



A Graph-Based Method for Optimal Active Electrode Selection in Cochlear Implants

Erin Bratu¹(✉), Robert Dwyer², and Jack Noble¹

¹ Department of Electrical Engineering and Computer Science, Vanderbilt University,
Nashville, TN 37235, USA

{erin.l.bratu, jack.noble}@vanderbilt.edu

² Department of Hearing and Speech Sciences, Vanderbilt University Medical Center,
Nashville, TN 37232, USA

Abstract. The cochlear implant (CI) is a neural prosthetic that is the standard-of-care treatment for severe-to-profound hearing loss. CIs consist of an electrode array inserted into the cochlea that electrically stimulates auditory nerve fibers to induce the sensation of hearing. Competing stimuli occur when multiple electrodes stimulate the same neural pathways. This is known to negatively impact hearing outcomes. Previous research has shown that image-processing techniques can be used to analyze the CI position in CT scans to estimate the degree of competition between electrodes based on the CI user's unique anatomy and electrode placement. The resulting data permits an algorithm or expert to select a subset of electrodes to keep active to alleviate competition. Expert selection of electrodes using this data has been shown in clinical studies to lead to significantly improved hearing outcomes for CI users. Currently, we aim to translate these techniques to a system designed for worldwide clinical use, which mandates that the selection of active electrodes be automated by robust algorithms. Previously proposed techniques produce optimal plans with only 48% success rate. In this work, we propose a new graph-based approach. We design a graph with nodes that represent electrodes and edge weights that encode competition between electrode pairs. We then find an optimal path through this graph to determine the active electrode set. Our method produces results judged by an expert to be optimal in over 95% of cases. This technique could facilitate widespread clinical translation of image-guided cochlear implant programming methods.

Keywords: Cochlear implants · Graph search · Image guided cochlear implant programming

1 Introduction

In the United States, it is estimated that 2 to 3 out of every 1000 children are born with some degree of hearing loss, and 37.5 million adults experience some degree of hearing loss [1]. The cochlear implant (CI) is a neural prosthesis that, over the last two decades, has become the standard treatment for severe-to-profound hearing loss [2]. As of 2012, an estimated 324,000 CIs have been implanted worldwide, and in the United

States, approximately 58,000 adults and 38,000 children have received a CI [1]. The CI is primarily used in cases of sensorineural hearing loss, where damage or defects affecting hearing are in the cochlea, a.k.a., the inner ear. In a subject without hearing loss, sounds reaching the cochlea would be transduced to electrical impulses that stimulate auditory nerve fibers. The nerve fibers are tonotopically organized, meaning that activation of nerve fibers located in different regions of the cochlea create the sensation of different sound frequencies. The frequency for which a nerve fiber is activated in natural hearing is called its characteristic frequency. As such, in natural hearing nerve fibers are activated when their characteristic frequencies are present in the incoming sound.

In a patient with hearing loss, sounds no longer properly activate auditory nerve fibers. The purpose of CI is to bypass the natural transduction mechanisms and provide direct electrical stimulation of auditory nerve fibers to induce hearing sensation. A CI consists of an electrode array that is surgically implanted in the cochlea (see Fig. 1a) and an external processor. The external processor translates auditory signals to electrical impulses that are distributed to the electrodes in the array according to the patient's MAP, which is the set of processor instructions determined by an audiologist in an attempt to produce optimal hearing outcomes. Tunable parameters in a patient's MAP include the active set of electrodes, the stimulation level of each electrode, and a determination of which electrodes should be activated when a particular frequency of sound is detected by the processor. Research indicates that the locations of electrodes within the cochlea impact the quality of hearing outcomes [3–9]. Most patients have less-than-optimal electrode array placement [5], so customizing the patient's MAP is critical for optimizing hearing outcomes. Previous studies have proposed methods for segmenting cochlear anatomy and electrode arrays from pre- and post-operative computed tomography (CT) images, permitting creation of 3D models of cochlear structures [10–12].

