Automatic Localization of Cochlear Implant Electrodes in CT

Yiyuan Zhao[1], Benoit M. Dawant[1], Robert F. Labadie[2], and Jack H. Noble[1]

[1] Dept. of Elect. Eng. and Comp. Sci., Vanderbilt University, Nashville, TN, USA
[2] Dept. of Otolaryngology – Head & Neck Surg., Vanderbilt University, Nashville, TN, USA

Abstract. Cochlear Implants (CI) are surgically implanted neural prosthetic devices used to treat severe-to-profound hearing loss. Recent studies have suggested that hearing outcomes with CIs are correlated with the location where individual electrodes in the implanted electrode array are placed, but techniques proposed for determining electrode location have been too coarse and labor intensive to permit detailed analysis on large numbers of datasets. In this paper, we present a fully automatic snake-based method for accurately localizing CI electrodes in clinical post-implantation CTs. Our results show that average electrode localization errors with the method are 0.21 millimeters. These results indicate that our method could be used in future large scale studies to analyze the relationship between electrode position and hearing outcome, which potentially could lead to technological advances that improve hearing outcomes with CIs.

Keywords: cochlear implant, electrode array, snake, segmentation.

1 Introduction

Cochlear Implants (CI) are surgically implanted neural prosthetic devices used to treat severe-to-profound hearing loss. In CI surgery, an electrode array is threaded into the cochlea. After surgery, a processor worn behind the ear sends signals to the implanted electrodes, which activate auditory nerve pathways inducing the sensation of hearing. Although CIs have been remarkably successful, a significant number of CI recipients experience marginal hearing restoration. Recent research has suggested that hearing outcomes with CIs are correlated with the location where the electrodes are placed [1-5]. However, without post-implantation imaging, the position of the electrodes is generally unknown since the array is blindly threaded into a small opening of the cochlea during surgery, with its insertion path guided only by the walls of the spiral-shaped intra-cochlear cavities.

In efforts to analyze the relationship between electrode location and outcome, several groups have proposed coarse electrode position measurements that can be visually assessed in CT images, e.g., whether all electrodes are within one of the two principal intra-cochlear cavities, depth of insertion of the first and last electrode, etc. [1-5]. Studies using these techniques have indicated that placement and outcome are indeed correlated, but it has not been possible to determine specific factors that affect outcome because dataset size was limited and because the electrode positions were never precisely quantified with these techniques. One factor that has limited the size

P. Golland et al. (Eds.): MICCAI 2014, Part I, LNCS 8673, pp. 331–338, 2014.

of the datasets in the studies is the amount of manual effort that must be undertaken to analyze the images. Our group has shown that knowledge of electrode location can be used to select better CI processor settings to significantly improve hearing outcomes compared to standard clinical results [6]. In the current work, we propose a fully automatic approach for localizing CI electrodes in CT images. An electrode localization approach that is automatic and accurate would be significant as it could facilitate precise quantification of electrode position on large numbers of datasets to better analyze the relationship between electrode position and outcome, which may lead to advances in implant design or surgical techniques. It could also automate the electrode localization process in systems designed to determine patient-customized CI settings such as the one proposed in [6], reducing the technical expertise required to use such technologies and facilitating transition to large scale clinical use.

Figure 1 shows an example of an electrode array in a CT slice. Localizing the electrodes in CT images is difficult because (a), as seen in the figure, the beam hardening artifacts caused by the metallic electrodes distort intensities in the region around the electrode array, leading to incorrect assignment of very high intensities during image reconstruction to nearby voxels that are not occupied by metal, thus making it difficult to segment electrodes via thresholding; and (b) the individual electrodes are so close that there is no contrast between them in standard CT images, even when acquired at very fine slice thickness and resolution. Our solution is to identify the centerline of the voxels occupied by the CI electrodes using a snake-based localization approach [7] and then to fit a 3D model of the electrode array to the extracted centerline. This is a similar approach to that which we proposed in [8]. However, the technique we presented in that paper leads to inaccurate results around the first and last electrodes due to curve shrinkage. This shrinking phenomenon is caused by the use of an intensity-based attraction function since the image intensity decreases mildly at the array endpoints relative to the rest of the array. Further, we found that the "forward energy," an external energy term designed to counteract endpoint shrinking errors by expanding the curve, became unstable and led to failures when applying the technique on clinical image datasets. As will be described in the following section, in this work, we propose a new technique to counteract the shrinking effect by localizing and fixing the endpoints prior to snake optimization. Our results, presented and discussed in Sections 3 and 4, will show that this fully automatic approach can reliably be applied to clinical images.

