Flaws in Spectral Ripple Stimuli for Listeners with Cochlear Implants

Matthew B. Winn¹ and Gabrielle O'Brien²

¹ Department of Speech-Language-Hearing Sciences, University of Minnesota, USA

² Department of Speech & Hearing Sciences, University of Washington, USA

Financial disclosures/conflicts of interest: None

All correspondence should be addressed to:

Matthew B. Winn, University of Minnesota Department of Speech-Language-Hearing
Sciences, 164 Pillsbury Drive SE, Minneapolis, MN 55455, USA. Email: mwinn@umn.edu

Abstract

Objectives

Spectral resolution is an important aspect of hearing speech accurately, and a known limiting factor in cochlear implants (CIs). The spectral ripple discrimination task is a psychophysical measure that has been found to correlate with word recognition, music perception, and various aspects of CI processing. It was hypothesized that the limited spectral sampling of CI processors would distort the output of spectral ripple stimuli in ways that could be problematic for their use in CI research.

Design

Idealized acoustic ripple spectra were analyzed through synthetic filters matched to CI processor channels, and the output was examined for spectral density and modulation depth cross a wide range of inputs. Existing data showing correlations between ripple discrimination and speech recognition were re-analyzed with special consideration of scores above the spectral aliasing limit.

Results

The analysis revealed numerous complicating factors in the transmission of spectral ripples through CIs, including spectral aliasing and modulation depth neutralization. These unintended distortions are nonlinear in nature, with the consequence that spectral ripple thresholds above a certain number cannot be ordered monotonically along a single parametric axis/dimension. For the Cochlear device, the critical limit for

spectral aliasing is around 2.1 RPO, and the critical limit for modulation saturation is about 3.5 RPO. Neither limitation is remediated by changing stimulus phase. Changes in modulation depth are limited by the single channel with the narrowest input filter octave-scaled bandwidth (channel 9 in the Cochlear device). Correlations between ripple discrimination and various measures of speech recognition are improved when high ripple scores are dismissed or winsorized. An additional difficulty with interpreting discrimination thresholds is that ripple stimuli often do not match the spectral characteristics of modulations found in actual speech sounds.

Conclusions

The authors conclude that the practice of identifying the highest spectral density at which ripple phase inversion can be detected should be abandoned because (1) it lacks the well-behaved mathematical properties a psychophysical experimenter requires, and (2) high spectral densities are a poor representation of ecologically-relevant cues. Alternative testing strategies (ripple depth detection, ripple phase shift, acoustic-phonetic cue perception) should be taken, with the specific consideration of real speech acoustics in their development.

INTRODUCTION

Despite the enormous success of the cochlear implant (CI) as a neural prosthesis to restore a sense of hearing, there remain significant limitations of the device that are the subjects of ongoing research. One of the more pressing issues in this line of work is the evaluation and improvement of spectral resolution, which is the ability to resolve differences in frequency patterns (or the spectral envelope) in sounds. In CI listeners, spectral resolution generally refers to the ability to resolve differences in spatial patterns of auditory nerve activation. It is widely hypothesized that poor spectral resolution in cochlear implant listeners may limit the effectiveness of their devices in transmitting complex spectral information relevant for speech perception (Friesen et al. 2001; Won et al. 2007; Litvak et al. 2007; Azadpour & McKay 2012; Winn & Litovsky 2015), although it is unknown to what extent individual differences in spectral resolution contribute to the wide variability in outcomes associated with implantation. In light of the numerous spectral cues involved in differentiating speech sounds that CI listeners are known to misperceive (Munson et al. 2003), it seems reasonable to suspect that spectral resolution is a chief limiting factor of the device. To better understand the mechanisms that determine cochlear implant efficacy, translational research has become increasingly concerned with the quantification of spectral resolution in CI listeners (Won et al. 2007; Won et al. 2011a; Drennan et al. 2010; Drennan et al. 2014; Shim et al. 2014; Drennan et al. 2015; Horn et al. 2017).

