Selecting electrode configurations for image-guided cochlear implant programming using template matching

Dongqing Zhang
Yiyuan Zhao
Jack H. Noble
Benoit M. Dawant

Selecting electrode configurations for image-guided cochlear implant programming using template matching

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

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

Abstract. Cochlear implants (CIs) are neural prostheses that restore hearing using an electrode array implanted in the cochlea. After implantation, the CI processor is programmed by an audiologist. One factor that negatively impacts outcomes and can be addressed by programming is cross-electrode neural stimulation overlap (NSO). We have proposed a system to assist the audiologist in programming the CI that we call image-guided CI programming (IGCIP). IGCIP permits using CT images to detect NSO and recommend deactivation of a subset of electrodes to avoid NSO. We have shown that IGCIP significantly improves hearing outcomes. Most of the IGCIP steps are robustly automated but electrode configuration selection still sometimes requires manual intervention. With expertise, distance-versus-frequency curves, which are a way to visualize the spatial relationship learned from CT between the electrodes and the nerves they stimulate, can be used to select the electrode configuration. We propose an automated technique for electrode configuration selection. A comparison between this approach and one we have previously proposed shows that our method produces results that are as good as those obtained with our previous method while being generic and requiring fewer parameters. © 2017 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JMI.5.2.021202](https://doi.org/10.1117/1.JMI.5.2.021202)]

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

Paper 17131SSPR received May 5, 2017; accepted for publication Aug. 16, 2017; published online Dec. 11, 2017.

1 Introduction

Cochlear implants (CIs) are one of the most successful neural prosthetics that have been developed over the last 20 years.¹ CIs are used to treat patients with severe-to-profound hearing loss. In surgery, an electrode array with 12 to 22 contacts is implanted in the cochlea. A processor worn behind the ear is connected to the array and sends signals to individual electrodes to stimulate the spiral ganglion (SG) nerves, i.e., the auditory nerves that are frequency-mapped within the cochlea. After implantation, the CI is programmed by an audiologist. Programming includes selecting contacts to deactivate and assigning sound frequency ranges and stimulation levels to each contact.

Recent studies have suggested that hearing outcomes with CIs are correlated with the spatial relationship between the electrode array and the nerves it stimulates.^{2–7} But, because the electrode array is blindly inserted into cochlea, its final position is generally not known. Recently, we have developed a series of algorithms that permit localizing both inner ear structures and the electrode array in CT images.^{8–13} We have shown that this information can be used to estimate how much overlap exists between neural areas each contact activates. Figure 1 shows the channel interaction or cross-electrode neural stimulation overlap (NSO) phenomenon, which is known to negatively affect hearing outcomes. We have shown in clinical studies that for both adults and pediatric populations, hearing outcomes improve significantly when stimulation overlap is detected and the configuration, i.e., the set of active electrodes, is adjusted to try to reduce it.^{14–16}

To estimate the amount of cross-electrode NSO, we have developed what we refer to as distance-versus-frequency curves (DVF's). We then use this information to guide programming, a process that we call IGCIP for image-guided CI programming.¹⁷ DVF's are 2-D plots, as shown in Fig. 2, that capture the patient-specific spatial relationship between electrodes and SG nerves. In this figure, the x axis is the position along the length of the SG nerves in terms of characteristic frequency, i.e., the sound frequency at which the local auditory nerves are activated in natural hearing, and the y axis is the distance to the SG nerves. The number of 2-D roughly parabolic curves in this plot is equal to the number of contacts in the electrode array. Each curve is labeled with the electrode number. The height of each individual DVF curve thus represents the distance from the corresponding electrode to the SG nerve pathways. An example is shown in Fig. 2. Here, the distance from electrode #2 to the SG nerves is larger than that from electrodes #1 and #3. We assume that electrodes that are farther away from the SG nerves activate a broader region than the electrodes that are close to the nerves. Here, this means that electrode #2 will interfere with electrodes #1 and #3. Electrodes #5, #6, and #7 are all far from the nerves and close to each other. It is thus also likely that they interfere with each other. Looking at the plot, an experienced user would deactivate electrodes #2 and #6. In Ref. 18 and its extension,¹⁹ Zhao et al. have proposed a method that captures the heuristics used by experts and translate them into a set of rules. Weights associated with each rule are estimated from a training set, and we have demonstrated that this method leads to

