Fragmented Skull Modeling Using Heat Kernels

Wei Yu [1], Maoqing Li [1], Xin Li* [2]

[1] Department of Automation, Xiamen University, China
[2] Department of Electrical and Computer Engineering, Louisiana State University, USA

Abstract

We develop a geometric reassembly algorithm that composes a 3D model from its fragments. One important application of this work is skull completion and modeling in archeology and forensics. Our reassembly algorithm employs a scale-space representation of shape based on the heat kernel, which only depends on the intrinsic geometry of the surfaces. Partial matching can then be conducted effectively. The entire assembly pipeline has three steps: (1) fragment-template matching based on heat-kernel; (2) matching refinement based on RANSAC and assembly computation; (3) assembly refinement using least square transformation error (LSTE) of break-curves. The main contribution of this paper is presenting novel algorithms for the first two steps. Experimental results on scanned skull fragments demonstrate the efficacy and robustness of our algorithm.

1 Introduction

This work studies geometric algorithms to reassemble fragmented 3D patches. Geometric reassembly algorithms can facilitate computer-aided data acquisition and completion when the model is not only incomplete but even fragmented. In this project, our goal is to explore reliable and effective algorithms to help recompose excavated skull fragments, which provide important information in archeology and forensic law enforcement. Anthropologists can reconstruct the face geometry from the excavated skull after they clean and complete the skull. For example, in law enforcement, such facial reconstruction is an important technique to help identify the deceased (when other ID information is not available due to environmental erosion or human activities) and has been demonstrated effective in many real cases [19]. The facial reconstruction is performed based on the statistical tissue depth and craniofacial anatomy. Before the tissue can be model, completion of the skull needs to be done, since it is very difficult to directly conduct tissue and craniofacial structure reconstruction when the skull is fragmented or has large regions missing. A geometric modeling and processing system that can complete the excavated skull reliably is therefore highly desirable and can greatly facilitate subsequent tasks.

With 3D scanning technologies, we can scan skull data into digital forms and use geometric processing technologies to repair the digital model. For the model that is in one piece and is just incomplete (e.g. it has holes and cracks), many model completion algorithms proposed in modeling/graphics literature [14, 18, 17] can be used. However, if the model is fragmented, we need to first reassemble these sub-pieces together, which is a challenging and less-explored problem. Two general strategies could be used for reassembly: we can rely on local properties of these fragment themselves, or we can use a complete template model for guidance.

In skull reassembly, a difficulty of using templates is that the geometry of skull is subtle and the template will not be exactly the same from the fragmented subject skull, so the registration between them might not be always accurate (especially when some fragment pieces are small). On the other hand, methods based on properties of fragments themselves are also nontrivial since some parts/pieces may be missing, different fragment scans have different scalings and samplings. In this work, we make use of both of these strategies to achieve a robust reassembly. The overall pipeline is illustrated in Fig. 1. The main contributions are two-folded.

1. We adopt a multi-scale surface descriptor based on the heat kernel, which has not been used for data reassembly before. We analyze its good properties in matching partial models to the complete model.
2. We integrate the developed methodologies into a three-step skull reassembly pipeline, and demonstrate that it is especially suitable for this task and is more effective and robust than existing techniques.

2 Background and Related Work

2.1 Template Matching and 3D Shape Descriptor

Template-based approaches are suitable for reassembling fragments whose sizes are not always big and whose boundary geometry is worn or partially damaged. When we use a template to guide the fragment reassembly, we need to solve a partial matching problem, which seeks a good mapping from a model F to a sub-region of another model M. Effective geometric feature extraction can greatly facilitate this partial matching problem.

A geometric descriptor is usually a function defined on each point of a model to describe local or global characteristic of this model. Local surface normals and curve curvatures are popular local descriptors in 3D object matching [29, 25]. This kind of descriptors is typically very easy to compute and compare. However, for a given point in the shape, there may be many points with the indistinguishable descriptor value. Therefore, they are usually used together with a voting scheme [1] or with an iterative alignment scheme to improve the discrimination ability [28].

Since templates and fragments might have different sizes and shapes, algorithms that are not sensitive to local geometric variance are desirable. One common strategy to derive point signatures is to summarize the shape distribution in the neighborhood of a point. For example, spin images [12] and shape contexts [3] are widely used point signatures that fall into this category. The spin image at a point is a 2D histogram of the counts of 3D points in a surrounding region, established using the local surface normal as a reference for invariance. This provides a discriminative local shape descriptor: points with the similar spin-image descriptors indicate highly likely similar (corresponding) regions. However, for two shapes that are not exactly the same, spin images could be too sensitive to local difference and may not correctly find corresponding feature points.

On the other hand, shape descriptors that capture the global geometric properties, such as Global Point Signature (GPS) [26, 23] are also studied. Eigenvalues of the Laplace-Beltrami operator are used together with the corresponding eigenfunctions to characterize the shape of models.

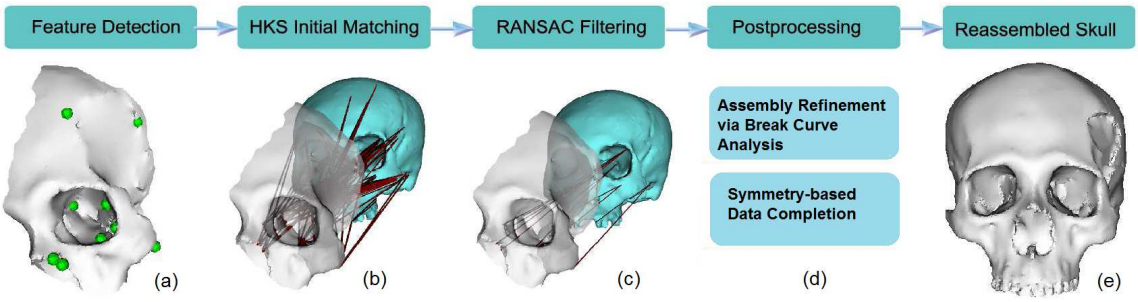

Figure 1: Fragmented Skull Completion Pipeline. First, extract multi-scale features on the fragments (a) and the template; second, perform HKS matching to obtain a superset of potential matches (b); then, use a RANSAC filtering to correct the matching (c) and get the rigid transformations on each fragment; then, perform a postprocessing (d) to complete the skull through a assembly refinement and model repair (hole filling) and finally obtain the completed skull (e).