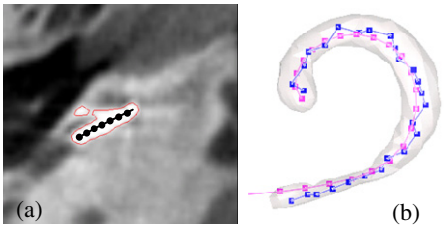

Fig. 1. Panel (a) shows a portion of an electrode array in an axial slice of a CT. Black dots indicate locations of individual electrodes. An isocontour around high intensity voxels is shown in red. Panel (b) shows a 3D isosurface of an electrode array with a manually determined centerline in purple. The blue curve is the coarse approximation to the centerline determined using our automatic initialization process discussed in Section 2.2.

2 Methods

The automatic segmentation method we propose is outlined in Figure 2. As can be seen in the figure, the first step (1) involves coarsely estimating the location of the region of interest (ROI), which is a local region ~1 cm^3 around the cochlea. This is done through registration with a known volume. The subsequent processing steps are then performed solely within the ROI. The next step (2) is to initialize our electrode array centerline localization. This is done by segmenting via thresholding the region of the image that contains the metallic electrodes and then computing the initialized centerline as the medial axis of the result. The thresholding step will produce a segmentation that includes electrode voxels as well as those that appear bright due to partial volume or beam hardening artifacts, but the medial axis extraction step is able to reliably and coarsely approximate the centerline of the electrode array. After initialization, the next steps (3-4) are to refine the centerline using a snake-based optimization approach [7]. In the third step, the curve endpoints are first localized within the neighborhood of their initialized positions using an endpoint detection filter we have designed. In the fourth step, the endpoints are fixed and the points in the rest of the curve are optimized. This is done using a snake with its external energy defined using the output of a vesselness filter that is applied to the original image to enhance the centerline of the electrode array [9]. By detecting and fixing the endpoints prior to snake optimization, curve shrinking effects discussed in the previous section are eliminated. The final step (5) is a straightforward resampling of the extracted centerline to determine individual electrode locations using *a priori* knowledge about the distance between neighboring electrodes. The following subsections detail this approach.

2.1 Data

The images in our dataset include images from 15 subjects acquired with a Xoran xCAT®. The images have voxel size 0.4 x 0.4 x 0.4 mm^3. As a pre-processing step, an ROI bounding the region around the electrode array in each target image is automatically localized by using a mutual information-based affine registration computed between the target image and a known reference image [10]. The ROI is then automatically cropped from the original target image and all subsequent steps are performed on the cropped image. Each cropped image includes approximately 30 × 30 × 30 mm^3. Each subject in this study was implanted with a Cochlear™ Contour Advance®. Thus, the methods presented are focused on segmenting this type of electrode array but could prove in future studies to be applicable to other implant models.

2.2 Centerline Initialization

The centerline is initialized by thresholding the region of the image that includes the electrode array and computing the medial axis of the result. We determine the threshold dynamically using a maximum likelihood estimation-based (MLE) threshold selection approach [11] since the best threshold can vary across subjects due to the relatively low signal-to-noise ratio (SNR) achieved using the low-dose acquisition

