SCIENCE CHINA
Information Sciences

• **RESEARCH PAPER** •
Progress of Projects Supported by NSFC

November 2013, Vol. 56 110901:1–110901:10
doi: 10.1007/s11432-012-4678-4

Multi-scale local features based on anisotropic heat diffusion and global eigen-structure

LI Shuai[1,2]*, QIN Hong[3] & HAO AiMin[1]

[1]*State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing* 100191*, China;*
[2]*School of Biological Science and Medical Engineering, Beihang University, Beijing* 100191*, China;*
[3]*Department of Computer Science, State University of New York at Stony Brook, Stony Brook, NY* 11794-4400*, USA*

Received June 7, 2012; accepted September 21, 2012; published online October 31, 2012

Abstract Multi-scale local feature detection enables downstream registration and recognition tasks in medical image analysis. This paper articulates a novel robust method for multi-scale local feature extraction on volumetric data. The central idea is the elegant unification of local/global eigen-structures within the powerful framework of anisotropic heat diffusion. First, the local vector field is constructed by way of Hessian matrix and its eigenvectors/eigenvalues. Second, anisotropic heat kernels are computed using the vector field's global graph Laplacian. Robust local features are manifested as extrema across multiple time scales, serving as volumetric heat kernel signature. To tackle the computational challenge for massive volumetric data, we propose a multi-resolution strategy for hierarchical feature extraction based on our feature-preserving down-sampling approach. As a result, heat kernels and local feature identification can be approximated at a coarser level first, and then are pinpointed in a localized region at a finer resolution. Another novelty of this work lies at the initial heat design directly using local eigenvalue for anisotropic heat diffusion across the volumetric domain. We conduct experiments on various medical datasets, and draw comparisons with 3D SIFT method. The diffusion property of our local features, which can be interpreted as random walks in statistics, makes our method robust to noise, and gives rise to intrinsic multi-scale characteristics.

Keywords volumetric heat kernel, anisotropic diffusion, multi-scale local feature, local/global eigen-structure, volumetric image

Citation Li S, Qin H, Hao A M. Multi-scale local features based on anisotropic heat diffusion and global eigen-structure. Sci China Inf Sci, 2013, 56: 110901(10), doi: 10.1007/s11432-012-4678-4

1 Introduction

Multi-scale local feature extraction of images is one of the most fundamental analytical tools, with widespread applications in feature-driven registration, region segmentation, object recognition, and illustrative visualization. In recent years, many powerful algorithms such as SIFT [1] and its variant [2] have been devised to support a suite of applications. Naively generalizing such 2D algorithms to handle 3D volumetric data have gained limited success, yet still encountering tremendous computational challenges

*Corresponding author (email: lishuaiouc@126.com)

when processing massive datasets in a desktop environment. In addition, 3D extensions of SIFT-like algorithms typically consider voxel intensity as scalar input while ignoring local/global inner structure. This is not enough due to the unpredictable external factors during data acquisition, such as severity of diseased organs, illumination, 3D camera viewpoint, equipment, etc. It will result in a large number of less salient or false alarm candidates; especially, ambiguities are unavoidable for the ones with low contrast or being poorly localized near an edge. Therefore, anisotropic and more material-aware method that can respect the local geometry structure and its orientations, must be considered during feature extraction.

Inspired by the scale-space theory and the use of heat diffusion in image analysis [3], one aim of this paper is to reveal the intrinsic structure of volumetric data to facilitate the local feature extraction. However, existing diffusion-based image analysis methods also tend to directly employ intensity information without discovering its local anatomical structure. When designing edge weights of the global graph Laplacian for anisotropic heat diffusion, how to incorporate the knowledge from the local anatomical structure into the feature-finding procedure remains elusive. Besides, existing diffusion methods such as heat kernel signature (HKS) [4] are more effective for locating large-scale features. Local features, on the other hand, are intrinsically linked to local details within certain smaller regions. To assist the heat kernel function to capture details more accurately in short-time scales, extra efforts are needed to design an initial meaningful heat field across the entire volume. Directly using voxel intensity is not a good solution towards robustness to both scale change and noisy input. At the time-performance front, heat-diffusion-based methods need to calculate the heat kernel using the eigen-spectrum of the Laplacian matrix or directly computing matrix exponentiation [5]. For massive volumetric data, such computation is formidable.

To overcome the aforementioned difficulties, our key idea is to employ the local eigen-structure of data as a high-dimensional vector input that governs the anisotropic heat diffusion, which automatically respects the anatomical structure during local feature extraction. As illustrated in Figure 1, the key steps of our algorithm include: local volumetric structure quantification via Hessian matrix computation and its eigen-decomposition, anisotropic heat diffusion in a feature-preserving and down-sampled volume, feature extraction at a coarse scale, and heat field design and local fine feature identification at a localized sub-volume. The novelty of our system lies at the powerful integration of local/global eigen-structures and anisotropic heat kernel computation towards multi-scale local feature extraction. Specifically, the contributions of this paper are as follows:

- We design an eigen-space local structure vectorization representation (ELSVR) for scalar volumetric image by way of Hessian matrix analysis and its eigenvector/eigenvalue concatenation, whose distance metric intrinsically measures the local material continuity of volumetric data.