However, pure global signature is also not suitable for handling fragment-template matching, because a sub-piece is globally different from a complete model and such a descriptor may not be applied to partial matching directly.

Based on heat diffusion process, the heat kernel function captures geometrical information around a point. The dissipating time of heat provides a natural notion of scale to describe the neighborhood of a point. Local shape properties can be characterized when observing this function within a small time range, while more global shape properties are encoded in a larger time range [30]. Since the heat kernel can capture surface geometry in a multi scale way, it is a powerful tool for data representation [33]. For instance, it is used for designing diffusion distance [16], isometry-invariant hierarchical segmentation [6], finding isometric matching [22][27], shape retrieval [21], and so forth.

2.2 Fragments Assembly

Assembling 2D fragments has been explored in computer vision fields [11, 15]. Papaioannou and Karabassi [24] developed a solution for assembly of 3D objects, which uses a complementary surface matching algorithm, in conjunction with facet boundary curve matching to compute the transformations. But a preprocessing is needed to ensure the fracture faces are nearly planar and they match each other completely.

Cooper et al. [5] proposed a framework of assembling 3D fragments mainly based on 3D measurement and matching of break-curve, sherd normal, etc. Willis et al. [35] then developed this approach by using Bayesian approach to assemble pots based on curve matching semi-automatically. Their experiments demonstrated these frameworks work very well on pottery assembly. Curves can also be matched through an optimization algorithms such as [32] that minimize the least square estimation of point patterns. Unlike assembly problems of objects which are simple in structure (like pottery), skull assembly may need different techniques. Without the pre-assumption on axially symmetric and contour smoothness, skulls have subtle facial geometry and details. On the other hand, unlike potteries who may have various distinct shapes, human skulls have similar global geometry, therefore, a template skull (from similar category of races and ages) can be used to assist the assembly. Yin et al.[37] presented an assembly and completion algorithm to first use a template skull to perform a rough assembly then perform assembly refinement based on break curve analysis in a reduced searching space. However, limited by the discrimination power of their descriptor, a very good template is needed to ensure the robustness of the reassembly.

3 Heat Kernel Shape Descriptor

In this section, we briefly recap the basic theory and properties of the heat kernel descriptor; then discuss why such a descriptor is ideal for this reassembly task.

3.1 Mathematical Background

Let M be a compact Riemannian manifold and $u(x,t)$ be the amount of heat at a point $x \in M$ at time t. The heat propagation over M is governed by the *heat diffusion equation*:

$$\begin{cases} \frac{\partial u(x,t)}{\partial t} = -\Delta u(x,t) \\ u(x,0) = f(x) \end{cases}$$

where Δ is the Laplace-Beltrami operator and $f(x)$ is the initial temperature defined on M. If M has boundaries, we additionally require u to satisfy the Dirichlet boundary condition $u(x,t) = 0$ for all $x \in \partial M$ at all t.

Given the initial function f, the solution to this heat equation at time t can be computed through the heat operator H_t:

$$u(x,t) = H_t f. \tag{1}$$

For any M, there exists a function $h_t(x,y)$ [8] that

$$u(x,t) = \int_M h_t(x,y) f(y) dy. \tag{2}$$

The $h_t(x,y)$ satisfying this equation is called the *heat kernel*, and its value can be thought as the amount of heat that is transferred from point x to point y during time t. For a compact surface M, the heat kernel has the following eigen-decomposition [13]:

$$h_t(x,y) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \Phi_i(x) \Phi_i(y) \tag{3}$$

where $\lambda_0, \lambda_1, ...$ are eigenvalues and $\Phi_0, \Phi_1, ...$ are the corresponding eigenfunctions of the Laplace-Beltrami operator, which satisfy $\Delta_M \Phi_i = \lambda_i \Phi_i$.

The *Heat Kernel Signature* (HKS) [30] is a powerful descriptor that characterizes local and global geometry of the surface patch centered at each point:

$$h_t(x) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \Phi_i(x)^2. \tag{4}$$

The HKS inherits many good properties from heat kernel and is therefore effective in describing shapes at different scales and identifying geometric features. For a piecewise linear surface mesh, HKS can be computered from the eigen-values and eigenvectors of the mesh Laplace opertor. This computation is detailed in [30]. We use the sparse eigen-solver in Matlab to compute them.

3.2 Properties of HKS for Geometric Reassembly

The heat kernel and heat kernel signature are desirable for the reassembly of fragmented objects. Specifically, for our fragmented skull assembly, we adopt it as a reliable descriptor for partial matching because it is multi-scale, informative (discriminative), and stable.

3.2.1 Multi-scale Property

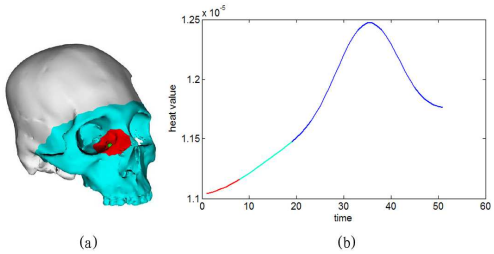

Figure 2: Multi-scale Property. The green point (a), considered in the fragment (red region) and in the whole model (cyan) has the overlapped signature curves (b).