*Address all correspondence to: Dongqing Zhang, E-mail: dongqing.zhang@vanderbilt.edu

Fig. 1 An illustration of the electrode interaction effect: gray spheres are electrodes and the spiral-shaped surface is the modiolus. Colored cones indicate neural areas the corresponding electrodes are stimulating.

Fig. 2 A synthetic 8-electrode DVF example.

satisfactory results, but it requires estimating the values of more than 10 parameters.

As the size of our DVF library increased, we hypothesized that it would become possible to find DVFs that are highly similar to each other. If this was the case, configurations that have been selected by an expert could be applied to new cases for which the configuration is unknown. But, with experience, we observed that matching complete DVFs is difficult and would require a very large library to be successful. We addressed this issue by decomposing entire DVFs into what we call patches and matching patches rather than entire DVFs. We then impose constraints to make the electrode configuration patterns compatible between subsequent patches. Our approach is tested on 20 cases for each of the three main CI manufacturers, and we compare the configurations generated by our method with those generated by Zhao's method as well as with configurations produced manually. The results we have obtained show that our method produces configurations that are generally as good as those obtained by Zhao's method, but it requires far

fewer parameters and could potentially improve further with larger library sizes.

2 Methods

CIs are manufactured by three main companies: Med-El (MD) (Innsbruck, Austria), Advanced Bionics (AB) (Valencia, California), and Cochlear (CO) (New South Wales, Australia). Implants distributed by these companies differ mainly by the number of electrodes in the array. To accommodate these differences, we create three libraries, one for each array type. In the study presented herein, the data set contains 58, 43, and 152 DVFs for MD, AB, and CO arrays, respectively. For each DVF in the libraries, the manual configuration, which we consider to be the ground truth, has been defined by Jack H. Noble (JHN) who is primarily responsible for creating configurations in the clinical studies we are conducting to evaluate our IGCIP approach. Each of these ground truth configurations has been used clinically. Both Dongqing Zhang (DZ) and Yiyuan Zhao have been trained by JHN to produce and evaluate deactivation solutions but did not produce deactivation solutions for the DVFs used in this study. They will be referred to as rater 1 and rater 2 in the remainder of this article. Each DVF in the libraries is decomposed into a series of patches that are defined as five consecutive electrode curves. Two consecutive patches overlap by four electrodes. Since the MD, AB, and CO implants have 12, 16, and 22 electrodes, respectively, this leads to 8, 12, and 18 patches per implant for the first, second, and third implant types. When computing the configuration automatically we use the same method for each CI type, but we adjust the number of electrodes to match the number of electrodes in the array of interest. When matching a new patient's DVFs to our libraries we also decompose them into patches. The matching process involves two steps. First, individual patches are matched and then patch-to-patch consistency is enforced. A flowchart of our algorithm is shown in Fig. 3. Without loss of generality, the description of our method is based on DVFs produced for CO implants.

To match individual patches, we define a feature vector as follows. For each electrode, eight features shown in Fig. 4 are computed. These are the minimum distance to the SG (1), the slopes of the curve at the intersection points with its left (2) and right (3) neighboring curves, the frequency ranges from the minimum to the two intersection points (4–5), the difference in the distance to the SG at the two intersection points (6–7), and the distance along a vertical line passing through the minimum between this minimum and the first intersection with one of the neighboring curves (8). For each individual feature, first- and second-order statistics are computed for all electrodes in the library. Each feature is subsequently normalized by subtracting its mean and dividing by its standard deviation. The

Fig. 3 The flowchart of our automatic electrode configuration algorithm.

