Automatic selection of the active electrode set for image-guided cochlear implant programming

Yiyuan Zhao
Benoit M. Dawant
Jack H. Noble

Automatic selection of the active electrode set for image-guided cochlear implant programming

Yiyuan Zhao,* Benoit M. Dawant, and Jack H. Noble

Vanderbilt University, Department of Electrical Engineering and Computer Science, 2301 Vanderbilt Place, Nashville, Tennessee 37235, United States

Abstract. Cochlear implants (CIs) are neural prostheses that restore hearing by stimulating auditory nerve pathways within the cochlea using an implanted electrode array. Research has shown when multiple electrodes stimulate the same nerve pathways, competing stimulation occurs and hearing outcomes decline. Recent clinical studies have indicated that hearing outcomes can be significantly improved by using an image-guided active electrode set selection technique we have designed, in which electrodes that cause competing stimulation are identified and deactivated. In tests done to date, an expert is needed to perform the electrode selection step with the assistance of a method to visualize the spatial relationship between electrodes and neural sites determined using image analysis techniques. We propose to automate the electrode selection step by optimizing a cost function that captures the heuristics used by the expert. Further, we propose an approach to estimate the values of parameters used in the cost function using an existing database of expert electrode selections. We test this method with different electrode array models from three manufacturers. Our automatic approach generates acceptable active electrode sets in 98.3% of the subjects tested. This approach represents a crucial step toward clinical translation of our image-guided CI programming system. © 2016 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JMI.3.3.035001](https://doi.org/10.1117/1.JMI.3.3.035001)]

Keywords: cochlear implant; stimulation strategy; image-guided cochlear implant programming.

Paper 16080PR received May 17, 2016; accepted for publication Aug. 30, 2016; published online Sep. 22, 2016.

1 Introduction

Over the last 20 years, cochlear implants (CIs) have become the most successful neural prosthesis and are used to treat severe-to-profound hearing loss.¹ In CI surgery, an array of electrodes is blindly threaded into the cochlea. After the surgery, the processor worn behind the ear sends signals to the implanted electrodes, which stimulate the auditory nerve pathways within the cochlea. After implantation, the CI is programmed by an audiologist. The CI programming begins with the selection of a general signal processing strategy, e.g., continuous interleaved sampling.² Then, the audiologist defines the “MAP,” i.e., the CI processor instructions that determine what signals are sent to the implanted electrodes in response to incoming sounds. The MAP is determined by selecting the electrode configuration, i.e., the active electrode set, by specifying stimulation levels for each active electrode based on measures of the user’s perceived loudness, and by selecting a frequency allocation table that specifies which electrodes will be activated when specific sound frequencies are detected. Electrode activation stimulates the spiral ganglion (SG) nerves, the nerve pathways that branch to the cochlea from the auditory nerve. In natural hearing, an SG nerve is activated when the characteristic frequency associated with that pathway is present in the incoming sound. The SG nerves, which are located within the modiolus of the cochlea, are tonotopically ordered by decreasing characteristic frequency along the length of the cochlea, and this precisely tuned spatial organization is well known [see Fig. 1(a)].^{3,4} The modiolar surface shown in Fig. 1(a) represents the interface between the

intracochlear cavities, where the electrodes are placed, and the modiolus, where the SG nerves that are stimulated by the electrodes are located. Recent research has suggested that hearing outcomes with CIs are correlated with the location at which the electrodes are placed in the cochlea.^{5–10} In surgery, the array is blindly threaded into the cochlea with its insertion path guided only by the walls of the spiral-shaped intracochlear cavities. The final position of the electrodes is not generally known in the traditional clinical workflow. However, we have developed techniques that accurately enable locating the electrodes using computed tomography images.^{11–13}

Recent research by our group^{11,14} has shown that the spatial relationship between the neural pathways and the electrodes can be used to estimate electrode interactions at the neural level, i.e., cross-electrode neural stimulation overlap [see Fig. 1(b)], which is a phenomenon known to negatively affect hearing outcomes.^{15,16} We have shown, in a large clinical study, that when stimulation overlap is detected and the configuration of active electrodes is adjusted to reduce that overlap, hearing outcomes are improved, and those improvements are statistically significant.¹⁷ Our goal now is to fully automate our system so that clinical translation of this technology is feasible.