For small values of t, the function $h_t(x,y)$ is mainly determined by a small neighborhood of x; its neighborhood grows to a bigger region as t increases. The multi-scale property of the heat kernel implies, in particular, that for small t, $h_t(x,y)$ only reflects local characteristic of the shape around point x, while for large values of t, $h_t(x,y)$ captures the global structure of M from the view of point x. This intuition can be formalized [10] as: *(a) For any smooth and relatively compact domain $D \subseteq M$, $\lim_{t\to 0} h_t^D(x,y) = h_t^M(x,y)$; and (b) if $D_1 \subseteq D_2 \subseteq \ldots \subseteq D_n$, $\cup_{D_i}^{D_n} = M$, then $\lim_{n\to\infty} h_t^{D_n}(x,y) = h_t^M(x,y)$ for any t.* The explicit relationship between time and the size of diffusion region is also discussed in [9]. Suppose W_x^t is the Brownian motion on M starting at point x, the heat kernel can be viewed as the transition density function of the Brownian motion, and is determined by the probability of W_x^t at time t (see [9, 20] for details).

In our fragment-template matching problem, a pure local descriptor can be easily affected by local noise and geometry disparity, while a global descriptor could not tolerate the intrinsic difference between a *complete* template and an *incomplete* fragment. Therefore, due to HKS's multi-scale property, unlike local descriptors such as spin images[12] or global descriptors such as GPS[23], the HKS allows us to perform multi-scale comparison between different neighboring regions of points on the same shape. Furthermore, the HKS of points on different shapes are commensurable, which allows us to perform the partial matching and registration even when the template and subject fragments have different subtle local geometries. Fig. 2 shows an example of HKS of the corresponding points on the template and the fragment.

3.2.2 Informative Property

Skull models have subtle and detailed geometry, a desirable shape descriptor needs to be discriminative. Different points from different skull locations have different HKS signatures. More rigorously, *let T: $M \to N$ be a surjective map between two Riemannian manifolds. If $h_t^M(x,y) = h_t^N(T(x),T(y))$ for any x, $y \in M$ and any $t>0$, then T is an isometry* [30].

The Varadhan's Lemma [10] indicates that for every pair of points $x, y \in M$, the geodesic distance $d^2(x,y) = -4\lim_{t\to 0} t \log h_t^M(x,y)$. The informative property implies that the heat equation contains all of the information about the intrinsic shape geometry and hence fully characterizes the shape.

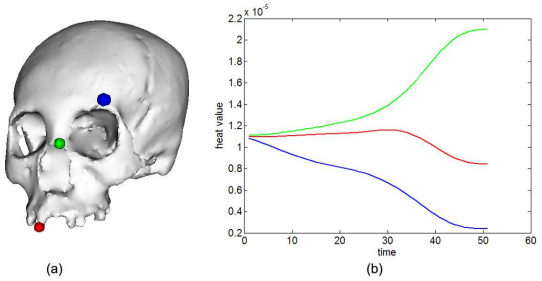

Figure 3: HKS are informative. Each point has a unique heat diffusion curves. Different points have different signatures.

HKS encodes the geometric information about the neighborhoods of a point x at various scales together. This property, together with the multi-scale property which is discussed earlier, make HKS a suitable discriminative descriptor (an example shown in Fig. 3) for our matching purpose .

3.2.3 Scaling-invariant HKS Matching

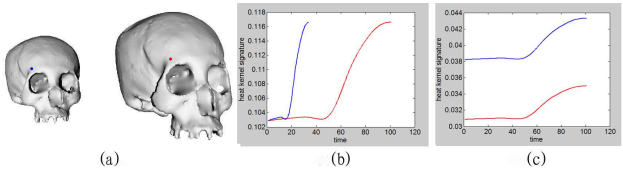

Figure 4: Normalization for Scaling Transformations. (a) shows one skull with two scaling, the right one is twice larger. (b) shows their HKS in the same coordinate, and (c) shows the result of normalization.

We consider the heat diffusion curve under a scaling transformation of an object. Given a model M and its scaled surface $M' = \beta M$, where β is the scaling factor, following eq (4), the new eigenvalues and eigenfunctions will satisfy $\lambda' = \beta^2\lambda$ and $\phi' = \beta\phi$. So we have the following equation:

$$h'_t(x) = \sum_{i=0}^{\infty} e^{-\lambda_i t\beta^2} \Phi_i(x)^2 \beta^2. \qquad (5)$$

This means HKS changes under the scaling transformation, relating the signature h' at time t for M' with the β^2 times of the signature h at time $\beta^2 t$ for M.

Fragments are scanned separately, so the scales of these digital models are usually inconsistent. Without a scaling-invariant descriptor and matching scheme, we need to preprocess the original skulls. A typical approach is placing markers and measure their distances, then using these distances to re-scale all digital fragment models into coherent sizes with respect to the template model. Such approach is tedious, error-prone, and could contaminate the original skull. So we want to seek for a geometric algorithm that can handle partial matching between models with inconsistent scaling.

Bronstein and Kokkinos [4] suggested a preprocessing to make HKS a scale invariant vision. It is based on a logarithmically sampled scale-space in which shape scaling corresponds, up to the multiplicative constant β, to a shift. And this shift is then undone by taking the discrete-time Fourier transform and Fourier transform modulus

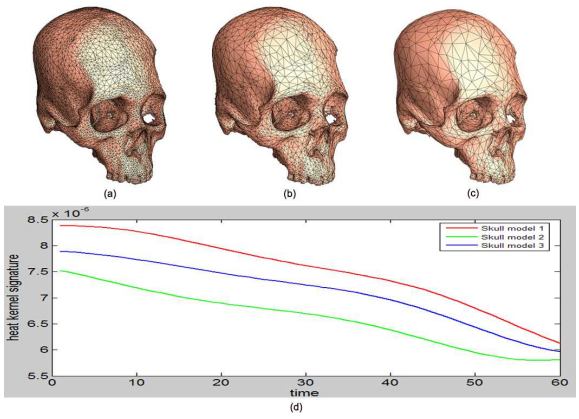

Figure 5: Signatures under Different Resolutions. (a,b,c) shows a skull with 35k, 20k, and 10k faces, respectively. Corresponding points in three skulls have very similar signature curves.

