} \nabla\varphi(c_{ijkl}), \tag{13}$$

that is, an average of the gradients of the adjacent cells, weighted by the inverse of the distance α^i_{jkl} between the central vertex $\mathbf{v}_i$ and the centroid of cell c_{ijkl}.

4.3 Pathline Integration

$V(t)$ is not a static field but varies with time t depending on the solution of $\nabla\varphi(t)$. Therefore, an iterative integration is employed based on each vertex. The time interval, $[0,1]$, is divided into K steps, and the time difference between the neighbor steps is $\delta t = 1/K$. Suppose the volume is uniformly distributed on the original manifold P, the volume density of each vertex ρ_{Pi} is 1 at all vertices. On Q, the density ρ_{Qi} equals $|\Omega_{Pi}|/|\Omega_{Qi}|$, the ratio of the voronoi volumes of that vertex on P and Q respectively. ρ_P and ρ_Q are constants through the iterative integration. Similarly, the density of the kth step $\rho_i(k)$ equals $|\Omega_{Pi}|/|\Omega_i(k)|$, where $|\Omega_i(k)|$ is the current voronoi volume. The initial state of the diffeomorphism g is an identity mapping, and each pathline $\mathbf{X}$ starts from the original place as $\mathbf{X}(0) = \mathbf{x}_0$. At the beginning of each step k, the Laplacian matrix L is calculated based the geometry of the current tetrahedral mesh, while b is based on the volume densities. $\varphi(k)$ is then solved with Eq. (11). The velocity field $V(k)$ is represented with $\nabla\varphi$ according to Eq. (13). The time integration $V(k)\delta t$ is applied to the current position $\mathbf{X}(k)$ in order to get a new $\mathbf{X}(k)$. Thus, we get the iteration

$$\mathbf{X}(k) = \mathbf{X}(k-1) + \frac{1}{K}V(k-1), \quad k = 1,2,...,K. \tag{14}$$

And the final diffeomophism g is represented with the sum of the displacements of all steps, as

$$g(\mathbf{x}) = \mathbf{x}_0 + \sum_{k=1}^{K} \frac{1}{K}\nabla\varphi(k-1). \tag{15}$$

In general, a larger number of iteration steps result in more accurate approximations. Empirically, K is set to 50, which satisfies most cases.

The movement of vertices is in principle determined by the designated volume changes. However, degenerate tetrahedrons may undermine the discrete computation due to inaccurate approximation of the Laplace-Beltrami operator. At the beginning of each step, we optimize the underlying tetrahedralization by performing local repairs to the combinatorial topology for the constrained Denaulay condition. Note that, the geometry of volumes is realized by their $\mathbf{R}^3$ embedding. Throughout this procedure, the discrete sampling of the shape remains unchanged, but is merely interpolated by a different tetrahedralization.

5 Volume-Preserving Registration

The previous sections showed an approach to find a volume-preserving mapping guided by the compressible flow of the volume forms. This section provides a solution to find a registration which can move feature points to their desired positions without breaking the original volume-preserving constraint. The problem can be defined as deforming the volume with a velocity field, such that all the salient feature points will be registered to predefined target positions and unconstrained part will follow to reach an alignment. Suppose there are N salient feature points $\{p_i\}$ and their target positions are $\{p_i'\}$. During the deformation, each point moves along its path $\gamma_i(t)$, such that $\gamma_i(0) = p_i$ and $\gamma_i(1) = p_i'$. The simplest path could be $\gamma_i(t) = p_i + (p_i' - p_i)t$, a straight line linking p_i and p_i' with constant speed. These paths give the solution to the registration of the finite feature points, but breaking the volume-preservation. In order to obtain volume-preserving registration, a velocity field $\mathbf{s}$ is required to deform the whole volume domain. To accomplish the volume registration, $\mathbf{s}$ must meet three criteria: First, each path γ_i is a pathline in $\mathbf{s}$; Second, $\mathbf{s}$ is divergence-free everywhere so that the volume elements are preserved under the deformation; Third, $\mathbf{s}$ on the boundary has no normal components, where no element moves in or out the original volume domain. Formally, the problem can be represented as finding a velocity field $\mathbf{s}$ defined on the volume domain M, such that

$$\begin{cases} \mathbf{s}(\gamma_i(t_0)) = \frac{d\gamma_i(t)}{dt}\Big|_{t=t_0}, \quad t_0 \in [0,1] \\ \nabla\cdot\mathbf{s} \equiv 0 \\ \mathbf{s}(x)\cdot\mathbf{n}(x) \equiv 0\Big|_{x\in\partial M}, \end{cases} \tag{16}$$

where $\mathbf{n}$ is the unit normal vector on the boundary ∂M.

To solve such an $\mathbf{s}$, a discretization of time and space is employed. Each pathline γ_i is represented as a sequence of path points from p_i to p_i' with $K+1$ samples, such as $\{\gamma_i(\frac{k-1}{K})\}, k = 1,2,...,K+1$. The boundary is represented with a set of evenly sampled points $\{b_j\}$ on it. Thus, the original problem of $\mathbf{s}$ is converted into

$$\begin{cases} \mathbf{s}(\gamma_i(\frac{k-1}{K})) = \gamma_i(\frac{k}{K}) - \gamma_i(\frac{k-1}{K}), \quad k = 1,2,...,K \\ \nabla\cdot\mathbf{s} \equiv 0 \\ \mathbf{s}(b_j)\cdot\mathbf{n}(b_j) \equiv 0, \quad b_j \in \partial M. \end{cases} \tag{17}$$

The algorithm is based on the discretization above with the following steps: First, a rigid registration is applied; Second, each pathline is generated; Then, a divergence-free interpolation is employed; Finally, the boundary constraint is enforced. The algorithm mainly focuses on the 3D sphere domain, but there is no difficulty in extending it to a lower or higher dimensional space with arbitrary shapes.