One step that has not yet been automated is the electrode configuration selection step. Thus far, electrode configurations have been manually selected based on the electrode distance-versus-frequency curves (DVF’s). The DVF is a technique developed by our group to facilitate the visualization of electrode interaction in individual patients.¹¹ It is a two-dimensional plot that captures important information about the patient-specific spatial

*Address all correspondence to: Yiyuan Zhao, E-mail: yiyuan.zhao@vanderbilt.edu

Fig. 1 Visualization of CI electrode activation patterns. In (a), the scala tympani (an intracochlear cavity) is shown with the modiolar surface, which represents the interface between the nerves of the SG and the intracochlear cavities and is color-coded with the tonotopic place frequencies of the SG in Hz. In (b), synthetic examples of stimulation patterns on the modiolar interface created by the implanted electrodes are shown in multiple colors to illustrate the concept of stimulation overlap.

relationship between the electrodes and the SG nerves, such as is shown in three-dimensional in Fig. 1(b). Figure 2(a) is an example of DVFs for a 7 electrode array. The horizontal axis represents position along the length of the modiolus in terms of the characteristic frequencies of adjacent SG nerves. Each DVF is labeled with a number representing its electrode number. The height of each DVF on the vertical axis represents the distance from the corresponding electrode to the frequency mapped modiolar surface. Thus, a DVF is constructed for a given electrode by finding the distance to that electrode from nearby, frequency-mapped sites on the modiolus. From this visualization technique, we can see that electrode 3 is ~ 1 mm from the modiolus, and the characteristic frequencies of the SG nerves closest to electrode 3 are around 1 kHz. Our current electrode configuration selection method is based on the assumption that if an electrode's DVF is not the closest DVF in the region around its minimum, it is likely that its stimulation region overlaps with other electrodes, and thus, it is negatively affecting hearing performance. As shown in Fig. 2, we can see that since the minimum of the DVF for electrode 4 is entirely above the DVF for electrode 3, it is likely that electrode 4 is stimulating the same neural region as electrode 3. Also, while the minimum of the DVF for electrode 6 falls below the other curves, its depth of concavity relative to the minimum envelope of the other neighboring DVFs is small, so it is likely that electrode 6 has an overlapping stimulation region with electrodes 5 and 7. Our active electrode set selection approach is to keep active the largest subset of electrodes that are not likely to cause stimulation overlap. Thus, in the example, we would remove electrodes 4 and 6 from the active electrode set. The DVFs of the updated electrode configuration are shown in Fig. 2(b).

As discussed above, we have shown in clinical studies that our manual approach for selecting active electrode set results in

significant improvement in hearing performance. While selecting the electrode set manually can usually be done relatively quickly (0.5 to 2 min), it requires specialized expertise, and training new individuals to become experts is time consuming. To develop an automated system that implements our approach and can be widely deployed for clinical use, we need an automated method that performs as well as an expert on average for selecting the electrode configuration. To solve this problem, we have developed a DVF-based cost function and proposed to optimize it using an exhaustive search method. The rest of this paper presents our approach.

2 Methods

Our dataset consists of DVFs and expert-defined optimal and acceptable electrode configurations for 96 cases. We divided the dataset into a training and a testing dataset. The training dataset contains 12 subjects implanted with arrays manufactured by Med-El (MD) (Innsbruck, Austria), 10 subjects implanted with arrays manufactured by Advanced Bionics (AB) (Valencia, California), and 14 subjects implanted with arrays manufactured by Cochlear (CO) (New South Wales, Australia). Our testing dataset contains 20 subjects of arrays manufactured by MD, 20 subjects of arrays manufactured by AB, and 20 subjects of arrays manufactured by CO. In our training dataset, we have 18 male and 18 female subjects. Subject age ranges from 18 to 84 with a mean age of 57.9 and standard deviation of 14.69 years. In our testing dataset, we have 28 male and 32 female subjects. The age range is 21 to 84 with a mean age of 58.1 and a standard deviation of 14.6 years.

Our approach is to develop a cost function that assigns a cost for a given electrode configuration, i.e., a particular set of "on" and "off" electrodes, for a subject based on the electrode DVFs. We then can use an exhaustive search method in which all

Fig. 2 Visualization of DVFs: (a) an example of a combination of the DVFs formed by 7 electrodes. Each single curve represents the distance from the corresponding electrode to the frequency mapped sites along the length of the modiolus. (b) the DVFs after electrode configuration adjustment.

possible configurations are generated, compute the cost for each configuration, and select the one with the minimum cost. In this work, we have chosen to design the cost function to be a linear combination of a set of DVF-based features that capture the heuristics we use for manually producing electrode configuration plans. The features aim at reducing the cross-electrode neural stimulation overlap as described in Sec. 1. We have defined a total of $N = 10$ feature cost terms. The weighted sum of the N feature cost terms is determined as the final cost value. The weights $\{w_i\}_{i=1}^N$ for the N feature cost terms are determined through a training process using the subjects in the training dataset. Each of the three electrode arrays type has a different number of electrodes (MD has 12, AB has 16, and CO has 22 electrodes) and a different geometry. Thus, they create different DVF patterns, which lead us to estimate the set of weights separately for each electrode type. After generating the estimates of the weights $\{w_i\}_{i=1}^N$, we apply the weights to the testing dataset for validation.