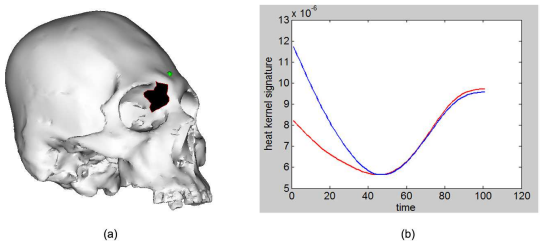

Figure 6: Robustness of HKS. The green point on an incomplete skull (a) has a similar signature (b, the blue curve) to the signature on the completed skull (b, the red curve).

(FTM). We adopt this approach. Fig. 4 illustrates the scale-invariant property on skulls of different scaling after the preprocessing.

3.2.4 Stability

Being scanned separately, different fragments may not only have different scaling, but also have different sampling and tessellations. Holes and local noise may exist due to occlusions, low reflectance, or other reasons. Therefore, a descriptor tolerating sampling resolution and geometric noise is desirable.

On discrete surfaces, [30] uses the mesh Laplace operator [2] to estimate the Laplace-Beltrami operator and shows that HKS is insensitive to the meshing and resolutions. This helps us to make our matching insensitive to different resolutions, as shown in Fig. 5.

Heat kernel is stable against local noise (e.g. small local geometric perturbation) due to the nature of heat diffusion process on the manifold. The heat kernel can be considered as a Brownian motion, which means that $h_t(x, y)$ is a weighted average possibility over all paths to reach y from x in time t [10]. Suppose W_x^t is the Brownian motion on manifold M, if we slightly perturb a subregion $M' \subset M$, only the paths passing through M' will get affected. The signature is therefore relatively stable. For a perturbation applied on a point x, the variation of HKS on x is significant for small t, but this decays as t increases. Fig. 6 shows an example of the robustness of HKS.

4 Skull Assembly Pipeline

Given a set of fragments $\{F_i\}$ (Fig. 7(a)), we shall compute transformations $\{T_i\}$ that transform these fragments to align with a given template M. Transformed fragments $\{T_i(F_i)\}$ will then roughly reassemble the original geometry of the subject skull (Fig. 7(b)). We call this step *coarse assembly*. After the coarse assembly, we further refine the assembly by analyzing and matching break boundary curves of fragments. The pipeline has 3 steps.

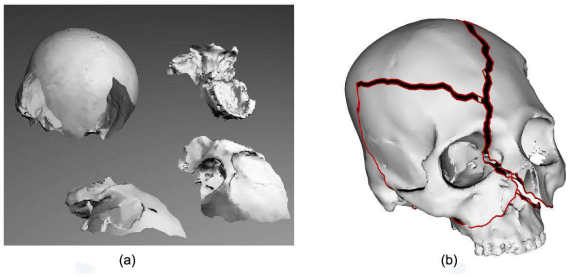

Figure 7: Coarse Assembly. (a) shows the skull fragments (S2); (b) shows the reassembled skull after Steps (1) and (2).

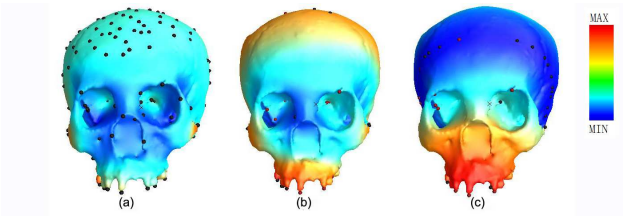

Figure 8: Multi-scale Feature Detection. The color indicates the heat value of the point, and features are extracted in different scales. (a) k = 0, (b) k = 60, (c) k = 100.

(1) Detect features by analyzing the heat field and heat diffusion;

(2) Extract a superset of initial matches by HKS, then compute a smaller set of final matches from these candidates using RANSAC, then compute the rigid transformations;

(3) Match the break curves, refine the assembled fragments, then complete the skull.

4.1 Feature Detection

On two given shapes M (template) and F (fragment), we firstly compute heat diffusion on M. A point v_i is called a feature in scale k, if $F^k(v_i)$ is the local maximum or minimum within its 2-ring region [34], where $F^k(v_i)$ means the heat value of v_i in scale k.

In our experiments, we compute the HKS signature curve for each model by uniformly sampling 100 points in the logarithmically scale over the time interval $[t_{min}, t_{max}]$ with $t_{min} = 0, t_{max} = 4 \log \frac{10}{\lambda_2}$, where λ_2 is the smallest eigenvalue and k is the step. For a small k, the detected features mainly encode local geometry while they characterize geometry more globally for a larger k. Therefore, long protrusions such as tooth tips can be extracted as local extremals in a large-scale k.

Fig. 8 shows the features detected in different scales. The features detected in the scale $k = 0$ are easily affected by noise. As k increases, the heat field becomes smoother gradually. The features detected in different scales can well depict the shape information in a multi-scale sense. In our experiments, we found $k = 60$ to be a suitable threshold (since HKS is not sensitive to the resolution, we found k values between 40 and 80 all provide relatively good and similar results).

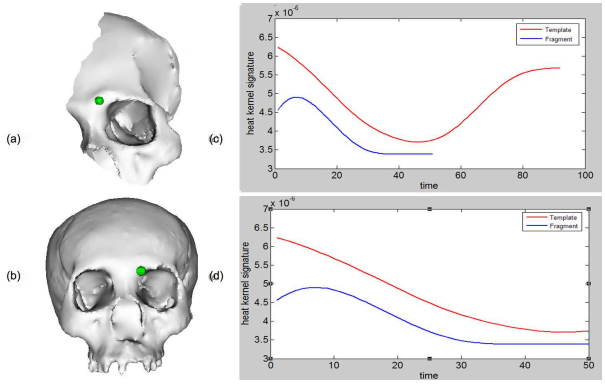

Figure 9: This figure shows the comparison process. (a) is the fragment and (b) is the template. (c) plots the HKS curves of the corresponding points on (a) and (b); (d) illustrates their common portion.