5.1 Rigid Registration

Although rigid registration is not required in the three criteria of $\mathbf{s}$, shorter paths are still desired to keep the local structure of the volume. The rigid registration can be considered as preprocessing. The rigid registration is a rotation matrix R such that

$$\arg\min_R \sum_{i=1}^{N} \| Rp_i - p_i' \| . \tag{18}$$

This problem can be convert into

$$\arg\max_R \sum_{i=1}^{N} p_i' p_i^T \cdot R, \tag{19}$$

which can be solved with Singular Value Decomposition (SVD) of the matrix $\sum_{i=1}^{N} p_i' p_i^T$. Once R is determined, the whole volume is rotated by R and $\{p_i\}$ is replaced by $\{Rp_i\}$.

5.2 Pathline Generation

In order to register p_i to p_i', the pathline is not unique. Constant straight line path is an obvious choice. In our applications, neuroimaging data is often registered within a 3D unit sphere domain. Some feature points can be near the sphere boundary or on the boundary. In these cases, simple straight line interpolation may destroy the local structure too much on the radial directions. To make the path as natural as possible near the boundary, the point moves with constant angular and radial speed towards to its target position. Such a radius path will preserve radial local structure near the boundary. It can be interpolated recursively with vector calculations. Consider the point pair p_1 to p_2 are two vectors in space. Neither of them is the origin. u_1 and u_2 are their unit and direction vectors, and $u_1 \neq -u_2$. Thus, the middle point p_m on the radius path from p_1 to p_2 must have its direction on $(u_1+u_2)/2$ and norm as $(\| p_1 \| + \| p_2 \|)/2$. Therefore, p_m of given p_1 and p_2 is represented as

$$p_m = \begin{cases} \frac{\|p_1\|+\|p_2\|}{4}(u_1+u_2) & \text{if } \| p1 \| > 0, \| p_2 \| > 0 \\ & \text{and } u_1 \neq -u_2 \\ \frac{p_1+p_2}{2} & \text{otherwise.} \end{cases} \tag{20}$$

Once any start and end points are given, the middle point can be interpolated. This means, the whole radius path can be interpolated recursively from start point p_i to the end points p_i' up to a desired resolution. In Fig. 1, the green pathline with circle marks is a radius path while the red one with triangle marks is the straight path. The feature point moves more naturally near the boundary along the radius path. In the more interior region, especially in the case that the origin is nearly in between of p_i and p_i', radius path is not that necessary. The straight path is more appropriate instead. Radius and straight path work better in their own regions. A global solution is to combine them linearly. In Fig. 1, the blue path with square marks is the combination of the radius and straight ones with the same weight. As discussed

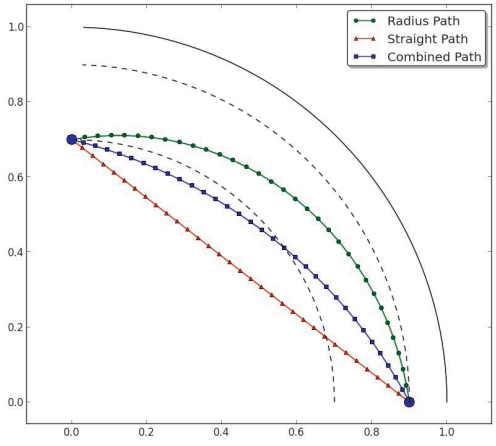

Fig. 1. Pathline interpolation between two points. The path with triangle markers is a straight line, the one with circle makers is a path with radius interpolation, while the one with square markers is a path combined with both straight line and radius interpolation.

above, given a feature point p_i and its target p_i', the radius and straight pathlines, γ_{ri} and γ_{si}, can be interpolated. Then, the final pathline γ_i is the weighted linear combination of these two pathlines, as

$$\gamma_i = (1-w)\gamma_{si} + w\gamma_{ri}, \tag{21}$$

where w is the combination weight and $w = (\| p_i \| + \| p_i' \|)/2$.

5.3 Divergence-Free Interpolation

The finite sample points are generated for N feature points p_i and K steps on each, which are represented with $\{\mathbf{x}_i, \mathbf{d}_i\}_{i=1}^{N\times K}$. Each of them moves with a constant speed, from its original location to the one to be registered. Based on this sample set, a velocity field $\mathbf{s}$ can be interpolated. Then, the whole volume is registered, following the registration of feature points. As aforementioned, $\mathbf{s}$ has to be divergence-free, which means $\nabla \cdot \mathbf{s} \equiv 0$, so that the volume-preserving is guaranteed during the whole registration procedure. Radial Basis Function (RBF) interpolation is powerful tool to generate a function which interpolates a discrete set of values. In our cases, the set of values is the displacement vector $\{\mathbf{d}_i\}$ for each constant time step. Therefore, instead of regular scalar RBF, vector-matrix RBF is employed, which is represented as

$$\mathbf{s}(\mathbf{x}) = \sum_{i=1}^{N\times K} \Phi(\| \mathbf{x} - \mathbf{x}_i \|)\mathbf{c}_i, \tag{22}$$

where Φ is an $N \times N$ matrix valued radial basis function in n dimension space, and $\mathbf{c_i}$ is a coefficient vector. The set of $\{\mathbf{c}_i\}_{i=1}^{N\times K}$ satisfies

$$\mathbf{s}(\mathbf{x}_i) = \mathbf{d}_i. \tag{23}$$

To interpolate a divergence-free field, we employ the radial basis function [20],

$$\Phi(\mathbf{x}) = (\nabla\nabla^T - \mathbf{I}\Delta)\psi(\mathbf{x}), \tag{24}$$

