

Automatic Detection of Anatomical Features on 3D Ear Impressions for Canonical Representation

Sajjad Baloch¹, Rupen Melkisetoglu¹, Simon Flöry², Sergei Azernikov¹,
Greg Slabaugh³, Alexander Zouhar¹, and Tong Fang¹

¹ Siemens Corporate Research, Princeton, NJ, USA

² Vienna University of Technology, Wien, Austria

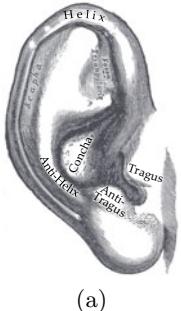
³ Medicsight PLC, London, UK

Abstract. We propose a shape descriptor for 3D ear impressions, derived from a comprehensive set of anatomical features. Motivated by hearing aid (HA) manufacturing, the selection of the anatomical features is carried out according to their uniqueness and importance in HA design. This leads to a *canonical ear signature* that is highly distinctive and potentially well suited for classification. First, the anatomical features are characterized into *generic* topological and geometric features, namely *concavities*, *elbows*, *ridges*, *peaks*, and *bumps* on the surface of the ear. Fast and robust algorithms are then developed for their detection. This indirect approach ensures the generality of the algorithms with potential applications in biomedicine, biometrics, and reverse engineering.

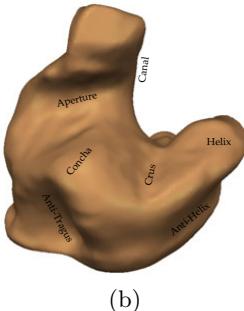
1 Introduction

Shape analysis typically involves the abstraction of complex structures by removing redundant details. It captures the essence of geometry via morphological descriptors that emphasize more on informative and distinctive features of the underlying surface. Such descriptors may be employed for classification in addition to guiding registration [1] and segmentation. The problem is very challenging due to the variability of organic surfaces. Moreover, it is typically not possible to consistently identify key features solely from geometric and topological information, and the need of anatomical features becomes imminent.

We focus on the surfaces representing the external and outer human ear and construct a canonical ear signature (CES). The idea is to derive a descriptor for subsequent modeling of the anatomy of the human ear. This will eventually lead to a 3D digital human ear atlas, and a framework for the design of implants and prosthetic devices for the ear. This is in line with the recent trends in medical image analysis, where the imaging technology is exploited for data modeling, intervention planning, and corrective treatment. Our approach herein is to derive the CES from the anatomy itself, which is quite well known in the medical community. The first comprehensive atlas dates back to Gray in 1858 (Fig. 1(a), with a 3D reconstruction in Fig. 1(b)). As shown, an ear impression consists of a spindle shaped *canal* that sits deep in the outer ear and a base that resides



(a)



(b)

Fig. 1. The human ear anatomy (courtesy [2]) and its 3D reconstruction

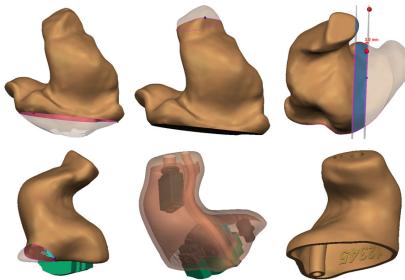


Fig. 2. Shell detailing and modeling. Last item is the finished shell

in the external ear. The two are separated by a narrow opening called *aperture*. Around the bottom of the external ear, there are two clamp like structures known as *tragus* and *anti-tragus*. The deepest point in the external ear is the *concha peak* that sits on the big bowl shaped *concha*. On the top, there is a narrow *helix* structure sandwiched between the *anti-helix* and the *crus*. These anatomical parts along with others are considered for the construction of the CES. Once a 3D impression of an ear is acquired, the problem then reduces to extracting its signature representation. To this end, we adopt a generalized approach, and first characterize the anatomical features into more generic topological and geometric features, such as *peaks*, *pits*, *concavities*, *elbows*, *ridges*, and *bumps*. Fast and robust algorithms are then constructed for the detection of generic features. As a result, the application of proposed algorithms is not limited to ears. For instance, the ridge detection algorithm may readily be used for tracing the nose in face recognition, or the crest lines on human brains.

There is substantial literature on surface feature detection [3,4,7]. However, not much work has been reported on ear shape analysis. We introduce three major contributions: (a) It is the first work that comprehensively analyzes the **outer** and **external** ear shape to automatically determine the CES. (b) Due to the varied nature of the features, a collection of algorithms are proposed to produce an overall system capable of their reliable detection. These methods enable automation in the HA manufacturing (HAM) [6] and have the potential of faster and more consistent design. (c) The resulting signature will serve as a foundation to a 3D digital atlas of the human ear. In comparison to [3,7], our algorithms are fast, and guarantee protection against broken ridges. We also consider the detection of planes, which has not been attempted previously.