Research has also shown that the spatial information garnered from these methods can be used to estimate channel interactions between electrodes [12, 13]. Channel interaction occurs when nerves, which naturally are activated for a finely tuned sound frequency, receive overlapping stimulation from multiple electrodes, corresponding to multiple frequency channels. This creates spectral smearing artifacts that lead to poorer hearing outcomes. Manipulating a subject's MAP to modify the active electrode set and the stimulation patterns can reduce these effects. The spatial relationship between electrodes and neural sites is a driving factor for channel interaction. Modeling an electrode as a point charge in a homogenous medium has been shown to yield similar electric field estimates to more sophisticated finite element models when using plausible tissue resistivity values within known ranges for the human cochlea [14]. Using a point charge model, Coulomb's law mandates that electric field strength, $E(\vec{x})$, at location \vec{x} is inversely proportional with squared distance between \vec{x} and the electrode location, \vec{c} .

$$E(\vec{x}) \propto \frac{1}{\|\vec{x} - \vec{c}\|^2} \quad (1)$$

Thus, as shown in Fig. 1b, when an electrode is close to neural sites (e.g., E8–E10), relatively little current is needed to activate nearby nerves, resulting in relatively little spread of excitation. However, when an electrode is distant to neural sites, neural activation requires broad stimulation patterns due to electrical current spread at greater

distance. When two electrodes are close together and both distant to neural sites (e.g., E5–6), they create substantial stimulation overlap, resulting in channel interaction.

One method of visualizing the spatial relationship between electrodes and neural sites to determine when channel interaction occurs is to use distance vs. frequency (DVF) curves. These curves represent the distance from the auditory nerve spiral ganglion (SG) cells, which are the most likely target of electrical stimulation, to nearby electrodes (see Fig. 1c). The characteristic frequencies of the nerve fiber SG sites are shown on the horizontal axis, and the distance to electrodes near to those neural sites are shown with the height of an individual curve for each electrode on the vertical axis. This simplifies the process of determining which nerve pathways are likely to be stimulated by a given electrode and where two electrodes might stimulate the same region. An electrode is most likely to stimulate the nerves it is closest to, as indicated by the horizontal position of the minimum of the curve. We refer to these nerves as having SGs located in the peak activation region (PAR) for the electrode. Determining which nerves are stimulated by multiple electrodes requires making additional assumptions about the spread of excitation of each contact. In this work, we use Eq. (1) to estimate electric field strength, and we assume the activation region for an electrode includes any nerves with SG sites \vec{x} that satisfy:

$$\frac{E(\vec{x})}{E(\text{PAR})} = \frac{\|\text{PAR} - \vec{c}\|^2}{\|\vec{x} - \vec{c}\|^2} \geq \tau, \quad (2)$$

which requires that the strength of the electric field in SGs must be greater than a certain fraction, τ , of the electric field in the PAR for those nerves to be considered active. This is equivalent to ensuring the ratio of squared distance from the PAR to the electrode to the square distance from another nerve SG site to the electrode is greater than τ . The DVF curves (see Fig. 1c) permit visually assessing the activation region of each contact. The activation region is defined by the width of the curve for which $\sqrt{\tau}$ times the curve height, $\sqrt{\tau}\|\vec{x} - \vec{c}\|$, is less than or equal to the minimum curve height, $\|\text{PAR} - \vec{c}\|$.

If substantial overlap of activation regions exists between neighboring electrodes, some electrodes may be selected for deactivation to reduce overlap. This is one approach for image-guided CI programming (IGCIP), i.e., a method that uses image information to assist audiologists optimize programming of CIs. The original technique for selecting the deactivated set required an expert to manually review each case and determine the optimal solution based on the information in the DVF curves. This process is not ideal for clinical translation as it can be time-consuming and requires expertise. Automated methods have been developed to eliminate the need for expert review, which either rely on an exhaustive search to optimize a cost function that relies on shape features of the DVF curves [15] or attempt to learn to replicate expert deactivation patterns using DVF curve template matching [16]. However, as presented below, the current state-of-the-art method [15] leads to optimal results in only 48% of cases, which is insufficiently reliable for widespread clinical translation. In this paper, we present an automated method for determining the active electrode set as a minimum cost path in a custom-designed graph. As our results will show, our method is significantly more robust in finding optimal deactivation plans compared to the state-of-the-art method and could facilitate automated clinical translation of IGCIP methods.




