Medical Image Analysis 52 (2019) 1–12

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Automatic graph-based method for localization of cochlear implant electrode arrays in clinical CT with sub-voxel accuracy

Yiyuan Zhao[a], Srijata Chakravorti[a], Robert F. Labadie[b], Benoit M. Dawant[a], Jack H. Noble[a,*]

[a] Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN 37235, USA
[b] Department of Otolaryngology – Head and Neck Surgery, Vanderbilt University Medical Center, Nashville, TN 37235, USA

ARTICLE INFO

Article history:
Received 26 March 2018
Revised 18 August 2018
Accepted 12 November 2018
Available online 13 November 2018

Keywords:
Cochlear implant
CI electrode array
Graph search
Segmentation

ABSTRACT

Cochlear implants (CIs) are neural prosthetics that provide a sense of sound to people who experience severe to profound hearing loss. Recent studies have demonstrated a correlation between hearing outcomes and intra-cochlear locations of CI electrodes. Our group has been conducting investigations on this correlation and has been developing an image-guided cochlear implant programming (IGCIP) system to program CI devices to improve hearing outcomes. One crucial step that has not been automated in IGCIP is the localization of CI electrodes in clinical CTs. Existing methods for CI electrode localization do not generalize well on large-scale datasets of clinical CTs implanted with different brands of CI arrays. In this paper, we propose a novel method for localizing different brands of CI electrodes in clinical CTs. We firstly generate the candidate electrode positions at sub-voxel resolution in a whole head CT by thresholding an up-sampled feature image and voxel-thinning the result. Then, we use a graph-based path-finding algorithm to find a fixed-length path that consists of a subset of the candidates as the localization result. Validation on a large-scale dataset of clinical CTs shows that our proposed method outperforms the state-of-art CI electrode localization methods and achieves a mean error of 0.12 mm when compared to expert manual localization results. This represents a crucial step in translating IGCIP from the laboratory to large-scale clinical use.

1. Introduction

Cochlear implants (CIs) are surgically implanted devices for treating severe-to-profound hearing loss (National Institute on Deafness and Other Communication Disorders, 2011). A CI uses an electrode array implanted within the cochlea to stimulate the spiral ganglion (SG) nerves to induce the sensation of hearing. The SG nerves are tonotopically ordered by decreasing characteristic frequency along the length of the cochlea (Greenwood, 1990; Stakhovskaya et al., 2007) (Shown in Fig. 1). A SG nerve is stimulated when the frequency associated with it exists in the incoming sound (Wilson and Dorman, 2008). After the CI surgery, an audiologist needs to program the CI to determine the stimulation level of each individual electrode based on perceived loudness, and to select a frequency allocation table to determine which individual electrodes are activated when the incoming sound contains specific frequencies. CIs lead to remarkable success in hearing restoration for the vast majority of recipients with average post-implantation sentence recognition rates over 70% correct for unilaterally implanted users and 80% correct for bilaterally implanted users, respectively (Gifford et al., 2014, 2008). However, there is a significant number of users experiencing only marginal benefits. Recent studies have demonstrated that there exists a correlation between hearing outcomes and the intra-cochlear locations of CI electrodes (Aschendorff et al., 2005; Rubinstein, 2004; Skinner et al., 2007; Verbist et al., 2005; Wanna et al., 2014, 2011). One factor that negatively affects hearing outcomes is electrode interaction (or channel interaction). Electrode interaction leads to nerve groups being activated in response to multiple frequency bands (Boëx et al., 2003; Fu and Nogaki, 2005). Electrode interaction can be alleviated by deactivating the electrodes that cause electrode interaction (Noble et al., 2014; Noble et al., 2016). In Fig. 1 we show the CI electrodes and their activation patterns for a subject. As can be seen, by deactivating some electrodes (labeled with red crosses), electrode interaction can be reduced.

Our group has developed methods for image-guided cochlear implant programming (IGCIP) (Noble et al., 2013) to assist audiologists with CI programming. IGCIP uses image processing techniques we have developed to analyze the spatial relationship between the CI electrodes and SGs for each individual recipients in order to estimate the occurrence of electrode interaction and select elec-

* Corresponding author.
E-mail addresses: yiyuan.zhao@vanderbilt.edu (Y. Zhao), jack.noble@vanderbilt.edu (J.H. Noble).

https://doi.org/10.1016/j.media.2018.11.005
1361-8415/

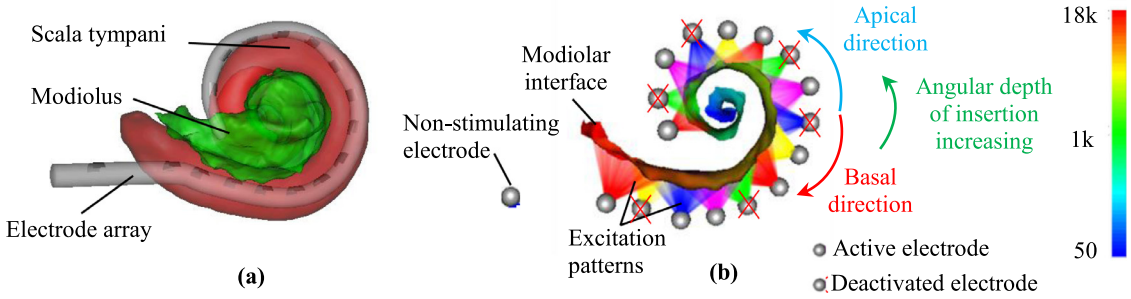

Fig. 1. Visualization of the intra-cochlear anatomy and CI electrode array. Panel (a) shows the scala tympani in red and the modiolus in green. The modiolus is the interface between the auditory nerves of the SG and the intra-cochlear cavities. Panel (b) illustrates the stimulation patterns produced by electrodes on one array. The modiolar surface is color-coded with the tonotopic place frequencies of the SG in Hz. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

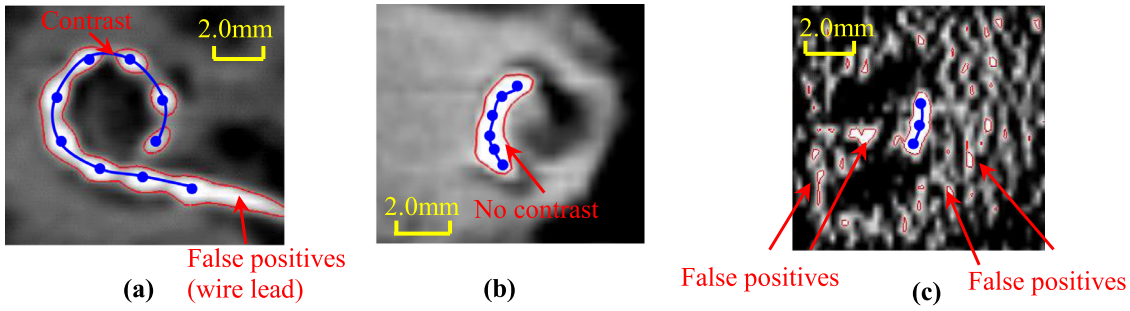

Fig. 2. Shown are examples of distantly (a) and closely (b) spaced arrays in eCTs. Panel (c) shows an example of a distantly-spaced array in a lCT.

trodes to deactivate to alleviate interaction. The major steps consist of (1) segmentation of the intra-cochlear anatomy, (Noble et al., 2011, 2012; Reda et al., 2014a, b, Wang et al., 2018), (2) localization of the implanted CI electrodes (Zhao et al., 2014; Noble and Dawant, 2015; Zhao et al., 2017; Zhao et al., 2018b), (3) analysis of the spatial relationship between the CI electrodes and the neural interface (Noble et al., 2013), and (4) automatic electrode configuration selection (Zhao et al., 2016, 2015; Zhang et al., 2017b). Clinical studies have shown that hearing outcomes are significantly improved when the CI electrode deactivation plans generated by IGCIP are adopted (Noble et al., 2014, 2016). A critical step in the generation of an IGCIP plan is the localization of the CI electrodes in a CT image. Since the end-users are audiologists who generally do not have experience with medical images, automatic electrode localization techniques would drastically reduce the barrier to adoption of the IGCIP approach by removing the need for audiologists to become experts in localizing electrodes in CT images. To fully translate IGCIP from the laboratory to large-scale clinical use, an accurate, automatic electrode localization method is needed. Automatic methods are also needed to conduct studies on large scale image databases to continue investigating the relationship between hearing outcomes and the intra-cochlear locations of CI electrodes.

Automating the electrode localization procedure is challenging. The first challenge is that the image quality of the clinical CTs is limited. The resolution of clinical CT images is usually coarse (resolution obtained nowadays is typically 0.2 x 0.2 x 0.3 mm^3) compared to the typical size of the CI electrodes which is on the order of 0.3 x 0.3 x 0.1 mm^3. Due to partial volume effects, it is difficult to localize the small-sized CI electrode array in clinical CTs. The image resolution is also coarse relative to the spacing between electrodes. This makes it difficult to separate the individual electrodes from each other (Shown in Fig. 2b). Further, because the electrodes are composed of radiodense platinum, beam hardening artifacts distort the intensities in the region around the electrode array, resulting in erroneous intensities assigned to voxels around the electrodes during reconstruction. This complicates the identification of individual electrodes in CTs. The second challenge is that even though the CI electrodes usually appear as high intensity voxel groups in CTs, it is difficult to select a threshold such that the thresholded image only contains voxels occupied by CI electrodes because voxels occupied by wire lead, receiver coils, and cortical bones are also usually assigned high intensity values. In this article, the non-electrode voxels with intensity values higher than a selected threshold are defined as “false positive” voxels. CT images are also reconstructed with different algorithms. In an image reconstructed with an “extended” Hounsfield Unit (HU) range (eCT), the metallic structures are assigned higher intensity values than the cortical bones. In an image reconstructed with a “limited” HU range (lCT), the maximum intensity is limited to the intensity of cortical bones. Thus, in an eCT, the false positive voxels are usually occupied by the metallic wire lead as shown in Fig. 2a. In an lCT, there are many more false positive voxels as shown in Fig. 2c. The third challenge is that there exist several models of electrode arrays, which lead to various intensity-based features in clinical CTs. A number of FDA-cleared models are produced by each of three manufacturers: Med-El® (MD) (Innsbruck, Austria), Advanced Bionics® (AB) (Valencia, California, USA), and Cochlear® (CO) (Sydney, New South Wales, Australia). Arrays differ by the number of electrodes, the size of electrodes, and the spacing between electrodes. Based on inter-electrode spacing, we classify CI electrode arrays into two broad categories: Closely-spaced and Distantly-spaced arrays. Closely-spaced arrays are such that individual electrodes cannot be resolved in the images and the set of electrodes thus form a single connected region as shown in Fig. 2b. We have proposed a centerline-based snake-based localization method (Zhao et al., 2014) to localize individual electrodes in this type of array. This method fails for distantly-spaced arrays because electrodes do not form a single connected region as shown in Fig. 2a. Other groups have also investigated methods for localizing CI electrodes in CTs (Braithwaite et al., 2016; Bennink et al., 2017). Bennink et al. proposed a method for localizing closely-spaced arrays by using the *a-priori* knowledge of the

Table 1
Specifications of different FDA-approved distantly-spaced CI electrode arrays in our dataset.