Fig. 4 Illustration of features 1 (green), 2 to 3 (black), 4 to 5 (purple), 6 to 7 (red), and 8 (orange).

feature vector for a patch is obtained by stacking the five individual feature vectors, leading to a 40-D feature vector. When determining the feature vector for new test patches not in the library, the new feature vector is normalized using the same mean and standard deviation computed from the library. The similarity score s between patches P_i and P_j with feature vectors $\mathbf{F}_i$ and $\mathbf{F}_j$ is then computed as

$$s(P_i, P_j) = \frac{\mathbf{F}_i \cdot \mathbf{F}_j}{\|\mathbf{F}_i\| \|\mathbf{F}_j\|}. \quad (1)$$

Figure 5 shows two exemplar patient patches and a series of library patches with their similarity score s . This figure shows

that our feature vector leads to a similarity metric that is reasonable, i.e., high scores correspond to patches that are visually similar.

With this similarity metric, a straightforward way to deactivate electrodes would be to find the best match for each test patch and to assign its central electrode's deactivation state to the test patch's central electrode. However, this oversimplifies the problem. When experts manually select the “on/off” state of one electrode, they not only consider the state of the central electrode, they also consider its consistency with the state of neighboring electrodes. An example can be seen in Fig. 2. Here, as discussed earlier, since electrode #6 is close to electrodes #5 and #7 but far away from the nerves, it is likely to interfere with electrodes #5 and #7. A good choice to reduce interaction would be to deactivate it but only if #5 and #7 are kept active. An alternative would be to deactivate #5 and #7. In that case, electrode #6 should be left on to avoid a big stimulation gap. However, it is possible that the best matches for each of #5, #6, and #7 all suggest turning off their middle electrodes but leaving neighbors on. This leads to deactivating all three electrodes. If this happens, electrode interaction is eliminated but too many members are turned off from the already limited number of electrodes in this region, and a stimulation gap is left between electrodes #4 and #8. This will result in excessive compression of the sound frequency spectrum, i.e., in fewer independent frequency channels than could be used and some neural areas may also not be stimulated at all. Both these issues would lead to suboptimal outcomes. So, without enforcing deactivation

Fig. 5 Two patient patches, A and B, each with three library patches are shown on the same row. The value of the similarity measure s increases as the visual similarity between patches increases.

Fig. 6 The directed graph we construct to find the optimal match sequence.

pattern consistency between patches, the strategy mentioned above would likely lead to bad configurations.

To tackle this problem, we impose patch-to-patch consistency as follows. We denote the i 'th ($i = 1, 2, 3, \dots, 18$) test patch as T_i . For T_i , instead of only selecting the most similar match, we store k matches with the highest similarity scores, forming a candidate set, denoted as $C_{i,j}$ with similarity scores $s_{i,j}$ ($j = 1, 2, 3, \dots, k$). For each test patch, our final goal is to find a match in the candidate set that is both similar to the test patch and leads to good patch-to-patch consistency. Specifically, for two neighboring patches, we define the inconsistency D using the number of "on/off" state disagreements in their two middle overlapping electrodes. With "1" meaning "electrode on" and "0" meaning "electrode off," suppose patches T_1 and T_2 in one set of DVFs consist of electrodes #1 to #5 and electrodes #2 to #6, respectively. Suppose also that for each patch we found a match, $C_{1,1}$ and $C_{2,1}$, respectively. Finally, suppose that the deactivation patterns of these matches are $a_{1,1} = [1, 1, 0, 0, 1]$ and $a_{2,1} = [1, 0, 1, 1, 0]$. The electrodes they both include are #2 to #5 and the two middle electrodes are #3 and #4. We define the disagreement as the Hamming distance²⁰ between the two state vectors for electrodes #3 and #4. In our example, it would be the Hamming distance between $[0, 0]$ and $[0, 1]$, i.e., $D(a_{1,1}, a_{2,1}) = 1$. The bigger the disagreement, the higher the inconsistency. Searching for the best global deactivation strategy involves minimizing a cost function containing both patch-match dissimilarity and patch-to-patch inconsistency terms. The simplest solution would be to do an exhaustive search but that would be time-consuming. Instead, we have formulated it as searching for the shortest path in a directed graph. As is shown in Fig. 6, every node in the graph is a match.