where the ∇ operator is represented with a vector operator $[\partial_{x_1}, \partial_{x_2}, ...\partial_{x_n}]^T$, Δ is a scalar valued Laplace operator, and ψ is a scalar valued radial basis function. To calculate the divergence of each column of Φ, ∇^T operator can be multiplied from the left. Then, there is

$$\begin{aligned} \nabla^T\Phi(\mathbf{x}) &= \nabla^T(\nabla\nabla^T - \mathbf{I}\Delta)\psi(\mathbf{x}) \\ &= ((\nabla^T\nabla)\nabla^T - \nabla^T\mathbf{I}\Delta)\psi(\mathbf{x}) \\ &= (\Delta\nabla^T - \nabla^T\Delta)\psi(\mathbf{x}) \\ &= \mathbf{0}_{1\times n}\psi(\mathbf{x}), \end{aligned} \tag{25}$$

where $\mathbf{0}_{1\times n}$ is a 1 by n zero operator. Eq. (25) verifies that the divergence of each column of Φ is 0. Therefore, $\mathbf{s}$, the finite combination of columns of Φ, is divergence-free. The next step is to solve the coefficient set $\{\mathbf{c}_i\}_{i=1}^{N\times K}$. Similar to a scalar radius basis function problem, the set of $\{\mathbf{c}_i\}_{i=1}^{N\times K}$ satisfies Eq. (23), which means

$$\sum_{i=1}^{N\times K} \Phi(\| \mathbf{x}_j - \mathbf{x}_i \|)\mathbf{c}_i = \mathbf{d}_j, \quad j = 1, 2, ..., N\times K. \tag{26}$$

These $N \times K$ independent linear equation arrays can be combined into one tensor equation:

$$T_\Phi \otimes m_c = m_d, \tag{27}$$

where $\Phi(\| \mathbf{x}_j - \mathbf{x}_i \|)$, $\{\mathbf{c}_i\}$ and $\{\mathbf{d}_j\}$ are filled into tensor T_Φ, matrix m_c, and matrix m_d. With the solution of this equation, $\mathbf{s}$ is interpolated, which is divergence-free and guarantees each sample point is along the pathline. In the experiments of this paper, the case of 3D unit sphere domain is employed. However, it is intuitive to extend to other dimensions. We choose Gaussian, $\psi(\mathbf{x}) = e^{-\lambda\|\mathbf{x}\|^2}$, as the scalar valued radius basis functions. Note that, theoretically, ψ could be other continuous functions. The matrix valued function Φ in 3D space is represented as

$$\Phi(\mathbf{x_0}) = \begin{vmatrix} -\partial_{yy} - \partial_{zz} & \partial_{xy} & \partial_{xz} \\ \partial_{xy} & -\partial_{xx} - \partial_{zz} & \partial_{yz} \\ \partial_{xz} & \partial_{yz} & -\partial_{xx} - \partial_{yy} \end{vmatrix} e^{-\lambda\|\mathbf{x}\|^2}|_{\mathbf{x}=\mathbf{x_0}}. \tag{28}$$

In the space, each column represents a vector field. The vector field of a single radial basis function is represented with the first column in the 3D space which is illustrated in Fig. 2. In this figure, the arrows represent the direction of the field while size and color for the magnitude.

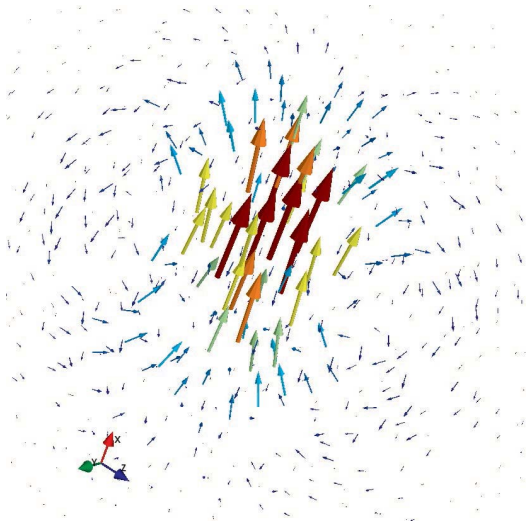

Fig. 2. The 3D vector field represented with the first column of the matrix function in Eq. (24).

5.4 Boundary Constraint

A divergence-free velocity field $\mathbf{s}$ is already obtained based on finite sampling of the registration pathlines, which guarantees volume is preserved during the registration. However, $\mathbf{s}$ on the boundary is not verified yet. If the volume deforms by following $\mathbf{s}$, the shape of the domain may change. For example, it can not be a unit sphere any more after the deformation. To keep the original shape of the domain, the velocity field must have zero normal component on the boundary, which means nothing is moving in or out of the original domain. Considering the normal components of the $\mathbf{s}$ at the discrete samplings on the boundary in Eq. (17), they are not guaranteed to be zero. To eliminate those normal component to zero, another velocity field $\mathbf{s}_c$ can be added to the existing $\mathbf{s}$. $\mathbf{s}_c$ is also divergence-free as $\mathbf{s}$, and has the same normal