Fig. 1. (a) A 3D representation of cochlea (red) and the electrode array (gray). (b) The modiolus containing the SG cells of the auditory nerve is shown in green with estimated spread of excitation from the CI electrodes multiple colors. (c) DVF curves for the same case, showing a deactivation plan. Active electrodes are represented by solid blue lines, and deactivated electrodes by dashed gray lines. (Color figure online)

2 Methods

The dataset in this study consists of 83 cases for which we have patient-specific anatomical data that is used to generate the DVF curves and electrode deactivation plans from the current state-of-the-art technique [15] and our proposed method. All cases use an implant from one of three manufacturers: MED-EL (MD) (Innsbruck, Austria), Advanced Bionics (AB) (Valencia, California), and Cochlear (CO) (New South Wales, Australia). Of these cases, 24 used an implant from MD, 32 from AB, and 27 from CO.

2.1 Graph Definition

We propose a graph, $G = \{N, E\}$, as a set of nodes N and edges E . Each node in N represents an electrode in the array. An optimal path resulting in minimum cumulative edge cost traverses edges connecting nodes corresponding to electrodes recommended for activation. Electrodes corresponding to nodes not in the path will be recommended for deactivation. Using this approach, we (1) select the start and end nodes of the path, (2) identify valid edges using hard constraints, (3) calculate edge costs using soft constraints, and (4) use Dijkstra’s algorithm [17] to find the globally optimal path.

The start node of the path is chosen to be the most apical contact (see Fig. 1a) because deactivating the most apical contact reduces stimulation of the lowest frequency nerves, creating perceived frequency upshifts that are generally bad for hearing outcomes [18]. Thus, it is desirable for this electrode to always be active. Similarly, the end node should be the electrode with PAR among the highest frequency nerves that can be effectively stimulated near the basal end of the cochlea. It is well known that electrodes outside the cochlea and those near the entrance of the cochlea are typically ineffective in stimulating auditory nerves. We use the active-shape-model based segmentation approach proposed by [19] to segment the cochlea and rely on the one-to-one point correspondence between the segmentation in the patient image and an atlas image to define a decision plane (see Fig. 1a). The plane corresponds to nerves with characteristic frequencies of 15 kHz. The first electrode apical to this plane is the end node of the path.

The edges E are defined to permit finding a minimum edge cost path from the start to the end node that represents the optimal set of active electrodes. Our proposed structure

of E is shown in Fig. 2. Edge e_{ij} is a directed edge connecting electrode i to electrode j with cost $C(e_{ij})$. Hard constraints (whether e_{ij} exists) and soft constraints (edge costs defined by a cost function $C(e_{ij})$) ensure the minimal path corresponds to the optimal active electrode set. Two necessary conditions must be met for e_{ij} to exist. First, e_{ij} only exists if $i < j$. This ensures the path traverses from the most apical electrode, E1, to a sequence of increasingly more basal neighbors until reaching the end node. As seen in Fig. 3, edges only exist connecting nodes to higher numbered nodes. Second, we encode a maximum allowable amount of activation region overlap between sequential active electrodes in the path as a hard constraint. We use Eq. 2 to define the activation region for each electrode and let $\tau = 0.5$ in our experiments. We found heuristically that this value of τ selected a similar rate of active electrodes as reported in studies of the number of effectively independent electrodes as a function of electrode distance [20], and it also matched behavior of experts when selecting electrodes for deactivation. We then define overlap acceptable if the activation region for electrode j does not include the PAR for electrode i and vice versa. Thus, if the region most likely to be activated by an electrode (its PAR) is also activated by another electrode, too much overlap is occurring, in which case e_{ij} does not exist. For example, in Fig. 1c the PAR for E3 falls within the activation region for E4, therefore $\nexists e_{34}$. An example of such a scenario is shown in our example graph in Fig. 2 with edge e_{12} . With n_1 and n_2 exhibiting too much overlap, $\nexists e_{12}$, and thus a path from n_1 must skip n_2 and instead traverse directly to n_3 or n_4 . One example allowable path in this graph is shown in red.