Manufacturer	Model	Total electrodes	Electrode spacing distance (mm)
Med-El	Standard (MD1)	12	2.4
	Flex28 (MD2)	12	2.1
Advanced Bionics	1 J (AB1)	17 (1 inactive electrode)	1.1 and 2.5
	Mid-Scala (AB2)	17 (1 inactive electrode)	0.95 and 3.0
	Helix (AB3)	18 (2 inactive electrodes)	0.85 and 3.0

Table 2
Datasets used in this study.

Dataset #	Purpose	Type of array	Number of eCTs	Number of lCTs	Total number of CTs
Dataset 1 (177 CTs)	Training (52 CTs)	AB1	15	0	15
		AB2	9	1	10
		AB3	3	0	3
		MD1	11	0	11
		MD2	12	1	13
	Validation (125 CTs)	AB1	19	6	25
		AB2	25	7	32
		AB3	4	0	4
		MD1	17	0	17
		MD2	36	11	47
Dataset 2 (28 CTs)	Robustness test	AB1	9	5	14
		AB2	9	5	14

CI array geometry. Braithwaite et al. proposed a method for localizing distantly-spaced arrays in CTs by using spherical measures. However, both of the two methods require manual initialization. Thus, they cannot be directly adopted by IGCIP for large-scale clinical use.

The graph-based path finding method (GP) we present in this article is designed to localize individual electrodes in distantly-spaced arrays. It builds upon and substantially improves a limited graph-based method (lGP) (Noble and Dawant, 2015) proposed by our group. In Section 2, we describe this method in detail. In Section 3, we present the localization results of GP, lGP and a preliminary implementation of GP (pGP) (Zhao et al., 2017) that does not produce sub-voxel accuracy. This is done on a large-scale dataset of clinically acquired CT images of subjects implanted with 4 different types of CI arrays. In Section 4, we analyze the results generated by GP, lGP, and, pGP. In Section 5, we summarize our work and discuss possible directions for extending it.

2. Methods

2.1. Dataset

In Appendix A we show the geometric models for three representative types of distantly-spaced electrode arrays, where distances between contacts were determined using publicly available documents provided by the manufacturers when available and otherwise were found by direct measurement using a surgical microscope and digital calipers. In Table 1, the specifications of the distantly-spaced CI electrode arrays included in this study are summarized. Table 2 lists the datasets we use in this study. Dataset 1 consists of whole head CTs of 177 patients from an IRB-approved adult cochlear implant imaging database at Vanderbilt University. Among these 151 are eCTs and the remaining 26 are lCTs. 144 of the 151 eCTs are acquired with a Xoran xCAT® flat panel scanner at the Vanderbilt University Medical Center (VUMC). The typical voxel size for Xoran eCTs is $0.4 \times 0.4 \times 0.4\text{mm}^3$. The remaining 7 eCTs are acquired with various scanners at various institutions. The typical voxel size for these 7 eCTs is $0.28 \times 0.28 \times 0.35\text{mm}^3$. The 26 lCTs are also acquired with various conventional scanners at various institutions. The typical voxel size for lCTs is $0.23 \times 0.23 \times 0.34\text{mm}^3$. The coarsest voxel size for lCT in our dataset is $0.46 \times 0.46 \times 0.50\text{mm}^3$. The scanners that are used to acquire the CTs above include Xoran xCAT, Siemens Somatom Definition AS, Siemens Somatom Force, Siemens Sensation 64, Siemens Somatom Emotion 16, Philips iCT 128, Philip Brilliance 64, Philips Mx8000 IDT16, Philips Comer-256, GE LightSpeed VCT, and GE Medical System BrightSpeed. Since our method includes several parameters, we randomly select 52 CTs from Dataset 1 that contain different types of electrode arrays for a parameter tuning process. The rest of the 125 CTs from Dataset 1 are used to validate the localization accuracy of our proposed method. Since we have only 7 eCTs and 26 lCTs acquired with scanners other than Xoran xCAT, we include 1 out of these 7 eCTs and 2 out of these lCTs in the training dataset. The rest of the 6 eCTs and 24 lCTs from other scanners are included in the testing dataset to confirm the ability of our proposed method to generalize to other scanners and acquisition parameters. In Dataset 1, an image processing expert with more than 4 years of experience with cochlear implant electrode localization manually localized all of the electrodes three times. Among the three sets of manual localization results, we randomly select two and average them to serve as the ground truth localization results. The remaining third manual localization result is used to estimate the rater's consistency error (RCE) defined as the distance between the ground truth and the third localization.

Dataset 2 consists of 14 CTs of a cochlear implant imaging phantom. We use Dataset 2 to directly evaluate the robustness of GP to various acquisition parameters (Chakravorti et al., 2017). The phantom was created using a cadaveric skull implanted with CIs in both left (AB1) and right (AB2) ears. For each side, we have acquired 14 CT scans with a range of acquisition parameters (the HU range, resolution, dose, and type of the implanted arrays) and with different scanners. For this data set, a high-dose, high-resolution, extended HU scan was used to create the ground truth. Three experts, including two image processing experts and one radiologist, provided a total of 10 sets of manual localization results for the high-resolution volume that were averaged to create the ground truth. In our previous study (Chakravorti et al., 2017), we found mean inter-rater differences were 0.06 ± 0.03 mm. Because we found in this multi-rater study that inter-rater variability in electrode localization is low we relied on a single rater to define the ground truth for Dataset 1. For each side in the 13 other CT images of Dataset 2, multiple (three or more) expert localiza-

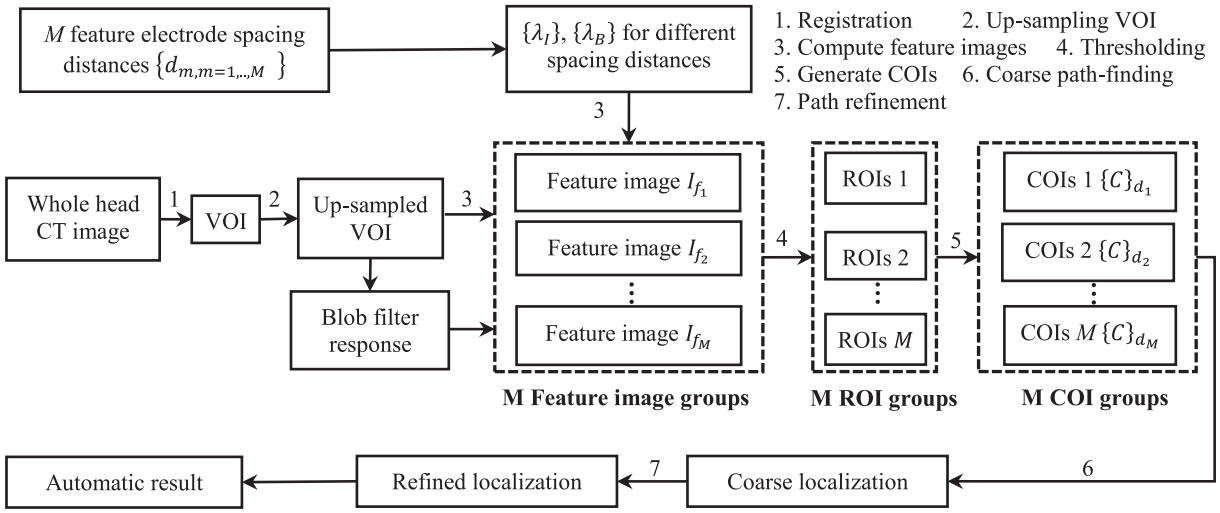

Fig. 3. Workflow of GP.

tions were averaged to produce a localization that represents the best possible localization achievable by an expert given the quality of that particular CT image. Comparing this expert localization to the ground truth permits measuring errors that are most likely due to the quality of the target image, i.e., the image-based localization error (IL), whereas comparing the automatic localization to the ground truth provides an estimate of the total error including errors due to image quality as well as to the algorithm. More details on this analysis can be found in our previous study (Chakravorti et al., 2017).

2.2. Overview

The workflow of GP, our proposed method, is outlined in Fig. 3. (1) We locate the volume of interest (VOI) that contains the cochlea region by registering the whole head CT to a reference image. (2) Next, we up-sample the VOI and the subsequent procedures are performed on the VOI. (3) Then, we determine the value of a set of parameters that will be used in the following steps using *a-priori* knowledge of the geometry of the array model. We call these parameters electrode spacing distance (ESD)-based parameters. As has been shown in Table 1, the distances between individual electrodes are known for each model. For a specific electrode array, we denote the distance between the centers of the ith and the $(i+1)$th electrodes as D_i and we define $\{D_i\}$ as the set of inter-electrode distances. We then define M sets of ESD-based parameters associated with this array corresponding with the $\{d_{m,\ m=1,..,\ M}\}$ unique distance values in $\{D_i\}$. For example, an AB2 array (shown in Appendix A) has $M = 2$ different ESDs, $d_1 =$ 2.5mm and $d_2 = 1.1$mm because $D_1 = 2.5$mm and $D_2 = D_3 = \ldots = D_{16} = 1.1$mm. ESD-based parameter values are used to tune filters or detection thresholds and produce M feature images, each optimized to detect electrodes separated by the corresponding distance d_i. (4) Next, we identify the regions-of-interest (ROIs) that contain voxels occupied by the CI electrodes by using the M feature images. (5) Then, we perform a voxel thinning method on each of the ROIs to extract the medial axis points as candidates of interest (COIs). At this stage, COIs consists of voxels occupied by electrodes and false positive voxels. (6) Once the COIs are extracted, we perform a coarse path-finding algorithm to find a fixed-length candidate path linking N COIs that minimizes a cost function to coarsely localize the electrodes. (7) Finally, we use a second path-finding algorithm to locally refine the location of each individual coarsely localized electrode. Each of these steps are detailed in the following subsections. In the remainder of this article the value of all the parameters denoted with Greek letters are determined through a parameter tuning process described in Section 2.6.