2 Canonical Ear Signature

First, we identify the anatomical constructs that uniquely identify the shape of an ear impression. To this end, we are inspired by the HAM application. HAs are generally custom made. Once a mesh is constructed from the 3D scan of an ear

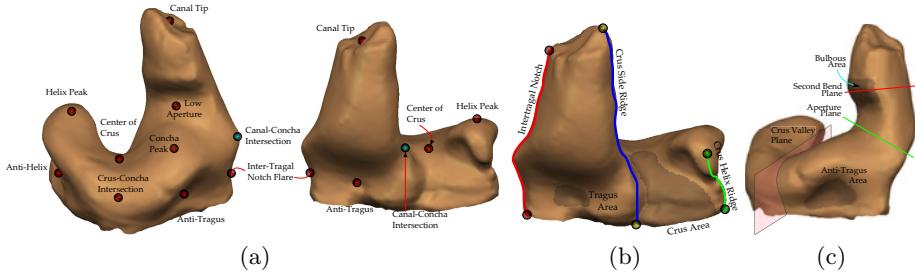


Fig. 3. Anatomical Features: (a) Points; (b) Curves and areas; (c) Planes and areas

impression, an operator carries out a sequence of surface modification operations (Fig. 2), to design an HA shell [6]. Typically, he relies on manual measurements (on sub-mm scale) implicitly using certain key anatomical features. Our hypothesis is that since an HA is designed to comfortably fit the ear(s) of a patient, the underlying features form a canonical representation of the anatomy. We have, therefore, identified a comprehensive set of features that are well known in the medical anatomy, and are implicitly used by HA designers (Fig. 3). Collectively, they capture the structure of an ear in a *Canonical Ear Signature*.

3 Generic Feature Detection

Our approach to canonical feature detection is based on generalization. We characterize various anatomical features via a set of generic features, *peaks* $\{\mathcal{F}_p\}$, *concavities* $\{\mathcal{F}_c\}$, *elbows* $\{\mathcal{F}_e\}$, *ridges* $\{\mathcal{F}_r\}$, and *bumps* $\{\mathcal{F}_b\}$. Although some anatomical features may not be represented by these geometric primitives, they may still be derived from the latter or other derived features. Algorithms are constructed for the generic features, thereby later providing a foundation for the derived features. An advantage of such a hierarchical approach is that the generic algorithms may readily be applied to other applications such as biometrics, and reverse engineering, without compromising on the complexity. The relationship between generic and anatomical features is tabulated in Table 1.

3.1 Problem Formulation

Let \mathcal{M} be a 2D manifold representing a 3D ear impression embedded in \mathbb{R}^3 locally parameterized as $\phi : \Omega \rightarrow \mathcal{M}$, where $\Omega \subset \mathbb{R}^2$ represents the parameter space. The problem is to detect a set of generic features $\mathcal{F} = \{\mathcal{F}_p \cup \mathcal{F}_c \cup \mathcal{F}_e \cup \mathcal{F}_r \cup \mathcal{F}_b\}$.

3.2 Peak Detection

A peak point is a prominent topological landmark on a surface. The basic idea for its detection emanates from *Morse theory*, and involves defining a smooth real valued function $f : \mathcal{M} \rightarrow \mathbb{R}$ on \mathcal{M} to abstract its shape via critical points. A point $p = \phi(u) \in \mathcal{M}, u \in \Omega$ is a *critical* point of f if the gradient of $f \circ \phi$ vanishes

Table 1. Generic Characterization of Anatomical Features

Characterization	Anatomical Features
Peak	Canal Tip, Concha Peak, Helix Peak
Concavity	Tragus, Anti-Tragus, Anti-Helix, Center of Crus
Elbow	First Bend, Second Bend
Ridge	Inter-Tragal Notch, Crus-Side Ridge, Crus-Helix Ridge
Bump	Canal Bulbous

at u . A critical point $p \in \mathcal{M}$ is regarded as *non-degenerate* if the Hessian $\nabla^2 f \circ \phi(u)$ is non-singular at $\phi(u)$. f is called *Morse*, if all of its critical points are non-degenerate. The problem, therefore, reduces to finding a suitable Morse function. For peak detection, we employ a *height function*, $h : \mathcal{M} \rightarrow \mathbb{R}$, that assigns to each point $p(x, y, z) \in \mathcal{M}$ a value equal to its height, $h(p) := h(x, y, z) = z$. For a non-degenerate surface, h is a Morse function and its critical points are the peaks, passes and pits of the surface. We use it for peak detection as in [4].