Soft constraints are encoded in an edge cost function,

$$C(e_{ij}) = \alpha d_i + (1 - \alpha)\beta^{(j-i-1)}, \quad (3)$$

where $d_i = \|\text{PAR}_i - \vec{c}_i\|$ is the distance from electrode i to its PAR, and α and β are parameters. The second term in the cost function ensures as many electrodes are active as allowable by the hard constraints, since when $j = i + 1$, no electrodes are deactivated, but when $j > i + 1$, some electrodes are skipped in the path, indicating they will be deactivated, and, assuming $\beta > 1$, a higher cost is associated with this. Further, a larger cost is assigned when deactivating multiple electrodes in sequence, i.e., when $j \gg i + 1$, to discourage deactivations that result in large gaps in neural sites where little stimulation occurs. Larger values of β result in greater values for this penalty. The first term in Eq. (3) rewards active electrodes that tend to have lower distance to SG sites. The parameter α controls the relative contribution of the two terms.




Fig. 2. Visualization of our graph design. Each node is indicated with a black circle and is labeled with a node number n_i . Dotted lines indicate invalid edges, and solid lines indicate valid edges. Red lines indicate an example path through the graph. (Color figure online)

From this graph, Dijkstra’s shortest-path algorithm can determine the global cost minimizing path. The resulting path represents the set of electrodes that should remain active, while electrodes not in the path will be recommended for deactivation.

2.2 Validation Study

Ideally, we would have an expert determine the optimal deactivation plan for each of the 83 cases in our dataset and measure the rate at which the algorithm produces the optimal plan. However, for a given case, it is possible there are multiple deactivation configurations that could be considered equally optimal, and it is difficult to determine a complete set of equally optimal plans. Thus, to assess the performance of our method, we instead implemented a masked expert review study to assess optimality of the results of our algorithm compared to the current state-of-the-art algorithm and control plans for each case. In this study, an expert reviewer was presented with a graphical representation of the DVF curves for each case, showing the planned active and deactivated electrodes, similarly to Fig. 1c. The reviewer was instructed to determine whether each plan was optimal, i.e., the reviewer would not adjust anything in the presented plan. Three sets of plans for each case were presented in this study. The first set consists of the results from our proposed method using parameters $\alpha = 0.5$ and $\beta = 4$. These parameter values were determined heuristically using DVF curves from 10 cases not included in the validation set. The second set is the deactivation result from the method described in [15]. The final set includes control plans, manually created by a second expert, where the active set is close to acceptable, but suboptimal. The inclusion of control plans is used to indicate if the reviewer has a bias toward rating plans as optimal, e.g., if numerous control cases are rated as optimal, the reviewer likely has such a bias. The three sets of plans were presented one at a time in random order. The reviewer was masked to the source of each plan in order to prevent bias towards any method.

2.3 Parameter Sensitivity Analysis

We performed a parameter sweep to assess the sensitivity of the parameters in our cost function across a set of values around the heuristically determined values of $\alpha = 0.5$ and $\beta = 4$ used above in the validation study. We used our proposed method to determine the active electrode set with parameter α in the range [0.1, 0.9] with step size of 0.1 and β in the range [2, 6] with a step size of 0.5. This resulted in 81 different parameter combinations for each case. We then used the Hamming distance metric to compare the resulting plan to the plan evaluated in the validation study. Large differences would indicate greater sensitivity of the method to the parameters.

3 Results and Discussion

The results of our validation study are shown in Table 1. Our reviewer judged 79 of the plans generated using our proposed method to be optimal, rejecting only four cases. Only 40 of the plans from the previous method described in [15] were rated optimal, and none of the control plans were marked optimal. Accepting none of the control plans

indicates that our expert reviewer is not biased toward accepting configurations and can accurately distinguish between optimal and close-to-optimal plans. We used McNemar mid-p tests to assess the accuracy of our plan to produce an optimal result versus that of the current state-of-the-art method in [15] as well as the control method. We found that the difference in success rates between the two methods and between the proposed and control method were highly statistically significant ($p < 10^{-9}$).