- We formulate novel anisotropic heat diffusion by integrating ELSVR distance metric into the global graph Laplacian over volumetric image, which naturally respects the local and global structure of volumetric data.

- We advocate volumetric heat kernel to guide multi-scale local feature extraction by employing the first eigenvalue of local Hessian as the initial heat field, of which robust local features, also called volumetric heat kernel signatures, are manifested as extrema across multiple time scales.

- We devise a hierarchical scheme to accelerate the volumetric heat kernel computation by introducing feature-sensitive multi-resolution analysis, of which large-scale features emerge in the down-sampled volume first and exact heat kernels are only computed locally to speedup the feature finding process.

2 Related work

We briefly review previous work in three categories: heat kernel, local image feature extraction, and diffusion-based image processing.

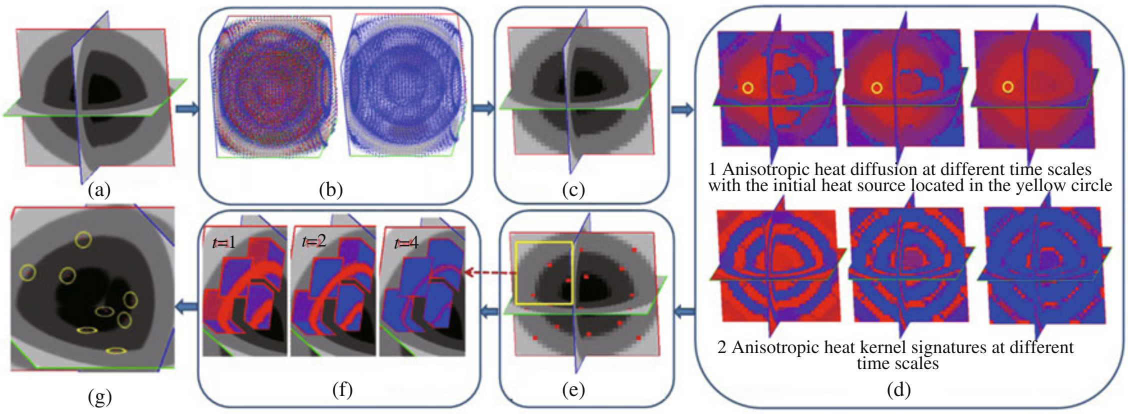

Figure 1 Algorithmic architecture of our new method. (a) Volumetric data; (b) eigenspace local structure vectorization representation; (c) feature-sensitive down-sampling; (d) multi-scale anisotropic heat kernels at coarse resolution; (e) multi-scale feature points at coarse resolution; (f) multi-scale diffusion of heat field at fine resolution; (g) extracted multi-scale feature points.

2.1 Heat kernel

In recent years, heat kernels and diffusion-driven utilities have gained momentum in multi-scale feature detection [4], scalar field smoothing [6] and shape retrieval [7] over manifold. In [8], a scale-invariant heat kernel signature is proposed using Fourier transformation to eliminate the scaling factor, and applied to shape retrieval. Vaxman et al. [5] proposed a multi-resolution approach to improve the computational efficiency of HKS for large meshes. The most significant advantage of such process is that, both diffusion and its kernel function afford robust multi-scale processing on manifold with intrinsic feature of scale-invariance and isometric deformation.

2.2 Local image feature extraction

Local image features are meaningful patterns, which are distinguishable from nearby regions. They typically have rich information content or are evidenced by large differential quantities. A comprehensive review of local image feature detectors can be found in [9]. The core idea of existing 2D feature extraction methods, such as Harris detector, Determinant of Hessian detector, Laplacian of Gaussian detector, and Difference of Gaussian detector, is to pinpoint local extrema by analyzing images' differential quantities. Naturally, these methods are inherently sensitive to scale changes, noisy inputs, and image deformation. To make improvements, Ref. [10] state that combining Harris or Hessian methods with scale selection is a good strategy, and it is recently generalized to a 3D setting [11]. At present, SIFT [1] remains to be one of the most successful algorithms to identify stable features in 2D images, and its generalization to high dimension is also available [12].

2.3 Diffusion-based image processing

During the past two decades, the diffusion theory was used in multi-scale image analysis by means of Gaussian kernel. The Gaussian kernel can mimic the diffusion kernel only in small regions on curved surfaces, but are not suitable for large regions. In essence, it is equivalent to convoluting the original image using classical isotropic heat diffusion. Later, bilateral filtering is proposed in [13] to smooth images while preserving edges. Since then, anisotropic diffusion based filtering has been ubiquitous in 2D image-processing applications, such as denoising, stylization, and optical-flow estimation. It is also worth noting that the idea of heat kernel signature [4] has been extended to isometry-invariant homogeneous solid geometry [14] of 3-manifold without considering material densities. This prompts us to investigate anisotropic heat kernels directly defined over volumetric medical data.