2.3. COI generation

The first step in our method is to identify the VOI that contains the cochlea region (a local region ~30 cm^3 around the cochlea). We achieve this by automatically registering (Zhang et al., 2017a) a reference image where the VOI bounding box is defined to the target CT. After determining the VOI, we up-sample it to a voxel size of $0.1 \times 0.1 \times 0.1$mm^3 and then compute a feature image I_f based on it. The feature image I_f is used for generating the ROIs and is computed as:

$$I_f(v) = \lambda_B(d_m)\frac{I_B(v) - T_B(\alpha_B\%)}{T_B(\alpha_B\%)} + \lambda_I(d_m)\frac{I(v) - T_I(\alpha_I\%)}{T_I(\alpha_I\%)} \quad (1)$$

where I is the intensity image of the VOI, I_B is the response to a blob filter applied to the VOI that is inspired by Frangi's vesselness filter (Frangi et al., 1998). As does Frangi, we use the value of the three eigenvalues (L_1, L_2 and L_3) of the 3×3 Hessian matrix at a voxel v to define the filter:

$$I_B(v) = \begin{cases} B_1(v) \cdot B_2(v) \cdot B_3(v), & L_1, L_2, L_3 < 0 \\ 0, & \text{otherwise} \end{cases}, \quad (2)$$

The three terms in Eq. (2) are defined as $B_1 = 1 - \exp\left(-\frac{\sum_{i=1}^{3} L_i^2}{S_1^2}\right)$, $B_2 = \exp\left(-\frac{r_{12}+r_{23}+r_{13}}{S_2}\right)$, and $B_3 = 1 - \exp\left(-\frac{L_{\min}}{S_3}\right)$, where $r_{ij} = |L_i - L_j|$, $L_{\min} = \min(-L_1, -L_2, -L_3)$, $S_1 = T_I(\alpha_I)$, $S_2 = 5000$, $S_3 = 40,000$. In Eq. (1), $T_I(\alpha_I\%)$ is a function which takes a percentage value $\alpha_I\%$ as input argument and generates an intensity threshold applied to I that corresponds to the top $\alpha_I\%$ of the cumulative histogram. S_2 and S_3 were empirically selected. The term B_1 enhances the voxels with high intensity. The terms B_2 and B_3 enhance the voxels that have spherical structures. The scales for our blob filter are selected as [0.2, 0.4] mm with a step of 0.04 mm, which is the typical range for the CI electrode radius. In Eq. (1), as is $T_I(\alpha_I\%)$, $T_B(\alpha_B\%)$ is a function that generates a threshold applied to I_B that corresponds to the top $\alpha_B\%$ of the cumulative histogram of I_B. $\lambda_I(d_m)$ and $\lambda_B(d_m)$ are functions of the ESD-based parameters d_m that return two weighting scalars. Because the weighting scalars returned by $\lambda_I(d_m)$ and $\lambda_B(d_m)$ are related to d_m, our method allows different weighting scalars to be assigned to the intensity and the blob filter response of the VOI depending on the spacing between electrodes. This is important because, for closer electrodes, heavier reliance on the blob filter

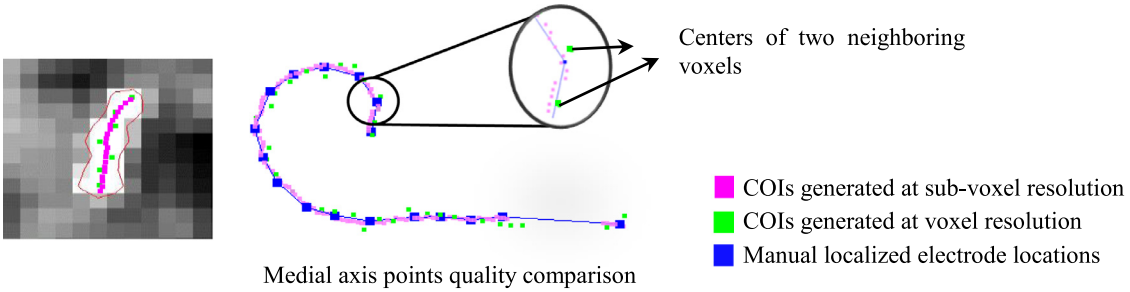

Fig. 4. Quality comparison between the COIs generated by our method on the feature image at sub-voxel resolution and at voxel resolution on the VOI. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

image is necessary to differentiate electrodes. For more distant electrodes, the more reliable intensity image can be emphasized in the cost function and the blob filter image is less important. Thus, λ_I and λ_B are defined as:

$$\lambda_I(d_m) = (-\kappa_I d_m + \beta_I)H(-\kappa_I d_m + \beta_I), \tag{3}$$

$$\lambda_B(d_m) = (\kappa_B d_m - \beta_B)H(\kappa_B d_m - \beta_B), \tag{4}$$

where β_I, κ_I, β_B, κ_B are positive weighting scalars. $H(\cdot)$ is the Heaviside function.

Each feature image I_f created with the corresponding d_m is then thresholded at 0. The thresholded regions are the ROIs for electrodes with a ESD value d_m. Next, we apply a voxel thinning method (Bouix et al., 2005) to the ROIs to generate the COIs associated with d_m. For each ROI, the voxel thinning method generates a series of points that are ordered sequentially as medial axis lines. Since we have up-sampled the VOI before generating feature images, ROIs, and COIs; the COIs we generate also have higher resolution than the COIs that would be generated by using the ROIs produced by the original VOI. Fig. 4 shows the difference between medial axis points generated by the voxel thinning method simply on the thresholded VOI without up-sampling and the medial axis points generated by our voxel thinning method on the up-sampled VOI. As can be seen, by up-sampling the VOI, our method permits to generate COIs with sub-voxel resolution. Among the COIs generated by using the up-sampled VOI (magenta), there exist candidate points that are closer to the actual locations of implanted electrodes (blue) than the COIs generated at voxel resolution (green). By up-sampling the VOI to a resolution higher than $0.1 \times 0.1 \times 0.1\text{mm}^3$, we could generate COIs with even higher resolution. However, we found empirically that the selected parameters lead to an adequate resolution for the COIs with an acceptable computational efficiency.

For a CT of a CI recipient implanted with an array with M ESD values $\{d_{m,\ m=1,\ldots,\ M}\}$, GP generates M sets of ROI groups, one for each ESD value. For each ROI, one set of COIs is generated. The complete set of COIs for the M ESD values are denoted as $\{C\}_{d_1}$, $\{C\}_{d_2}$, $\ldots$, $\{C\}_{d_M}$. We denote a COI that is the kth medial axis point on the medial axis line of the jth ROI in the mth ROI group as $c_m^{j,k}$. These COIs serve as the candidate nodes in a graph search problem used to coarsely localize the individual electrodes. In the following descriptions, we note p as a candidate path, p_i as the ith COI on the path p, and $\{C\}_{d_m}^{j}$ as the set of COIs that are on the medial axis line of the jth ROI in the mth ROI group associated with d_m.

2.4. Coarse path-finding algorithm

The coarse path-finding algorithm aims at finding a fixed-length path of N COIs representing the electrodes on the array, where N is the number of the electrodes on the array. While a standard technique such as Dijkstra's algorithm (Dijkstra, 1959) is typically used for path-finding problems because it guarantees finding a globally optimum solution, we instead use a custom path-finding algorithm that provides no such guarantee because it permits using non-local geometric-based constraints during the search. At each iteration of our proposed path-finding algorithm, a grow stage and a prune stage are included. At the first iteration, the algorithm uses every node in $\{C\}_{D_1}$ as a seed COI representing a candidate path that are each of length 1 in a candidate path group $\{p\}$. The candidate path group $\{p\}$ is used to store the candidate paths during the path-finding algorithm. At the i^{th} iteration, in the grow stage, each candidate path in $\{p\}$ is grown into a new set of candidate paths by connecting each of the "reachable" COIs in $\{C\}_{D_i}$ to it. Reachability of a COI to a candidate path is defined in Appendix B. The new set of candidate paths replace $\{p\}$ before the prune stage. Because the number of candidate paths in $\{p\}$ would grow exponentially at each iteration and the problem would become computationally intractable if left unchecked, we use a prune stage to reduce the set of candidate paths after the grow stage. This is done by computing at each iteration the value of a candidate path cost function and keeping the $\eta_{\max}$best candidate paths in $\{p\}$ in the prune stage. The cost function consists of a shape-based cost term and an intensity-based cost term, which capture the geometric and intensity features of the electrode arrays in clinical CTs. After N-1 iterations, $\{p\}$ consists of candidate paths of length N, and each node in these paths corresponds to a candidate electrode position. Node positions of the path with the lowest cost are used as coarse electrode positions. Fig. 5 shows a grow stage step for one candidate path with 3 reachable COIs. Among the three reachable COIs for path p, the path formed by adding c^2 leads to the lowest overall cost.