Table 1. Validation study results. Each row indicates the number of optimal and non-optimal plans generated by the given method in each column.

	Proposed	Algorithm from [14]	Control
Optimal	79	40	0
Non-optimal	4	43	83

Inspecting the four cases where the proposed deactivation plan was rejected, the reviewer noted that the plans for these cases were actually optimal, and the rejection in each case was due to erroneous reading of the DVF curves when the amount of activation region overlap between electrodes was very close to the acceptable overlap decision threshold. DVF curves for one such case are shown in Fig. 3a, along with the deactivation plan suggested by [15] in panel (b). The plan from the proposed algorithm in (a) was rejected because the reviewer mistakenly believed the PAR for E5 (green) fell within the activation region for E4 (red). Note that the hard constraints imposed by our proposed method guarantee plans that are free of this type of error, which is a significant benefit of graph-based, compared to other optimization methods. The plan from algorithm [15] in (b) is also borderline but was correctly judged to be unacceptable because the PAR for E2 (green) falls just outside the activation region for E3 (red), meaning E2 should be active.




Fig. 3. (a) A rejected plan generated using our proposed method. (b) A rejected plan generated using the method from [15], showing the same case as (a).

The results of our parameter sensitivity study are shown in Fig. 4. We found that our method was relatively insensitive to low values for α and high values for β , i.e., the deactivation plan did not change from the $\alpha = 0.5$, $\beta = 4$ solution used in the validation study in this region of the parameter space. However, large numbers of plans changed when α was high or β was low. Since our validation study revealed that $\alpha = 0.5$, $\beta = 4$ produced optimal solutions, changes in many plans indicates that those configurations

likely produce sub-optimal results. This finding is reasonable since, when β is low or α is high, deactivating numerous electrodes in sequence is not properly penalized in the cost function.

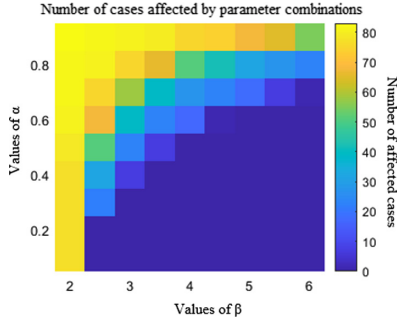


Fig. 4. Parameter sensitivity test results. Higher values indicate a greater number of plans with differences from the originally generated plans, where $\alpha = 0.5$ and $\beta = 4$.

4 Conclusion

In this study, we presented an automated graph-based approach for selecting active electrode sets in CIs. Automated selection methods reduce the time required to develop a patient-specific plan and remove the necessity for an expert reviewer to manually select the active electrodes from a set of DVF curves. Clinical translation of IGCIP techniques requires that our developed methods be robust and reliable to maximize positive hearing outcomes in patients. Our approach utilized spatial information available from previous techniques for segmenting cochlear structures and electrode arrays. We used this information to develop a graph-based solution for selecting an optimal active electrode set. To validate our results, we asked an expert reviewer to rate electrode configurations as optimal or non-optimal, where for a plan to be considered optimal, the reviewer would make no changes to that plan. 95.2% of plans created from our method were accepted as optimal, compared to only 48.2% of plans generated using the current state-of-the-art technique. Further, post-evaluation review revealed that the four rejected plans from our proposed method were actually optimal. These results suggest that our method is significantly more robust than the current state-of-the-art method and could facilitate widespread, automated clinical translation of IGCIP methods for CI programming. In the future, we plan to evaluate our method in a clinical study to confirm that the results of our method produce improved hearing outcomes for CI recipients in practice. This study would examine improvements in hearing outcomes for subjects relative to their current implant configuration over the course of several weeks by collecting data before reprogramming and again after a 3 to 6-week adjustment period to the new electrode configuration. Following successful clinical confirmation of our method, we will perform a multi-site study to assess clinically translating this method to other institutions.