Table 1 p -values of McNemar mid- p tests between our method and Zhao's method at different ranks.

p -value	"Better than manual" rate	"At least equivalent as manual" rate	Acceptance rate
Rater #1	0.73	0.5	1
Rater #2	0.3	0.23	0.51

Matches in the same column are the candidates for the same test patch. We put two extra nodes in the graph as the starting and ending nodes. Arbitrary similarity scores are assigned to them and their disagreements with neighboring matches are set to 0. Edges only exist from patch $C_{i,j}$ to $C_{i+1,l}$, ($i = 0, 1, 2, \dots, 18$ and $j, l = 1, 2, \dots, k$), and the cost along an edge is defined as

$$\text{cost}(C_{i,j}, C_{i+1,l}) = \lambda[1 - s(T_{i+1}, C_{i+1,l})] + (1 - \lambda)D(a_{i,j}, a_{i+1,l}). \quad (2)$$

Here, we use the first term $1 - s(T_{i+1}, C_{i+1,l})$ to represent the patch-match dissimilarity between T_{i+1} and $C_{i+1,l}$. $a_{i,j}$ and $a_{i+1,l}$ are the deactivation patterns of patch $C_{i,j}$ and $C_{i+1,l}$. The second term $D(a_{i,j}, a_{i+1,l})$ is the deactivation pattern inconsistency. λ ($0 < \lambda < 1$) balances these two terms. The shortest path from the starting node to the ending node gives the optimal match sequence. To find the shortest path, we used Dijkstra's algorithm.²¹

Fig. 7 Validation study results. (a)–(c) The results of validation studies performed by rater 1 and (d)–(f) rater 2 on our proposed automatic method, Zhao's automatic and control electrode configurations.

After the optimal match sequence is found, the on/off state of each optimal match's central electrode is used for the corresponding contact in the new patient's DVFs. The state of the first two and of the last two electrodes are assigned to be the

same as the state of these electrodes in the first and last matches of the sequence, respectively.

Note that with the proposed algorithm, the final electrode configurations depend on the value of two parameters, k and λ .

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

The next section will present the results we have obtained. In Sec. 4, we discuss how these parameters were chosen as well as the sensitivity of our results to these parameter values.

3 Results

At the time of writing, we have evaluated our technique on DVFs generated from images of 20 patients/brand for a total of 60 patients for whom manual solutions are available. We automatically generated the electrode configurations using the method proposed herein and Zhao's method. To evaluate the solutions, human rater 1 and rater 2 were asked to compare automatic and manual configurations as well as automatic and control configurations. The control configurations were generated by manually producing a configuration that is not "acceptable" but "close" to acceptable for all test subjects. This was done by DZ who changed the on/off state of 1 to 3 electrodes. All the validation solutions were created before the comparative study started to minimize the chance that particular solutions could be remembered. This control configuration was added to avoid evaluation bias, i.e., to avoid having the human rater biased toward rating solutions as acceptable/comparable if they knew that they were always presented with a human-generated and computer-generated solution. Specifically, for each subject, three comparisons were performed: the manual configuration versus automatic configuration #1 (generated by our proposed method), the manual configuration versus automatic configuration #2 (generated by Zhao's method), and the manual configuration versus the control configuration. For each comparison, the raters were presented with the two configurations blind to their origin and the ordering was randomized. For each comparison, the two raters were asked to rank the solutions in terms of quality and to rate whether each was acceptable or not. "Acceptable" means that in the rater's opinion, the configuration can be used for CI programming and is likely to lead to hearing outcomes that are comparable to those achieved using the best possible configuration.