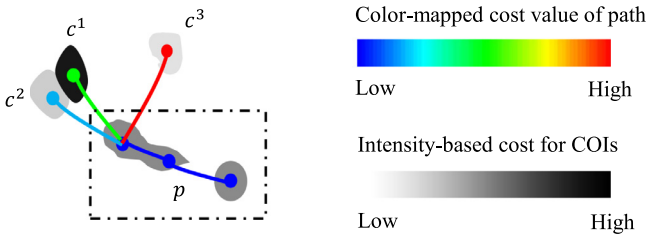

Fig. 5. A simplified example of the coarse path-finding algorithm in GP. At the ith iteration, the existing path p consisting of i-1 nodes has 3 reachable COIs c^1, c^2, and c^3. The path-finding algorithm computes the shape-based cost and intensity-based cost for the three COIs, and adding c_i^2 to the existing path will result in the lowest cost. Compared to c^2, c^1 has acceptable shape-based features but its intensity-based cost is high. Although c^3 has the lowest intensity-based cost, the sharp turn formed by c^3 and p makes its shape-based cost high. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

At the ith iteration, candidate path p has $i-1$ COIs. The cost for adding a new COI c into path p is:

$$C_{O1}(c, p) = \rho C_{I1}(c) + C_{S1}(c, p) \tag{5}$$

where ρ is a weighting scalar to specify how much we rely on the intensity-based term $C_{I1}(c)$ relative to the shape-based cost term $C_{S1}(c, p)$. N is the total number of electrodes in the array. The intensity based cost term $C_{I1}(c)$ is defined as:

$$C_{I1}(c) = w \cdot \left(\mu_I \frac{I_{\max} - I(c)}{I_{\max}} + \mu_B \frac{I_{\mathrm{Bmax}} - I_{\mathrm{B}}(c)}{I_{\mathrm{Bmax}}} + \mu_V \frac{I_{\mathrm{Vmax}} - I_{\mathrm{V}}(c)}{I_{\mathrm{Vmax}}}\right), \tag{6}$$

where $I(c)$, $I_{\mathrm{B}}(c)$, and $I_{\mathrm{V}}(c)$ are the image intensity, blob filter response, and vesselness filter response at COI c; and $I_{\max}$, I_{Bmax}, I_{Vmax} are the maximum values of those images for all the COIs, respectively. The blob filter is as described in Eq. (2). The vesselness filter is Frangi's vesselness filter (Frangi et al., 1998) with a scale of 0.25 mm. $\mu_I = 1$, $\mu_B = \lambda_B$, and $\mu_V = \lambda_I$ are weighting scalars. We include the image intensity because voxels occupied by metallic electrodes are usually assigned high intensity. The blob filter response is included because the electrodes often have a blob-like appearance. When d_m increases, I_{B} becomes more reliable and μ_B increases. We also include the vesselness filter response because the electrodes sometimes have a tubular appearance if there is not much contrast between them in CT images. When d_m decreases, I_{V} becomes more reliable and μ_V increases. w is a multiplier we use to punish solutions for which the first electrode is selected as a COI with low blob filter response. We do so to capture the fact that the first electrode usually has a high blob filter response because it only has a neighbor in one direction. At the i^{th} iteration, w is defined as:

$$w = \begin{cases} 100, & i = 1 \text{ and } I_B(c) < T_B\left(\alpha_B'\%\right) \\ 1, & \text{otherwise} \end{cases}, \tag{7}$$

where $T_B(\alpha_B'\%)$ is a function that defines a threshold value applied to I_B that corresponds to the top $\alpha_B'\%$ of the cumulative histogram of the blob filter response. The value w at the first iteration is empirically determined with the heuristic that the most basal electrode is usually inserted in a shallow position of the cochlea, which makes its intensity-based feature obvious. Next, the shape-based cost term $C_{S1}(c, p)$ evaluates the geometric features of a candidate path p when a COI c is added. It is defined as:

$$C_{S1}(c, p) = \mu_{\mathrm{d}} C_{\mathrm{d}}(c, p_{i-1}) + \mu_{\mathrm{s}}(C_{\mathrm{a}}(c, p_{i-1}, p_{i-2}) + C_{\mathrm{ins}}(c, p_{i-1})) \tag{8}$$

where $C_{\mathrm{d}}(\cdot)$, $C_{\mathrm{a}}(\cdot)$, and $C_{\mathrm{ins}}(\cdot)$ are the distance-based, smoothness-based, and the angular depth of insertion (DOI) based cost terms, respectively. The first term $C_{\mathrm{d}}(c, p_{i-1})$ is defined as:

$$C_{\mathrm{d}}(c, p_{i-1}) = |\mathrm{dist}(c, p_{i-1}) - D_{i-1}|, \tag{9}$$

$$\mu_{\mathrm{d}} = \begin{cases} \mu_{\mathrm{d1}}, & if\ \mathrm{dist}(c, p_{i-1}) < D_{i-1} \\ \mu_{\mathrm{d2}}, & if\ \mathrm{dist}(c, p_{i-1}) \geq D_{i-1} \end{cases} \tag{10}$$

where $\mathrm{dist}(c, p_{i-1})$ is the Euclidean distance between a COI c to the endpoint of a candidate path p. Eqs. (9) and (10) punish the candidate path from growing an edge that is shorter or longer than the expected distance. $C_{\mathrm{a}}(c, p_{i-1}, p_{i-2})$ is determined as:

$$C_{\mathrm{a}}(c, p_{i-1}, p_{i-2}) = \left(\angle(c, p_{i-1}, p_{i-2}) - \tilde{\angle}_{i-1}\right) H\left(\angle(c, p_{i-1}, p_{i-2}) - \tilde{\angle}_{i-1}\right), \tag{11}$$

where $H(\cdot)$ is the Heaviside function, $\angle(c, p_{i-1}, p_{i-2})$ is the bending angle formed by adding c to the last two endpoints p_{i-1}, p_{i-2} of an existing candidate path p and is defined as:

$$\angle(c, p_{i-1}, p_{i-2}) = 1 - \frac{(c - p_{i-1}) \cdot (p_{i-1} - p_{i-2})}{\mathrm{dist}(c, p_{i-1}) \cdot \mathrm{dist}(p_{i-1}, p_{i-2})} \tag{12}$$

and $\tilde{\angle}_{i-1}$ is a heuristically selected threshold bending angle value. Eq. (11) punishes paths with bending angles that are sharper than the threshold value. From the ground truth localization results in our training dataset, we observed that (1) the electrodes inserted deeper in the cochlea have a sharper bending angle than the electrodes that are inserted shallower because the curvature of the cochlea increases with increasing DOI, and (2) arrays from the MD family have sharper bending angles than arrays from the AB family due to a larger spacing distance between electrodes for MD arrays. Thus, we determine $\tilde{\angle}$ values for arrays from AB ($\tilde{\angle}_{\mathrm{AB}}(\cdot)$) and MD ($\tilde{\angle}_{\mathrm{MD}}(\cdot)$) families separately. $\tilde{\angle}_{\mathrm{AB}}(\cdot)$ and $\tilde{\angle}_{\mathrm{MD}}(\cdot)$ are set as:

$$\tilde{\angle}_{\mathrm{AB}}(i) = \begin{cases} 0.30\ , & i \leq E_{\mathrm{Half}} \\ 0.59\ , & i > E_{\mathrm{Half}} \end{cases}, \tag{13}$$

$$\tilde{\angle}_{\mathrm{MD}}(i) = \begin{cases} 0.56 & i \leq E_{\mathrm{Half}} \\ 1.27\ , & i > E_{\mathrm{Half}} \end{cases}, \tag{14}$$

where $E_{\mathrm{Half}} = \frac{N}{2}$ is used to empirically distinguish the electrodes that are inserted deeply versus shallowly in the cochlea. The values in Eqs. (13) and (14) were selected as 130% of the maximum bending angles observed among training AB and MD arrays when $i \leq E_{\mathrm{Half}}$ and $i > E_{\mathrm{Half}}$. The DOI cost $C_{\mathrm{ins}}(c, p_{i-1})$ is defined as:

$$C_{\mathrm{ins}}(c, p_{i-1}) = (H(\mathrm{DOI}(p_{i-1}) - \mathrm{DOI}(c)) + H(|\mathrm{DOI}(c) - \mathrm{DOI}(p_{i-1})| - 180^\circ)) \tag{15}$$

where $\mathrm{DOI}(c)$ is the angular depth of insertion value for COI c. As the cochlea has a spiral shape with 2.5 turns, the depth of any position within the cochlea can be quantified in terms of an angle from 0 to 900°. To obtain the $\mathrm{DOI}(\cdot)$ values, we register a pre-implantation CT, in which the intra-cochlear anatomy is segmented, to our post-implantation target CT. For recipients that do not have pre-implantation CTs, our group also has developed a method to segment the intra-cochlear anatomy from post-implantation CTs directly (Reda et al., 2014). These two methods are used to generate a DOI map for each individual voxel in the post-implantation CT. The first term in Eq. (15) punishes paths in which a newly added COI c has a $\mathrm{DOI}(c)$ value that is smaller than the endpoint p_{i-1} on the path p. The second term in Eq. (15) punishes adding a COI c into an existing path p when the COI c is more than a half turn (180°) ahead or behind the endpoint of p. The DOI term constrains the candidate path to grow in the correct direction and to not cross two turns of the cochlea. We have observed that in some training cases, several electrodes fall on the boundary of two turns of the cochlea and using the nearest neighboring voxel of those COIs for DOI estimations can be inaccurate. We have designed a solution for the computation of the DOI cost term that is robust to boundary COIs and these details are presented in Appendix C.

With the cost function defined above, GP runs the first path-finding algorithm to coarsely localize the location of the electrodes. Then, the candidate path with the lowest overall cost is selected as the coarsely localized electrode array.