The feature cost terms $\{f_i\}_{i=1}^N$ are defined as follows. First

$$f_1 = \begin{cases} 0 & \text{If the most apical electrode} \in \text{active set} \\ 1 & \text{If the most apical electrode} \notin \text{active set} \end{cases}, \quad (1)$$

which assigns a zero cost to configurations whose most apical electrode, i.e., the deepest electrode in the cochlea [see Fig. 1(b)], is activated and a nonzero cost otherwise. f_1 is included because deactivating the most apical electrode, which stimulates nerves with lower characteristic frequencies, can result in an up-shift in perceived sound frequency. This affects hearing quality and is usually not preferred. Next

$$f_2 = \frac{1}{K_a}, \quad (2)$$

where K_a is the number of electrodes that are active in the configuration. While other terms below are designed to deactivate electrodes to increase channel independence, f_2 captures the heuristic that keeping more electrodes active is desirable because it results in less frequency compression and better

outcomes if those electrodes provide independent stimulation. Next

$$f_3 = \left(\sum_{i=1}^K e^{-\text{AreaTerm}_i} \right) / K, \quad (3)$$

where K is the total number of electrodes, and

$$\text{AreaTerm}_i = \begin{cases} \frac{T(D_i)}{\text{Area}_i} & \text{if electrode } i \in \text{active set} \\ \frac{\text{Area}_i}{T(D_i)} & \text{if electrode } i \notin \text{active set} \end{cases}, \quad (4)$$

where Area_i is a term that captures the channel independence of electrode i by measuring the area above the DVF for electrode i and below the envelope of the other DVF curves, and $T(D_i)$ is a term that defines the value of Area_i at which activating or deactivating electrode i is equally desirable as a function of the distance D_i between electrode i and modiolus. Equation (3) is designed to assign a lower cost for activating (deactivating) electrodes with DVFs, whose Area_i is larger (smaller) than the threshold value $T(D_i)$. Figure 3(a) shows qualitatively the term Area_i for several DVF curves. In this example, $\text{Area}_2 > \text{Area}_3$, $\text{Area}_2 > T(D_2)$, and $\text{Area}_3 < T(D_3)$, which leads to electrode 2 having a small cost for being active and 3 having a large cost for being active. This will favor configurations with electrode 2 being activated and electrode 3 being deactivated. Optimal electrode configurations will thus tend to consist of electrodes with DVF curves that have larger Area_i values. To compute Area_i , we sum the squared distances measured between the DVF for the i 'th electrode and the envelope of the other DVFs at discrete positions sampled along the frequency axis. We found empirically that defining Area_i as the sum of the squared distances between the curves is better than a sum of direct distances for describing expert-perceived channel overlap because the sum of squared distances is larger for DVFs that have at least some regions that lie relatively far below the envelope of the other DVFs. $T(\cdot)$ is a function that is defined using a subset of electrodes in our training dataset as follows. Figure 4(a) shows a scatter plot of electrodes-of-interest

Fig. 3 Visualization of three DVF-based features.

Fig. 4 The visualization of (a) area-distance and (b) depth of concavity-distance relationship and the empirical separation line for electrodes in the training dataset.

(EOIs), which are a subset of electrodes from our training dataset for which the expert identified that the decision to keep them active or not was driven by the DVF area. Area_i is shown on the y-axis and the electrode distance to the modiolus, and D_i , is shown on the x-axis. As observed in the plot, the activation decision is a function of both D_i and Area_i . This is because when D_i is larger, the electrode is farther from the modiolus, and we expect wider spread of excitation. Thus, we would require a greater Area_i to obtain adequate channel independence and to want to keep the electrode active. Thus, we define $T(\cdot)$ as a polynomial function of modiolar distance that best separates the active and inactive EOIs from the training dataset in this plot in a least-squares sense

$$T(D_i) = -0.2660 + 1.4125D_i + 0.5398D_i^2. \quad (5)$$

T_i is shown as the green curve in Fig. 4(a). The coefficients and the order of the polynomial function are determined with our training dataset. First, we randomly separate the EOIs into 90 training EOIs and 20 validation EOIs. Next, we investigated first-, second-, and third-order polynomials as candidate functions. The coefficients of each polynomial are chosen so that the polynomial best separates the active and inactive training EOIs in a least-squares sense. Next, we evaluated each of the three candidate polynomials with the validation EOIs. We found that the second-order polynomial correctly classified the largest number of testing EOIs. Thus, we chose to use the second-order polynomial as $T(\cdot)$ and this is shown as the green curve in Fig. 4(a). Next