The results are shown in Fig. 7. According to rater 1, across all 60 subjects, Zhao's method generated "better than the manual" configurations for 5 subjects, "at least equivalent as the manual" ("better than the manual" + "replicate the manual" + "equally good as the manual") configurations for 35 subjects and unacceptable configurations for 2 subjects. Our proposed method generated "better than the manual" configurations for 4 subjects, which was one less than Zhao's method, "at least equivalent as the manual" configurations for 40 subjects, 5 more than Zhao's method, and unacceptable configurations for the same number of subjects as Zhao's method. According to rater 2, Zhao's method generated "better than the manual" configurations for 10 subjects, "at least equivalent as the manual" configurations for 40 subjects, and unacceptable configurations for 5 subjects. Our proposed method generated "better than the manual" configurations for 6 subjects, which was 4 less than Zhao's method, "at least equivalent as the manual" configurations for 46 subjects, 6 more than Zhao's method, and unacceptable configurations for only 3 subjects, 2 less than Zhao's method. For statistical analysis, we performed a McNemar mid- p test between groups of configurations. The McNemar mid- p test is a statistical test used for binary matched-pairs data.²² We found that the differences between acceptance rates of our method and control configurations are statistically significant for both rater 1 ($p = 1.8 \times 10^{-13}$) and rater 2 ($p = 4.4 \times 10^{-16}$). We did not find statistically significant difference

when comparing the acceptance rate of our method across the two raters ($p = 0.63$). Neither did we find statistically significant difference when comparing the acceptance rate of Zhao's method across the two raters ($p = 0.22$) and acceptance rate of control configurations across the two raters ($p = 0.21$). To summarize, our proposed method generated more configurations of high quality (at least equivalent as the manual configurations) and has a higher acceptance rate than Zhao's method, although it does not outperform the manual configurations as often as Zhao's method. McNemar mid- p tests between our method and Zhao's method in terms of the number of configurations reaching different ranks, e.g., differences in the number of configurations deemed "better than the manual configuration" between the two methods by each rater are shown in Table 1. None of them are found to be significantly different.

In Fig. 8, we show the DVFs for automatically determined electrode configurations for four example cases. The blue curves represent DVFs for electrodes that are kept active in the configuration and the red dashed curves represent DVFs for electrodes that are deactivated. In Fig. 8(a), a result for an AB case is identified as better than the manual configuration that is shown in Fig. 8(e) because electrode #10 is likely to interfere with electrode #12 in the manual configuration and it is turned off in the automatic one. In Fig. 8(b), the automatic configuration for an AB case is rated as equally good as the manual configuration. The automatic configuration deactivates electrode #10, but the manual configuration keeps it. Though deactivating electrode #10 reduces some interaction, because the gap between electrode #8 and #11 is fairly big, it will be offset by an increase in frequency spectrum compression artifacts we mentioned above. In Fig. 8(c), the automatic configuration for an MD case is identified as only acceptable because turning off electrode #11 will probably result in slight frequency spectrum compression artifacts but it is not a serious issue here. In Fig. 8(d), the automatic configuration for a CO case is not acceptable because there will be serious electrode interaction between electrode #1, #2, and #3 if electrode #2 is not deactivated.

We also conducted an experiment to evaluate the sensitivity of our proposed method to the size of the library. The CO library is the only one that is large enough to conduct such a study. For this array type, we used 10, 20, 40, and 80 DVFs randomly selected from the complete set of 152 DVFs. We then produced four different configurations, one for each library size and asked rater 1 to evaluate them. In this experiment, we did not compare our method to Zhao's method, but a control configuration was included for each subject. This results in five comparisons. Each comparison involves the manual configuration versus one of the four automatic or the control configurations. Again, the trials are presented in a blinded fashion with the ordering randomized. Figure 9 shows a plot of the number of "at least equivalent

Fig. 9 Evaluation results when different library sizes are used.

Fig. 10 Evaluation results when different λ s are used ($k = 5, 10$, and 15 , respectively).

as the manual" configurations and "unacceptable" configurations for different library sizes. The result obtained with the entire library (152 DVFs) is also shown. This figure shows that results tend to improve as the library size increases, at least up to 80 DVFs. With the number of DVFs currently at our disposal we cannot assess whether we have reached a plateau or whether results could improve further with a larger library.

