



Published in final edited form as:

Proc SPIE Int Soc Opt Eng. 2014 March 21; 9034: 90342V-. doi:10.1117/12.2043260.

An artifact-robust, shape library-based algorithm for automatic segmentation of inner ear anatomy in post-cochlear-implantation CT

Fitsum A. Reda^a, Jack H. Noble^a, Robert F. Labadie^b, and Benoit M. Dawant^a

^aDept. of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN 37235, USA

^bDept. of Otolaryngology–Head and Neck Surgery, Vanderbilt University Medical Center, Nashville, TN 37232, USA

Abstract

A cochlear implant (CI) is a device that restores hearing using an electrode array that is surgically placed in the cochlea. After implantation, the CI is programmed to attempt to optimize hearing outcome. Currently, we are testing an image-guided CI programming (IGCIP) technique we recently developed that relies on knowledge of relative position of intracochlear anatomy to implanted electrodes. IGCIP is enabled by a number of algorithms we developed that permit determining the positions of electrodes relative to intra-cochlear anatomy using a pre- and a post-implantation CT. One issue with this technique is that it cannot be used for many subjects for whom a pre-implantation CT was not acquired. Pre-implantation CT has been necessary because it is difficult to localize the intra-cochlear structures in post-implantation CTs alone due to the image artifacts that obscure the cochlea. In this work, we present an algorithm for automatically segmenting intra-cochlear anatomy in post-implantation CTs. Our approach is to first identify the labyrinth and then use its position as a landmark to localize the intra-cochlear anatomy. Specifically, we identify the labyrinth by first approximately estimating its position by mapping a labyrinth surface of another subject that is selected from a library of such surfaces and then refining this estimate by a standard shape model-based segmentation method. We tested our approach on 10 ears and achieved overall mean and maximum errors of 0.209 and 0.98 mm, respectively. This result suggests that our approach is accurate enough for developing IGCIP strategies based solely on post-implantation CTs.

Keywords

Cochlear implant (CI) surgery; CI programming; intra-cochlear anatomy; segmentation; registration

1. INTRODUCTION

A cochlear implant (CI) is a device that restores hearing by directly stimulating the auditory nerve using an electrode array that is surgically placed in the cochlea. After placement, the CI is programmed by an audiologist who determines a number of device programming parameters that define how signals are sent to the implanted electrodes to try to optimize

hearing outcome. We recently developed and are currently testing an image-guided cochlear implant programming (IGCIP) technique that relies on knowledge of the spatial relationship between the electrodes and intra-cochlear anatomy [1]. Figure 1a shows the surfaces of the two principal intra-cochlear structures, the scala tympani (ST) and the scala vestibuli (SV), and a surface representing the spiral ganglion (SG), which is the group of nerves targeted for stimulation by the electrode array. Figure 1b shows an example surface model of the electrode array inserted into the cochlea, and Figure 1c shows a surface of the active region (AR), which is the interface between (1) the SG and (2) the union of the ST and SV. This is the region where electrical activation of nerves is most likely.

IGCIP is enabled by a number of algorithms we have developed that permit determination of the position of the electrodes relative to intra-cochlear anatomy using a pre- and a post-implantation CT [2]-[7]. In a preliminary study with over thirty subjects, we have shown that IGCIP can significantly improve hearing outcomes [1]. One issue with our current technique is that it has not been possible to localize intra-cochlear structures in post-implantation CTs directly due to the image artifacts caused by the electrode array that obscure the cochlea in the image (see Figure 1d and 1e). Thus far, the shape of the cochlea and intra-cochlear anatomy has been determined using a pre-implantation CT, which we then register to a post-implantation CT. For unilateral CI recipients where only a post-implantation CT with both ears in the field of view (FOV) is available, we have also developed a technique that permits estimating the shape of the cochlea in the implanted ear using information from the contralateral normal ear [8]. However, the approaches we have developed thus far cannot be used for many CI recipients for whom a pre-implantation CT of neither ear is available. In this paper, we present a technique by which we can localize the intra-cochlear structures in post-implantation CTs directly, without the need for a pre-implantation CT of either ear, despite the substantial artifacts present in these images. This new approach capitalizes on the physical relationship between the cochlear anatomy and the labyrinth, i.e., the rest of the inner ear.

2. METHODS

Our approach consists of two main steps. First, (A) we localize the entire labyrinth, which we use as a landmark structure, by (1) coarsely estimating its position using a shape chosen from a library of labyrinth shapes, (2) automatically creating a statistical shape model that is specific to the subject, and (3) refining the coarse estimate by performing a statistical shape model-based segmentation. The labyrinth is a structure that shares the external wall of the cochlea with the intra-cochlear anatomy and that also includes the semi-circular canals. Figure 2a and 2b show a surface of the labyrinth and the intracochlear structures that it externally bounds. Next, (B) we segment the SOIs by fitting the subset of SOI model points, that represent the external wall of the cochlea (see Figure 2d and 2e), to the part of the labyrinth that represent the same (see Figure 2c). To do the fitting, we established offline a one-to-one point correspondence between the subset of SOI model points and the subset of labyrinth model points that represent the exterior of the cochlea. The exterior region of the cochlea, which we use to fit the SOIs model, is the only portion of the SOIs that (a) has contrast in CT and (b) can also be localized by the labyrinth.