2.5. Path refinement

The process described in sub-Section 2.4 coarsely localizes the electrodes. The second path finding procedure is used to refine the coarse result in a local region around each coarsely localized electrode. In this step, the method defines a set of COIs $\{c\}^i$ around each coarsely localized electrode p_i by sampling a fine rectangular

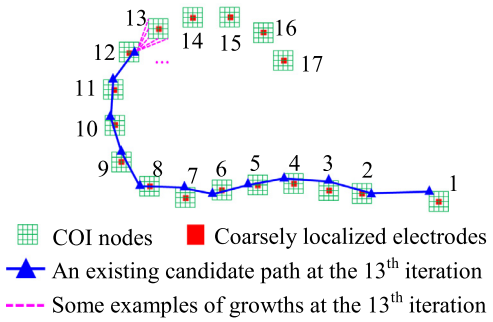

Fig. 6. Visualization of the path-refinement process at iteration 13 for an existing candidate path. This path grows by adding all the COI nodes (the re-sampled rectangular grids) around the 13th electrode to it. The prune step keeps only $\eta_{\max 2}$ candidate paths with lowest costs for the next iteration. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

grid of points (Shown in Fig. 6). The set of candidate COIs around the ith coarsely localized electrode is defined as:

$$\{c\}^i = \{p_i + \varphi_q[x, y, z]\}_{x,y,z\in[-\varphi_r,\ \varphi_r]} \tag{16}$$

In the path refinement algorithm, our method aims at localizing N electrodes after N iterations. We use a candidate path group $\{p\}$ which is similar to the one being described in sub-Section 2.4 to store the candidate paths during the path finding algorithm. At the first iteration, all the nodes in $\{c\}^1$ are treated as seed nodes which represent candidate paths with length 1. At the ith iteration, the method grows the candidate paths by adding the candidate nodes $\{c\}^i$ to the existing candidate paths in the candidate path group (Shown in Fig. 6). Then the method prunes the candidate path group by keeping only $\eta_{\max 2}$ paths with the lowest cost in the group after each iteration. The cost function to evaluate the quality of a new candidate path constructed by adding a COI c to an existing candidate path p consists of an intensity-based cost term and a shape-based cost:

$$\mathrm{Cost}_2(c, p) = \mathrm{C}_{\mathrm{I2}}(c) + \mathrm{C}_{\mathrm{S2}}(c, p) \tag{17}$$

The intensity-based cost term $C_{\mathrm{I2}}(c)$ is defined as:

$$C_{\mathrm{I2}}(c) = -(\varphi_I G_\sigma(I(c)) + \varphi_B I_B(c)) \tag{18}$$

where $G_\sigma(I(c))$, and $I_B(c)$ are the Gaussian filter response, and the blob filter response at c, respectively. σ is the scale of the Gaussian filter. The shape-based cost is defined as:

$$C_{\mathrm{S2}}(c, p) = |\mathrm{dist}(c, p) - D_{i-1}| \cdot \begin{cases} \varphi_{\mathrm{d1}}, & dist(c, p) < D_{i-1} \\ \varphi_{\mathrm{d2}}, & dist(c, p) \geq D_{i-1} \end{cases} \tag{19}$$

where dist(c, p) is the Euclidean distance between node c to the endpoint electrode on path p. After N iterations, the path with the lowest overall cost is selected as the final localization result generated by GP.

2.6. Parameter tuning for GP

The parameter tuning process was performed by using the CTs in our training dataset. The initial values of the parameters were heuristically determined. Then, parameters were optimized sequentially and iteratively until a local optimum was reached for each parameter with respect to the mean localization errors in the training dataset. The parameters used in the coarse localization step were optimized first and then the parameters used in the refinement step were optimized. After determining the optimized values of all the parameters, we fixed those parameter values and performed validation study of the GP on the testing dataset.

Table 3
The selected values for parameters in GP.

Coarse path-finding algorithm		Path refinement algorithm	
$\eta_{\max}$	1200	$\eta_{\max 2}$	500
α_I (%)	0.048 (%)	φ_q	0.03
α_B(%)	0.028 (%)	φ_r	3
β_I	2.72	σ	0.275
κ_I	1.82	φ_I	32
β_B	1.14	φ_B	16
κ_B	1.21	φ_{d1}	0.6
γ_1	0.6	φ_{d2}	2.5
γ_2	1.2		
ρ	2.0		
α'_B(%)	0.007 (%)		
μ_{d1}	10		
μ_{d2}	6		
μ_S	450		

2.7. Evaluation of methods

In our validation study, we compared the performance of the proposed method GP with the baseline method lGP (Noble and Dawant, 2015) and a preliminary implementation of GP (pGP) (Zhao et al., 2017) on our testing dataset in Dataset 1 with 125 clinical CTs implanted with different types of distantly-spaced electrode arrays. The baseline method lGP relies solely on image intensity of the VOI to generate ROI and COIs. Because of this limitation, we expect it to generate less accurate results for most eCTs and unacceptable results for most lCTs because the false positive COIs in lCTs are assigned the same maximum intensity as the true positive COIs. pGP is a preliminary implementation of GP. It uses a set of two fixed weighting scalars (λ_B and λ_I in Eq. (1)) to generate feature images for ROIs and COIs generation. For lCTs, to reduce false positives among COIs pGP performs image opening on the ROIs with an empirically selected kernel size, which may accidentally remove true positive COIs. With GP, a cost function term is used as a soft-constraint so that true positive COIs are not eliminated. In contrast to lGP and pGP, GP generates COIs with sub-voxel resolution, which should permit more accurate results with the subsequent path-finding algorithms. To evaluate the difference between electrode localization methods, we performed a paired t-test with Bonferroni correction on the mean localization errors generated by lGP, pGP, GP and RCE on all the possible pairs. We visually confirmed that these data approximately form normal distributions with different means to ensure the usage of t-test is appropriate.

We used Dataset 2 to evaluate the localization robustness of all the automatic electrode localization methods to different CT acquisition parameters and the different models of CI electrode arrays. We performed paired t-tests with Bonferroni correction on the mean localization errors generated by lGP, pGP, GP, and IL on all the possible pairs to evaluate the statistical differences among them.

3. Results

3.1. Parameter tuning

Table 3 lists the parameter values after the tuning process. To show the effectiveness of the parameters we selected, the parameter sweeping procedure is visualized in Fig. 7 with respect to the mean localization errors in log-scale. Each parameter was swept from 0 to the double of its final selected value with uniform step size. Two exceptions are $\eta_{\max}$ and $\eta_{\max 2}$. For $\eta_{\max}$ and $\eta_{\max 2}$, we start by setting them as 1 because the two path-finding algorithms need to store at least one candidate path.

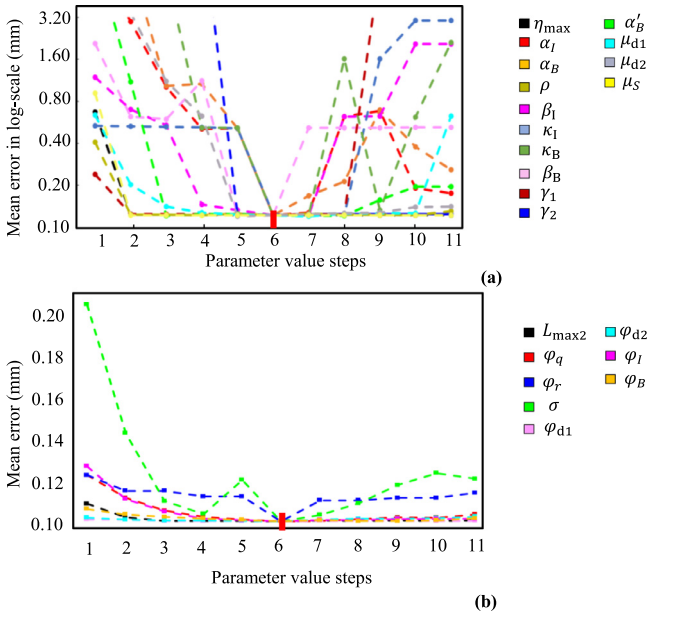

Fig. 7. Visualization of errors when testing each parameter used in the coarse path-finding algorithm (a) and the path refinement (b) in GP. Each parameter is tested over a range from 0 to the double of the optimal values. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.2. Electrode localization accuracy

We run GP, lGP, pGP on a standard Windows Server PC [Intel (R), Xeon (R) CPU X5570, 2.93 GHz, 48GB Ram]. The average running time for GP from CT registration to electrode localization is ~40 s, which is longer than pGP (~8 s) and lGP methods (~5 s). GP has a longer run time because it up-samples VOIs to generate COIs with sub-voxel resolution. This COI generation process takes ~32 s. The two path-finding algorithms in GP take ~8 s.

We define a "failure" as a case for which a method fails to find a fixed-length path from the COIs it generates or for which the method generates a solution that has a maximum error that is larger than one voxel diagonal. Among 125 clinical CTs in our testing dataset, lGP, pGP, and GP fails for 74, 47, and 5 cases, respectively. One major reason for the methods to fail is that COIs cannot be produced for one or more electrodes, and thus the subsequent coarse path-finding algorithm is not able to find a fixed-length path with N COIs representing the electrodes on the array that obeys the hard reachability constraints. Among the failures by lGP, pGP, and GP, 6, 13, and 2 cases are due to this reason, respectively. Fig. 8a shows the quantitative analysis of the localization results generated by lGP, pGP, GP, and the rater's consistency errors (RCEs) in boxplots. Besides the cases for which no path could be found, the proposed method generates coarse localization results with a mean error of 0.15 mm and final localization results with a mean error of 0.12 mm. The mean of GP's final localization error is close to the mean RCE error, which is 0.09 mm. Fig. 8b shows the distribution of the number of cases that have localization errors that fall into the intervals [0, 25%), [25%, 50%), [50%, 75%), [75%, 100%), and larger than or equal 100% of the voxel diagonal, as well as the failure cases. As can be seen from Fig. 8b, GP generates 120 out of 125 (96%) localization results that have maximum errors within one voxel diagonal, which is close to the RCE (100%) and outperforms pGP (58%) and lGP (41%). Example results of the GP, lGP, and pGP methods are shown in Figs. 9 and 10.

The t-test values obtained when comparing the mean localization errors of GP, lGP, pGP and RCE are as follows: 8.36×10^{-12} for lGP-pGP, 1.07×10^{-15} for lGP-GP, 3.21×10^{-16} for lGP-RCE, 2.20×10^{-8} for pGP-GP, 8.16×10^{-9} for pGP-RCE, and 1.24×10^{-1} for GP-RCE. These results show that the results generated by GP are significantly different from those obtained with lGP and pGP but are not significantly different from RCE.