$$f_4 = \sum_{i=1}^K e^{-\text{DepthTerm}_i} / K, \quad (6)$$

where K is the total number of electrodes

$$\text{DepthTerm}_i = \begin{cases} \frac{\text{Depth}_i}{R(D_i)} & \text{if electrode } i \in \text{active set} \\ \frac{R(D_i)}{\text{Depth}_i} & \text{if electrode } i \notin \text{active set} \end{cases}, \quad (7)$$

$$\text{Depth}_i = \min(C_{iL}, C_{iR}), \quad (8)$$

where C_{iL} and C_{iR} are the depth of concavity of the i 'th electrode DVF relative to its left and right neighbors, Depth_i is the overall depth of concavity for the curve, and $R(D_i)$ is the value

of Depth_i for which activating and deactivating the electrode are equally desirable as a function of the distance D_i between electrode i and modiolus. Equation (7) is designed to assign a lower cost for activating (deactivating) electrodes with DVFs whose depth of concavity Depth_i is larger (smaller) than the threshold value $R(D_i)$. This term captures the property that optimal configurations consist of electrodes whose DVFs have large depth of concavity. Figure 3(b) shows an example of the depth of concavity measurement. In this example, $\text{Depth}_2 = c_{2R} < R(D_2)$ and $\text{Depth}_3 = c_{3L} > R(D_3)$, which leads to a large cost for activating electrode 2 and a small cost for activating electrode 3. This will favor solutions in which electrode 3 is activated and electrode 2 is deactivated. $R(\cdot)$ is a polynomial function that is defined using a subset of electrodes selected from our training dataset in a manner identical to T as

$$R(D_i) = 0.0328 + 0.005D_i + 0.0351D_i^2. \quad (9)$$

Figure 4(b) shows a scatter plot of EOIs in our training dataset for which the expert decision to keep them active or not was driven by the depth of concavity. Depth_i is shown on the y-axis, the electrode distance to the modiolus D_i is shown on the x-axis, and R is shown in green. As observed in the plot, the activation decision is a function of both electrode distance D_i and Depth_i . This is because, similarly to Eq. (4) above, when D_i is larger and we expect wider spread of excitation, we would require larger Depth_i to indicate adequate channel independence and to keep the electrode active. Next

$$f_5 = \sum_{i=1}^K (D_i - M_i)u(D_i - M_i), \quad (10)$$

where D_i is the distance from the electrode i to the modiolus and M_i is a linear function defined as

$$M_i = -17.29 \log_{10}(\text{Freq}_i) + 72.81, \quad (11)$$

where Freq_i is the place frequency of the nerves closest to electrode i and $u(\cdot)$ is the unit step function. f_5 is designed to assign a cost to electrodes that fall above the line defined by Eq. (11). This line is shown in Fig. 3(c), which shows a small, zoomed in portion of the plot shown in Fig. 3(b). Since the line is steep, electrodes located above it are located in the very high-frequency region (>13 kHz) near the entrance of the cochlea.

These electrodes are often deactivated clinically because they are outside or nearly outside the cochlea or provide abnormal perception due to loss of neural survival that is common in this region. Thus, f_5 is used to indicate that electrodes in this region are less desirable. As shown in Fig. 3(c), Eq. (12) was designed by finding the least-squares fit line that separates the groups of electrodes in the training electrode configurations that were set as activated (red) and deactivated (blue). Also shown are distances D_i and M_i for one electrode (magenta). Next

$$f_6 = K_I/K_a, \quad (12)$$

where K_a is the number of active electrodes in the configuration and K_I is the number of DVFs that have a minimum that falls above the envelope of other electrodes' DVFs [see electrode 4 in Fig. 2(a)]. When this term is larger than 0, it is a strong indicator that one or more electrodes are stimulating the same frequency range as other electrodes but less effectively since they are located further away from the modiolar surface. Next

$$f_7 = \left(\sum_{i=1}^{K_s} e^{-\text{ratio}_{S(i)}} \right) / K_s, \quad (13)$$

where S is the set of K_s active electrodes that have active neighbors on both the left and right side

$$\text{ratio}_i = \min(A_{iL}, A_{iR}) / \max(A_{iL}, A_{iR}), \quad (14)$$

where A_{iL} and A_{iR} indicate the left and right half area terms of the DVF curve of electrode i [see Fig. 4(a)]. A_{iL} and A_{iR} are defined as the sum of the distances measured between the DVF curve for the i 'th electrode and the envelope of the other DVFs at the discrete positions sampled along the frequency axis to the left and right of the minimum, respectively. Equation (14) assigns a low cost to the configurations with symmetric DVFs, and a high cost to the configurations with one or more highly asymmetric DVFs. Finally, $f_8 = \sqrt{f_3}$, $f_9 = \sqrt{f_4}$, and $f_{10} = \sqrt{f_7}$. These terms were included after testing all combinations of squares and square roots of f_{1-7} and finding that including these terms led to better results.