The following subsections detail our methods. In Section 2.1, we present the datasets we use in this study. The image-to-image registration process and active shape segmentation process we use at various steps throughout our work are detailed in Section 2.2 and 2.3, respectively. Next, in Section 2.4, we present the process we use to create our shape library. Finally, in Section 2.5, we present the multi-step process we propose to localize the labyrinth and use its position to estimate the position of the intra-cochlear structures.

2.1 Data

In this study, we use several groups of head CT scans, which are presented in Table I. The scans were acquired from several conventional scanners and a low-dose flat-panel volumetric CT (fpVCT) scanner (Xoran Technologies xCAT® ENT). Conventional CTs of 70 subjects are used for creating SOI and labyrinth shapes library as discussed in section 2.4, conventional CTs of 25 subjects are used for creating an active shape model (ASM) of the labyrinth as discussed in section 2.4, fpVCT scans of 14 subjects are used for creating an intensity model for each point on an ASM of the labyrinth as discussed in section 2.5.2, and CT-fpVCT pairs of 8 subjects are used for validating our segmentation results as discussed in 2.5.3. Our validation dataset (dataset 5) is constructed such that it allows us to (1) generate automatic segmentations on post-implantation CTs using the approach we propose (2) register the post-implantation CTs to the corresponding preimplantation CTs and (3) validate our results by comparing registered automatic segmentations to ground truth segmentations established on the pre-implantation CTs. Typical voxel size for conventional CTs is $0.25 \times 0.25 \times 0.3 \text{ mm}^3$; for flat-panel CTs it is $0.4 \times 0.4 \times 0.4 \text{ mm}^3$.

2.2. Image registration methods

In this subsection, we present the image-to-image registration process that we use at various steps throughout our study. Given a “fixed” image, i.e., an atlas or reference image, and a “floating” image, i.e., the target image, we use the process outlined in Figure 3 to register them. First, we affinely register the entire but downsampled images using an intensity-based affine registration method [9]-[10]. Next, we refine this registration by performing intensity-based affine registration at full image resolution on a pre-determined region that encompasses the ear structures. Finally, we further refine the registration by performing intensity-based non-rigid registration on the ear region [11].

2.3. Active shape model (ASM)-based segmentation

Various processes we describe in the following subsections rely on the creation of an active shape model (ASM) and performing active shape segmentation. Thus, in the following subsections we describe the general ASM framework that we use.

2.3.1. Active shape model (ASM) creation—Given a reference surface and a set of training surfaces of a structure with a one-to-one point correspondence between the points on the reference surface and the points on each training surface, we perform the following steps to create an ASM of a structure. First, we register each training surface to the reference surface with a 7-DOF (three translations, three rotations, one isotropic scaling) transformation that minimizes the root-mean-squared (RMS) distance between the surfaces. Next, we use the registered surfaces to build the structure ASM according to the procedure

described by Cootes in [12]. Finally, we store the ASM in the reference image space. The ASM is represented by the mean shape $\{\bar{x}_i\}_{i=0}^{N-1}$, with N being the number of points in the shape, and L eigenvectors $U = [\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{L-1}]$ that are corresponding to the largest eigenvalues $\lambda_0, \lambda_1, \dots, \lambda_{L-1}$. Mathematically,

$$\{\lambda_l, \mathbf{u}_l\}_{l=0}^{L-1} : \lambda_l \mathbf{u}_l = X \mathbf{u}_l, \quad (1)$$

where X is the covariance matrix of the points on registered surfaces.

2.3.2. Active shape segmentation—Active shape segmentation of the structure is performed by fitting the ASM to an initial estimate of the shape. This process consists of three main steps: (1) *Shape initialization*: We determine an initial coarse estimate of the shape $\{x_i\}_{i=0}^{N-1}$ by projecting the mean shape $\{\bar{x}_i\}_{i=0}^{N-1}$ from the reference image space to the target image space using the registration transformation that registers the two images. This registration transformation is computed using the image-to-image registration process described in Section 2.2. (2) *Shape adjustment*: We adjust the initial shape by iteratively finding a candidate position for each i th point in the initial shape and fitting the shape model to these candidate positions in a weighted least squares sense. The candidate position x'_i for each initial point x_i is determined along the surface normal \hat{n}_i in the interval $[-1.5, 1.5]$ mm, equivalently,

$$x'_i = x_i + \Delta d \cdot k_{\min} \cdot \hat{n}_i, \quad (2)$$

where $d = 0.15$ mm, and k_{\min} is chosen as,

$$k_{\min} = \underset{k}{\operatorname{arg\,min}} C_i(k) : k \in \{-10, -9, \dots, 10\}, \quad (3)$$

i.e., the candidate position for the i th point is the position at which the cost function $C_i(\bullet)$ is the smallest cost value in the interval $[-1.5, 1.5]$ mm along \hat{n}_i . The cost function is tailored to the type of image we use as described in the following section. We then fit the shape

model to the candidate points $\{x'_i\}_{i=0}^{N-1}$ to obtain an adjusted shape $\{x''_i\}_{i=0}^{N-1}$, given by

$$x''_i = \psi^{-1} \left(\bar{x}_i + \sum_{l=0}^{L-1} b_l \mathbf{u}_{l,i} \right), \quad (4)$$

where \mathbf{b} , a vector of parameters that defines the shape, is given by

$$\mathbf{b} = [b_0, b_1, \dots, b_{L-1}] = (U^T W^T W U)^{-1} U^T W^T W \mathbf{d}, \quad (5)$$

and $\mathbf{d} = \{d_i\}_{i=0}^{N-1}$, defined as

$$d_i = \psi(x'_i) - \bar{x}_i \text{ for } i \in \{0, N-1\}, \quad (6)$$

is the residual between the mean shape and candidate points, after they are registered to the mean shape with a 7-DOF (three translation, three rotation, and one isotropic scaling) transformation ψ , computed as,

$$\psi = \arg \min_{\psi} \sum_{i=0}^{N-1} w_i^2 \left\| \psi(\mathbf{x}'_i) - \bar{\mathbf{x}}_i \right\|^2. \quad (7)$$