4.2 Coarse Matching

The coarse matching in step 2 of our assembly pipeline contains two sub-steps, which will be discussed in the following subsections.

4.2.1 Initial Matching

On both template M and fragments set $F = \{F_i\}$, where F_i are the fragments pieces, we first put all the features detected by HKS into candidate sets, denoted as $\mathcal{P} \subseteq F$ and $\mathcal{Q} \subseteq M$, respectively.

Then for each point $p_i \in \mathcal{P}$, we search for its corresponding features $q_i \in \mathcal{Q}$. We correspond two points if their HKS curves match well. We use $\Phi^F_{p_i}$ to denote the heat diffusion curve p_i on F and use $\Phi^F_{p_i}(k)$ to indicate the heat value in scale k.

Unlike isometric matching [34][22], corresponding points in the fragment and template may have different initial HKS and different diffusion time, therefore, simple distance such as $\sum_{k=1}^{n} ||\Phi^F_{p_i}(k) - \Phi^M_{q_i}(k)||$ is not suitable for our problem.

We modify the matching of HKS and develop different metrics at a given range of scales. Firstly, to make the time t meaningful and stable for different surfaces, we normalize the HKS of every match into a same coordinate system, and compare the common portion. We then evaluate the quality of the match (p_i, q_i) by computing the difference between $\Phi_{(p_i)}{}^F$ and $\Phi_{(q_i)}{}^M$ in their common diffusion time:

$$E_{(p_i,q_i)}(k) = ||\Phi^F_{p_i}(k) - \Phi^M_{q_i}(k)||, \quad (6)$$

where $E_{(p_(i,),q_i)}(k)$ is a function to evaluate the matching (p_i, q_i) and $E_{(p_(i,),q_i)}(k)$ indicates the difference in scale k. And finally we use the variance $D(E)$ to evaluate the difference between two HKS so as to find the closest curve or the best match for a given point. Fig. 9 demonstrates this process.

After the computation of feature points and descriptors, we can build an initial matching graph, correlating each feature point on M and its n most similar matches on F. In our experiments, we set $N = 5$. We use such a conservative threshold (which may leads to quite a few false positives, but usually guarantee at least one reliable matching), and we will refine this matching in the next section. Fig. 10(a) shows an example of the result of this step.

4.2.2 RANSAC Refinement

The initial matching provides a conservative many-to-many correspondence, where the correct match for each point is included but many incorrect matching are also involved. Given such a candidate matching graph (typically, 4 out of 5 matches are wrong), we need a filter to eliminate these wrong matches.

We develop such a filtering scheme based on the RANSAC strategy, which is a classical techniques for parameter estimation that can optimize the fitting of a functional description to presented data. Rather than using as much as possible of the data to obtain an initial solution then attempting to eliminate the invalid data points, RANSAC uses a small initial data set and enlarges this set with consistent data when possible [7][36]. Therefore, only a small number of random guesses to find a candidate set is needed to start the process. Intuitively, although we could skip the initial matching step discussed in the previous section and enumerate all possible matches between template and fragments and evaluate all their validity, this will lead to an inefficient process with big time consumption and poor accuracy.

By using HKS, the approximate features can be obtained on the template and fragment, it is the result of the global similarity of skulls. Tevs et al. [31] developed a RANSAC subgraph extraction method. We compute the Euclidean distances of every correct matching here (which is enough for our problem), and the aggregate of these distance errors, for a correct match, yields a Gaussian distribution:

$$p(l_1, ..., l_{n(i,j)}|m_{i,j}) = \frac{1}{\sigma^m (2\pi)^{m/2}} \prod_{k=1}^{m} exp(\frac{l_k - l_k^{(0)}}{2\sigma^2}), \quad (7)$$

where l_k denotes the actual distance observed and $l_k^{(0)}$ the correct distance. Then, Bayes rule is employed to compute the probability of match $m_{i,j}$ being correct. $p(m_{i,j})$ is the prior probability of match, which in our case is given by the descriptor matching.

$$p(m_{i,j}|l_1, ..., l_{n(i,j)}) = \frac{p(l_1, ..., l_{n(i,j)}|m_{i,j})p(m_{i,j})}{\sum_{i,j} p(l_1, ..., l_{n(i,j)}p(m_{i,j}))} \quad (8)$$

Based on the above discussion, we use a RANSAC randomized sampling algorithm, judged by the probability, to evaluate how well the matches preserve the Euclidean distance. The algorithm is formulated as follows. Fig. 10 shows an example.

In: fragment F, template T, the max iteration number k, and a matching candidate super-set C;

Out: the refined matching set S;

0) $S = \emptyset$;

1) Select 3 matches from C randomly (which can be used to solve a rigid transformation) and stored in temperate set D;

2) Enumerate each of the rest candidates in $c_i \in C$ and evaluate whether for this match, eq 8 exceeds a threshold ϵ, if so, add this match c_i into D;

3) Check whether the size of D is larger than the size of S, if so, $S \leftarrow D$;

4) $D \leftarrow \emptyset$; Go to step 1, unless the max iteration number k is reached.

A threshold ϵ that is not too strict is usually desirable. In our experiments, we set $k = 500$ and we found $\epsilon = 90\%$ always leads to good results.

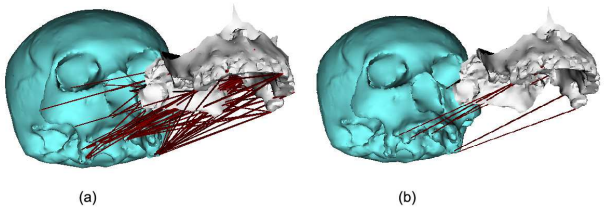

Figure 10: This figure shows input and output of the RANSAC filter, (a) is the superset of matches which includes many wrong matches and (b) is the final matches sifted by the filter.