A linear combination of the cost terms is used to define an overall cost function for a given configuration, i.e.,

$$C = \sum_{i=1}^N w_i f_i, \quad (15)$$

Because current electrode arrays have ~ 22 electrodes, it is practical to find the globally optimal configuration through

an exhaustive search that evaluates all possible configurations. The values for the set of scalar coefficients $\{w_i\}_{i=1}^N$ used to weigh each of the cost terms in Eq. (16) are estimated using a training set of existing manually selected electrode configurations and a least-squares technique.

Our methods are summarized in Fig. 5. As can be seen in the figure, there is a training stage and a testing stage. The training stage is used to determine the parameter (weight) values $\{w_i\}_{i=1}^N$ that control the contribution of each feature term in the overall cost function. Input 1 is the DVF-based feature set. Using this feature set, a cost term is computed for each feature for all the possible electrode configurations in the set of training cases. The resulting cost terms (output 2) are passed to the least-squares solver, which solves equations of the form

$$\left\{ \sum_{i=1}^N f_i^{m,o} w_i + \delta = C^{m,o} \right\}_{m=1, o=1}^{M,O}, \quad (16)$$

where $\{f_i^{m,o}\}$ is the set of N cost terms for each of the M electrode configurations for the O subjects in our training dataset, $\{C^{m,o}\}$ is the set of cost estimates for each configuration, and δ is a constant. We compute $\{C^{m,o}\}$ using a piecewise function defined as

$$C^{m,o} = \begin{cases} 0 & e_{m,o} = e_{\text{opt},o} \\ \frac{1}{2} & e_{m,o} \in \{e_{\text{acc},o}\} \\ \text{dist}(e_{m,o}, e_{\text{opt},o}) & \text{otherwise} \end{cases}, \quad (17)$$

where $e_{\text{opt},o}$ is the electrode configuration chosen manually by an expert for the o 'th subject, $\{e_{\text{acc},o}\}$ is a set of other electrode configurations that were identified by the expert as being acceptable for the o 'th subject, and $\text{dist}(e_{m,o}, e_{\text{opt},o})$ is an electrode configuration distance metric we have defined on all other electrode configurations. $\text{dist}(\cdot, \cdot)$ needs to capture the difference in quality between configurations and is thus a critical element of our method. A straightforward approach would be to use the hamming distance between the electrode configurations. However, we found this to be suboptimal as certain configuration patterns, such as on-off-on-off versus off-on-off-on would be assigned the highest possible distance value even though this often does not lead to very different stimulation patterns. To address this issue, $\text{dist}(e_{m,o}, e_{\text{opt},o})$ is computed in this work in two steps as shown in Fig. 6: (1) The activation status of each electrode in $e_{m,o}$ is compared with the corresponding electrode in $e_{\text{opt},o}$. For each j 'th electrode $e_{m,o,j}$ in $e_{m,o}$ that does not match $e_{\text{opt},o,j}$, we compute the distance, in terms of the number of electrodes, to the nearest electrode in $e_{\text{opt},o}$ that does match $e_{m,o,j}$. This results in an array of distances, $d = \{d_j\}$, where

Fig. 5 The workflow of the automatic electrode configuration selection method.

Fig. 6 The distance metric between electrode configuration patterns (marker +: electrodes activated and marker -: electrodes deactivated). Both configurations have five differences in the electrode activation patterns. With the optimal distance metric, configuration $e_{b,o}$ is assigned with larger distance compared to configuration $e_{a,o}$ to the optimal configuration $e_{opt,o}$.

$d_j = |j - k|$ is the distance from $e_{m,o,j}$ to $e_{opt,o,k}$, the closest electrode in $e_{opt,o}$ that matches $e_{m,o,j}$. (2) We then compute $\text{dist}(e_{m,o}, e_{opt,o})$ as the sum of the local maxima in $\vec{d}$. This metric is designed to assign a higher cost to configurations that have more distant mismatches, which indicates greater disagreement with the optimal configuration. In summary, our approach assigns higher values to $C^{m,o}$ for less desirable electrode configurations and lower values to $C^{m,o}$ for more desirable electrode configurations.

The set of weights $\{w_i\}_{i=1}^N$ can be determined by solving Eq. (17) once offline using a constrained least-squares linear system solver in MATLAB[®] 2014b (Mathworks, Inc. Natick, Massachusetts), with the constraint $w_i \geq 0 \forall i = [1, N]$. This constraint represents an additional piece of *a-priori* knowledge that captures the fact that the cost function should increase when feature terms increase since, as designed, the value of the features increases for less desirable electrode configurations. We have found that this constraint leads to better results. Once the weights are defined by using the training dataset, the optimal electrode configuration for a new subject is determined automatically by finding the global minimum of the cost function through an exhaustive search.