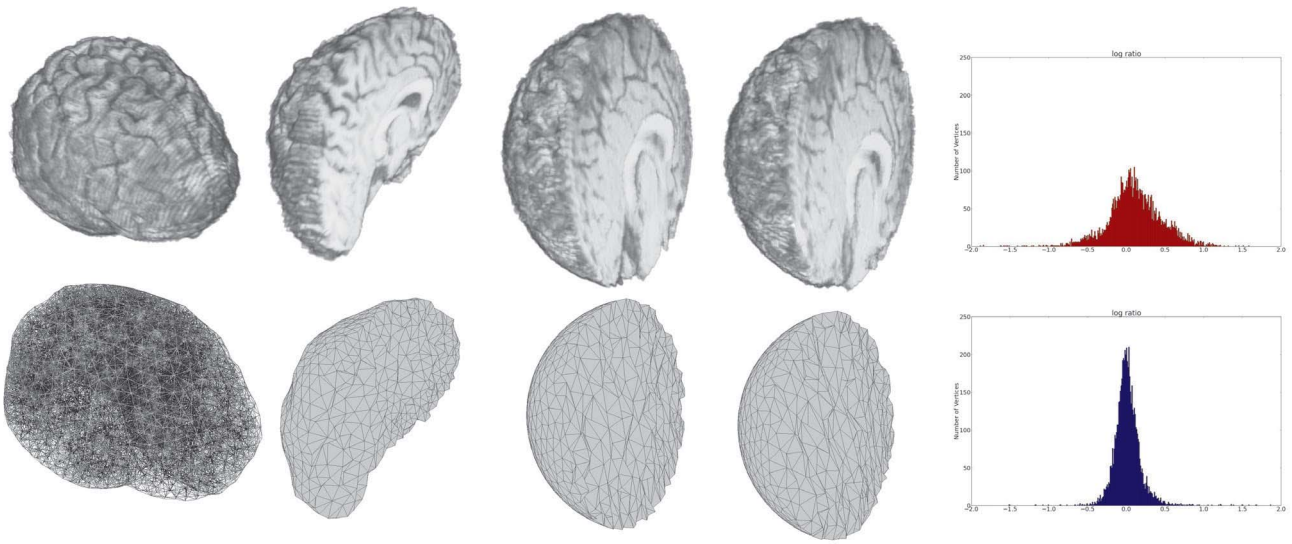

Fig. 3. Volume-preserving mapping pipeline. The volume is represented with a 3D tetrahedral mesh. The volume is first mapped to a 3 dimension sphere with harmonic mapping as shown in the third column. Then, the volume-preserving mapping is applied to preserve the local volume element as shown in the fourth column. During the mapping procedure, the mapping casts on the whole volume simultaneously. The volume is cut by a virtual plane for the interior visualization. The histograms on the right show the log volume ratios on the vertices from the original volume to harmonic (top) and volume-preserving mapping (bottom) ones, respectively.

components at those sampling points on the boundary, but with opposite normal directions. Note that, $\mathbf{s}$ is supposed to register the volume based on those pathlines, so $\mathbf{s}_c$ has to be zero at the pathline samplings in order not to affect the registration. Thus, there is another negative sampling set for $\mathbf{s}_c$

$$\begin{cases} \mathbf{s}_c(\gamma_i(\frac{k-1}{K})) \equiv 0, & k = 1,2,...,K \\ \mathbf{s}_c = -\mathbf{s}(b_j)\cdot\mathbf{n}(b_j), & b_j \in \partial M. \end{cases} \tag{29}$$

With these samplings, the divergence in the previous section can be applied again to generate $\mathbf{s}_c$. Then, adding $\mathbf{s}_c$ to $\mathbf{s}$ leads to the final registration velocity field $\mathbf{s}_r = \mathbf{s} + \mathbf{s}_c$.

Note that, the procedure generating the velocity field $\mathbf{s}_r$ is guaranteed to satisfy the three criteria in Eq. (17). The registration deforms the volume under $\mathbf{s}_r$ with K steps. The salient feature points will be registered to their target locations, and the other parts of the volume will also follow the continuous flow to reach an alignment. Both volume element and local topological structure will be preserved as well.

6 Experimental Results and Applications

In this section, we present the experimental results on the performance and accuracy of our volume-preserving mapping and registration. And an example application will be demonstrated to show how our method can facilitate collective data visualization of 3D multimodal brain imaging database.

6.1 Experimental Results and Validation

The typical volume-preserving mapping pipeline and results are demonstrated in Fig. 3. A closed connected volume is initially mapped onto a 3D sphere. Then, our volume-preserving mapping corrects the local volume, which squeezes the log ratios of volume near to 0. Figure 4 shows another example, where the spherical volume map of the input data is cut to show both the boundary and interior, and the color map indicates the volume distortion rate. The sequence shows the process of the correction of volume distortion from its initial mapping to its 15th and 30th iterations.

To evaluate the registration accuracy, we synthesize incompressible deformations to simulate physical shape changes of a brain, relying on the fact that 75-78% of brain is water, as it is shown in Fig. 5. Then, the volume-preserving mapping is employed to pre-

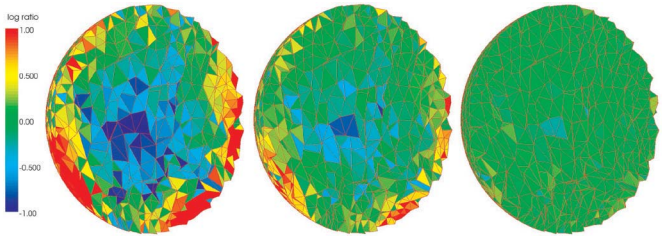

(a) (b) (c)

Fig. 4. Volume distortion correction process, where the spherical volume map is cut to show both boundary and interior, and the color indicates the volume distortion rate. The red means increasing volume while blue means shrinking volume. Green means no volume distortion. (a) Its initial mapping with large volume distortions; (b) and (c) show volume-preserving correction at the 15th and 30th steps, respectively.