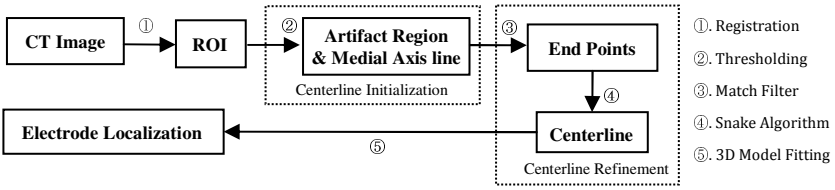

Fig. 2. Flow chart of the electrode array centerline localization process

protocols on a flat panel scanner. We would also expect that a dynamic threshold would account for differences between scanners, but this was not tested in this study. The MLE approach we have designed is to fit a model, defined as the sum of two Gaussian distributions, to the ROI image histogram and compute a threshold based on this result. One distribution $G(\mu_1, \sigma_1)$ corresponds to soft tissue and another $G(\mu_2, \sigma_2)$ corresponds to bony tissue. While air and metal are present in the ROI image, their relatively small volumes contribute little to the shape of the histogram, and thus these intensity classes are ignored in the histogram fitting. Once the distributions are estimated, the threshold is selected based on the upper tail of the Gaussian that models the intensity distribution of bone to be $\mu_2 + 5\sigma_2$, which was empirically determined to lead to good results. We chose to use this MLE-based approach, rather than a simpler percentile-based approach, because this approach is not sensitive to differences in ROI volume or differences in volume of metal present in the ROI, which can vary across subjects. After a threshold is determined, the medial axis of the resulting thresholded volume is computed using the medial axis extraction techniques presented by Bouix et al. [12]. The resulting curve provides a close but coarse approximation to the centerline of the electrode array. An example result of this process is shown in blue in Figure 1b.

2.3 Centerline Refinement

After the curve is initialized, we refine its position using a snake-based algorithm. The traditional snake algorithm localizes a contour by minimizing the energy equation

$$E = \int_0^1 \rho_1 \|x'(s)\|^2 + \rho_2 \|x''(s)\|^2 + E_{ext}\big(x(s)\big)\, ds, \tag{1}$$

where $x(s)$ is the position of the parameterized curve at s, ρ_1 and ρ_2 are the tension and rigidity weighting terms, and E_{ext} is the external energy term. In our experiments, we set $\rho_1 = 0.03$ and $\rho_2 = 0.08$ as these values were empirically determined to lead to good results, and we define E_{ext} to be the output of a vesselness response filter applied to the ROI image [9]. We apply the filter at scales $\sigma = \{0.08, 0.16, \ldots, 0.8\}$ mm and set the other internal parameters to be $\alpha = 0.5$, $\beta = 0.5$, and $\gamma = 500$. Vesselness response, rather than, for example, a direct function of image intensity is used as an external energy because the high intensity voxels in the region around the electrode array can be noisy, and voxels with intensity that is locally maximal often do not fall on the centerline of the homogeneous bright region in the image (see Figure 1). Since the electrode array has the appearance of a tubular structure, a vesselness response filter is a natural choice to enhance the centerline of the electrode array.

The robustness of the vesselness filter in detecting the centerline of the electrode array is high along the length of the array but diminishes at the endpoints. Thus, with no additional information, optimizing the snake would result in a shrinking of the curve at the endpoints. To address this, we determine the endpoint positions using an endpoint detection filter and fix them during the snake optimization. The endpoint detection filter we have constructed, $M_{\hat{v}}(\boldsymbol{\omega})$, is a match filter. For the sake of simplicity, we define $M_{\hat{v}}(\boldsymbol{\omega})$ such that $\boldsymbol{\omega} = \mathbf{0}$ lies at the center of the filter (see Figure 3a). We also orient the filter using $\hat{v}$, which represents the orientation of the centerline of the electrode array at the endpoint. To define $M_{\hat{v}}(\boldsymbol{\omega})$, we first define $M'_{\hat{v}}(\boldsymbol{\omega})$ as

$$M'_{\hat{v}}(\boldsymbol{\omega}) = \begin{cases} r^2 - \|\boldsymbol{\omega}\|^2 & \boldsymbol{\omega} \cdot \hat{v} \geq 0 \\ r^2 - \|\boldsymbol{\omega} - (\boldsymbol{\omega} \cdot \hat{v})\hat{v}\|^2 & \boldsymbol{\omega} \cdot \hat{v} < 0 \end{cases}. \tag{2}$$