We performed a validation study to show the robustness of our method. To evaluate our method on our testing dataset, we asked two electrode configuration selection experts (JHN and YZ) who currently verify all the configurations used in our clinical studies to perform a blinded and randomized evaluation of the automatic configurations against control configurations. To do this, for the 60 subjects in our testing dataset, we generated three sets of electrode configurations: manual, automatic, and

control electrode configurations. The manual electrode configurations were manually selected by JHN and have been implemented in patients in our previous clinical research studies. The automatic electrode configurations were generated by running our proposed method on the subjects in testing dataset. Control electrode configurations were constructed for each subject in the testing set by the experts by manually selecting a configuration that is not “acceptable” but “close” to acceptable for all testing subjects. An electrode configuration is judged as “acceptable” when the expert believes it can be used for CI programming and is likely to lead to hearing outcomes that are nearly as good as those that would be achieved using the best possible configuration. For each test subject, two tests were done in which each expert was presented with a pair of electrode configurations and asked to rank them in terms of quality and rate whether each configuration was acceptable. In one test, the pair of configurations consists of the automatic and manual plan. In the other test, the control and the manual plan are ranked and rated. The ordering of all tests across all test subjects was randomized and the expert was masked to the identity of each configuration. The control configurations used for tests with one expert were generated by the other expert. Masking the identity of all the configurations, including control configurations, and randomizing the order of tests were steps done to minimize the potential for the experts to be biased toward evaluating all configurations as acceptable and so that the presence of such a bias could be detected in the results. Rating a significant portion of the control plans as acceptable would be indicative of such a bias. Two experts were included so that interrater variability could be characterized.

3 Results

The parameter training process was implemented in MATLAB[®] (Mathworks Inc., Natick, Massachusetts), and the electrode configuration selection algorithm was implemented in C++. The training process is an offline process, which generates the feature cost term weights in 1, 4 min, and 7 h 40 min on a standard Windows Server PC [Intel (R), Xeon (R) CPU X5570, 2.93 GHz, 48 GB RAM] for MD, AB, and CO arrays, respectively. The electrode configuration selection algorithm required 15s, 30 s, and 2 min for MD, AB, and CO arrays, respectively. Compared to the manual selection done by an expert (requires 0.5 to 2 min), our automatic electrode configuration selection algorithm is comparable but does not require any specialized training. The feature weight values computed for the MD, AB, and CO arrays from the training dataset are shown in Table 1. As can be seen from the table, the feature that prevents deactivating the most apical electrode (f_1) and one of the channel interaction features (f_6) were assigned the highest weight values for all three types of implants. For the AB and the CO arrays, the other feature that was assigned a high weight value is the term punishing electrodes falling around the entrance of the cochlea (f_5). The other features were assigned relatively low weight values. For MD, the term punishing activating electrodes that fall around the entrance of the cochlea (f_5) and the term favoring a large area for each DVF curve (f_3) were assigned moderately high weight values. The remaining features were assigned weights with very low magnitude ($\leq 10^{-13}$). In experiments on the MD training set, we found that removing the features that were assigned the very low weights produced identical electrode configurations. This confirms that the features with low magnitude weights ($\leq 10^{-13}$) do not play a

Table 1 The feature cost terms generated for MD, AB, and CO arrays.

Feature cost terms	MD	AB	CO
1	0.68	0.28	1.09
2	0	1.55×10^{-9}	1.62×10^{-4}
3	3.72×10^{-3}	8.89×10^{-5}	8.95×10^{-4}
4	0	6.02×10^{-9}	5.09×10^{-5}
5	3.23×10^{-4}	8.32×10^{-2}	4.89×10^{-4}
6	2.28	1.37	5.43
7	0	6.75×10^{-10}	4.75×10^{-5}
8	0	3.58×10^{-9}	5.55×10^{-5}
9	0	5.07×10^{-10}	4.98×10^{-5}
10	0	1.18×10^{-9}	4.80×10^{-5}

significant role in achieving the best results and can be ignored. Thus, for MD, we only kept f_1 , f_3 , f_5 , and f_6 and removed the other features by setting their weights as 0.

The results of our validation study on our testing set are shown in Fig. 7. As can be seen from Fig. 7(a), according to expert 1 (JHN), across the 60 subjects in our testing dataset, 14 of the automatically generated electrode configurations were found to be better than the manually selected configuration. In these tests, the manual configuration was found to be acceptable, but not as good as the automatic configuration. In the remaining tests, 33 automatic configurations were found to be equivalent to or exactly the same as the manual configurations, 12 were found to be not as good as the manual configuration but still acceptable, and only 1 was found to be not acceptable. None of the control configurations was evaluated as equivalent to or better than the manual configuration, 8 were evaluated as acceptable and 52 were evaluated as not acceptable. For expert 2 (YZ), 24 automatic configurations