We assign a reliability weight $w_i \in [0, 1]$ for each candidate point. The reliability weight computation, as we will explain in the following sections, is tailored for the type of image we want to segment. The weight matrix $W = \text{diag}([w_0, w_1, \dots, w_{N-1}])$, with $w_i = [w_i, w_i, w_i]$, in Eqn. (5) is designed so that candidate points with high reliability have more influence on the least squares model fitting. (3) *Iterate shape adjustment*: We iterate the shape adjustment process until the RMS distance between the adjusted shape and the initial shape at that

iteration is small, specifically until $\left(1/N \sum_{i=0}^{N-1} \|\mathbf{x}''_i - \mathbf{x}_i\|^2\right)^{1/2} < \varepsilon$ is satisfied, where ε is empirically set to 0.01 mm.

In summary, given an ASM of a structure and its initial shape estimate, we iteratively fit the ASM to segment the structure. At each iteration, we determine a candidate position for each i th point using Eqn. (2), we re-compute a weight for each i th candidate point, and finally we determine an adjusted shape by fitting the ASM to the candidate points in a weighted least squares sense using Eqn. (4).

2.4. Shape library creation

Our segmentation approach, as we will discuss in the next subsection, relies on first determining a coarse estimate of the labyrinth, which we use as a landmark, and of the intracochlear structures, which are the SOIs, by mapping surfaces of the labyrinth and SOIs chosen from a library of labyrinth and SOI surfaces to the subject's images. To enable this approach we create a library of internal ear structures segmented in a number of subjects' pre-implantation CTs. Specifically, we use the set of CTs in dataset 2 (see Table I) to create a library of surfaces that represent the labyrinth and SOIs of each subject's left or right ear, chosen randomly.

To produce a surface of the labyrinth in each CT in this dataset, we perform an ASM-based segmentation of the labyrinth in the CT using the active shape segmentation process described in Section 2.3.2. The labyrinth ASM we use in this step is created offline, using the reference CT in dataset 1 and the set of training CTs in dataset 3, according to the process we previously reported in [8] for the same purpose. When segmenting an image with this model, the cost $C_i(k)$ we use for candidate position selection in Eqn. (3) is given by

$$C_i(k) = - |I(\mathbf{x}_i + \Delta d \cdot (k+1) \cdot \hat{\mathbf{n}}_i) - I(\mathbf{x}_i + \Delta d \cdot (k-1) \cdot \hat{\mathbf{n}}_i)|, \quad (8)$$

where $I(\bullet)$ is image intensity in the CT at a given point. It is thus designed such that a candidate position for the i th point is chosen to be the position with the largest intensity gradient over the interval $[-1.5, 1.5]$ mm along $\hat{\mathbf{n}}_i$. We use two different approaches for selecting candidate positions, one for *contrasted* points $C \subset \{\mathbf{x}_i\}_{i=0}^{N-1}$, which are the subset

of points we know *a priori* are well contrasted in CT, and one for *non-contrasted* points $C' = \{x_i\}_{i=0}^{N-1} \setminus C$, which are the rest of the labyrinth surface points. Points that belong to C and C' have been labeled at the time the models were created. For each $x_i \in C$, we determine a candidate position x'_i using Eqn. (2) and assign a reliability weight of $w_i = 0.99$, while for each $x_i \in C'$ we use its original initial position determined via image registration as a candidate position and we assign a reliability weight of $w_i = 0.01$. A relatively high weight is thus assigned to the candidate positions for C so that the shape fitting is influenced more by those points with contrast in the CT. Although the results obtained with this techniques are generally satisfactory, there are cases where mis-segmentation is observed. This is likely caused by the limited number of shapes we use to create our ASM, which may not be able to capture enough variability to segment accurately all the images in our library. To deal with this issue, at each iteration, we determine the final adjusted point for the i th point using the equation

$$x_i^a = \alpha \cdot x_i'' + (1 - \alpha) \cdot x'_i, \quad (9)$$

which is a weighted combination of the position of the fitted model position x_i'' , given by Eqn. (4), and the candidate position x'_i , given by Eqn. (2), controlled by the weight parameter α . We set α to be 0.8 initially and we perform our iterative shape adjustment while decrementing α by 0.1, at the end of each iteration, for the first six iterations and use the final value of α for the remaining iterations. The value of α is set such that we largely rely on the model at the beginning. As we iteratively obtain better estimates of the shape, we gradually rely more on the candidate points which are likely to be positions with strong image gradient.

Finally, after the labyrinth is segmented algorithmically we manually adjust the segmentation to correct for any visually identifiable error. We then rely on the segmented labyrinth surface and an ASM of the SOIs, which we previously created and reported in [2], to segment the SOIs. To do this, we first establish offline a one-to-one point correspondence between the model points of the SOIs and the model points of the labyrinth. The SOI model points are then fitted to the corresponding points on the segmented labyrinth.