The algorithm for detecting the critical points of h follows from *Morse deformation lemma*, and analyzes the level sets of h (horizontal planes) for topological changes. By gradually increasing $h \in [0, H]$ in K steps, we find the intersections of the surface with corresponding planes. Intersections are subsequently analyzed for changes in topology between two successive planes. If a change in topology is detected, we notice the existence of a critical level between them, and zoom in to analyze the surface with larger K (effective $K \leftarrow K^2$, since the interval is smaller). The process is repeated until convergence to a critical point.

3.3 Concavities

Concavities are marked by depressions on a surface. For their detection, we utilize orthogonal scans on a surface to generate a surface profile that is composed of the intersection contours. Spline representations of individual contours are analyzed for variations in signed curvature, where the negative sign identifies a concavity. First, a profile in one direction is considered, and subsections of contours with negative curvature are identified. For these subsections, the points of least curvature are found, with their average computed as a seed point. This seed point is corrected by a scan, orthogonal to the previous scan, shifting it towards the lowest curvature point. Consequently, the seed point is pushed deeper in the valley. The process is repeated iteratively to achieve the absolute local minimum similar to minimization by alternating variables. Once the center of concavity is identified, region growing based on negative curvature is used to determine the concave region.

3.4 Elbows

The detection of elbows in the presence of noise and bumps is quite challenging, which limits the use of only the curvature. Our approach, first identifies points of high curvature on a surface followed by a selection/rejection strategy to fit a plane along the elbow. This leads to a robust algorithm when compared with computations based solely on the skeleton.

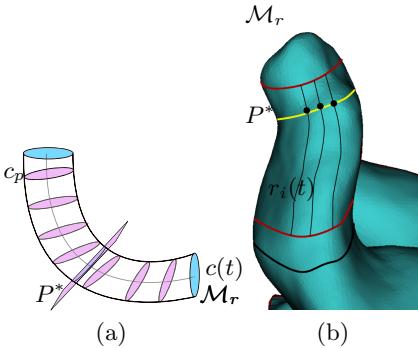


Fig. 4. Elbow Detection: (a) Scanning; (b) Two planes (shown red) define the ROI \mathcal{M}_r in the computation of t_c -sensitive elbow (yellow). Black radial lines represent the spline representations r_i .

The proposed algorithm relies on a predefined region of interest (ROI) $\mathcal{M}_r \subset \mathcal{M}$. First, \mathcal{M}_r is scanned with planes P_t oriented along the centerline $c(t) : t \in [0, 1]$, to generate its x-sectional profile $\{c_p\} = \{\mathcal{M} \cap P_t, \forall t \in [0, 1]\}$ (Fig 4(a)). The profile contours, thus, capture information about bends on the surface. Correspondences are then established among these contours along the radial direction (Fig 4(b)), and each radial contour is eventually represented by its spline parametrization $r_i(t)$ for robustness to noise. Once radial contours are parameterized, we proceed to identify points $Q = \{q_i\}$ of maximal curvature along these contours. However, not all points reliably represent an elbow, due to potential presence of bumps. The set Q is, hence, pruned via a point rejection strategy for plane fitting similar to the deterministic RANSAC to increase robustness.

3.5 Ridge Detection

We define a ridge $r : [0, 1] \rightarrow \mathcal{M}$ as a geodesic on \mathcal{M} characterized by points of high curvature, $\kappa : \mathcal{M} \rightarrow \mathbb{R}$. Note that local extrema of κ alone are not sufficient for its detection, as isolated points may exhibit high curvature due to noise. This is a major limitation of existing methods [3] in addition to their inherent inability to maintain the integrity of a ridge if not prominent, or if broken.

At times, it is also required to identify a ridge in addition to its detection. We, therefore, develop a method that first detects its starting and ending points p_s and p_e respectively, using either the geometric information and/or some a priori knowledge. A geodesic is then run between the two points. We are specifically interested in a geodesic g that minimizes the cost of going from p_s to p_e , where the cost is defined as a weighted combination of the geodesic distance and surface curvature: $\mathcal{C}(g) = \int_0^1 w(\kappa(t))g(t)dt$, where $g(0) = p_s$ and $g(1) = p_e$ and $w(\kappa(t))$ is selected as a decreasing functional of curvature. Hence, the ridge is a minimizer of \mathcal{C} , and may easily be computed for a triangulated mesh, through Dijkstra's algorithm with curvature weighted edge lengths. Consequently, the resulting