There are numerous challenges in evaluating spectral resolution in a meaningful way. First, using natural speech sounds as a test stimulus is typically fraught with confounding factors such as the involvement of non-spectral cues (such as amplitude

envelope cues, linguistic context cues, familiarity with the test stimulus), and the tendency for newly-implanted patients to increasingly adapt to the new sound of their device over the first 6 to 18 months of use (Blamey et al. 2013). While that acclimatization period is beneficial for the patient, it is problematic for the experimenter who desires a measure of peripheral auditory ability independent from experience, i.e. without having to wait upwards of a year to stabilize. Furthermore, it would be ideal to have a test of spectral resolution that is largely independent of development or cognitive status, and which is not sensitive to learning effects. A test with these constraints could prove useful in the prognosis of an implanted listener in the earliest days post-activation without a long period of uncertainty about their development of device use.

Description of the Spectral Ripple test in CI users

The spectral ripple test is a popular psychophysical method that is generally considered a probe of spectral resolution in individual listeners, and which has been argued to satisfy the constraints mentioned above (Won et al. 2007; Jones et al. 2013). Often, the test is presented in a discrimination paradigm, in which listeners distinguish two phases of a sinusoidally modulated spectral envelope (the spectral ripple) presented at equal modulation depths. It is also presented as a task in which the listener discriminates a spectrally modulated envelope from a spectrally flat envelope, i.e. detection of spectral modulations rather than discrimination of different spectral modulations. The carrier signal being spectrally modulated is typically filtered white noise, pink noise, or a phase-randomized harmonic complex. In some studies, the ripple density (the number of spectral modulations per octave) is increased until an adaptive

tracking procedure converges on the highest ripple density at which the listener can reliably complete the task (Won et al. 2007). Other studies have fixed the ripple density and tracked the smallest modulation depth that can be reliably discriminated (Litvak et al. 2007; Saoji et al. 2009). One could also imagine a case where at fixed density and fixed modulation depth, a task where listeners are probed for the smallest *phase shift* that is detectable.

One reason for the spectral ripple test's popularity is that spectral ripple thresholds have been demonstrated to correlate with speech perception tests, including word, consonant, and vowel recognition (Henry et al. 2005), and speech recognition in noise (Won et al. 2007; Lawler et al. 2017; although note that Anderson et al. 2011, did not observe robust correlations with any speech perception measure). Higher (better) ripple discrimination thresholds are also associated with better musical perception (Won et al. 2010). The test appears to be sensitive to physical and psychophysiological properties of the implant, including differences between processing strategies (Drennan et al. 2010), and indices of electrical current spread or channel interaction (Jones et al. 2013; Won et al. 2014; Scheperle & Abbas 2015). These correlations support the hypothesis that excellent spectral resolution *enables* effective perception of speech and other sounds. They could also be taken as support for the appealing idea that increasing spectral resolution in a listener via an alternative stimulation strategy may improve their performance. Enthusiasm for the spectral ripple test is therefore understandable from both a clinical and academic perspective.

Limitations of the Spectral Ripple test in CI users

Despite their correlations with scores on various speech tasks, several concerns about the spectral ripple test have been raised that suggest thresholds may not be easily interpretable. For example, Azadpour and McKay (2012) suggest that listeners could attend to differences in loudness, spectral centroid, and changes to the spectral edges rather than differences in broadband spectral density. They further suggest that processing of the stimulus through the CI speech processor exacerbate problems of edge detection and further deteriorate the stimulus if peak-picking strategies are used (which they commonly are, in Cochlear devices). One potential confound is that a listener capable of attending to a frequency region containing less than a single ripple could use a local loudness cue to discern ripple phases. Additionally, listeners could use the highest and lowest audible frequencies or the spectral centroid as cues. To mitigate these issues, Aronoff and Landsberger (2013) developed the modified spectral ripple test, which introduces temporal modulations (in the form of drifting ripple phase) that reduce local loudness cues and limit movement of spectral centroid. In the sections that follow, we attempt to quantify the extent of any possible distortions of the spectral ripple stimulus so that they can be better understood.

Goals of the current study

In addition to the concerns raised by Azadpour and McKay (2012), several potentially serious concerns remain unaddressed, which are the topic of the current article. Here we address the problems of 1) spectral aliasing, 2) decline and eventual saturation of spectral modulation depth, 3) the potential role of experience, learning or

cognitive factors other than spectral auditory perception, and 4) the lack of any meaningful correspondence between spectral ripples and the acoustics of speech.