3.3. Robustness to acquisition parameters

Table 4 lists the mean localization errors of the automatic methods along with the image-based localization error for Dataset 2. Five results from lGP and three results from pGP were not included in these results because they were too inaccurate and could lead to spurious inferences. From Table 4, we find that among GP, lGP, and pGP, GP generated the electrode localization results with the lowest mean localization error for all groups of CT images. The error of GP is close to the image-based localization error IL. From the t-test results, we found that both lGP and pGP add significant algorithmic errors beyond the image-based localization errors for both AB1 and AB2, which is not unexpected. However, the auto-

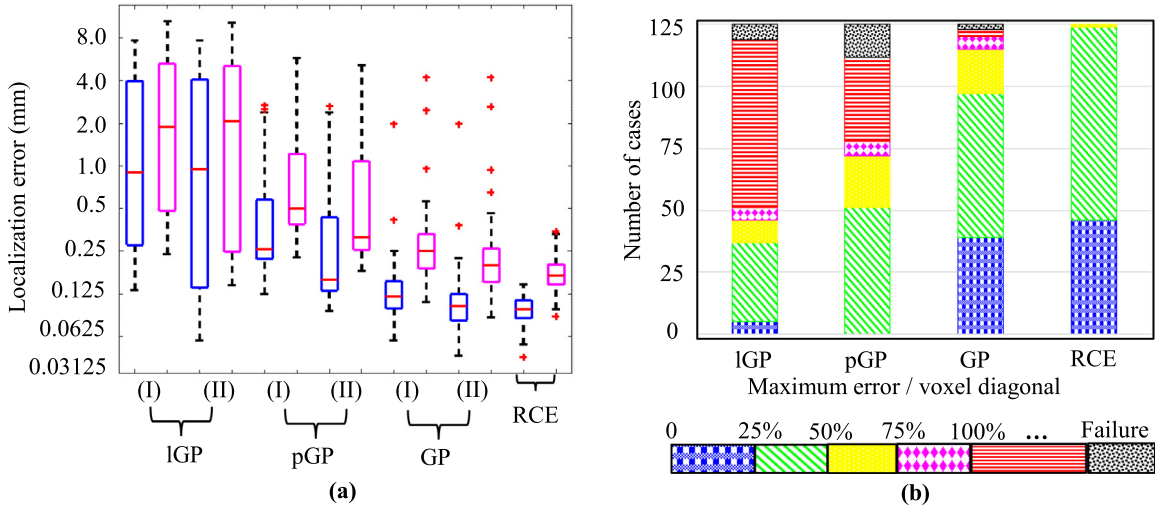

Fig. 8. Panel (a) shows the boxplot (in log-scale) of mean (blue) and maximum (magenta) coarse (I) and refined (II) localization errors of the automatically generated results by lGP, pGP, GP and the rater's consistency errors (RCEs). Panel (b) shows the bar plot of the number of cases on which lGP, pGP, GP, and RCE achieves maximum final localization errors lower than 25% (blue), lower than 50% (green and blue), lower than 75% (magenta, green, and blue), lower than 100% (yellow, magenta, green, and blue), over 100% (red) the length of a voxel diagonal. Failure cases where no result could be generated are shown in black. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

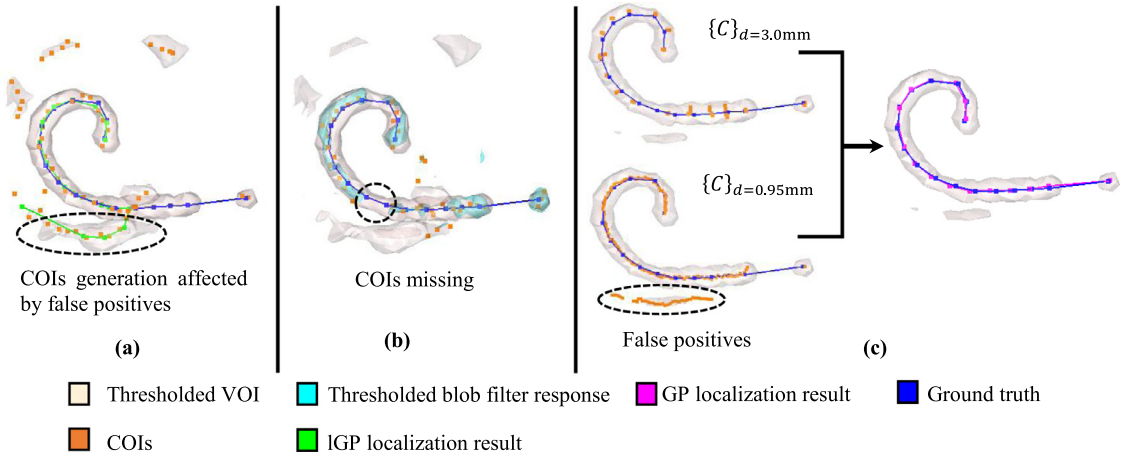

Fig. 9. Visualization of localization results generated by (a) lGP and by (c) GP. In (b), pGP fails to generate a fixed-length path as final localization result because the COIs are missing around two electrodes. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

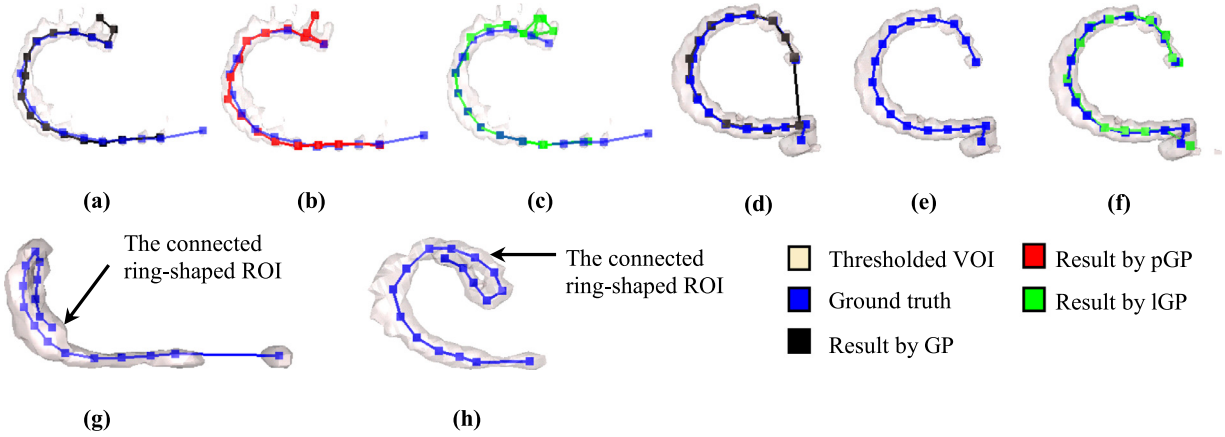

Fig. 10. Panel (a)–(c) and (d)–(f) show localization results generated by GP, lGP, and pGP for two cases, respectively. Panel (g) and (h) show two failure cases for GP. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 4
Mean localization errors for each image group in mm.

	HU range		Resolution		Dose		Array	
	lCTs	eCTs	Low	Mid	Mid	High	AB1	AB2
lGP	1.59 ± 1.97	0.14 ± 0.10	1.59 ± 1.97	0.15 ± 0.10	1.22 ± 1.83	0.18 ± 0.11	0.50 ± 1.17	0.56 ± 1.22
pGP	0.40 ± 0.42	0.19 ± 0.06	0.20 ± 0.07	0.33 ± 0.36	0.38 ± 0.38	0.17 ± 0.06	0.26 ± 0.27	0.45 ± 0.43
GP	**0.13 ± 0.06**	**0.08 ± 0.05**	**0.13 ± 0.06**	**0.08 ± 0.05**	**0.12 ± 0.06**	**0.10 ± 0.06**	**0.10 ± 0.06**	**0.11 ± 0.07**
IL	0.11 ± 0.05	0.07 ± 0.04	0.10 ± 0.06	0.07 ± 0.04	0.10 ± 0.05	0.08 ± 0.04	0.08 ± 0.05	0.14 ± 0.15

matic localization error obtained with GP is not significantly different from the image-based localization errors (IL).

4. Discussion

4.1. Sensitivity of parameters

From Fig. 7a, we observe that every parameter contributes to the coarse localization step and setting any of them to 0 increases the mean localization error. This indicates that the cost terms we have designed are useful, and the parameters we selected are effective in achieving low localization errors. Among the parameters in the coarse path-finding algorithm, α_I, α_B, β_I, β_B, κ_I, κ_B, and μ_{d2} are sensitive because adjusting them from their selected values results in much larger errors. Aside from μ_{d2}, the other sensitive parameters are all related to feature image construction and COIs generation, which shows that the COI generation step plays a crucial role for the following path-finding algorithms to localize the array. η_{max}, ρ, α'_B, μ_{d1}, γ_1, γ_2, and μ_S are not sensitive around the selected values. However, using the selected values for those parameters leads to a local mean localization error minimum on the training dataset in our parameter tuning process. From Fig. 7b, we can observe that the refined localization errors are relatively flat around the selected values of each individual parameters. The most sensitive parameter is the scale σ for the Gaussian blur filter. The other parameters are not sensitive around their selected values. However, setting any parameter as 0 increases the mean localization error on the training dataset.

4.2. Electrode localization accuracy

Fig. 9 shows an example case (eCT implanted with an AB2 array) in which the GP localization method (panel c) was able to very accurately localize the electrode locations among the correct set of COIs, even when some false positive COIs are present. For the same case, the lGP method generates an inaccurate result (panel a). This

is because the threshold selected for generating the ROIs and COIs is not high enough to eliminate many false positive voxels in the VOI, and lGP only relies on the image intensity for COIs generation. We also performed pGP on the same case but it fails to generate a result (Fig. 9b). pGP uses the blob filter response to enhance the high intensity blob-like structures in the VOI, but because pGP uses a single set of fixed weighting scalars for image intensity and blob filter response to construct a feature image rather than ESD-based parameters, the method removes some ROIs that have less blob-like features because they are relatively closely-spaced. Consequently, the method fails to find a fixed length path with 17 COIs representing all the electrodes on the array. GP generates two sets of COIs for two ESD values. As can be seen, for $d = 0.95$mm, the COIs generation relies more on the image intensity, which results in more false positives but is less likely to miss electrodes. For $d = 3.0$mm, the COIs generation step relies on the blob filter response, which enhances the distantly-spaced electrodes that have a more obvious blob-like shape in the CT image. GP also upsamples the VOI, which permits the generation of more accurate COIs.