4 Parameter Selection

The algorithm we present herein involves two parameters: k , which is the number of candidate matches for each test patch, and λ , which is the weight of the patch-match dissimilarity term in the cost function. To find the best values for these parameters, we created a validation test that contained 10 randomly selected CO subjects. The remaining 142 DVFs were used as library for these validation cases. We generated the configurations for this validation set and asked one of the raters to do the evaluation. The evaluation was done as described in the library size sensitivity study discussed in Sec. 3. Since the evaluation is a laborious process, instead of doing a grid search for the two parameters, we adopted the following strategy: we first heuristically chose a few values of k and evaluated the results when different λ s were used. Following this experiment, we fixed the value of λ and evaluated the method on a larger range of k values.

Specifically, we started with k set to 5, 10, and 15 and we varied λ from 0 to 1 and counted both the number of configurations that were at least equivalent to manual configurations and the number of unacceptable configurations. The results are shown in Fig. 10. As can be seen, the results are largely insensitive to the value of λ , but combining patch-match similarity and patch-to-patch consistency is nevertheless important.

Based on the results discussed above, we chose $\lambda = 0.5$, which is roughly in the middle of the [0.0001 and 0.99] interval, and varied k between 1 and 80, starting at one and iteratively doubling its value. Figure 11 shows that the results improve

Fig. 11 Evaluation results when different k s are used ($\lambda = 0.5$).

as k increases from 1 to 10 and reaches a plateau when $k = 10$. Based on these observations, we chose $\lambda = 0.5$ and $k = 10$ to run our experiments on the testing set that was used to generate the results presented in Sec. 3.

5 Conclusions

This paper presents a generic algorithm for automatic electrode configuration selection for CIs. This method produces results that are comparable to those obtained by Zhao's method described in Refs. 17 and 18. Specifically, in the large-scale study we have conducted, this new method generates more configurations of high quality (at least equivalent to manual configurations) than Zhao's. It also generates fewer unacceptable solutions, but it does not outperform the manual configurations as often as Zhao's method. We have also observed that performance tends to increase with library size. As more data become available, we will assess whether our technique reaches a plateau or keeps improving. Finally, this method only requires a few parameters (λ and k , i.e., the number of patches kept when matching the library to a new patch). In contrast, the previous method relies on heuristics used by experts to come up with a deactivation strategy. Capturing a complete set of heuristics, translating them into equations, assigning weights to each rule, and dealing with unusual cases are difficult. The weights would also need to be reestimated when the training set is modified to, for instance, include new electrode arrays. Here, we rely on deactivation plans that have been vetted by experts, and we match new cases to known ones. In terms of computation, the proposed algorithm takes around 1 s/case. Zhao's method takes 2 to 3 s for MD and AB implants and ~35 s for CO implants, which have more contacts. We note, however, that computation time is not a critical issue at this point.

Disclosures

This work was performed under an IRB protocol for human data approved by Vanderbilt University. No conflicts of interest, financial or otherwise, are declared by the authors.

Acknowledgments

This research has been supported by grants NIH R01 DC014037 and R01 DC014462 from the National Institute of Deafness and Other Communications Disorders. The content is solely the responsibility of the authors and does not necessarily represent the official views of this institute. An earlier version of this paper was printed in the proceedings of the SPIE Conference on Medical Imaging 2017: Image-Guided Procedures, Robotic Interventions, and Modeling, 101350Q.