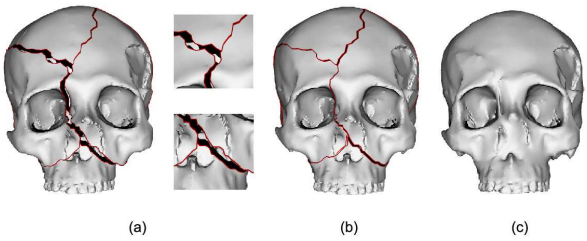

Figure 11: Postprocessing: assembly refinement and skull completion. (a) shows the reassembled skull (S2) after rough assembly; (b) shows the result after break curve matching and assembly refinement; (c) is the final completed skull.

4.2.3 Local Registration

After the correct matching is computed, we can compute the rigid transformation T_i for each fragment by solving an over-determined system:

$$T_i \begin{pmatrix} p_i^1 \\ p_i^2 \\ \dots \\ p_i^n \end{pmatrix} = \begin{pmatrix} q_i^1 \\ q_i^2 \\ \dots \\ q_i^n \end{pmatrix},$$

where T_i is the rigid transformation on F_i. As long as we have more than 3 pairs of corresponding features, we can solve the transformation and thus reassemble the fragments. Fig. 7(b) shows a coarse assembly example.

4.3 Postprocessing and Skull Completion

With the help of template skull, the fragments are roughly assembled after the above steps. However, such a coarse reassembly may not be very accurate. Gaps or intersections between fragments may exist in the above assembling result. Sometimes this is due to the geometric difference between the (fragmented) subject skull and the template skull. We further refine the reassembly by analyzing the local geometry of these fragments, using the method introduced in [37]. The rough assembly result is refined through an optimization of the least square transformation error (LSTE) of break-curves. Finally, when the reassembly is finished. The missing or damaged regions (e.g. holes) are repaired using the inherent symmetry [37] of the skull. Fig. 11 shows the refine and completion process.

5 Experimental Results

In our experiments, all the skulls are scanned and provided by forensic anthropologists. We used our assembly algorithm to repair four sets (two male skulls $s1$, $s2$ and two female skulls $s3$, $s4$) of fragmented skulls. In practice, a template M is selected from an organized skull database, in which skulls are classified by the sex, race, age, cranial form, and etc. Given an incomplete subject

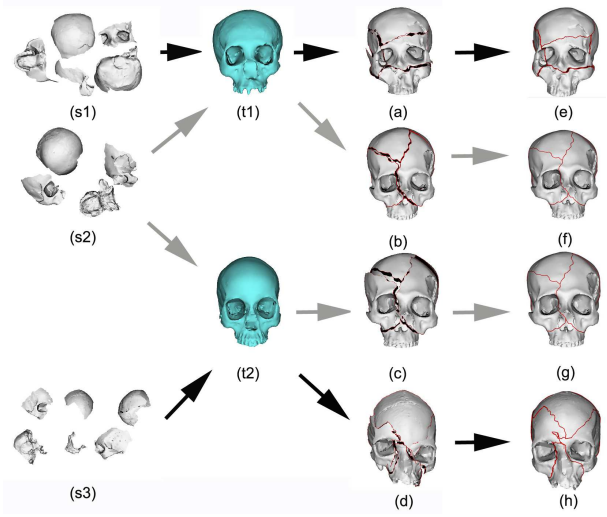

Figure 13: (s1 - s3) are the fragmented skulls. (t1) and (t2) are the templates. The different coarse reassembling results are shown in (a) - (d); (e) - (h) show the results after the refinement guided by break curve matching.

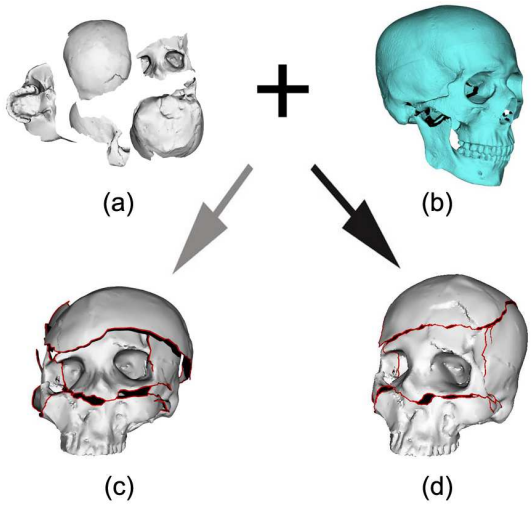

Figure 14: Comparison of our proposed reassembly and the algorithm of [37]. The fragments (a) are assembled using the template (b). (c) is the result from [37] and (d) is the result of our method.

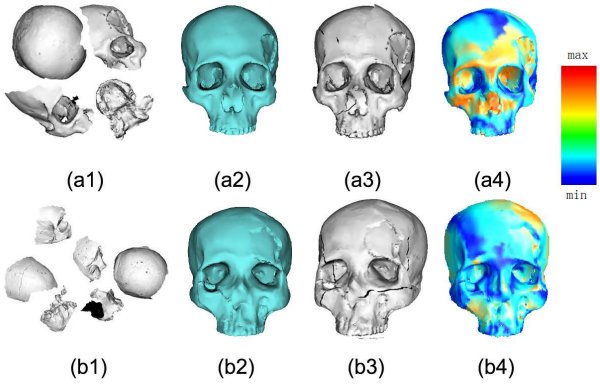

Figure 15: Reassembly of Simulated Fracturing. (a1) and (b1) are fragments partitioned from the complete models (a2) and (b2) respectively. (a3) and (b3) are reassembled skulls; (a4) and (b4) color-encode the reassembly errors. The red indicates larger deviation while the blue indicates smaller errors ($Min = 0.0$ and $Max = 3\%$ of the length of the bounding box).