We produce the surfaces such that there is a one-to-one, across subject, point correspondence between the points composing the surfaces. For the purpose of segmentation strategy, which we will explain, we divide the points on each labyrinth included in the library into two groups: *near* points, which are points that may be close to implanted electrodes, and *far* points, which are the rest of points. Figure 2f shows a surface of the labyrinth with the two point groups rendered with different colors.

2.5. Labyrinth and SOIs segmentation

Our approach for segmenting both the labyrinth and SOIs in a target CT is to first determine a coarse estimate of the structures and then refine this coarse estimate. To identify a coarse estimate of the structures we map surfaces of the structures, chosen from our shape library, that best localize the structures in the target CT. We then create ASM models for the structures using the subset of shapes chosen, from our shape library, based on their

similarity to the corresponding structure shapes in the target CT. Finally, we refine the coarse estimate using a standard weighted ASM-based segmentation method. The following subsections detail our shape library-based coarse shape estimation and our shape model-based segmentation refinement steps.

2.5.1. Shape library-based segmentation initialization—We coarsely localize the labyrinth with a labyrinth surface chosen from our shape library. The surface is chosen such that its *far* points (see Figure 2f) best approximate the *far* points portion of the labyrinth in the target image. This process includes several steps. First, we determine the *far* points of the labyrinth in the target image by fitting the *far* points of the labyrinth ASM (see Section 2.3.1) following the segmentation process described in Section 2.3.2. The *far* points are likely to be far from implanted electrodes. They are thus unlikely to have been affected by implant-related artifacts in the image. Next, we register each labyrinth surface in our shape library to the target image using the transformation that minimizes the RMS distance between the *far* points on the library surface and the *far* points localized in the target image in the previous step. Finally, we compute a dissimilarity quantity for each registered surface as the residual RMS. The registered surface with the smallest dissimilarity quantity k_s is used as the coarse segmentation, with k_s defined as

$$k_s = \arg \min_k \left(\frac{1}{N^f} \sum_{i=0}^{N^f-1} \|\mathbf{x}_i - T_k(\mathbf{x}_{ki})\|^2 \right) : k \in \{0, 1, \dots, M-1\}, \quad (10)$$

in which M is the number of subjects in the library, $\{\mathbf{x}_i\}_{i=0}^{N^f-1}$ is the set of *far* points localized in the image, $\{\mathbf{x}_{ki}\}_{i=0}^{N^f-1}$ is set of *far* points in the k th shape in the library, N^f is the number points in the *far* portion of the labyrinth, and T_k is the 6-DOF (three rotations, three translations) transformation that registers the two *far* point sets, computed as,

$$T_k = \arg \min_T \frac{1}{N^f} \sum_{i=0}^{N^f-1} \|\mathbf{x}_i - T(\mathbf{x}_{ki})\|^2. \quad (11)$$

The value of the dissimilarity term is low when the shape represented by the *far* points localized in the image closely matches the shape represented by the *far* points in the k th surface. As we will show in the results section, the *far* portion of the labyrinth can be used as a good landmark for predicting the position of the labyrinth. A coarse segmentation of the SOIs is obtained by projecting the k_s -th subject's SOI surfaces to the target image through T_{k_s} .

2.5.2. Shape model-based segmentation refinement—To refine the coarse segmentations, we first segment the labyrinth by performing a weighted active shape segmentation and then segment the SOIs by fitting their ASM to the segmented labyrinth. This process is summarized in Figure 4.

First, we create two ASMs, one for the labyrinth and another for the SOIs using a subset of surfaces from our shape library. These are chosen as the five (a number chosen

experimentally) surfaces with the smallest dissimilarity quantity. The ASMs we create are thus specific to each target image. Next, using this target-specific ASM, we refine the initial labyrinth segmentation as discussed in Section 2.3.2. In this process, we use the coarse labyrinth localized in section 2.5.1 as the initial shape. We then iteratively refine it by first finding candidate position \mathbf{x}'_i (see Eqn. (2)) for each i th point \mathbf{x}_i and then fitting the ASM to the candidate positions in a weighted least squares sense (see Eqn. (4)). The cost function $C_i(\bullet)$ we use for candidate position selection in Eqn. (3) is a function of an intensity model of the image at that point. To build the intensity model, we rely on a set of manually segmented labyrinth surfaces obtained from dataset 4. For each j th training surface $\{\mathbf{x}_{ji}\}_{i=0}^{N-1}$ an intensity profile $\mathbf{p}(\mathbf{x}_{ji})$ is extracted at each i th point along the normal $\hat{\mathbf{n}}_{ji}$ using the equation

$$\mathbf{p}(\mathbf{x}_{ji}) = [I_j(\mathbf{x}_{ji} - \Delta d \cdot 9 \cdot \hat{\mathbf{n}}_{ji}), I_j(\mathbf{x}_{ji} - \Delta d \cdot 8 \cdot \hat{\mathbf{n}}_{ji}), \dots, I_j(\mathbf{x}_{ji} + \Delta d \cdot 9 \cdot \hat{\mathbf{n}}_{ji})]^T, \quad (12)$$