In Fig. 10, we show the only 2 cases for which GP fails to generate localization results (Shown in Fig. 10a–f) and 2 example cases for which GP generate localization results with maximum errors more than one voxel diagonal (Shown in Fig. 10g–h). These are complicated cases for which all the automatic methods fail. Panels (a)–(c) and panels (d)-(f) show three sets of localization results generated by GP, pGP, and lGP for two cases implanted with AB1 arrays. In Case 1, the CT has abnormal intensity features due to beam hardening artifact. Around the most apical electrodes, several false positive voxels are assigned similar high intensity values as the voxels occupied by the actual electrodes. Meanwhile, the inactive electrode has low intensity and blob filter response. This causes all three localization methods to miss the inactive electrode and wrongly select one of the false positives as the most apical electrode. In Case 2 shown in Fig. 10d–f, the inactive electrode lies much closer than usual to the first active electrode because the array is kinked between the electrodes. This leads to poor localization results generated by all the three methods. Fig. 10g–h shows 2 cases on which GP fails to generate a path. This is because the electrode arrays in these two cases are folded. The ROIs generated by GP could not distinguish the electrodes that are pushed together. These 4 cases indicate our method is not robust to extreme cases with severe imaging artifacts or where the array is kinked or folded.

4.3. Robustness to CT acquisition parameters

As is presented in sub-Section 3.2, the mean localization errors of GP and IL are close and the Bonferroni corrected p value of the paired t-test between them is larger than 0.05. This indicates that with GP, we have achieved the best level of localization accuracy that can be reasonably expected with these images given the imaging technique employed. The accuracy also isn't unduly affected by the variations of the parameters when compared to the image-based localization errors. Unlike lGP and pGP, GP does not produce poor localizations in case of low resolution or low dose images. The proposed method is thus highly accurate and robust to changes in CT acquisition parameters.

4.4. Limitations and future works

One limitation of this study is that the generation of COIs is highly dependent on intensity-based features. As is shown in Fig. 10, abnormal intensity-based features could lead to failure or inaccurate localization results. Another limitation of our proposed method is that it is not robust to electrode arrays that are kinked or folded. Future work will be aimed at addressing these limitations.

The proposed GP method fully automates the localization of distantly spaced electrodes for IGCIP and achieves accurate localization results on most of the CTs in our dataset. This enables the translation of IGCIP methods for clinical use by end-users such as audiologists who in general are not trained in medical imaging. However, the sensitivity of the improvement in hearing outcomes when using IGCIP with respect to the electrode localization error remains unknown. In the future, we plan to conduct a prospective clinical study to evaluate the sensitivity of IGCIP and hearing outcomes with respect to different levels of electrode localization errors (Zhao et al., 2018a,c).

5. Conclusion

In this study, we propose an automatic graph-based method for localizing distantly-spaced CI electrode arrays in clinical CTs with sub-voxel accuracy. We use a method to generate candidate voxels of interests that are around electrodes at a sub-voxel resolution and use two path-finding algorithms to find a fixed-length path whose nodes represent electrodes on the array. We perform a parameter tuning process for our proposed method on a training dataset with clinical CTs implanted with different types of distantly-spaced arrays. The results of the validation studies on a large-scale testing dataset including 125 clinical CTs, and 28 phantom CTs show the accuracy and robustness of our proposed method. When compared with the other two previously developed methods, our proposed GP achieves the lowest mean localization error of 0.12 mm and fails to generate localization results with maximum errors within one voxel for only 5 cases. Our proposed automatic method generates localization results that are not significantly different from the localization results generated by an expert. The validation study on 28 CTs acquired from a cochlear implant imaging phantom indicate that our proposed method is robust to CT acquisition parameters. The overall localization errors of GP are significantly different from the errors of the previously developed methods and are close to the rater consistency errors.

To our best knowledge, GP is the most accurate automatic CI electrode localization method available. This represents a crucial step for fully automating our IGCIP techniques and translating IGCIP into clinical use. It also enables us to conduct comprehensive large scale studies on the correlation between hearing outcomes and the intra-cochlear locations of CI electrodes.

Conflict of interest

There are no known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome.

Acknowledgments

This work was supported in part by grants R01DC014037, R01DC008408 and R01DC014462 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 these institutes.

Appendix A. Three types of distantly-spaced CI electrode arrays provided by the two major manufacturers

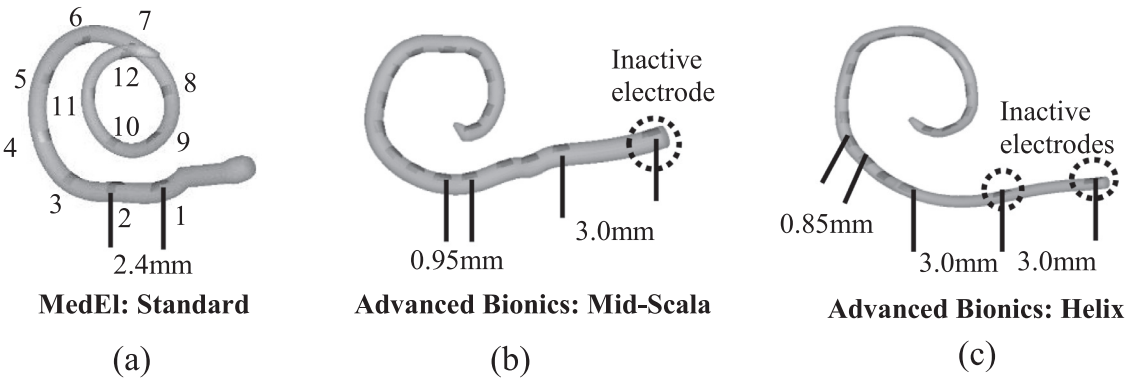

Appendix B. Reachability of a COIs to a candidate path

At the ith iteration, a candidate path p consists of $i-1$ COIs. In the grow stage, a COI $c_m^{j,k}$ is considered "reachable" for a candidate path p if it obeys the following 5 hard constraints. First, it should be such that $\gamma_1 D_{i-1} < \text{dist}(p_{i-1}, c_m^{j,k}) < \gamma_2 D_{i-1}$. In this equation, $\text{dist}(p_{i-1}, c)$ is defined as the Euclidean distance between a COI $c_m^{j,k}$ and the endpoint p_{i-1} of the candidate path p. This constrains the distance between the current endnode of the path and the candidate node to be close to the expected *a-priori* distance D_{i-1}. The second constraint imposes that $c_m^{j,k}$ is only reachable for p if $D_{i-1} = d_m$. This constrains the candidate node to belong to the corresponding ESD value. The third constraint imposes that $c_m^{j,k} \notin p$, which forbids to add a COI to a path if the COI is already in the path, keeping the path from looping back upon itself. The fourth constraint imposes that if $p_{i-1} \notin \{C\}_{D_{i-1}}^{j}$ (the ROI for $c_i^{j,k}$), then it is only permitted to add $c_i^{j,k}$ to p if $p_z \notin \{C\}_{D_{i-1}}^{j}$, $\forall z \in [1, i-2]$. This constraint does not allow the path to return to the ROIs that the candidate path p has already visited. The last constraint imposes that if $p_{i-1}, p_{i-2} \in \{C\}_{D_{i-1}}^{j}$ and $d_m = D_{i-1}$, $c_m^{j,k}$ is only reachable for p if p_{i-2}, p_{i-1}, and $c_m^{j,k}$ are monotonically ordered in the medial axis line $\{C\}_{D_{i-1}}^{j}$. This constraint prevents the path from looping back within an ROI, since COIs belonging to a ROI should be ordered identically to the ROI's medial axis.

Appendix C. Solution for the computation of the DOI cost term of the COIs on the boundary of two turns of the cochlea

During the COIs generation step, due to (1) the registration errors between pre- and post-implantation CTs, (2) the localization errors for intra-cochlear anatomy segmentation, and (3) the limited accuracy of voxel thinning method for generating COIs from ROIs, the DOI values of the COIs close to the boundary of two turns of cochlea could be inaccurately estimated. Fig. C.1 shows one example that is implanted with a MD2 array. In Fig. C.1, we label the DOI$(\cdot)$ values for each individual ground truth location of the electrodes. As we can see, the 11th electrode is on the boundary between the second turn and the first turn of cochlea. The DOI of the COIs for the 11th electrode would be estimated incorrectly if using a nearest neighbor approach. In the path-finding algorithm, this will lead to a large cost value when growing a path from the 10th electrode to the COIs for the 11th electrode. To solve this issue, for each COI, we find the maximum and minimum (DOI_{Max} and DOI_{Min}) among the DOI values for each voxel in a $3 \times 3 \times 3$ neighborhood around its nearest neighbor voxel. If $\text{DOI}_{\text{Max}} - \text{DOI}_{\text{Min}} \geq 180°$, the COI is near a border and so we create an additional "phantom" COI for the original COI at the same location in the image. The DOI values of the phantom COI and the COI are assigned DOI_{Max} and DOI_{Min}, respectively. Aside from DOI values, the phantom COI has the same information as the original COI. Thus, the path-finding algorithm has equal chance to visit the phantom COI and the original COI and evaluate the cost value for the candidate path with two estimates of the DOI values.

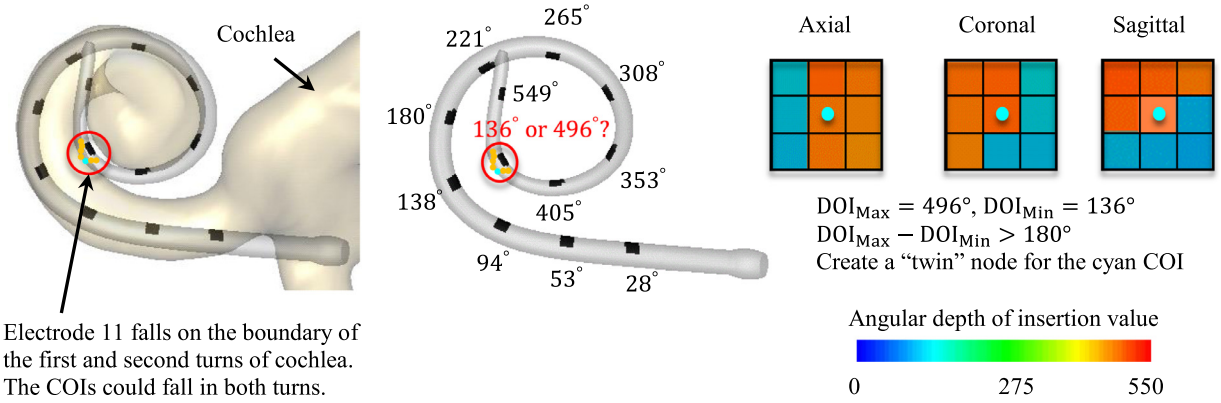

Fig. C.1. One example of the boundary problem when computing DOI$(\cdot)$. The 11th electrode falls on a boundary between two turns of the cochlea. On the right side, the color-coded angular DOI map around the electrode is shown. The DOI map is generated by resampling a $3 \times 3 \times 3$ voxel rectangular grid around the closest voxel to the 11th electrode.