were found to be better than the manual configurations, 26 were found to be equivalent to or exactly the same as the manual configurations, 9 were found to be acceptable, and only 1 was evaluated as not acceptable. The same automatic case was rated as not acceptable by both experts. None of the control configurations was rated as equivalent to or better than manual configurations, only 6 were evaluated as acceptable and the remaining 54 were rated as not acceptable. These results show that, with the exception of one unacceptable result, our method performs similarly to an expert. On average, the automatic method slightly outperforms an expert since more automatic plans are ranked better than manual plans than vice versa. Two-tailed paired-sign tests were used to compare the acceptance rate for control versus automatic plans and showed that the rate at which the automatic plans are judged to be acceptable was significantly better for both expert 1 ($p = 1 \times 10^{-15}$) and expert 2 ($p = 1 \times 10^{-16}$). No statistically significant differences were found when comparing ratings of the automatic electrode configurations across the two raters ($p = 1$).

In Fig. 8, we show the DVFs for automatically determined electrode configurations for several cases. The blue dotted curves represent DVFs for electrodes that are removed from the active electrode configuration and the red solid curves represent DVFs for electrodes that are active. The electrode numbers are in increasing order from the left to the right. To facilitate interpretation, we label the EOIs in the figure. In Fig. 8(a), a result for an AB case is shown that is identified as better than the manual configuration because the second and the fourth electrodes are deactivated in the automatically generated configuration. Deactivating those electrodes is good because they are likely to interfere with electrode 3. Figure 8(b) presents a result for an MD case that is identified as equivalent to the manual configuration. The automatic plan deactivates electrode 5 while the manual plan keeps it. The plans are judged to be equivalent because it is hypothesized that reducing channel interaction artifacts by turning off electrodes will be offset by an increase in frequency compression artifacts resulting in equivalent outcomes. Figure 8(c) presents a result for a CO case that is judged to be not optimal compared with the manual configuration but still acceptable. The 11-14-17 configuration in the automatic plan is not as good as the 10-12-15-18 because the

Fig. 7 Validation study results. (a) The results of validation studies performed by expert 1 (JHN) and (b) expert 2 (YZ) on automatic and control electrode configurations.

Fig. 8 Visualization of automatically selected (a–d) and corresponding manual (e–h) electrode configurations for several cases. An automatic AB plan that was judged as better than the manual plan is shown in (a). An automatic MD plans judged to be equivalently good are shown in (b). An automatic CO plans judged as acceptable are shown in (c). An automatic MD result that was judged as not acceptable is shown in (d).

minimum of electrode 11 and 17 in the automatic plan is very close to the curves of the neighbor electrodes 10 and 18. Thus, the 11-14-17 configuration in the automatic plan does not adequately address the channel interaction problem between electrodes 10 and 11 and electrodes 17 and 18. Figure 8(d) presents the only automatic configuration for an MD case that is judged to be not acceptable. In Fig. 8(d), the automatic configuration deactivates electrode 2 and 9. This is not desirable because of the relatively large distances between electrodes 1 and 3 and 8 and 10. This plan is likely to cause frequency compression artifacts.

4 Conclusion

In this study, we propose the first approach for automatic selection of electrode configurations for image-guided CI programming (IGCIP). This is a crucial step toward clinical translation of our IGCIP system that has been shown in clinical studies to lead to significant improvement in outcomes. Our approach is to design a DVF-feature-based cost function and to train its parameters using existing electrode configuration plans that we have accumulated in our database. Our validation study has shown that our method generalizes well on a

large-scale testing dataset and can produce acceptable electrode configurations in the vast majority of cases. In the validation tests with implant models from the three major CI manufacturers, our automatic method produces acceptable configurations for 98.3% of the arrays tested. According to the evaluation results from two experts in our group, around 83% of the configurations produced by our automatic method were ranked as at least equivalent to the manual configurations. Around 33% of the configurations produced by our automatic method were ranked as better than the manual configurations, whereas only 17% of the manual configurations were ranked as better than the automatic. These results suggest that our method is a viable approach for automatically selecting electrode configurations for IGCIP with similar performance to a trained expert. While the best approach to assess our IGCIP system would be to analyze a collection of hearing outcomes data from CI recipients before and after using IGCIP with the automatic and the manual electrode configuration selection methods, such data are difficult to obtain. This is because it would require subjects to come back once for reprogramming and again to re-evaluate outcomes 3 to 6 weeks after reprogramming. In the future, we plan to perform such a study with a limited number of recipients who live in close proximity to our institution. While our method generates acceptable configurations for the vast majority of cases tested, it is still capable of producing unacceptable configurations. Thus, in future work, we will investigate developing an automatic method to evaluate the quality of the electrode configuration generated by our method. This would enable our IGCIP system to notify the user that expert intervention might be needed to select the electrode configuration when our automatic method fails.