where $d = 0.15$ mm, and $I_j(\bullet)$ is the intensity for the j th training image at a given point. The intensity model at the i th point is given by $\{\mathbf{p}(\mathbf{x}_{ji})\}_{i=0}^{M-1}$, where M is the number of training surfaces. Finally, the cost function is designed as

$$C_i(k) = \min_j \|\mathbf{p}(\mathbf{x}_i + \Delta d \cdot k \cdot \hat{\mathbf{n}}_i) - \mathbf{p}(\mathbf{x}_{ji})\| : j \in [0, 1, \dots, M-1], \quad (13)$$

which defines the cost for selecting $\mathbf{x}_i + \Delta d \cdot k \cdot \hat{\mathbf{n}}_i$ as the candidate position for \mathbf{x}_i as the minimum Euclidean distance between the intensity profile at $\mathbf{x}_i + \Delta d \cdot k \cdot \hat{\mathbf{n}}_i$ and all the M intensity profiles contained in the set of model profiles at the i th point. The reliability $w_i \in [0, 1]$ we assign for each i th point is based on the intensity profile extracted at the i th point in I , and is given by,

$$w_i = \frac{\#\{k \in \{-10, -9, \dots, 10\} : I(\mathbf{x}_i + \Delta d \cdot k \cdot \hat{\mathbf{n}}_i) < R\}}{21} \quad (14)$$

where R is an intensity threshold that separates the bright metallic artifact from the rest of the structures. This weight is high when the set of intensity values in a given profile are below R , which indicates that the extracted profile is far from the image artifact in the image and is thus more likely to be reliable. To determine this threshold, the maxima along all the intensity profiles extracted along the surface normals at the points composing the initial shape are first computed. The threshold is then chosen experimentally to be the 90th percentile of the distribution of maxima. It is thus adapted to each image. Finally, we segment the SOIs by fitting the points on the target-specific SOIs' ASM to their corresponding points on the segmented labyrinth, as discussed in Section 2.4.

2.5.3. Segmentation validation—We validate our method by automatically segmenting the ST, SV, SG and labyrinth in the post-implantation CTs in dataset 5 (see Table I) using the approach we propose and by measuring the resulting segmentation errors. The gold-standard surfaces that we use for comparison were created in the corresponding pre-implantation CTs by manually editing surface points on segmentations that are automatically initialized by pre-implantation CT segmentation techniques we previously developed [2],

[8]. For each structure, we measure a distance from each point on its automatically generated surface to the corresponding point on its gold-standard surface, and report the mean, standard deviation, median and maximum of the distances we measure over all points on the surface. To quantify the improvement afforded by our refinement method, we measure the same segmentation error when we only use the segmentation initialization step.

3. RESULTS AND DISCUSSIONS

In Table II, we present errors obtained in segmenting the entire SOIs. The mean, standard deviation, median and maximum errors are 0.254, 0.128, 0.224, 0.76 mm, respectively, for results achieved using our segmentation initialization approach alone and 0.209, 0.128, 0.181, 0.98 mm, respectively, for results achieved by refining the initial results using our segmentation refinement approach. Table III presents the same information for the labyrinth. Overall SOI average segmentation error is close to half the voxel size in the segmented CT and errors are all sub-millimetric (<1 mm). In Figure 5, we show renderings of segmented SOI surfaces colormapped with segmentation errors for all testing ears. For the majority of the cases, SOI segmentation errors are sub-voxel, except for one (1L). For ear 1L we observed that the labyrinth surface chosen from the shape library using Eqn. (10) does not localize the SOIs as well as it does the *far* points region of the labyrinth, i.e., the relative position of the *far* points and the *near* points in this particular subject is different than in the selected library shape. This is also the case for the subset of surfaces selected to build the target-specific ASM for performing our segmentation refinement step. Because of this, the initial SOI segmentation errors are relatively large and they get worse when using our segmentation refinement step. We conducted an experiment for 1L where we perform our segmentation refinement step using an ASM built with all the shapes, rather than a subset of shapes, in our library. The final errors in segmenting the SOIs we obtain by doing so are 0.37, 0.12, 0.36, 0.78, which are smaller than the initial errors. This indicates that for ear 1L our target-specific ASM did not capture the target structure shapes well. The same phenomenon has been observed for ear 7R. However, for 80% of the test ears (8 out of 10), performing our segmentation refinement step has led to a reduction in initial segmentation errors. We do not use all shapes in the library to create the target specific ASMs because experiments we conducted show that the smallest segmentation error overall ten ears is obtained when we use target specific ASMs built with the five most similar shapes.

A method we previously developed for segmenting the SOIs in pre-implantation CTs [2] resulted in mean and maximum segmentation errors of 0.15 and 1.6 mm, respectively. This shows that the segmentation errors we achieve in post-implantation CTs are close to those that are achievable in pre-implantation CTs, despite the presence of metallic artifacts in the image that obscure the structures. Figure 6 shows qualitative results for a case whose average error is close to the overall average error (7R). Both the gold-standard (yellow) and automatic (red for ST, blue for SV and green for SG) contours are overlaid on both the pre-implantation CT and the registered post-implantation CT. As can be seen in the figure, although the structures are obscured by the bright artifact introduced by the implant, there is a good agreement between the two contours along the length of the structures.