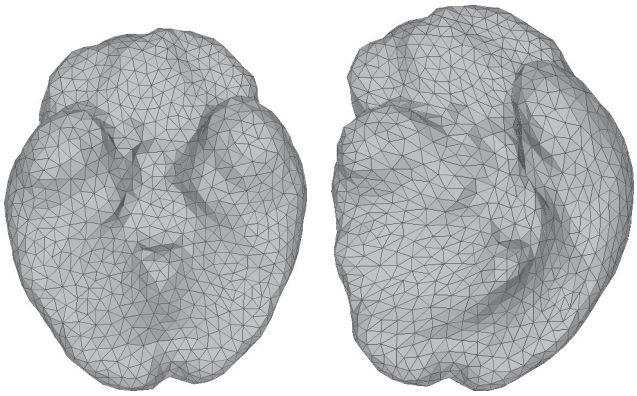

Fig. 5. Deformation synthesis. The left one is the original brain while the right one is the incompressible deformation under a divergence-free field.

dict the interior change. The incompressible deformation is generated with a divergence-free field. The reference and deformed shapes are all mapped and registered to the same unit sphere with both harmonic and volume-preserving mappings, respectively, for the purpose of comparison. For the synthetic data, the vertex correspondence is known. The deformed volumes and their references will have slightly different vertex distributions in the unit sphere. The accuracy is evaluated with the average vertex displacements from the deformed ones to their references within the unit sphere. As shown in table 1, the conformal mapping reduces the vertex displacement to 3.09% and introduces 32.3% volume distortions. Our volume-preserving mapping gives the displacement as 2.32% and keeps the distortion within 12.9%, which is significantly better.

Table 1. Comparison of Predictions

	Conformal	Volume-Preserving
Vertex Displacement	3.09%	2.32%
Volume Distortion	32.3%	12.9%

In many applications, it is desired that some salient features can be registered so that analysis can be performed across subjects with volume-preserving constraints. Fig. 6 demonstrates a divergence-free velocity field during the registration of salient points at one iteration step. To better visualize the velocity field and movement, we select only 3 feature points for registration. The directions shows the moving direction at the current locations to the target positions. Please note that, the boundary condition restricts the field within the sphere domain. The salient points from different shapes are registered to their reference positions. The whole volumes are also registered globally, following the incompressible flow. Fig. 7 shows a slice view of the volume-preserving registration process, where (a) is the reference brain, (b) is the brain to be registered to the reference brain, (c) shows an intermediate result and (d) shows the final registration result. Note that, in real applications, we use many more feature points in order to achieve a registration for analysis purposes. In our experiments, we do not choose any features as registration constraints in the caudates or thalamus structures inside brain. In order to evaluate the registrations of these structures, we calculate the dice coefficients based on our database of 20 MRI scans. Then, they are compared with the results of the HAMMER algorithm [26] and the flow one [13], as shown in table 2. Note that, these two algorithms focus only on registration, but our method, with the volume-preserving constraint, is still comparable to their registration accuracy. In table 2, we also compare the volume distortions of HAMMER [26], the flow method [13] and ours. The results show HAMMER [26] introduced an average volume distortion of 13.3% during the elastic registration, the flow method [13] 14.2%, while ours volume-preserving mapping and registration limits the distortion within 3.0%.

Our volume preserving mapping and registration are both iteration algorithms. Each iteration is to solve a linear system. It is linear to the number of the vertices with the preconditioned bi-conjugate gradient method. With constant K steps, the global complexity is still linear. We have performed our algorithm on a laptop with an i5 2.53 GHz CPU and 8G memory. The algorithm is implemented with python and C++. The average numbers of vertices and tetrahedrons of our brain data used above are 5200 and 26052, respectively. It takes 521 seconds to compute each entire volume-preserving mapping pipeline and 180 seconds for each pair registration, on average.

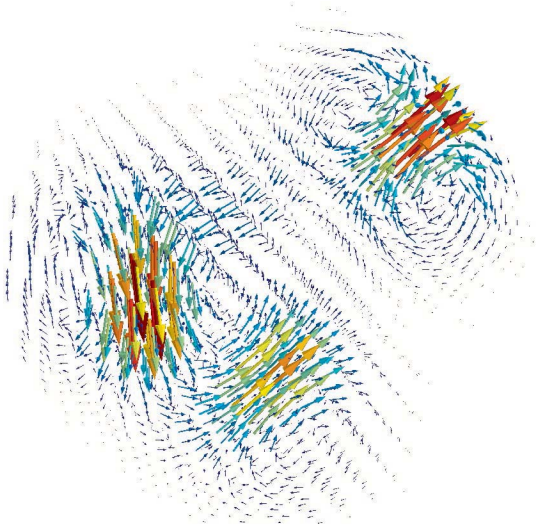

Fig. 6. Divergence-free velocity field in a 3D sphere. Red color and long arrow demonstrate high velocity while blue color and short arrow for low velocity. For the clear visualization, there are only 3 salient points involved.

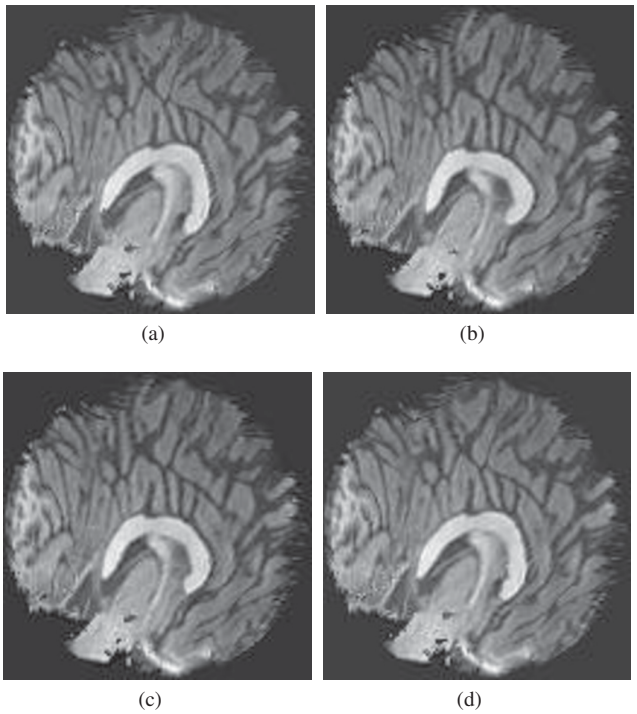

Fig. 7. Slice view of the volume-preserving registration process. (a) shows the reference brain, (b) shows the brain to be registered to the reference brain, (c) shows an intermediate registration deformation result and (d) shows the final registered result.