This equation defines $M'_{\hat{v}}(\boldsymbol{\omega})$ such that when $\boldsymbol{\omega} \cdot \hat{v} \geq 0$, i.e., in the $\hat{v}$ direction from the origin as seen in Figure 3a, $M'_{\hat{v}}(\boldsymbol{\omega})$ matches a semispherical structure, whereas in the opposite direction where $\boldsymbol{\omega} \cdot \hat{v} < 0$, the filter matches a tubular structure. The radius, r, of the sphere and tube are set to be 0.3 mm, which is approximately the radius of the electrode arrays in our images. The final form of the filter is defined as $M_{\hat{v}}(\boldsymbol{\omega}) = M'_{\hat{v}}(\boldsymbol{\omega})\left(\rho_3 H\left(M'_{\hat{v}}(\boldsymbol{\omega})\right) + (1-\rho_3)H\left(-M'_{\hat{v}}(\boldsymbol{\omega})\right)\right)$, where $H(\cdot)$ is the Heaviside function and $\rho_3 = 0.97$ is a parameter we chose empirically to optimize results and tunes the weighting between the fore- and background regions of the filter.

To find each endpoint using this filter, we set $\hat{v}$ to be the orientation of the central axis of the electrode array as estimated by the vesselness response at $\boldsymbol{x}_e^i$, the location that the endpoint was initialized using the methods described in Section 2.2, and then compute the endpoint location $\boldsymbol{x}_e$ as

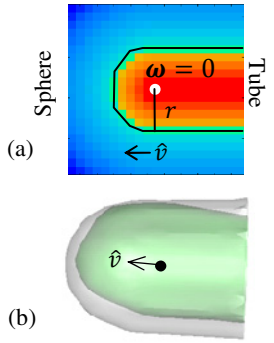

Fig. 3. (a) shows a slice of $M'_{\hat{v}}(\boldsymbol{\omega})$ with $M'_{\hat{v}}(\boldsymbol{\omega}) = 0$ isocontour in black and $\boldsymbol{\omega} = \mathbf{0}$ shown as white dot. (b) shows the 3D isosurface of $M'_{\hat{v}}(\boldsymbol{\omega})$ (white) aligned with the tip of an electrode array (green).

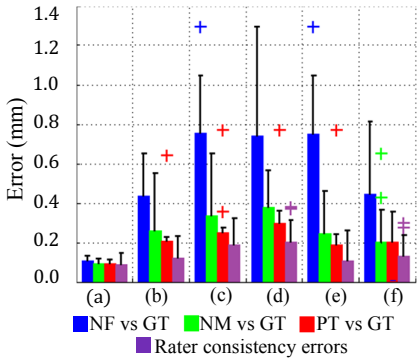

Fig. 4. Barplots of mean (a) and max (c) curve distances; mean (b) and max (d) electrode distances; and tip (e) and base (f) endpoint distances.

$$\boldsymbol{x}_e = \text{argmax}_{\boldsymbol{x}\in N(\boldsymbol{x}_e^i)} \sum_{\boldsymbol{y}\in L(\boldsymbol{x})} I(\boldsymbol{y})\, M_{\hat{v}}(\boldsymbol{y}-\boldsymbol{x}). \quad (3)$$

$N(\boldsymbol{x}_e^i)$ is a neighborhood function that we define as the set of 16 x 16 x 16 points uniformly sampled in a 1.2 x 1.2 x 1.2 mm^3 box surrounding $\boldsymbol{x}_e^i$, I is the ROI image, and $L(\boldsymbol{x})$ is a neighborhood function defined as the set of 21 x 21 x 21 points uniformly sampled in a 1.2 x 1.2 x 1.2 mm^3 box oriented in the $\hat{v}$ direction surrounding $\boldsymbol{x}$. In summary, Eqn. (3) selects the endpoint as the point in a local region around the initial endpoint that maximizes the response of the endpoint enhancement filter, and the filter response should be maximized when it is aligned with and centered on the tip of the electrode array.