Table IV presents, for each testing ear, mean, standard deviation, median and maximum surface distance errors in mm for the AR, which is the part of the SOIs most likely to be stimulated by implanted electrodes and thus the most important surface for our programming application. The overall mean, standard deviation, median, and maximum errors in segmenting the AR are 0.202, 0.131, 0.169 and 0.98 mm, respectively. The corresponding errors we obtain in segmenting the AR using our segmentation initialization step alone, prior to performing our segmentation refinement step, are 0.258, 0.127, 0.225 and 0.70 mm, respectively. Figure 7 shows renderings of the segmented AR surface for each testing ear that are colormapped with surface distance errors in mm. As can be seen from these, errors are sub-voxel (< 0.4 mm) for the majority of AR for all testing ears, except for one (1L).

5. CONCLUSIONS

The IGCIP strategies we recently developed and are currently testing require accurate localization of the position of implanted electrodes relative to intra-cochlear anatomy. So far, we have made this possible for subjects for whom a CT has been acquired prior to implantation, where we segment the SOIs in the pre-implantation CT, identify the electrodes in the post-implantation CT, and register the two CTs to determine the spatial relationship between the implanted electrodes and the SOIs. We have also recently presented a technique that makes IGCIP possible for subjects with no pre-implantation CT but who are implanted unilaterally. For this population of subjects, we determine the SOIs in the implanted ear using information extracted from the normal ear in the post-implantation CT. In this study, we have presented a shape library-based algorithm that does not require a pre-implantation CT of either ear to segment the SOIs.

Our approach for segmentation relies on first approximating the shape of the labyrinth by mapping a labyrinth surface that is selected from a library of such surfaces, and then refining this shape by performing a weighted active shape segmentation with an ASM built to be specific for the target image. We then segment the SOIs by fitting their shape model to the external wall of the cochlea established on the segmented labyrinth. As the results we present show, we achieve sub-millimetric errors at all points on the surfaces, and overall SOI segmentation error averages 0.209 mm. This average error is 0.202 mm for the AR, the “important” part of the SOIs. These results, which we achieve on post-implantation CTs, are comparable to those that are achievable on pre-implantation CTs and this suggests that our approach is accurate enough for use in position-based sound processing strategies. It is of note that our approach achieves this level of accuracy on an imperfect dataset composed of low-dose fpVCT images. We speculate that our approach could produce even more accurate segmentations when employed on post-implantation CTs acquired with standard CT scanners.

The segmentation results we achieve using our segmentation initialization approach alone are very close to the results we achieve by refining the initial results using our segmentation refinement approach. This suggests that the segmentation initialization approach alone can be used to achieve the task of segmenting the SOIs, particularly, in cases where the external wall of the cochlea in the image is completely obscured by the implant, which could prevent

our segmentation refinement approach from improving the initial results. Future work will focus on exploring techniques for automatically determining when to use the segmentation initialization step alone to achieve the segmentation task.

ACKNOWLEDGMENTS

This work was supported by NIH grants R01DC008408, R21DC012620, and R01DC010184 from the National Institute on Deafness and Other Communication Disorders. The content is solely the responsibility of the authors and does not necessarily represent the official views of this institute.

5. REFERENCES

- [1]. Noble JH, Labadie RF, Gifford RH, Dawant BM. Image-guidance enables new methods for customizing cochlear implant stimulation strategies. *Neural Systems and Rehabilitation Engineering*. Sept.2013 21(5):820–829. IEEE Transactions on. [PubMed: 23529109]
- [2]. Noble JH, Labadie RF, Majdani O, Dawant BM. Automatic Segmentation of Intracochlear Anatomy in Conventional CT. *Biomedical Engineering*. Sept.2011 58(9):2625–2632. IEEE Transactions on.
- [3]. Noble, Jack H.; Schuman, Theodore A.; Wright, Charles G.; Labadie, Robert F.; Dawant, Benoit M. Automatic identification of cochlear implant electrode arrays for post-operative assessment. *Proc. SPIE 7962, Medical Imaging 2011: Image Processing*; 796217
- [4]. Schuman TA, Noble JH, Wright CG, Wanna GB, Dawant B, Labadie RF. Anatomic Verification of a Novel, Non-rigid Registration Method for Precise Intrascalar Localization of Cochlear Implant Electrodes in Adult Human Temporal Bones Using Clinically-available Computerized Tomography. *The Laryngoscope*. 2010; 120(11):2277–2283. [PubMed: 20939074]
- [5]. Wanna GB, Noble JH, McRackan TR, Dawant BM, Dietrich MS, Watkins LD, Rivas A, Schuman TA, Labadie RF. Assessment of electrode positions and audiological outcomes in bilateral cochlear implant patients. *Otology & Neurotology*. 2011; 32(3):428–432. [PubMed: 21283037]
- [6]. Noble JH, Gifford RH, Labadie RF, Dawant BM. Statistical Shape Model Segmentation and Frequency Mapping of Cochlear Implant Stimulation Targets in CT. *Medical Image Computing and Computer-Assisted Intervention–MICCAI*. 2012; 2012:421–428.
- [7]. Noble JH, Dawant BM, Gifford RH, Labadie RF. Automatic, Image-based Cochlear Implant Electrode-to-Spiral Ganglion Position Analysis: Implications for Programming. Presented at the American Otological Society. Apr.2012
- [8]. Reda, FA.; Dawant, BM.; McRackan, TR.; Labadie, RF.; Noble, JH. Automatic segmentation of intra-cochlear anatomy in post-implantation CT. *Proc. SPIE 8671, Medical Imaging 2013: Image-Guided Procedures, Robotic Interventions, and Modeling*; 86710I
- [9]. Maes F, Collignon A, Vandermeulen D, Mrchal G, Suetens P. Multimodality image registration by maximization of mutual information. *IEEE Trans. Med. Imaging*. 1997; 16:187–198. [PubMed: 9101328]
- [10]. Wells WM III, Viola P, Atsumi H, Nakajima S, Kikinis R. Multi-modal volume registration by maximization of mutual information. *Med. Image Anal*. 1996; 1:35–51. [PubMed: 9873920]
- [11]. Rohde GK, Aldroubi A, Dawant BM. The adaptive bases algorithm for intensity-based nonrigid image registration. *IEEE Trans.Med. Imag*. Nov.2003 22(no. 11):1470–1479.
- [12]. Cootes TF, Taylor CJ, Cooper DH, Graham J. Active shape models - Their training and application. *Comp. Vis. Image Understanding*. 1995; 61(no. 1):39–59.