3 Eigenspace local structure vectorization representation

Hessian matrix H is a powerful means to encode local shape information in volumetric data, and it intuitively depicts how the normal (i.e., gradient) of an iso-surface changes. As a real-valued and symmetric matrix, H has real-valued eigenvalues ($|\lambda_1| \geqslant |\lambda_2| \geqslant |\lambda_3|$) and corresponding eigenvectors ($\boldsymbol{e}_1, \boldsymbol{e}_2, \boldsymbol{e}_3$). In some sense, the directions corresponding to maximal eigenvalue represent the most direct transition from one material to other neighboring materials, while the direction corresponding to minimal eigenvalue shows the material flow of certain material, and the maximal and minimum eigenvalues represent the principal curvatures. Thus, from the standpoint of anisotropic heat diffusion, considering both eigenvalues and eigenvectors simultaneously overcome the limitation of isotropic Laplacian filter, our goal is to extract local structure and quantify material property change among neighboring voxels.

In [15], Alfiansyah et al. documented the relationship between different combinations of the eigenvalues of Hessian matrix and their corresponding 3D local structures (sheet, blob, and noise) as follows:

$$\begin{aligned} R_{\text{sheet}} &= (|\lambda_3| + |\lambda_2|)/|\lambda_1|, \\ R_{\text{blob}} &= (2|\lambda_1| - |\lambda_2| - |\lambda_3|)/2|\lambda_1|, \\ R_{\text{noise}} &= \sqrt{\lambda_3^2 + \lambda_2^2 + \lambda_1^2}. \end{aligned} \tag{1}$$

According to [15], the eigenvalues for voxels on a tube surface will be: $|\lambda_1| = 0$ and $|\lambda_2| = |\lambda_3| >> 0$; thus we can homoplastically define a ratio to depict the tube-like local structure as

$$R_{\text{tube}} = (|\lambda_3| + |\lambda_2| - |\lambda_1|)/|\lambda_1|. \tag{2}$$

The values of the four ratios become 0 when their corresponding local structures emerge. We now set a stage to construct a high-dimensional eigen-space by concatenating eigenvectors, eigenvalues, and their ratios computed from local Hessian matrices. Our eigen-space local structure vectorization representation is formulated as ELSVR= $(\boldsymbol{e}_1^{\mathrm{T}}, \boldsymbol{e}_2^{\mathrm{T}}, \boldsymbol{e}_3^{\mathrm{T}}, \lambda_1, \lambda_2, \lambda_3, R_{\text{sheet}}, R_{\text{blob}}, R_{\text{tube}}, R_{\text{noise}})$.

Next, ELSVR is employed to govern the anisotropic heat diffusion via initial heat field design and global graph Laplacian assembly.

4 Anisotropic heat diffusion and global graph Laplacian

4.1 Heat kernel

Heat diffusion over high-dimensional ELSVR manifold is governed by a heat equation. Let $(M;\mu)$ be a compact Riemannian manifold with metric μ. The heat diffusion over $(M;\mu)$ is controlled by

$$\frac{\partial u(x,t)}{\partial t} = -\triangle_\mu u(x,t), \tag{3}$$

where $\triangle_\mu$ denotes the Laplace-Beltrami operator on $(M;\mu)$, and $h_t(x,y)$: $R^+ \times M \times M \longrightarrow R^+$ is called the heat kernel if, for all $f \in L^2(M;\mu), t > 0, x \in M$,

$$H_t f(x) = \int_M h_t(x,y) f(y) \mathrm{d}\mu(y). \tag{4}$$

The heat kernel h_t is the integral kernel of the operator H_t. Assume the Laplace-Beltrami operator $\triangle_\mu$ of M has the eigen-decomposition $\{\lambda_k, \phi_k\}_{k=1}^\infty$. Then it can be formulated as a series of eigenfunctions:

$$h_t(x,y) = \sum_{k=1}^{\infty} \mathrm{e}^{-t\lambda_k} \phi_k(x)\phi_k(y). \tag{5}$$

The physical meaning of heat kernel is the heat diffused from x to y at any time t, with a unit heat source placed at x. Heat kernels have many attractive properties. They are symmetric, isometric-invariant, multi-scale, informative, and stable. They also satisfy the semigroup identity: $h_{t+s}(x,y) = \int_M h_t(x,z)h_s(y,z)\mathrm{d}z$.