After the endpoints are determined, they are fixed and the positions of the remaining points in our curve are optimized by iterating the standard snake update equations [7] until convergence or until reaching 100 iterations. Once the final curve is localized it is straightforward to resample the curve to identify the location of individual electrodes based on *a priori* knowledge of the distance between electrodes in the array.

2.4 Validation

We quantified the accuracy of our automatic electrode array extraction technique in a dataset of fifteen head CT images by comparing centerlines computed automatically using the proposed technique (PT) to ground truth (GT) curves, which were created by averaging of three sets of curves independently defined by an expert. Metrics used to characterize distance between two curves include mean and max curve distance (mean and max of the distances computed from each point on curve 1 to the closest point on curve 2 and vice versa), mean and max electrode distance (distance between each electrode location in curve 1 to the corresponding electrode in curve 2 after determining electrode locations along the curves as described in 2.3), and distance between corresponding endpoints in curves 1 and 2. To show the benefit our matched filter provides, we also report quantitative errors that result from computing the curve when (a) endpoints are fixed at their initialization position without the matched filter update (NM) and (b) when the endpoints are not fixed but optimized with the snake method similarly to the rest of curve (NF).

To assess whether the PT produces acceptable results, we conducted a second study in which an expert was asked to select between the GT and PT endpoints, blind to their identity. We focused on the endpoints because, as our results will show, this is the area in which there are the largest discrepancies between GT and PT curves.

3 Results

The quantitative comparisons between the GT and PT centerlines for all the datasets are shown in Figure 4 in red, and Figure 5 shows visualizations of two cases. In Figure 4, for each barplot, the height of the bars, crosses, and black whiskers denote the mean, outlier data, and maximum non-outlier value. Data are considered outliers if they fall above $q_3 + 1.5(q_3 - q_1)$, where q_3 and q_1 are the 25$^{\text{th}}$ and 75$^{\text{th}}$ percentiles

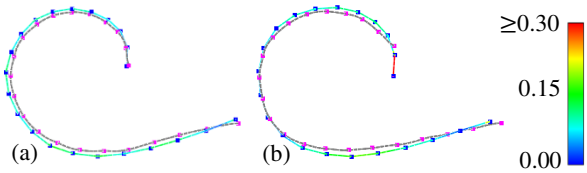

Fig. 5. 3D renderings of GT (colormapped with curve distance in mm) and PT (shown in transparent black) curves for our best (a) and worst (b) case errors. Points indicate electrode locations along curves determined by distance priors

of the dataset. As can be seen in the figure, our proposed method results in mean curve errors of 0.09 mm (0.13 of a voxel diagonal) and average maximum curve errors of 0.25 mm (0.36 of a voxel diagonal) with an overall maximum of 0.80 mm. Our method extracts a much more accurate centerline compared to prior work in which we achieved mean curve errors of 0.2 millimeters [8]. Further, the mean electrode localization error with our currently proposed method is only 0.21 mm. The utility of fixing the endpoints and optimizing them with our matched filter is also apparent in Figure 4 as NF and NM lead to much larger electrode and endpoint localization errors. This difference is not as pronounced in mean curve errors since curve distances along the length of the curve are not sensitive to errors at the endpoints. The mean tip and base endpoint errors with PT are 0.19 mm and 0.2 mm. These quantities are slightly higher for NM and substantially higher for NF. The outlier values for PT that fall above 0.6 mm all correspond to the case shown in Figure 5b, where the tip of the array was localized incorrectly due to lower than normal SNR in the image. We also show in purple in Figure 4 rater consistency errors computed among the three sets of curves manually delineated by an expert. We find mean and overall maximum consistency curve errors of 0.09 and 0.35 mm, suggesting that except for the outlier case, errors in our PT are close to the level of rater repeatability.