6.2 Collective Data Visualization for Large 3D Brain Imaging Database

Quantitative multimodal imaging across a large number of human subjects can provide complementing information to assess the many aspects of brain processes in the normal and diseased states. We have applied our method to a database of 300 human brain images, including Magnetic resonance imaging (MRI), Positron emission tomography (PET) and Diffusion Tensor Imaging (DTI) data. The feature extraction from the multimodal data in the large number of subjects in the database can now be indexed and corresponded based on regularized local volumetric element for accurate and efficient data sampling, access and visualization. In brief, our system first applies volume-preserving mapping and registration to 3D data collected from different individuals and brings each individual brain data to a registered volumetric canonical space. And then, an icosahedron subdi-

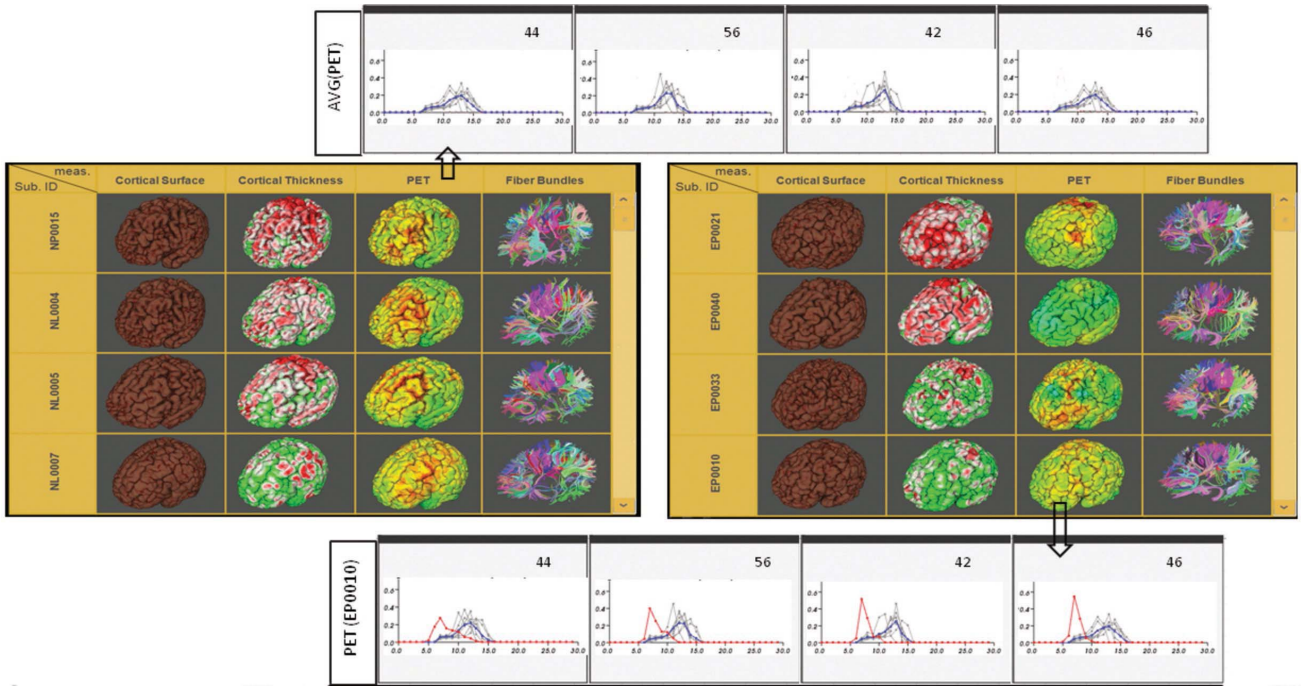

Fig. 8. The left table is the patient table while the right table is the normal's table. The table is generated after the query of finding salient features in a certain diseased group based on the indexed volume-preserving and cross-subject registered elements. The normal distribution was generated based on the same query on normal subjects.

Table 2. Comparison of Dice Coefficients

Structure	HAMMER	Flow Method	Our Method
Left Caudate	0.5820	0.6912	0.6371
Right Caudate	0.8107	0.6474	0.6638
Left Thalamus	0.7365	0.8463	0.7676
Right Thalamus	0.8719	0.8291	0.7560
Average Volume Distortion	13.3%	14.2%	3.0%

vision scheme is employed to produce volumetric elements. As we have shown, establishing a volume-preserving mapping and registration from one individual to volumetric canonical space can guarantee that the same number of samples is used for feature computation and the correspondence is generated in all the subdivided index elements. Subsequently, multimodal feature extraction can be operated in each subjects volumetric space with extracted features indexed in a common template.

With the above volume-preserving mapping and registration process, the simultaneous exploration and collective data visualization across a large number of human subject can now be facilitated upon the features computed on the correspondent elements. The exploration of salient features to certain disease typically starts from a single subject. As features confirm an abnormal variance when compared to the normal population group, this set of salient features becomes an important insight that demands further scrutiny. To confirm this insight and graduate it to a potential hypothesis, the system will formulate this set of features as a formal query and search the database within the current subject's disease group to identify whether similar insight can be found within it. If a substantial number of patients are found having a similar profile, statistical parametric tests are initiated to confirm whether the hypothesis is true. Since confirmed hypothesis leads to treatment decisions, scientists will usually verify these insights for each patient queried from the database and exclude cases which do not make sense. To verify the insights, we provide an interactive multiview coordinated 3-D table display that lists each subject as rows in the table. Fig. 8 illustrates the result obtained after the formal query is executed. The figure shows two separate tables. The table on the right lists the patients as sets of rows, with the columns listing the modalities from where the salient features were found for each patient. Each display cell in the table, when in focus, will open a new window that will visualize its salient feature distribution as a 2-D plot and the corresponding normal distribution. Similarly, the table on the left is a table of normal subjects with the same features as columns so that an analyst will also know the source of the normal distribution the patient is compared against. Clicking on the column heading will show the feature distributions of all the subjects in a single window. Because the cells are interactive and coordinated, rotating the brain within a table display cell will rotate the brain in all the other display cells so that the same anatomical area will be in focus.