4.2 ELSVR-weighted global graph Laplacian

Over ELSVR manifold, the behavior of anisotropic heat diffusion is determined by its graph Laplacian that encodes neighborhood information. Consider volumetric data as an undirected graph $G = (V, E)$, where V represents all voxels and E satisfies $E \in V \times V$, and an edge $E_{ij} \in E$ exists if and only if the corresponding voxel pair satisfies the one-ring connectivity requirement (we denote the neighboring voxel set by N_i). In order to favor heat diffusion across similar material and suppress its conduction among different materials, the key idea is to design a global graph Laplacian which is ELSVR-weighted. We first measure to what extent two neighboring voxels V_i and V_j belong to the same material:

$$S_{i,j} = \begin{cases} 0, & \text{if } R^i_{\text{noise}} = 0 \text{ or } R^j_{\text{noise}} = 0, \\ \mathrm{e}^{-d_r}, & \text{otherwise}, \end{cases} \tag{6}$$

where $d_r = \sum_{R_k \in R}(R^i_k - R^j_k)^2$, $R = \{R_{\text{sheet}}, R_{\text{blob}}, R_{\text{tube}}\}$. The edge weight $W(i,j)$ can be defined as

$$W_{i,j} = \begin{cases} \mathrm{e}^{\frac{-(1-S_{i,j})d_{i,j}d^{\alpha}_{i,j}}{\sigma^2}}, & V_j \in N_i, \\ 0, & \text{otherwise}, \end{cases} \tag{7}$$

where $d_{i,j}$ is the Euclidean distance between V_i and V_j with sampling intervals d_x, d_y, d_z, $d^{\alpha}_{i,j} = (\arccos(\boldsymbol{e}^i_1 \cdot \boldsymbol{e}^j_1) + \arccos(\boldsymbol{e}^i_2 \cdot \boldsymbol{e}^j_2) + \arccos(\boldsymbol{e}^i_3 \cdot \boldsymbol{e}^j_3))$, and σ is a control parameter and proportional to the variance of structure distance of two voxels. For highly noisy images, we must find a way to suppress noise influence. One feasible method is to smooth the image with 3D Gaussian kernel first. Thus the graph Laplacian is defined as

$$L_{i,j} = \begin{cases} W_{i,j}, & V_j \in N_i, \\ -\sum\limits_{j \in N_i} W_{i,j}, & i = j, \\ 0, & \text{otherwise}. \end{cases} \tag{8}$$

Based on (4), our volumetric anisotropic heat kernel $h^v_t(V_i, V_j)$ (VAHK) can be approximated by using a finite number of eigenfunctions of this matrix. The ELSVR-weighted graph Laplacian enables new heat kernels that take local structure information into consideration and thus is more material-aware and structure-sensitive (see Figure 1 for its anisotropic behavior).

4.3 Initial heat field design

One naive way to initialize the heat field is to directly use voxel intensity. Yet, this simple approach ignores rich local structure encoded by ELSVR. As shown in Figure 2, eigenvalue components in ELSVR are potential candidates for feature-sensitive initial heat field design, because: 1) the first eigenvector represents the direction of the maximum curvature; 2) the eigenvalues are rotation-invariant; 3) the first and third eigenvalues approximately exhibit a negative symmetry in boundary areas, while the second eigenvalue often lacks good contrast. We choose the first eigenvalue to initiate the heat field.

We now conduct heat kernel convolution with multiple time scales in a local region and determine the feature location by searching local extrema. The initial heat field $f_0(\lambda^i_1)$ after one-step convolution should be updated as

$$f_t(\lambda^i_1) = \sum_{V_j \in \Omega_i} h^v_t(V_j, V_i) f_0(\lambda^j_1), \tag{9}$$

where Ω_i is the local convolution region of V_i.

Rewriting (8) in matrix form, we can get $F_t(\lambda_1) = A_t F_0(\lambda_1)$. By using the semigroup identity of heat kernel, the heat field after n-step convolution can be obtained as $F_{nt}(\lambda_1) = A_{t_0} F_{(n-1)t}(\lambda_1) = \cdots = A^{n-1}_{t_0} F_t(\lambda_1) = A^n_{t_0} F_0(\lambda_1)$. By fixing time $t = t_0$ for the one-step diffusion, as n increases, the diffusion can be approximated by a sparse matrix and its multiplication; thus a GPU acceleration employing CUDA is readily available towards performance improvement.

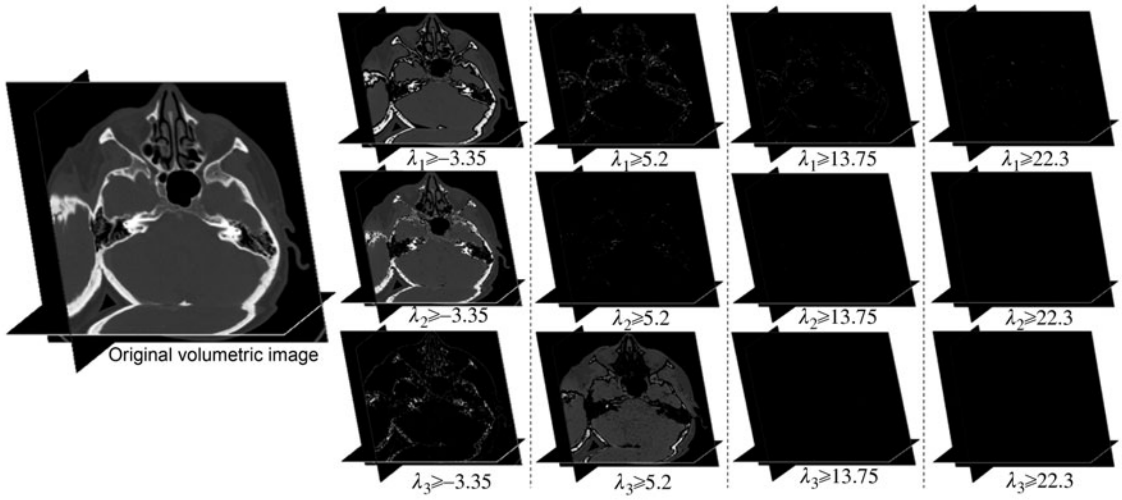

Figure 2 A set of images corresponding to different value ranges of λ_1.