References

1. NIDCD Fact Sheet: Cochlear Implants, "National institute on deafness and other communication disorders," NIH Publication No. 11-4798, 2011, <https://www.nidcd.nih.gov/sites/default/files/Documents/health/hearing/FactSheetCochlearImplant.pdf> (30 August 2017).
2. A. Aschendorff et al., "Quality control after insertion of the nucleus contour and contour advance electrode in adults," *Ear Hear.* **28**, 75s–79s (2007).
3. C. C. Finley et al., "Role of electrode placement as a contributor to variability in cochlear implant outcomes," *Otol. Neurotol.* **29**, 920–928 (2008).
4. L. K. Holden et al., "Factors affecting open-set word recognition in adults with cochlear implants," *Ear Hear.* **34**, 342–360 (2013).
5. G. B. Wanna et al., "Impact of intrascalar electrode location, electrode type, and angular insertion depth on residual hearing in cochlear implant patients: preliminary results," *Otol. Neurotol.* **36**, 1343–1348 (2015).
6. K. F. Nordfalk et al., "Scalar position in cochlear implant surgery and outcome in residual hearing and the vestibular system," *Int. J. Audiol.* **53**, 121–127 (2014).
7. 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).
8. J. H. Noble and B. M. Dawant, "Automatic graph-based localization of cochlear implant electrodes in CT," *Med. Image Comput. Comput. Assisted Intervention* **9350**, 152–159 (2015).
9. F. A. Reda et al., "Automatic segmentation of intra-cochlear anatomy in post-implantation CT of unilateral cochlear implant recipients," *Med. Image Anal.* **18**, 605–615 (2014).
10. F. A. Reda et al., "An artifact-robust, shape library-based algorithm for automatic segmentation of inner ear anatomy in post-cochlear-implantation CT," *Proc. SPIE* **9034**, 90342V (2014).
11. J. H. Noble et al., "Automatic segmentation of intracochlear anatomy in conventional CT," *IEEE Trans. Biomed. Eng.* **58**, 2625–2632 (2011).
12. Y. Zhao et al., "Automatic localization of cochlear implant electrodes in CT," in *Int. Conf. on Medical Image Computing and Computer-Assisted Intervention*, Vol. **17**, pp. 331–338 (2014).
13. Y. Zhao, B. M. Dawant, and J. H. Noble, "Automatic localization of cochlear implant electrodes in CTs with a limited intensity range," *Proc. SPIE* **10133**, 101330T (2017).
14. J. H. Noble et al., "Initial results with image-guided cochlear implant programming in children," *Otol. Neurotol.* **37**, e63–e69 (2016).
15. J. H. Noble et al., "Clinical evaluation of an image-guided cochlear implant programming strategy," *Audiol. Neurotol.* **19**, 400–411 (2014).
16. R. F. Labadie et al., "Results of postoperative, CT-based, electrode deactivation on hearing in prelingually deafened adult cochlear implant recipients," *Otol. Neurotol.* **37**, 137–145 (2016).
17. 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).
18. Y. Zhao, B. M. Dawant, and J. H. Noble, "Automatic electrode configuration selection for image-guided cochlear implant programming," *Proc. SPIE* **9415**, 94150K (2015).
19. Y. Zhao, B. M. Dawant, and J. H. Noble, "Automatic selection of the active electrode set for image-guided cochlear implant programming," *J. Med. Imaging* **3**(3), 035001 (2016).
20. R. W. Hamming, "Error detecting and error correcting codes," *Bell Syst. Tech. J.* **29**(2), 147–160 (1950).
21. E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik* **1**, 269–271 (1959).
22. M. W. Fagerland, S. Lydersen, and P. Laake, "The McNemar test for binary matched-pairs data: mid-p and asymptotic are better than exact conditional," *BMC Med. Res. Method.* **13**(1), 91 (2013).

Dongqing Zhang received his BS degree in electrical and information engineering from the University of Science and Technology of China, Hefei, China, in 2013. Currently, he is a PhD student in the Department of Electrical Engineering and Computer Science, Vanderbilt University. His research interest is image-guided surgery, medical image processing, and machine learning.

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.

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.

Benoit M. Dawant received his MSEE degree from the University of Louvain, Louvain-La-Neuve, Belgium, in 1983, and his PhD from the University of Houston, Houston, Texas, in 1988. He is currently a Cornelius Vanderbilt Professor in 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.