References

Aschendorff, A., Kubalek, R., Turowski, B., Zanella, F., Hochmuth, A., Schumacher, M., Klenzner, T., Laszig, R., 2005. Quality control after cochlear implant surgery by means of rotational tomography. Otol. Neurotol. 26, 34–37.

Bennink, E., Peters, J.P.M., Wendrich, A.W., Vonken, E., van Zantan, G.A., Viergever, M.A., 2017. Automatic localization of cochlear implant electrode contacts in CT. Ear Hear. 38 (6), e376–e384. doi:10.1097/AUD.0000000000000438.

Boëx, C., de Balthasar, C., Kós, M.-I., Pelizzone, M., 2003. Electrical field interactions in different cochlear implant systems. J. Acoust. Soc. Am. 114, 2049–2057.

Bouix, S., Siddiqi, K., Tannenbaum, A., 2005. Flux driven automatic centerline extraction. Med. Image Anal. 9, 209–221.

Braithwaite, B., Kjer, H.M., Fagertun, J., Ballaster, M.A.G., Dhanasingh, A., Mistrik, P., Gerber, N., Paulsen, R.R., 2016. Cochlear implant elecrode localization in post–operative CT using a spherical measure. Biomedical Imaging (ISBI), 2016 IEEE 13th International Symposium on Prague, Czech Republic.

Chakravorti, S., Bussey, B.J, Zhao, Y., Dawant, B.M., Labadie, R.F., Noble, J.H., 2017. Cochlear implant phantom for evaluating computed tomography acquisition parameters. J. Med. Imaging 4 (4), 045002. doi:10.1117/1.JMI.4.4.045002, (2017).

Dijkstra, E.W, 1959. A note on two problems in connexion with graphs. Numer. Math. 1, 269–271.

Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A., 1998. Multiscale vessel enhancement filtering. In: Wells, W.M., Colchester, A., Delp, S. (Eds.), Medical Image Computing and Computer-Assisted Interventation—MICCAI'98: First International Conference Cambridge, MA, USA, October 11–13, 1998 Proceedings. Springer, Berlin Heidelberg, pp. 130–137.

Fu, Q.-J., Nogaki, G., 2005. Noise susceptibility of cochlear implant users: the role of spectral resolution and smearing. J. Assoc. Res. Otolaryngol. 6, 19–27.

Gifford, R.H., Shallop, J.K., Peterson, A.M., 2008. Speech recognition materials and ceiling effects: considerations for cochlear implant programs. Audiol. Neuro-otol. 13, 193–205.

Gifford, R.H., Hedley-Williams, A., Spahr, A.J., 2014. Clinical assessment of spectral modulation detection for adult cochlear implant recipients: a non-language based measure of performance outcomes. Int. J. Audiol. 53, 159–164.

Greenwood, D.D., 1990. A cochlear frequency-position function for several species – 29 years later. J. Acoust. Soc. Am. 87, 2592–2605.

National Institute on Deafness and Other Communication Disorders, 2011. NIDCD Fact Sheet: Cochlear Implants, NIH Publication No. 11-4798.

Noble, J.H., Labadie, R.F., Majdani, O., Dawant, B.M., 2011. Automatic segmentation of intra-cochlear anatomy in conventional CT. IEEE Trans. Bio-med. Eng. 58, 2625–2632.

Noble, J.H., Gifford, R.H., Labadie, R.F., Dawant, B.M., 2012. Statistical shape model segmentation and frequency mapping of cochlear implant stimulation targets in CT. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (Eds.). Medical Image Computing and Computer-Assisted Intervention – MICCAI 2012. MICCAI 2012. Lecture Notes in Computer Science, 7511. Springer, Berlin, Heidelberg.

Noble, J.H., Labadie, R.F., Gifford, R.H., Dawant, B.M., 2013. Image-guidance enables new methods for customizing cochlear implant stimulation strategies. Neural Syst. Rehabil. Eng. IEEE Trans. 21, 820–829.

Noble, J.H., Gifford, R.H., Hedley-Williams, A.J., Dawant, B.M., Labadie, R.F., 2014. Clinical evaluation of an image-guided cochlear implant programming strategy. Audiol. Neurotol. 19, 400–411.

Noble, J.H., Dawant, B.M., 2015. Automatic graph-based localization of cochlear implant electrodes in CT. In: Navab, N., Hornegger, J., Wells, M.W., Frangi, F.A. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part II. Springer International Publishing, Cham, pp. 152–159.

Noble, J.H., Hedley-Williams, A.J., Sunderhaus, L., Dawant, B.M., Labadie, R.F., Camarata, S.M., Gifford, R.H., 2016. Initial results with Image-guided Cochlear Implant Programming in Children. Otol. Neurotol. 37 (2), e63–e69.

Reda, F.A., McRackan, T.R., Labadie, R.F., Dawant, B.M., Noble, J.H., 2014a. Automatic segmentation of intra-cochlear anatomy in post-implantation CT of unilateral cochlear implant recipients. Med. Image Anal. 18, 605–615.

Reda, F.A., Noble, J.H., Labadie, R.F., Dawant, B.M., 2014b. An artifact-robust, shape library-based algorithm for automatic segmentation of inner ear anatomy in post-cochlear-implantation CT. In: Proceedings of SPIE–the International Society for Optical Engineering 9034, 90342V.

Rubinstein, J.T., 2004. How cochlear implants encode speech. Curr. Opin. Otolaryngol. Head Neck Surg. 12, 444–448.

Skinner, M.W., Holden, T.A., Whiting, B.R., Voie, A.H., Brunsden, B., Neely, J.G., Saxon, E.A., Hullar, T.E., Finley, C.C., 2007. In vivo estimates of the position of advanced bionics electrode arrays in the human cochlea. Ann. Otol. Rhinol. Laryngol. 116, 2–24.

Stakhovskaya, O., Sridhar, D., Bonham, B.H., Leake, P.A., 2007. Frequency map for the human cochlear spiral ganglion: implications for cochlear implants. J. Assoc. Res. Otolaryngol. 8, 220–233.

Verbist, B.M., Frijns, J.H., Geleijns, J., Van Buchem, M.A., 2005. Multisection CT as a valuable tool in the postoperative assessment of cochlear implant patients. Am. J. Neuroradiol. 26, 424–429.

Wanna, G.B., Noble, J.H., McRrackan, T.R., Dawant, B.M., Dietrich, M.S., Watkins, L., Rivas, A., Schuman, T.A., Labadie, R.F., 2011. Assessment of electrode placement and audiologic outcomes in bilateral cochlear implantation. Otol. Neurotol. 32, 428.

Wanna, G.B., Noble, J.H., Carlson, M.L., Gifford, R.H., Dietrich, M.S., Haynes, D.S., Dawant, B.M., Labadie, R.F., 2014. Impact of electrode design and surgical approach on scalar location and cochlear implant outcomes. Laryngoscope 124, S1–S7.

Wang, J., Zhao, Y., Noble, J.H., Dawant, B.M., 2018. Conditional Generative Adversarial Networks for Metal Artifact Reduction in CT Images of the Ear. In: Frangi, A., Schnabel, J., Davatzikos, C., Alberola-López, C., Fichtinger, G. (Eds.). Lecture Notes in Computer Science, Medical Image Computing and Computer Assisted Intervention – MICCAI 2018. MICCAI 2018, vol. 11070. Springer, Cham https://link.springer.com/chapter/10.1007/978-3-030-00928-1_1.

Wilson, B.S., Dorman, M.F., 2008. Cochlear implants: current designs and future possibilities. J. Rehabil. Res. Dev. 45, 695–730.

Zhang, D., Liu, Y., Noble, J.H., Dawant, B.M., 2017a. Localizing landmark sets in head CTs using random forests and a heuristic search algorithm for registration initialization. J. Med. Imaging 4 (4), 044007.

Zhang, D., Zhao, Y., Noble, J.H., Dawant, B.M., 2017b. Selecting electrode configurations for image-guided cochlear implant programming using template matching,. J. Med. Imaging 5 (2), 021202.

Zhao, Y., Dawant, B.M., Labadie, R.F., Noble, J.H., 2014. Automatic localization of cochlear implant electrodes in CT. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2014: 17th International Conference, Boston, MA, USA, September 14–18, 2014, Proceedings, Part I. Springer International Publishing, Cham, pp. 331–338.

Zhao, Y., Dawant, B.M., Noble, J.H., 2015. Automatic electrode configuration selection for image-guided cochlear implant programming. 94150K.

Zhao, Y., Dawant, B.M., Noble, J.H., 2016. Automatic selection of the active electrode set for image-guided cochlear implant programming. J. Med. Imaging 3 (3), 035001.

Zhao, Y., Dawant, B.M., Noble, J.H., 2017. Automatic localization of cochlear implant electrodes in CTs with a limited intensity range. In: Proceedings of SPIE – The International Society for Optical Engineering, 101330T.

Zhao, Y., Labadie, R.F., Dawant, B.M., Noble, J.H., 2018a. Validation of cochlear implant electrode localization techniques. In: Proceedings of SPIE – The International Society for Optical Engineering 10576, 105761U.

Zhao, Y., Dawant, B.M., Labadie, R.F., Noble, J.H., 2018b. Automatic localization of closely spaced cochlear implant electrode arrays in clinical CTs. Med. Phys. 45 (11), 5030–5040. doi:10.1002/mp.13185.

Zhao, Y., Labadie, R.F., Dawant, B.M., Noble, J.H., 2018c. Validation of automatic cochlear implant electrode localization techniques using μCTs. J. Med. Imaging 5 (3), 035001.