Acknowledgements. This work was supported in part by grants R01DC014037, R01DC014462, and R01DC008408 from the NIDCD and by training grant T32EB021937 from the NIBIB. This content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

1. NIDCD Quick Statistics About Hearing. National Institute on Deafness and Other Communication Disorders. <https://www.nidcd.nih.gov/health/statistics/quick-statistics-hearing>
2. NIDCD Fact Sheet: Cochlear Implants. 2011 National Institute on Deafness and Other Communication Disorders. NIH Publication No. 11-4798 (2011)
3. Holden, L.K., et al.: Factors affecting open-set word recognition in adults with cochlear implants. *Ear Hear.* **34**(3), 342–360 (2013)
4. Finley, C.C., Skinner, M.W.: Role of electrode placement as a contributor to variability in cochlear implant outcomes. *Otol Neurotol.* **29**(7), 920–928 (2008)
5. Chakravorti, S., et al.: Further evidence of the relationship between cochlear implant electrode positioning and hearing outcomes. *Otol Neurotol.* **40**(5), 617–624 (2019)
6. Aschendorff, A., et al.: Quality control after cochlear implant surgery by means of rotational tomography. *Otol Neurotol.* **26**(1), 34–37 (2005)
7. Wanna, G.B., et al.: Impact of electrode design and surgical approach on scalar location and cochlear implant outcomes. *Laryngoscope* **124**(S6), S1–S7 (2014)
8. Wanna, G.B., et al.: Assessment of electrode placement and audiologic outcomes in bilateral cochlear implantation. *Otol Neurotol.* **32**(3), 428–432 (2011)
9. O’Connell, B.P., et al.: Electrode location and angular insertion depth are predictors of audiologic outcomes in cochlear implantation. *Otol Neurotol.* **37**(8), 1016–1023 (2016)
10. Zhao, Y., et al.: Automatic graph-based method for localization of cochlear implant electrode arrays in clinical CT with sub-voxel accuracy. *Med. Image Anal.* **52**, 1–12 (2019)
11. Noble, J.H., Dawant, B.M.: Automatic graph-based localization of cochlear implant electrodes in CT. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9350, pp. 152–159. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24571-3_19
12. 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. Neural Syst. Rehabil. Eng.* **21**(5), 820–829 (2013)
13. Noble, J.H., Gifford, R.H., Labadie, R.F., Dawant, B.M.: Statistical shape model segmentation and frequency mapping of cochlear implant stimulation targets in CT. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012. LNCS, vol. 7511, pp. 421–428. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33418-4_52
14. Rattay, F., Leao, R.N., Felix, H.: A model of the electrically excited human cochlear neuron. II. Influence of the three-dimensional cochlear structure on neural excitability. *Hear Res.* **153**(1–2), 64–79 (2001)
15. Zhao, Y., Dawant, B.M., Noble, J.H.: Automatic selection of the active electrode set for image-guided cochlear implant programming. *J. Med. Imaging* **3**(3), 035001 (2016)
16. Zhang, D., Zhao, Y., Noble, J.H., Dawant, B.M.: Selecting electrode configurations for image-guided cochlear implant programming using template matching. *J. Med. Imaging* **5**(2), 021202 (2018)
17. Dijkstra, E.W.: A note on two problems in connexion with graphs. *Numer. Math.* **1**, 269–271 (1959)

18. Stakhovskaya, O., Sridhar, D., Bonham, B.H., Leake, P.A.: Frequency map for the human cochlear spiral ganglion: implications for cochlear implants. *J. Assoc. Res. Otolaryngol.* **8**(2), 220–233 (2007)
19. Noble, J.H., Labadie, R.F., Majdani, O., Dawant, B.M.: Automatic segmentation of intracochlear anatomy in conventional CT. *IEEE Trans. Biomed. Eng.* **58**(9), 2625–2632 (2011)
20. Berg, K.A., Noble, J.H., Dwyer, R.T., Labadie, R.F., Gifford, R.H.: Speech recognition as a function of the number of channels in perimodiolar electrode recipients. *J. Acoust. Soc. Am.* **145**(3), 1556–1564 (2019)