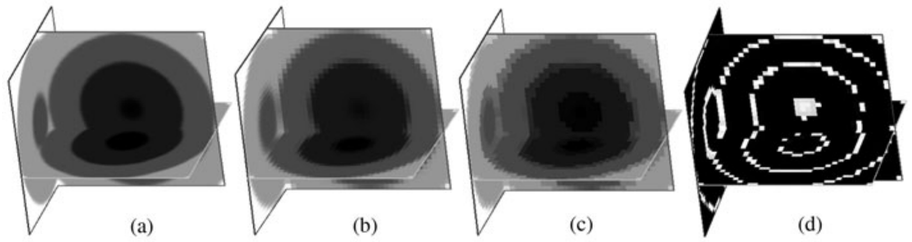

Figure 3 Comparison of down-sampling methods. (a) Original volume, 128^3 voxels; (b) average down-sampling method, 32^3 voxels; (c) our method, 32^3 voxels; (d) difference of (b) and (c).

5 Multi-resolution strategy and numerical techniques

The biggest challenge for volumetric data is its extremely high computational complexity. In practice, the multi-resolution approach is much more desirable for performance speedup. The main idea is to organize volumetric data in a hierarchical fashion, so that we process a much smaller dataset at a coarser grid, and then refine our feature extraction procedure in a localized sub-volume at a finer grid. Since any down-sampling strategy might result in important information loss, we propose a feature-preserving down-sampling method:

$$\widetilde{I(V_i)} = \frac{1}{\sum_{V_j \,\in\, \Omega_i} \mathrm{e}^{|\lambda_1^j|}} \sum_{V_j \,\in\, \Omega_i} \mathrm{e}^{|\lambda_1^j|} I(V_j), \tag{10}$$

where $\widetilde{I(V_i)}$ is the down-sampled voxel intensity at a coarser grid, Ω_i is the sampling region corresponding to $\widetilde{I(V_i)}$. The comparison between our down-sampling method and another method is illustrated in Figure 3, where the down-sampled volumes using the average method, our method and the subtraction result are illustrated respectively from left to right. It can be noted that, our method is sensitive to local curvature information in original volumetric data, and this down-sampling method takes advantage of already-computed ELSVR and reuses its curvature information encoded in ELSVR.

After down-sampling, the discrete anisotropic heat kernel is approximated by the first 300 smallest eigenvalues and the corresponding eigenfunctions. As noted in [4], HKS is very robust to large scale features. In the same spirit, the down-sampled volume retains large scale features because of its feature-preserving property. Consequently, it is ideal to compute the volumetric anisotropic heat kernel signature (VAHKS) using

$$h_t^v(V_x, V_x) = \sum_{k=1}^{n} \mathrm{e}^{-t\lambda_k} \phi_k(V_x)\phi_k(V_x). \tag{11}$$

The important parameter is time scale t, which is intrinsically relevant to the size of local heat diffusion region for each voxel. The smallest t in down-sampled volume should be set to $D_{\text{size}} \times \min(d_x, d_y, d_z)$, and D_{size} is the down-sampling interval. After getting multi-scale VAHKS, we can approximate locations of feature points by choosing voxels that always occur as local extremum of VAHKS under multiple time scales.

We then refine feature locations in the original volumetric data. In practical implementation, we employ Matlab function *eigs* to conduct Laplacian matrix eigen-decomposition, according to our experience, which will be very time-consuming when the matrix size is larger than 40000×40000. Besides, the largest down-sampling interval in our experiments is 16-voxel long. Therefore, corresponding to each feature location at a coarser grid, it suffices to choose the size of sub-volume as $\Omega_{\text{refine}} = 32 \times 32 \times 32$ during VAHKS recomputation. In each sub-volume, we conduct heat field diffusion using the same method as described in Subsection 4.3. Similar local extremum searching methods can be adopted to extract local multi-scale feature points.

However, a remaining problem about the relationship between local diffusion region and the diffusion time must be addressed. Heat kernel $h(V, \cdot)$ can be viewed as the transition density function of the Brownian motion, whose influence range is determined by a geodesic ball $B(V, O\sqrt{c_n t \lg t})$ [16]. Therefore, around a voxel V_i in the original volume, it is justifiable and acceptable to define a heat region as

$$\Omega_t^i = \{V_j | V_j \in \Omega_{\text{refine}} \quad \text{and} \quad h_t^v(V_j, V_i) > C(t)\}, \tag{12}$$

where $C(t)$ is a threshold, which is set to $\min(d_x, d_y, d_z)/(1 + t)$. Moreover, as features located near the boundary of the sub-volume are not stable, we should discard them. In our experiments, the number of extracted feature points for a volumetric image is around 100–1500, which is large enough to support the subsequent feature-based applications.