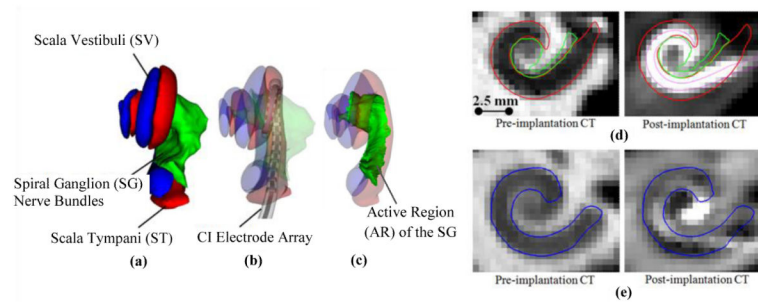


Figure 1.

Shown in (a) and (b) are surfaces of the ST (red), the SV (blue), and the SG (green). In (b), a surface model of a CI electrode array inserted into the ST is shown. Panel (c) shows surfaces of the AR (green), the ST (transparent red), and the SV (transparent blue). Panel (d) shows contours of the ST (red), the SG (green), and the electrodes (purple) in the coronal view of a pre-implantation CT and a corresponding post-implantation CT. Shown in (e) are contours of the SV (blue) in the coronal view of a pre-implantation CT and a corresponding post-implantation CT.




Figure 2.

Shown in (a) are surfaces of a labyrinth (transparent orange) and of the intra-cochlear anatomy (ST (transparent red), SV (transparent blue), and SG (transparent green)). Panel (b) shows the same structures in a different orientation. In (c) the set of points that represent the external wall of cochlea and that are used to fit the SOI model to the labyrinth model is shown on the surface of the labyrinth. In (d) and (e) the same set of points is shown on the SOI surfaces. Panel (f) shows a labyrinth surface with near points in yellow and far points in purple.

- 1) Affine registration
 - a. Affinely register the floating image to the fixed image, after downsampling both images by a factor of four in each dimension.
 - b. Crop the ear region from the affinely registered floating image.
 - c. Affinely register the floating ear region to the fixed ear region image at full image resolution.
- 2) Non-rigid registration
 - a. Non-rigidly register the floating ear region to the fixed ear region.

Figure 3.
Image registration process.

- 1) Build a target-specific ASM for both the labyrinth and SOIs
- 2) Perform ASM segmentation of the labyrinth
 - a. Initialize the shape
 - b. Adjust the shape
 - c. Iterate (b) until convergence
- 3) Segment SOIs by fitting their ASM to the labyrinth determined in (2).

Figure 4.
Segmentation refinement process.




Figure 5. Surfaces of intra-cochlear structures colormapped with segmentation errors viewed on the coronal plane (top row) and sagittal plane (bottom row).

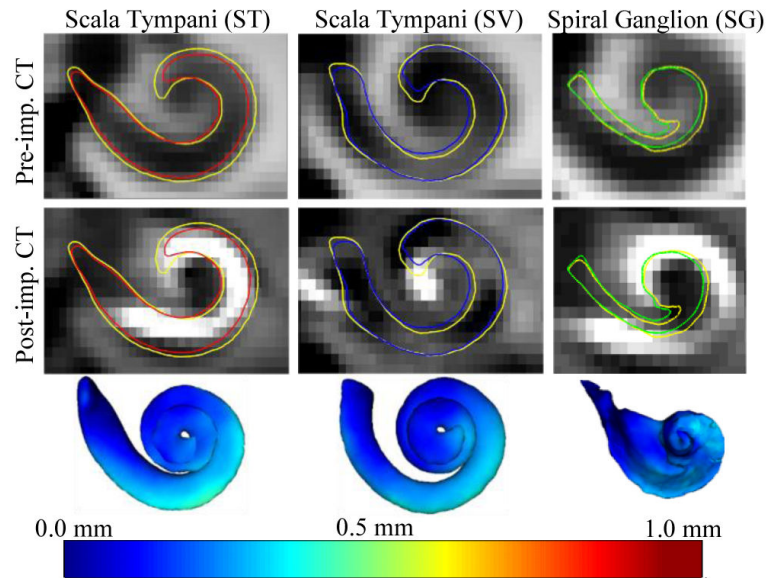


Figure 6. Results for a case with average error close to the overall average error. The contours shown are the ST (left panel), SV (middle panel), and SG (right panel). Contours for gold-standard ST (red), SV (blue), SG (green) surfaces and contours for automatic surfaces