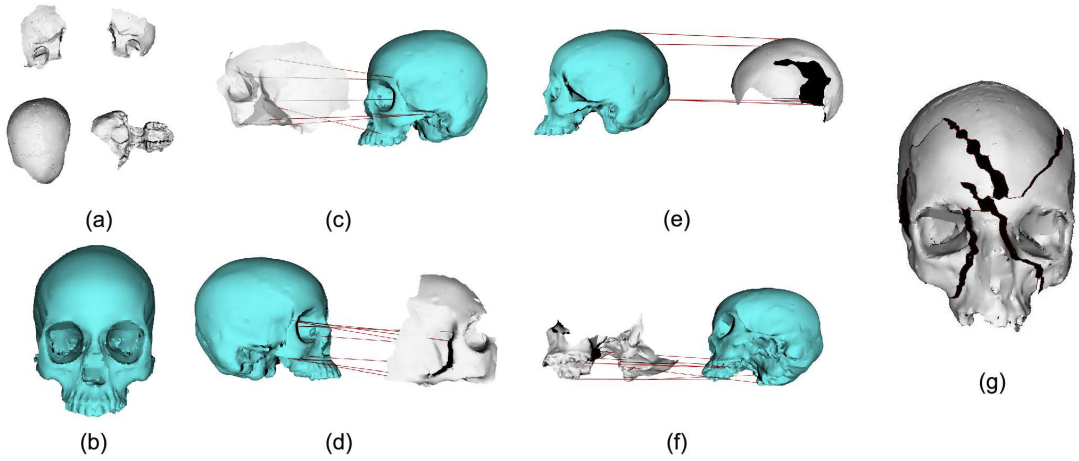

Figure 12: The Assembly of Skull $s4$. (a) shows the fragments of $s4$ and (b) is the template. (c-f) show the matching from fragments to the template. (g) is the coarse reassembly result (before refinement).

Table 1: Runtime Table: The fragmented skulls to be completed and the template skulls are listed in the columns *Subject* and *Template*, respectively. $\#\Delta$(K): the number of thousand triangles in the mesh; $\#F$: number of fragments; T_{HKS}: the time of computing HKS in seconds; T_{RAN}: the time of RANSAC process with 500 iterations. T_{Com}: the time of post-processing and skull completion. Experimental time is measured in seconds.

Skulls	Temp	$\#\Delta$(K)	$\#F$	T_{HKS}	T_R	T_{Com}
S_1	T_1	35.1	4	328.2	6.9	27.8
S_1	T_2	35.1	4	341.6	7.2	28.1
S_2	T_1	37.2	5	321.9	7.3	28.4
S_2	T_2	37.2	5	333.7	7.5	29.0
S_3	T_1	52.4	6	492.3	6.2	40.3
S_3	T_2	52.4	6	489.1	6.3	38.8
S_4	T_1	43.8	6	400.6	6.1	36.3
S_4	T_2	43.8	6	402.2	6.2	37.9

skull, anthropologists analyze their anthropometry features (lengths/ratios of distances between features), predict which category this skull belongs to, and select a template skull from the same category as the template. We perform our assembly algorithms on a 2.4GHz desktop with 2GB RAM. The runtime table is given in Table. 1. The assembling results of the skulls $s1$, $s2$ are shown in Fig. 1 (e) and Fig. 11 (c), and the assembly of the female skull $s3$ is shown in Fig. 13 (d). Fig. 12 illustrates the assembly of skull $s4$,

In Fig. 13, we show our experiments to reassemble fragments using different templates. For $s2$, we use two templates $t1$ and $t2$ for the matching. Partial matching against different templates leads to slightly different coarse assembly, but after the refinement, the results are similar. Our entire pipeline therefore is not very sensitive to the selection of templates.

Fig. 14 shows the comparison between our proposed method and the algorithm of [37]. Due to relatively significant difference between the subject skull and the template, the coarse assembly of [37] fails and can not be fixed successfully by the break curve matching. In contrast, our proposed approach demonstrates better reliability against the not-similar template and successfully recompose the subject skull.

In the above experiments, we only have the scanned skull fragments without the ground truth complete skull. We are not able to quantitatively measure the assembly error. We also develop experiments to evaluate our algorithms by intentionally breaking our complete skulls into pieces. We then reassemble these fragments and measure the deviation between the result and the original model. The numerical errors illustrating this accuracy are plotted in Fig. 15. The reconstruction error is smaller than 3%.

6 Conclusion

We introduce an geometric reassembly algorithm for fragmented skull completion. We employ a multi-scale descriptor based on heat kernel and analyze the its several desirable properties in geometric reassembly and in our task. Then we develop a partial matching algorithm based on this descriptor. We integrate our developed algorithm into the skull assembly pipeline, which mainly consists of three components: matching computation between the fragments and template, rough assembly computation, and matching refinement. The new scheme improves the efficacy of the first two steps, and our rough assembly results therefore outperform [37].

A limitation of the current assembly algorithm is in handling tiny fragmented pieces. If a tiny fragment does not have enough salient geometric features, then its matching with template is difficult and unreliable. Right now we give up reassembling such fragments and rely on the subsequent hole filling algorithm to repair the corresponding missing region. Also, refinement guided by the break curve matching is sometimes not reliable if the boundary is worn or partially damaged. Currently after the refinement of each fragment, if the matching error of break curves with adjacent parts is bigger than a threshold, we skip the refinement on this fragment and restore the assembly suggested by the template-subject matching. We will explore a better refinement strategy to deal with this issue.

7 Acknowledgements

The skull models are from LSU Forensic Anthropology and Computer Enhancement Services (FACES) lab, provided by Warren Waggenspack from LSU Mechanical Engineering Department. This project is partially supported by Louisiana Board of Regents (LA-BOR) Research Competitiveness Subprogram (RCS) LEQSF(2009-12)-RD-A-06, LA-BOR PFund:NSF(2011)-PFund-236, and National Natural Science Foundation of China No. 61170323. We

thank the anonymous reviewers for their constructive comments; and we thank Zhao Yin and Li Wei for valuable discussions.