6 Experimental results and evaluation

We conduct experiments on various data sets and make comparison with 3D SIFT method. Our method is implemented using C++ on a computer with Intel Core (TM) i7 CPU (1.6 GHz, 4 cores) and 4G RAM. To improve the time performance, we employ CUDA to conduct parallel ELSVR computation of original volumetric image. Table 1 documents the statistics and time performance of key parts in our experiments. The 4th and 5th columns show the timing (in seconds) for coarse level feature locating (CLFL, also including down-sampling, VAHKS computation), and per sub-volume feature refining (PSFR, including local VAHKS computation, convolution and multi-scale feature extraction). It indicates that the time expense is comparable with that of 3D SIFT method.

Figure 4(a) illustrates experimental results and their comparison with 3D SIFT over volumetric MRI images of human head. Figure 4(b) is a down-sampled version of Figure 4(a) with $32 \times 32 \times 9$ resolution, where VAHKS shown is computed at $t = 6$. Figure 4(c) shows the refined VAHKS when $t = 0.6$ and $t = 1.2$ in a sub-volume of original volumetric images, which corresponds to a coarse feature location. It shows that the VAHKS are more informative, since features extracted using VAHKS are mostly located at the boundaries of different materials or regions with high curvature changes, while features from 3D SIFT (see Figure 4(e)) lack such a property and occur irregularly. To examine the robustness of our method, we add 5% and 10% (of average intensity) random noise to original images. The results are shown in Figures 4(f)–(i). Comparing Figures 4(f), (h) with Figure 4(d) and comparing Figures 4(g), (i) with Figure 4(e), we see that most of VAHKS feature points are stable and only a small portion of them change when adding noisy perturbation (see points circled by red rings in Figures 4(f) and (h)). The results from 3D SIFT show obvious changes (note that the features completely disappear in Figure 4(g) and feature points in Figure 4(i) change drastically).

Figure 5 and Figure 6 show experimental results on volumetric CT images of human pelvis (Figure 5(a)) and human abdomen (Figure 6(a)). In Figure 5, Figure 5(b) and Figure 5(c) respectively show the feature points extracted from original volumetric images by using our method and 3D

Table 1 Time performance of experiments (s)

Data	Resolution	ELSVR (CUDA)	Down-sampling	CLFL	PSFR
Figure 4(a)	$512 \times 512 \times 147$	36.4	$32 \times 32 \times 9$	139	487
Figure 5(a)	$512 \times 512 \times 174$	45.8	$32 \times 32 \times 10$	164	506
Figure 6(a)	$512 \times 512 \times 256$	79.1	$32 \times 32 \times 16$	436	484

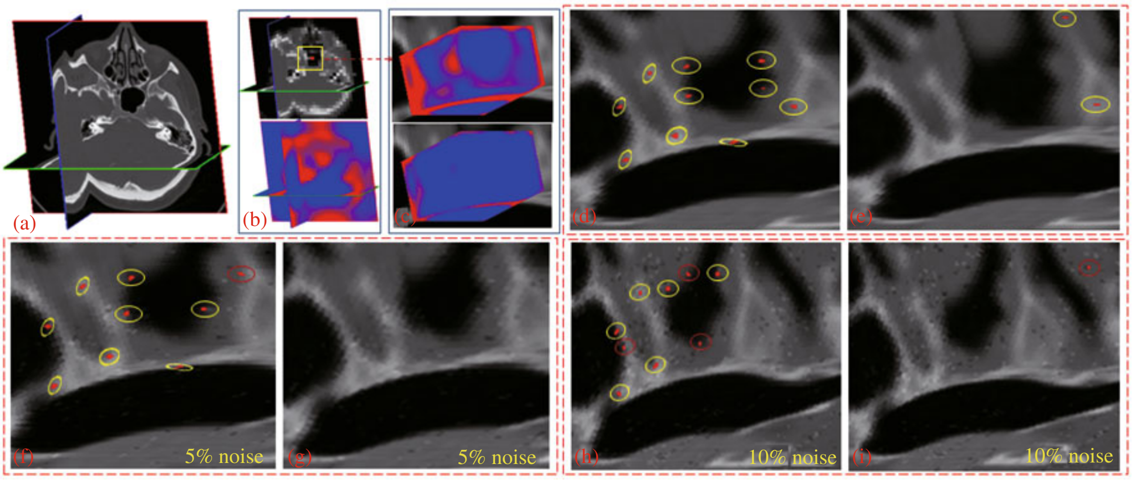

Figure 4 Experimental results and their comparison with 3D SIFT over MRI head volume. (a) Input images; (b) VAHKS computed when $t = 6$ in the down-sampled images; (c) refined VAHKS when $t = 0.6$ and $t = 1.2$ in original images; (d) features extracted by VAHKS from original images; (e) features extracted by SIFT from original images; (f) features extracted by VAHKS from 5% noise perturbed images; (g) features extracted by SIFT from 5% noise perturbed images; (h) features extracted by VAHKS from 10% noise perturbed images; (i) features extracted by SIFT from 10% noise perturbed images.