In particular, a group of doctors from neurosurgery department have applied and evaluated our collective data visualization system and method for the analysis of epilepsy and normal brain activities using Positron Emission Tomography (PET). The detailed anatomy structure is obtained from Magnetic Resonance Imaging (MRI). Fig. 9 illustrates how the brains of different subjects are registered to a common canonical sphere domain. The abnormality in the patient's brain can be visualized directly, which is statistically different from the normal one according to the cross-subject comparison of the PET values among the aligned and volume-preserved local volume elements. For the quantitative analysis, they have employed PET and MRI data of 50 normal brains. First, all the brain volumes are mapped and registered to a unit 3D sphere. A reference brain in terms of PET is calculated based on the average PET values. When a patient brain is undergoing analysis, the mapping and registration processes are applied as well. Epileptic abnormal foci can be discovered with the comparison of PET values among the aligned and volume-preserved local volume elements, as shown in Fig. 10. The detected foci have been compared to the current gold standard, intracranial electroencephalogram (EEG), which is an invasive method for recording brain electrical signals with the patient skull openned. The accuracy in terms of detected foci volume overlap between the cross-subject PET analysis and EEG

measures reaches 94.2%, which is significantly improved as compared to the traditional individual image analysis (only 82.2%).

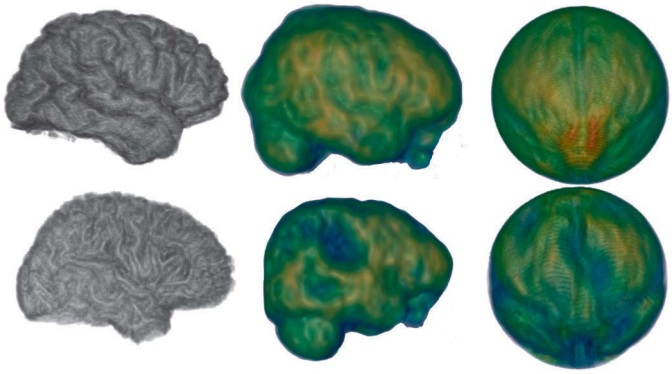

Fig. 9. Volume-preserving mapping of brains. The columns represent MRI data, PET data, and mapped PET data, respectively, while rows represent the normal brain and one from a patient.

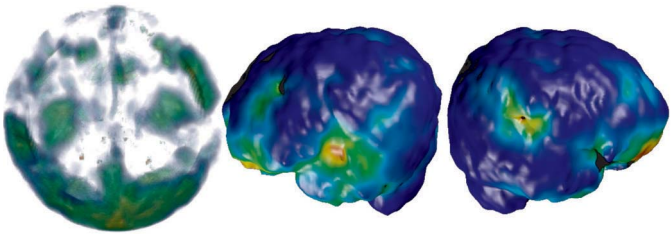

Fig. 10. Map the PET difference back to the original shape. The PET difference shows the abnormalities of the patient's brain. With the bijective mapping, the difference is mapped back to the original brain shape, indicating the epileptogenic region.

7 Conclusion

We have presented a volume mapping methodology which preserves the volume form among general manifolds together with salient feature registration. Given an initial mapping between two closed connected manifolds with the same total volume, a volume-preserving mapping can be derived from a compressible flow of the volume form on the target manifold. Then, salient feature based registration is performed by constructing a divergence-free vector field within the domain. All the features move towards to the desired positions as the manifold keeps moving along the field. The divergence-free field guarantees no volume distortion will be involved in the registration. Different from optimization approaches, a volume-preserving mapping and a registration are guaranteed at the final state of the mathematical integration without the concern of convergence. The volume-preserving registration facilitates the alignment of a class of volume data with similar geometric structures, such as patient brains, to the canonical domain. As the local volume element is preserved and indexed correspondingly across subjects, the quantitative analysis can be uniformly performed upon the same volume size or based on the same number of samples. The volume-based statistics analysis framework is more accurate due to the very minimal distortions and accurate registration. Hence, it facilitates an effective map-based collective data analysis and visualization over a large number of volume objects.

Note that, the discrete algorithm suffers from errors in numerical iterations. Our experiments show that increasing the resolution of the tetrahedral meshes will improve the result. However, high resolution meshes also require more computing power. Better numerical solvers will be investigated in future work. In addition, our current method has to process the boundary and interior separately. We believe that it is a drawback which causes volume preserving and registration errors. A more unified method which simultaneously processes both boundary and interior area/volume preserving constraints could lead to an improved solution.

Acknowledgments

We would like to thank the reviewers for their valuable comments. This work is supported in part by the following research grants, including NSF grants IIS-0915933 and IIS-0937586, NIH grants 1R01NS058802-01A2 and 2R01NS041922-05A1, as well as NSFC grant 61228103.

References

[1] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis. On area preserving mappings of minimal distortion. In T. Djaferis and I. Schick, editors, *System Theory: Modeling, Analysis, and Control*. Springer US, Kluwer, Holland, 1999.

[2] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vision*, 61(2):139–157, 2005.