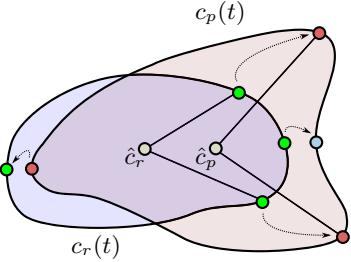


Fig. 5. Dotted lines indicate correspondence between c_p and c_r . Corresponding points (q_p, q_r) are compared via their distances from individual centroids (\hat{c}_p, \hat{c}_r) . $q_p \in c_p$ is bumpy if $\|q_p - \hat{c}_p\|_2 > \|q_r - \hat{c}_r\|_2$. Red points are bumpy but not the blue.

ridge does not suffer from the limitations outlined earlier, and, its accuracy depends solely on robust detection of the end points. Curvature weighting ensures that the geodesic passes through the high curvature ridge.

3.6 Bumps

We are interested in bumps on somewhat tubular subsurfaces. A bump is characterized with gentle slope, instead of pointy high curvature area. The problem is significant in CAD, for removing unwanted bumps, or in medical image analysis, for instance, for detecting soft-plaques in CT angiography.

Our approach is to analyze a tubular subsurface via x-sectional scans in an ROI, \mathcal{M}_r , to generate a set of profile contours $\{c_p\}$. The problem is, hence, reduced to (1) identifying the contours that correspond to a bump, followed by (2) determining “bumpy points” of a “bump contour”. A reference contour c_r is defined, with which all c_p are compared, starting at one end of \mathcal{M}_r and gradually moving towards the other. \mathcal{M}_r is selected to be sufficiently large to avoid the boundary of \mathcal{M}_r falling on a bump. Hence, the x-section at a boundary of \mathcal{M}_r may be selected as the initial c_r , which is later continuously updated with the last contour found without a bump.

For the identification of bump contours, the area of each contour c_p is compared to that of c_r . If it exceeds the reference area by a certain threshold t_P , c_p is considered to be on a bump. Once a bump contour is identified, the challenge lies in determining its bumpy part. A bump contour is projected to the plane of c_r . In this 2D subspace of the reference plane, a correspondence is found between the projected points and c_r . Corresponding points are then compared by way of their distances from individual centroids. If the distance of a projected point to the centroid of c_p , \hat{c}_p , exceeds the distance of the corresponding reference point to the reference centroid, \hat{c}_r , the point is marked bumpy. Such a criterion takes into account possible shifts of centroids and gives accurate estimates for bump points. After all x-sections are examined, bump points are binned into areas of connected points through recursion within 1-ring neighborhood.

4 Anatomical Feature Detection

We now detect the CES. The inter-tragal notch and the crus-side ridge require the detection of their end points. We use PCA of the x-section to identify two points on the canal as candidates for the inter-tragal notch top and the crus-side ridge top. The bottom of the inter-tragal notch is detected by considering the convexity of the bottom contour between tragus and anti-tragus. Geodesics are run from the ridge tops to the bottom point according to Section 3.5. The shorter geodesic is classified as the inter-tragal notch. The other ridge top is regarded the crus-side ridge top, and the shortest geodesic from this point to the bottom contour is the crus-side ridge. For detecting bulbous areas, and first and second bends, the canal was selected as the ROI. The canal-concha (or canal-crus) intersection is detected as an intersection of two geodesics, one running along a canal ridge, while the other traced from the concha peak (or crus).

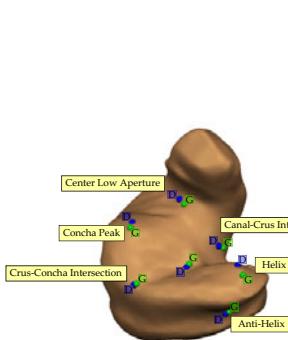


Fig. 6. Automatic detection of Point features: G = GT, D = detected features

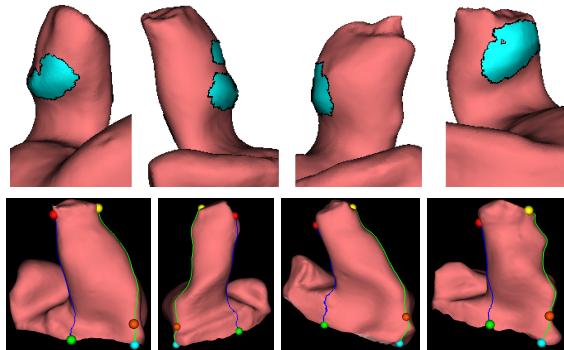


Fig. 7. Automatic detection of: (Top) Canal Bulbous; (Bottom) Inter-tragal notch, crus-side ridge, and inter-tragal notch flare

Helix ridge is detected as the shortest curvature weighted geodesic between the helix and the shell boundary. Crus area is computed as the area enclosed by appropriately weighted geodesics run between the following feature pairs: (1) center crus-helix ridge bottom; (2) center crus-crus ridge bottom; and (3) the boundary contour. Crus-concha intersection is detected by analyzing the tangential profile of the intersection of the shell with the crus valley plane.