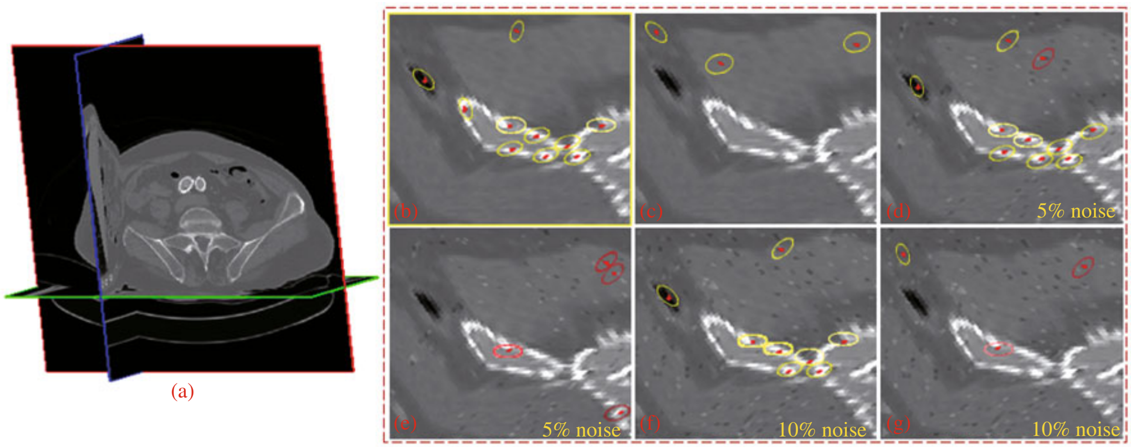

Figure 5 Experimental results and their comparison with 3D SIFT over CT pelvis volume. (a) Input images; (b) features extracted by VAHKS from original images; (c) features extracted by SIFT from original images; (d) features extracted by VAHKS from 5% noise perturbed images; (e) features extracted by SIFT from 5% noise perturbed images; (f) features extracted by VAHKS from 10% noise perturbed images; (g) features extracted by SIFT from 10% noise perturbed images.

SIFT method, Figures 5(d), (e) and Figures 5(f), (g) are the corresponding results when adding 5% and 10% random noise to original pelvis images. In Figure 6, Figures 6(b), (c) and Figures 6(d), (e) respectively demonstrate the experimental results on original abdomen images and the perturbed ones with 10% random noise. For both our method and 3D SIFT method, the points circled by yellow rings represent stable feature points which are not affected by noise. Such results further verify the robustness of our method. Thus, our method is not sensitive to the modalities of volumetric images, and it is multi-

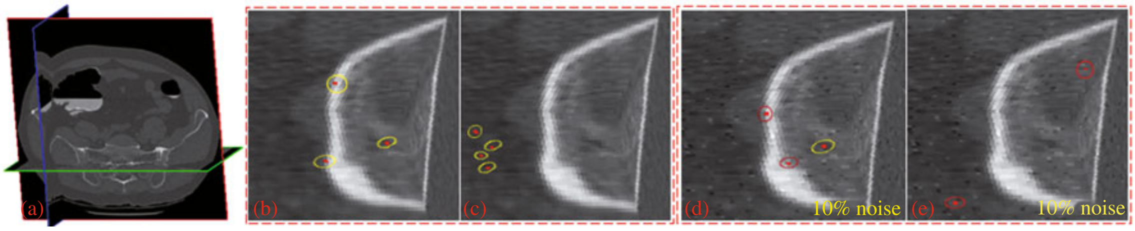

Figure 6 Experimental results and their comparison with 3D SIFT over CT abdomen volume. (a) Input images; (b) features extracted by VAHKS from original images; (c) features extracted by SIFT from original images; (d) features extracted by VAHKS from 10% noise perturbed images; (e) features extracted by SIFT from 10% noise perturbed images.

scale, informative and robust since it directly inherits such excellent properties from heat kernel and random walk.

Our volumetric heat kernel based feature extraction can facilitate many down-stream applications, such as feature-based sparse registration, feature-constrained dense registration, feature-preserving denoising/smoothing, feature-centric segmentation, and feature-driven illustrative visualization. Especially for registration, Ovsjanikov et al. [17] have employed heat kernel to map isometric shapes by defining injective heat kernel map (HKM) for surface points over the 2D manifold. Although the original HKM still suffers from some limitations in injection property and discriminating power, inspired by HKM, we can further define a simple and robust descriptor by means of the proposed VAHKS at multiple different time scales, which can be used to match the sparse local features between two homologous volumetric medical images. Furthermore, we can take the sparsely matched feature points as anchors and extend HKM from one source to multiple sources, since multiple sources can offer direction and distance information during feature matching, and thus have stronger discriminating power. This mechanism offers good promise to further support the dense registration of volumetric images. Our immediate efforts will be geared towards broadening the application scope of our method.