[3] M. Ben-Chen, C. Gotsman, and G. Bunin. Conformal flattening by curvature prescription and metric scaling. *Comput. Graph. Forum*, 27(2):449–458, 2008.

[4] M. Desbrun, M. Meyer, and P. Alliez. Intrinsic parameterizations of surface meshes. *Comput. Graph. Forum*, 21:209–218, 2002.

[5] A. Dominitz and A. Tannenbaum. Texture mapping via optimal mass transport. *IEEE Trans. on Vis. Comput. Graph.*, 16(3):419–433, 2010.

[6] M. fei Li, S. Liao, and R. Tong. Facial hexahedral mesh transferring by volumetric mapping based on harmonic fields. *Computers and Graphics*, 35(1):92–98, 2011.

[7] X. Gu, S. J. Gortler, and H. Hoppe. Geometry images. *ACM Trans. Graph.*, 21(3):355–361, 2002.

[8] X. Gu and S.-T. Yau. Global conformal surface parameterization. In *Proc. the 2003 Eurographics symp. on geometry processing*, pages 127–137, 2003.

[9] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, G. Sapiro, and M. Halle. Conformal surface parameterization for texture mapping. *IEEE Trans. on Vis. Comput. Graph.*, 6(2):181–189, 2000.

[10] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent. Optimal mass transport for registration and warping. *Int. J. Comput. Vision*, 60(3):225–240, 2004.

[11] J. Hua, Z. Lai, M. Dong, X. Gu, and H. Qin. Geodesic distance-weighted shape vector image diffusion. *IEEE Trans. on Vis. Comput. Graph.*, 14(6):1643–1650, 2008.

[12] M. Jin, Y. Wang, S.-T. Yau, and X. Gu. Optimal global conformal surface parameterization. In *Proceedings of VIS '04*, pages 267–274, 2004.

[13] A. Joshi, R. Leahy, A. W. Toga, and D. Shattuck. A framework for brain registration via simultaneous surface and volume flow. In *IPMI Proceedings*, pages 576–588, 2009.

[14] L. Kharevych, B. Springborn, and P. Schröder. Discrete conformal mappings via circle patterns. *ACM Trans. Graph.*, 25(2):412–438, 2006.

[15] V. Kraevoy and A. Sheffer. Cross-parameterization and compatible remeshing of 3d models. *ACM Transactions on Graphics*, 23(3):861–869, 2004.

[16] B. Lévy, S. Petitjean, N. Ray, and J. Maillot. Least squares conformal maps for automatic texture atlas generation. *ACM Trans. Graph.*, 21(3):362–371, 2002.

[17] X. Li, X. Guo, H. Wang, Y. He, X. Gu, and H. Qin. Harmonic volumetric mapping for solid modeling applications. In *Proc. ACM symp. on Solid and physical modeling*, pages 109–120, 2007.

[18] X. Li, X. Guo, H. Wang, Y. He, X. Gu, and H. Qin. Meshless harmonic volumetric mapping using fundamental solution methods. *IEEE Trans. on Automation Science and Engineering*, 6, 2009.

[19] X. Li, H. Xu, S. Wan, Z. Yin, and W. Yu. Feature-aligned harmonic volumetric mapping using mfs. *Computers and Graphics*, 34, 2010.

[20] S. Lowitzsch. *Approximation and interpolation employing divergence-free radial basis functions with applications*. PhD thesis, Texas A&M University, 2002.

[21] T. Mansi, X. Pennec, M. Sermesant, H. Delingette, and N. Ayache. ilogdemons: A demons-based registration algorithm for tracking incom-

pressible elastic biological tissues. *Int. J. Comput. Vision*, 92(1):92–111, 2011.

[22] M. Meyer, M. Desbrun, P. Schröder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. In H.-C. Hege and K. Polthier, editors, *Visualization and Mathematics III*, pages 35–57. Springer-Verlag, Heidelberg, 2003.

[23] J. Moser. On the volume elements on a manifold. *AMS Trans.*, 120(2):286–294, 1965.

[24] P. V. Sander, J. Snyder, S. J. Gortler, and H. Hoppe. Texture mapping progressive meshes. In *Proceedings of SIGGRAPH '01*, pages 409–416, 2001.

[25] P. V. Sander, Z. J. Wood, S. J. Gortler, J. Snyder, and H. Hoppe. Multi-chart geometry images. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symp. on Geometry processing*, pages 146–155, 2003.

[26] D. Shen and C. Davatzikos. Hammer: Hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. Med. Imaging*, 21:1421–1439, 2002.

[27] O. Sorkine, D. Cohen-Or, R. Goldenthal, and D. Lischinski. Bounded-distortion piecewise mesh parameterization. In *Proceedings of VIS '02*, pages 355–362, 2002.

[28] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache. Symmetric log-domain diffeomorphic registration: A demons-based approach. In *MICCAI*, pages 754–761, 2008.

[29] W. von Funck, H. Theisel, and H.-P. Seidel. Vector field based shape deformations. *ACM Transactions on Graphics*, 25(3):1118–1125, 2006.

[30] S. Wan, Z. Yin, K. Zhang, H. Zhang, and X. Li. A topology-preserving optimization algorithm for polycube mapping. *Computers and Graphics*, 35, 2011.

[31] E. Zhang, K. Mischaikow, and G. Turk. Feature-based surface parameterization and texture mapping. *ACM Trans. Graph.*, 24(1):1–27, 2005.

[32] L. Zhu, S. Haker, and A. Tannenbaum. Flattening maps for the visualization of multibranched vessels. *IEEE Trans. Med. Imaging*, 24(2):191–198, 2005.

[33] G. Zou, J. Hu, X. Gu, and J. Hua. Authalic parameterization of general surfaces using Lie advection. *IEEE Trans on Vis. and Comp. Graphics*, 17(6):2005–2014, 2011.