References

[1] G. Barequet and M. Sharir. Partial surface and volume matching in three dimensions. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 19:929–948, 1994.

[2] M. Belkin, J. Sun, and Y. Wang. Discrete laplace operator on meshed surfaces. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, SCG '08, pages 278–287. ACM, 2008.

[3] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(4):509 –522, apr 2002.

[4] M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, 2010.

[5] D. B. Cooper, A. Willis, S. Andrews, J. Baker, Y. Cao, D. Han, K. Kang, W. Kong, F. F. Leymarie, X. Orriols, S. Velipasalar, E. L. Vote, M. S. Joukowsky, B. B. Kimia, D. H. Laidlaw, and D. Mumford. Assembling virtual pots from 3d measurements of their fragments. In *Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, VAST '01, pages 241–254. ACM, 2001.

[6] F. de Goes, S. Goldenstein, and L. Velho. A hierarchical segmentation of articulated bodies. In *Proceedings of the Symposium on Geometry Processing*, SGP '08, pages 1349–1356, 2008.

[7] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. In M. A. Fischler and O. Firschein, editors, *Readings in computer vision: issues, problems, principles, and paradigms*, pages 726–740. 1987.

[8] Y. E. Gliklikh. Stochastic analysis on manifolds. In *Global and Stochastic Analysis with Applications to Mathematical Physics*, Theoretical and Mathematical Physics.

[9] A. Grigor′yan. Escape rate of brownian motion on riemanian manifolds. *Applicable Analysis*, 71(1-4):63–89, 1998.

[10] A. Grigor′yan. *Heat kernels on weighted manifolds and applications*. Amer Mathematical Society, 2006.

[11] K. Hori, M. Imai, and T. Ogasawara. Joint detection for potsherds of broken earthenware. In *Computer Vision and Pattern Recognition (CVPR).*, volume 2, 1999.

[12] A. Johnson. *Spin-Images: A Representation for 3-D Surface Matching*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, August 1997.

[13] P. W. Jones, M. Maggioni, and R. Schul. Manifold parametrizations by eigenfunctions of the laplacian and heat kernels. *Proceedings of the National Academy of Sciences of the United States of America*, 105(6):pp. 1803–1808, 2008.

[14] T. Ju. Robust repair of polygonal models. *ACM Trans. Graph.*, 23:888–895, 2004.

[15] W. Kong and B. Kimia. On solving 2d and 3d puzzles using curve matching. In *CVPR*, volume 2, pages II–583 – II–590, 2001.

[16] S. Lafon, Y. Keller, and R. Coifman. Data fusion and multicue data matching by diffusion maps. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 28(11):1784 –1797, nov. 2006.

[17] X. Li, Z. Yin, L. Wei, S. Wan, W. Yu, and M. Li. Symmetry and template guided completion of damaged skulls. *Computers and Graphics*, 35(4):885 – 893, 2011.

[18] P. Liepa. Filling holes in meshes. In *Proc. Symp. Geometry processing*, pages 200–205, 2003.

[19] M. Manhein, G. Listi, R. Barsley, R. Musselman, N. Barrow, and D. Ubelaker. In vivo facial tissue depth measurements for children and adults. *J Forensic Sci.*, 45(1):48–60, 2000.

[20] F. Memoli. Spectral gromov-wasserstein distances for shape matching. In *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, 2009.

[21] M. Ovsjanikov, A. M. Bronstein, L. J. Guibas, and M. M. Bronstein. Shape google: a computer vision approach to invariant shape retrieval. In *Proc. NORDIA*, 2009.

[22] M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. J. Guibas. One point isometric matching with the heat kernel. *Comput. Graph. Forum*, 29(5):1555–1564, July 2010.

[23] M. Ovsjanikov, J. Sun, and L. Guibas. Global intrinsic symmetries of shapes. In *Proceedings of the Symposium on Geometry Processing*, SGP '08, pages 1341–1348, 2008.

[24] G. Papaioannou and E.-A. Karabassi. On the automatic assemblage of arbitrary broken solid artefacts. *Image Vision Comput.*

[25] S. Ruiz-Correa, L. Shapiro, and M. Melia. A new signature-based method for efficient 3-d object recognition. In *Proc. CVPR*, volume 1, pages I–769 – I–776, 2001.

[26] R. M. Rustamov. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the fifth Eurographics symposium on Geometry processing*, pages 225–233. Eurographics Association, 2007.

[27] A. Sharma and R. Horaud. Shape matching based on diffusion embedding and on mutual isometric consistency. In *CVPR Workshops (CVPRW)*, pages 29 –36, 2010.

[28] G. Sharp, S. Lee, and D. Wehe. Icp registration using invariant features. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(1):90 –102, jan 2002.

[29] F. Stein and G. Medioni. Structural indexing: Efficient 3-d object recognition. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14:125–145, 1992.

[30] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proceedings of the Symposium on Geometry Processing*, SGP '09, pages 1383–1392, 2009.

[31] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel. Isometric registration of ambiguous and partial data. In *CVPR*, 2009.

[32] S. Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13(4):376 –380, 1991.

[33] A. Vaxman, M. Ben-Chen, and C. Gotsman. A multi-resolution approach to heat kernels on discrete surfaces. *ACM Trans. Graph.*, 29:121:1–121:10, July 2010.

[34] S. Wang, T. Hou, Z. Su, and H. Qin. Multi-scale anisotropic heat diffusion based on normal-driven shape representation. *The Visual Computer*, 27:429–439, 2011.

[35] A. Willis and D. Cooper. Bayesian assembly of 3d axially symmetric shapes from fragments. In *Proc. CVPR*, volume 1, pages I–82 – I–89, 2004.

[36] S.-W. Yang, C.-C. Wang, and C.-H. Chang. Ransac matching: Simultaneous registration and segmentation. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 2010.

[37] Z. Yin, L. Wei, M. Manhein, and X. Li. An automatic assembly and completion framework for fragmented skulls. In *International Conference on Computer Vision (ICCV) 2011*, 2011.