In the expert endpoint selection test, among the 30 endpoints in the 15 cases, 8 PT endpoints were judged to be equally accurate to GT, and 29 of 30 PT endpoints were judged to be acceptable. The lone exception was the tip endpoint shown in Figure 5b.

4 Conclusions

In this work, we have designed an automatic cochlear implant electrode array centerline extraction method. Our experiments show that our method is highly accurate, even when applied to clinical images. Compared to our prior method reported in [8], the method we propose here achieves results with errors that are half as large on average. This improvement is due in large part to the use of our matched filter, which leads to better endpoint localization. Our approach requires approximately 3 minutes of computation time on a standard PC.

Our method did result in unacceptably large errors for one of fifteen images. Future studies will involve developing techniques to detect and handle such errors. Additionally, we plan to test our method with images acquired with different scanners and of subjects with different implant models. We also plan to apply our method to large

numbers of datasets to facilitate studying how the location of individual electrodes correlates with outcomes with the goal of developing technologies that can improve hearing outcomes with CIs.

Acknowledgements. This research has been supported by NIH grants R01DC014037, R01DC008408, and R21DC012620. The content is solely the responsibility of the authors and does not necessarily represent the official views of this institute.

References

1. Verbist, B.M., Frijns, J.H.M., Geleijns, J., van Buchem, M.A.: Multisection CT as a Valuable Tool in the Postoperative Assessment of Cochlear Implant Patients. Am. J. Neuroradiol. 26, 424–429 (2005)
2. Aschendorff, A., Kubalek, R., Turowski, B., Zanella, F., Hochmuth, A., Schumacher, M., Klenzner, T., Laszig, R.: Quality control after cochlear implant surgery by means of rotational tomography. Otol Neurotol 26, 34–37 (2005)
3. Skinner, M.W., Holden, T.A., Whiting, B.R., Voie, A.H., Brundsen, B., Neely, G.J., Saxon, E.A., Hullar, T.E., Finley, C.C.: In vivo estimates of the position of advanced bionics electrode arrays in the human cochlea. Annals of Otology, Rhinology and Laryngology Supplement 197, 2–24 (2007)
4. Wanna, G.B., Noble, J.H., McRacken, T.R., Dawant, B.M., Dietrich, M.S., Watkins, L.D., Schuman, T.A., Labadie, R.F.: Assessment of electrode placement and audiologic outcomes in bilateral cochlear implantation. Otol Neurotol 32, 428–432 (2011)
5. Wanna, G.B., Noble, J.H., Carlson, M.L., Gifford, R.H., Dietrich, M.S., Haynes, D.S., Dawant, B.M., Labadie, R.F.: Impact of electrode design and surgical approach on scalar location and cochlear implant outcomes. Laryngoscope (in press, 2014)
6. Noble, J.H., Labadie, R.F., Gifford, R.H., Dawant, B.M.: Image-guidance enables new methods for customizing cochlear implant stimulation strategies. IEEE Trans. on Neural Systems and Rehabilitiation Engineering 21(5), 820–829 (2013)
7. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active Contour Models. Int'l Jour. of Computer Vision, 321–331 (1988)
8. Noble, J.H., Schuman, T.A., Wright, C.G., Labadie, R.F., Dawant, B.M.: Automatic Identifi-cation of Cochlear Implant Electrode Arrays for Post-Operative Assessment. In: Dawant, B.M., Haynor, D.R. (eds.) Medical Imaging 2011: Image Processing, Proc. of the SPIE Conf. on Med. Imag., vol. 7962, p. 796217 (2011)
9. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: Wells, W.M., Colchester, A., Delp, S. (eds.) MICCAI 1998. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998)
10. Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P.: Multimodality image registration by maximization of mutual information. IEEE Trans. Med. Imag. 16, 187–198 (1997)
11. Duda, R.O., Hart, P.E.: Pattern Classification and Scene Analysis. A Wiley-Interscience Publication, pp. 192–202. Wiley, New York (1973)
12. Bouix, S., Siddiqi, K., Tannenbaum, A.: Flux driven automatic centerline extraction. Medical Image Analysis 9, 209–221 (2005)