7 Conclusion and discussion

We have detailed a novel robust method for multi-scale local feature extraction in volumetric data. Our newly-proposed solutions include many technical elements, including ELSVR, ELSVR-weighted graph Laplacian, anisotropic volumetric heat kernel computation in a hierarchical fashion, and initial heat field design using the first eigenvalue of local Hessian matrices. The center piece of our method is the unified treatment of local/global eigen-structures and their participation in the anisotropic heat diffusion process. As a result, we are capable of locating multi-scale local features, supporting multi-resolution analysis in volumetric images. Our experiments demonstrate good properties of this volumetric anisotropic heat kernel, especially its robustness to noisy input.

Based on the theoretical analysis and experimental results shown above, we also discuss some limitations and open problems of our method. In general, the volumetric dataset may have thousands of millions of voxels, which can't afford to conduct global spectral analysis, since the corresponding Laplacian matrix will be a square matrix with many billions of entries. While our multi-resolution method aims to tackle this challenge, the full spectrum eigen-decomposition of the low-resolution volume (or sub-volume) still remains very time-consuming. We will have to adopt a brute-force truncation method to restrain the heat diffusion in a sub-volume, where the uncertainty will gradually increase when the heat diffusion start to approach boundaries of the sub-volume. As a result, the feature points close to the boundary tend to be less trustworthy, which should be abandoned, so we may miss some feature points. In practice, however, we typically need hundreds of sparse feature points, and our method can extract enough feature points to enable the down-stream feature-based applications. Meanwhile, the feature-sensitive down-sampling method might still cause information loss, which remains a challenge to quantitatively analyzing the relationship between the image feature preservation and the multi-resolution representation of images.

Acknowledgements

This research is supported in part by National Natural Science Foundation of China (Grant Nos. 61190120, 61190121, 61190125) and National Science Foundation of USA (Grant Nos. IIS-0949467, IIS-1047715, IIS-1049448).

References

1 Lowe D G. Distinctive image features from scale-invariant keypoints. Int J Comput Vis, 2004, 2: 91–110
2 Osada K, Furuya T, Ohbuchi R. SHREC'08 entry: local volumetric features for 3D model retrieval. In: International Conference on Shape Modeling and Applications, New York, 2008. 245–246
3 Zhang F, Hancock E R. Graph spectral image smoothing using the heat kernel. Pattern Recognit, 2008, 41: 3328–3342
4 Sun J, Ovsjanikov M, Guibas L. A concise and provably informative multi-scale signature based on heat diffusion. In: Proceedings of the Symposium on Geometry Processing, Berlin, 2009. 1383–1392
5 Vaxman A, Ben-Chen M, Gotsman C. A multi-resolution approach to heat kernels on discrete surfaces. ACM Trans Graph, 2010, 29: 1–10
6 Patanè G, Falcidieno B. Multi-scale feature spaces for shape processing and analysis. In: Proceedings of the 2010 Shape Modeling International Conference, Washington, DC, 2010. 113–123
7 Bronstein A M, Bronstein M M, Guibas L J, et al. Shape Google: geometric words and expressions for invariant shape retrieval. ACM Trans Graph, 2011, 30: 3328–3342
8 Bronstein M M, Kokkinos I. Scale-invariant heat kernel signatures for non-rigid shape recognition. In: Proceedings of Computer Vision and Pattern Recognition, San Francisco, 2010. 1704–1711
9 Tuytelaars T, Mikolajczyk K. Local invariant feature detectors: a survey. Found Trends Comput Graph Vis, 2008, 3: 177–280
10 Moreels P, Peron P. Evaluation of features detectors and descriptors based on 3d objects. Int J Comput Vis, 2007, 73: 263–284
11 Glomb P. Detection of interest points on 3D data: extending the Harris operator. Computer Recognition Systems 3. Berlin: Springer, 2009. 103–111
12 Cheung W, Hamarneh G. N-SIFT: N-dimensional scale invariant feature transform for matching medical images. In: Proceedings of 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Arlington, VA, 2007. 720–723
13 Aurich V, Weule J. Non-linear gaussian filters performing edge preserving diffusion. In: Proceedings of Mustererkennung 1995, 17. DAGM-Symposium, Bielefeld, 1995. 538–545
14 Raviv D, Bronstein M M, Bronstein A M, et al. Volumetric heat kernel signatures. In: Proceedings of the ACM Workshop on 3D Object Retrieval, New York, 2010. 39–44
15 Alfiansyah A, Ng K H, Lamsudin R. Bone segmentation based on local structure descriptor driven active contour. In: 13th International Conference on Biomedical Engineering, Singapore, 2009. 1077–1081
16 Grigor'yan A. Escape rate of brownian motion on riemannian manifolds. Appl Anal, 1999, 71: 63–89
17 Ovsjanikov M, Merigot Q, Memoli F, et al. One point isometric matching with the heat kernel. Comput Graph Forum, 2010, 29: 1555–1564