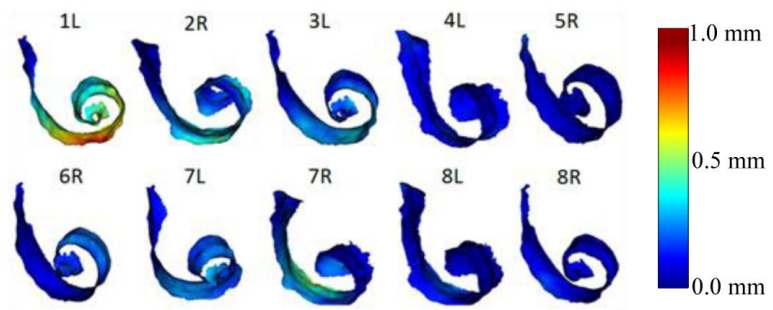


Figure 7.
Surface of the active region colormapped with segmentation errors for each testing ear.

Table I

Datasets used in this study

Dataset #	Purpose	Dataset Size	Acquisition			CI electrodes	
			Xoran fpVCT	Conventional	No CIs	One CI	Two CIs
1	Reference (atlas)	1		×	×		
2	Shape library creation	70		×	×		
3	Labyrinth ASM creation	25		×	×		
4	Intensity model creation	14	×		×		
		6		×	×		
5	Segmentation validation		×			×	
		2		×	×		
			×				×

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table II

Mean, standard deviation, median and maximum intra-cochlear anatomy (SOIs) segmentation errors in millimeters for the segmentation initialization and for the segmentation refinement steps. L is for left and R is for right ear.

Ear	Intra-cochlear anatomy							
	Initial Error				Final Error			
	Mean	Std. Dev.	Median	Maximum	Mean	Std. Dev.	Median	Maximum
1L	0.39	0.12	0.40	0.76	0.44	0.18	0.44	0.98
2R	0.41	0.11	0.41	0.72	0.26	0.05	0.27	0.39
3L	0.27	0.09	0.26	0.52	0.24	0.10	0.23	0.53
4L	0.16	0.07	0.15	0.38	0.13	0.05	0.12	0.31
5R	0.18	0.07	0.17	0.50	0.13	0.04	0.13	0.27
6R	0.16	0.06	0.16	0.37	0.17	0.06	0.17	0.31
7L	0.35	0.09	0.35	0.66	0.25	0.09	0.25	0.54
7R	0.18	0.06	0.18	0.35	0.23	0.08	0.23	0.50
8L	0.15	0.05	0.15	0.27	0.14	0.05	0.14	0.29
8R	0.30	0.09	0.28	0.54	0.11	0.05	0.10	0.33
Overall	0.254	0.128	0.224	0.76	0.209	0.128	0.181	0.98

Table III

Mean, standard deviation, median and maximum labyrinth anatomy segmentation errors in millimeters for the segmentation initialization and for the segmentation refinement steps.

Ear	Labyrinth Anatomy							
	Initial Error				Final Error			
	Mean	Std. Dev.	Median	Maximum	Mean	Std. Dev.	Median	Maximum
1L	0.32	0.16	0.34	0.76	0.32	0.18	0.26	0.95
2R	0.32	0.18	0.27	0.88	0.24	0.07	0.24	0.54
3L	0.24	0.12	0.20	0.72	0.24	0.12	0.23	0.75
4L	0.19	0.10	0.17	0.52	0.16	0.07	0.14	0.47
5R	0.15	0.09	0.13	0.56	0.12	0.06	0.12	0.42
6R	0.15	0.08	0.14	0.49	0.13	0.06	0.13	0.39
7L	0.28	0.15	0.25	0.79	0.21	0.10	0.19	0.49
7R	0.15	0.07	0.15	0.46	0.19	0.10	0.18	0.50
8L	0.15	0.07	0.14	0.42	0.16	0.08	0.15	0.44
8R	0.28	0.12	0.27	0.77	0.14	0.08	0.12	0.45
Overall	0.223	0.139	0.185	0.88	0.192	0.131	0.169	0.95

Table IV

Mean, standard deviation, median and maximum active region (AR) segmentation errors in millimeters for the segmentation initialization and for the segmentation refinement steps. L is for left and R is for right ear.

Ear	Active Region							
	Initial Error				Final Error			
	Mean	Std. Dev.	Median	Maximum	Mean	Std. Dev.	Median	Maximum
1L	0.41	0.09	0.42	0.69	0.45	0.19	0.46	0.98
2R	0.42	0.10	0.41	0.70	0.28	0.04	0.28	0.39
3L	0.25	0.08	0.24	0.44	0.21	0.10	0.20	0.53
4L	0.15	0.07	0.14	0.35	0.11	0.05	0.10	0.28
5R	0.18	0.06	0.18	0.32	0.12	0.03	0.12	0.22
6R	0.16	0.05	0.17	0.33	0.17	0.07	0.18	0.30
7L	0.36	0.08	0.37	0.62	0.24	0.09	0.23	0.50
7R	0.18	0.05	0.18	0.34	0.22	0.09	0.23	0.44
8L	0.15	0.05	0.16	0.27	0.14	0.04	0.15	0.25
8R	0.32	0.08	0.31	0.54	0.10	0.05	0.09	0.29
Overall	0.258	0.127	0.225	0.70	0.202	0.131	0.169	0.98