5 Experiments

Statistical validation of the detected features was carried out over a dataset of 198 shells, for which an expert annotated the Ground Truth (GT). Features were then detected automatically, and compared with GT (e.g., Fig. 6).

For quantitative validation, point features were compared by way of their distance from the GT. The similarity measure for plane features considered their orientation and location. Orientations were compared via inner product between plane normals. Deviation of plane locations was the average distance between the two planes. Sensitivity and specificity were used for validating area features. The mean and standard deviations of these measures are given in Fig. 8. The results indicate good agreement with the GT. For the crus area, the mean *sensitivity* was found to be 0.83 with a mean *specificity* of 0.93. Since bulbous is not always present on an impression, only its qualitative evaluation was considered (Fig. 7(Top)). The ridges were analyzed qualitatively and in more than 90% of the cases the experts were satisfied with the performance. Examples in Fig. 7(Bottom) precisely follow the actual ridges on the canal. Inter-tragal notch flare is also at the correct location, precisely where the saddle is formed.

In the second set of experiments, an expert provided acceptable error tolerances for various features (3mm for points and plane locations, 15° for plane orientations). The *success rate* was then computed as the percentage of cases that resulted in features within the acceptable range. Mean success rate was

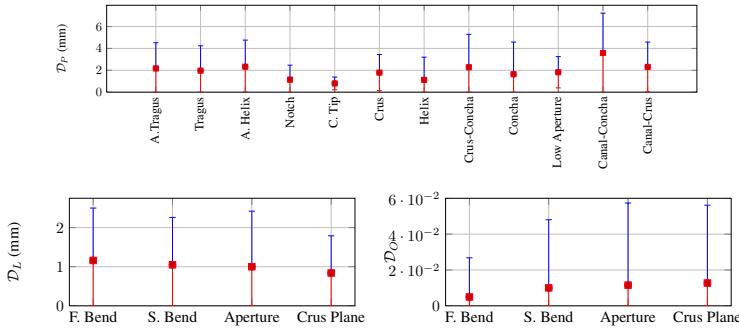


Fig. 8. Feature Validation – Comparison with GT: (Top) Point Features – Mean and standard deviation of the point deviations. (Bottom) Plane Features – Mean and standard deviation of deviations between plane (Left) locations; (Right) orientations.

87% for points (excluding Canal-Concha Intersection) and 82% for the planes. On average, all individual features exhibited acceptable performance (overall average deviation of ~ 2 mm and $\sim 10^0$) with the exception of the Canal-Concha Intersection and the Second Bend. Computationally, on average it takes 1.7s to compute all features for a dense mesh (79ms per feature) of ~ 23 k vertices.

6 Conclusions

We have proposed a comprehensive set of features for canonical representation of human ears, in addition to fast and robust algorithms for their detection. The representation may be used for automatic design of HA shells [5], feature guided registration, and classification. Due to their robustness, they have already been introduced for the automatic design of HAs by a major HA manufacturer. We emphasized that the utility of the proposed algorithms is not limited to the HAs. They have been designed with generality in mind, thereby providing a very powerful tool for medical image analysis and CAD.

References

1. Cohen-Or, D., et al.: Three-dimensional distance field metamorphosis. *ACM Trans. Graph.* 17(2), 116–141 (1998)
2. Henry, G.: *Gray's Anatomy: Descriptive and Surgical* (1858)
3. Hildebrandt, K., et al.: Smooth feature lines on surface meshes. In: *SGP 2005* (2005)
4. Shinagawa, Y., Kunii, T.L.: Constructing a Reeb graph automatically from cross sections. *Computer Graphics and Appl.* 11(6), 44–51 (1991)
5. Sickel, K., Baloch, S., et al.: Semi-automatic manufacturing of customized hearing aids using a feature driven rule-based framework. In: *VMV 2009* (2009)
6. Slabaugh, G., et al.: 3D Shape Modeling for Hearing Aid Design. *IEEE Signal Processing Magazine* 5(25), 98–102 (2008)
7. Stylianou, G., Farin, G.: Crest lines for surface segmentation and flattening. *IEEE Visualization and Computer Graphics* 10(5), 536–544 (2004)