Acknowledgments

This work was supported in part by Grant Nos. R21DC012620, R01DC014037, 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 this institute. This research was also supported by a fellowship from the Vanderbilt Institute in Surgery and Engineering. An earlier version of the paper was printed in the Proceedings of the SPIE Conference on Medical Imaging 2015: Image-Guided Procedures, Robotic Interventions, and Modeling, 94150K.

References

1. NIDCD Fact Sheet: Cochlear Implants, "National Institute on Deafness and Other Communication Disorders," NIH Publication No. 11-4798, <https://www.nidcd.nih.gov/sites/default/files/Documents/health/hearing/FactSheetCochlearImplant.pdf> (2011).
2. B. S. Wilson et al., "Better speech recognition with cochlear implants," *Nature* **352** (6332), 236–238 (1991).
3. D. D. Greenwood, "A cochlear frequency-position function for several species—29 years later," *J. Acoust. Soc. Am.* **87**(6), 2592–2605 (1990).
4. O. Stakhovskaya et al., "Frequency map for the human cochlear spiral ganglion: implications for cochlear implants," *J. Assoc. Res. Otol.* **8**(2), 220–223 (2007).
5. A. Aschendorff et al., "Quality control after cochlear implant surgery by means of rotational tomography," *Otol. Neurotol.* **26**(1), 34–37 (2005).
6. J. T. Rubinstein, "How cochlea implants encode speech," *Curr. Opin. Otolaryngol. Head Neck Surg.* **12**(5), 444–448 (2004).
7. M. W. Skinner et al., "In vivo estimates of the position of advanced bionics electrode arrays in the human cochlea," *Ann. Otol. Rhinol. Laryngol.* **116** (4), 2–24 (2007).
8. B. M. Verbist et al., "Multisection CT as a valuable tool in the postoperative assessment of cochlear implant patients," *AJNR Am. J. Neuroradiol.* **26**(2), 424–429 (2005).
9. G. B. Wanna et al., "Impact of electrode design and surgical approach on scalar location and cochlear implant outcomes," *Laryngoscope* **124**(Suppl. 6), S1–S7 (2014).
10. G. B. Wanna et al., "Assessment of electrode placement and audiologic outcomes in bilateral cochlear implantation," *Otol. Neurotol.* **32**(3), 428–432 (2011).
11. J. H. Noble et al., "Image-guidance enables new methods for customizing cochlear implant stimulation strategies," *IEEE Trans. Neural Syst. Rehabil. Eng.* **21**(5), 820–829 (2013).
12. Y. Zhao et al., "Automatic localization of cochlear implant electrodes in CT," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2014*, pp. 331–338, Springer, Orlando, Florida (2014).
13. J. H. Noble and B. M. Dawant, "Automatic graph-based localization of cochlear implant electrodes in CT," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*, pp. 152–159, Springer, Munich, Germany (2015).
14. J. H. Noble et al., "Statistical shape model segmentation and frequency mapping of cochlear implant stimulation targets in CT," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2012*, pp. 421–428, Springer, Nice, France (2012).
15. C. Boëx et al., "Electrical field interactions in different cochlear implant systems," *J. Acoust. Soc. Am.* **114**(4), 2049–2057 (2003).
16. Q.-J. Fu and G. Nogaki, "Noise susceptibility of cochlear implant users: the role of spectral resolution and smearing," *J. Assoc. Res. Otol.* **6**(1), 19–27 (2005).
17. J. H. Noble et al., "Clinical evaluation of an image-guided cochlear implant programming strategy," *Audiol. Neurotol.* **19** (6), 400–411 (2014).

Yiyuan Zhao received his BS degree in telecommunication engineering from Xidian University, Xi'an, Shaanxi, China, in 2011. He is currently a PhD student of electrical engineering from Vanderbilt University, Nashville, Tennessee. His current research project is automatic techniques for cochlear implant CT image analysis. His research interests include medical image processing, image segmentation and registration, and computer vision.

Benoit M. Dawant received his MSEE degree from the University of Louvain, Leuven, Belgium, in 1983, and his PhD from the University of Houston, Houston, Texas, in 1988. He is currently a professor at the Department of Electrical Engineering and Computer Science, Vanderbilt University. His primary research areas include medical image processing, segmentation, and registration. He currently focuses on applying these methods to the development of novel techniques for guiding surgical and interventional procedures.

Jack H. Noble received his BE, MS, and PhD degrees in electrical engineering from Vanderbilt University, Nashville, Tennessee, in 2007, 2008, and 2011, respectively. He is currently a research assistant professor at the Department of Electrical Engineering and Computer Science, Vanderbilt University. His primary research interests include medical image processing, image segmentation, registration, statistical modeling, and image-guided surgery techniques.