MATERIALS AND METHODS

The approach to spectral ripple analysis was purely analytical in nature. Idealized acoustic ripple spectra were filtered through a synthetic CI speech processor. We used the channel-frequency allocations obtained from the fitting software for the devices made by the Cochlear Corporation (), and analyses for other devices are available in the Supplemental Digital Content. For any particular ripple density and phase, Spectral energy was binned within each channel; the intensity of all the acoustic frequency components were averaged into a single number representing the output of each channel. To simplify the interpretation, no peak-picking analysis was done. Figure 1 illustrates the acoustic and corresponding electric representations of a small number of ripples to verify the process.

162

[[Figure 1 here]]

163

164 Fig 1. Spectral ripples analyzed through a synthetic CI processor. The gray sinusoidal
 165 spectra are idealized acoustic ripples. Banded rectangular shapes are the simulated
 166 electric channel output, consisting of the average energy of the acoustic energy within
 167 each channel.

168

169

RESULTS

170

171

172

173

174

175

176

Figure 1 illustrates how at low spectral densities (e.g. 0.5 RPO), the electric representation / processor output represents the acoustic input spectrum rather faithfully. However, at higher densities, the electric spectrum is unable to represent the spectral phase changes. Figure 2 illustrates a wide range of finely-sampled spectral densities along with their corresponding electric representations. It is evident from Figure 2 that at some RPO density between 2 and 3, the electric representation of the spectrum loses correspondence with the intended input. We conclude that the lack of

177 correspondence is a potentially serious flaw for an experimenter who wants to
178 rigorously vary a single parameter – in this case, spectral density.

179 Spectral aliasing

180 The spectral sampling limit of a cochlear implant must necessarily be limited by
181 the number of channels on an electrode array, and the spectral envelopes employed in
182 the ripple test may be subject to aliasing effects. This is a simple result of the Nyquist
183 sampling theorem in the spectral domain. In short, for a rippled spectrum of N ripples
184 per octave, there must be at least $N*2$ frequency channels *per octave* in order to encode
185 the spectral envelope (just as there must be a time sampling rate of $f*2$ in order to
186 encode frequency f).

187

188

[[Figure 2 here]]

189

190 Fig 2. Extended sampling of spectral ripple densities following the same pattern as

191 Figure 1. Around 2.1 ripples per octave (RPO), the simulated electric output no longer

192 faithfully encodes the intended spectral modulations.

193

194

195 Anderson et al. (2011) argued that spectral aliasing is not necessarily
196 problematic, as the spectral ripple task does not depend on faithfully replicating an
197 acoustic stimulus and can be performed simply by detecting any change in the phase-
198 inverted spectrum (but see Won et al. 2011b for counterargument). As long as even a
199 single electrode itself can encode a dense-enough spectral modulation, the task can be
200 completed. While the mathematical premise of Anderson et al.'s comment is sound,
201 there is a disconnect between it and the intended goal of the test as a broadband sound
202 whose whole spectral envelope is described with one number (RPO), and indeed whose
203 broadband spectral density is the essential defining feature that is used to correlate it
204 with speech perception scores. It should be reasonable for an experimenter to assume
205 that a stimulus with broadband periodic spectral ripples contains those rippled
206 components relatively equally across the whole spectrum. Otherwise, the stimulus is not
207 much more useful than single-electrode psychophysics than the spectral ripple test is
208 intended to supplant.

209

210 Chaotic spectral components

211 It is worth understanding what sort of spectral envelopes are being tested in the
212 range of test parameters where spectral undersampling occurs. Without fully
213 characterizing the stimuli that are commonly used as probes of spectral resolution, the
214 interpretation of this measure remains incomplete and potentially misleading. For
215 higher-density ripples are sent through the CI processor, the question is if we don't

216 produce the intended spectral modulations, *what spectral modulations densities are*
217 *actually present?* To analyze this, we produced log-spectrum frequency modulation
218 spectra for each of the simulated electrical outputs illustrated in Figure 2. Figure 3 (left
219 column) shows these modulation spectra, which we describe here. Ideally each of these
220 spectral modulation spectra would have a distinct peak at the intended RPO density
221 (indicated on the figure as a gray line). This goal is satisfied for lower-density spectra, up
222 to roughly 2.1 RPO. Past that, there are two problems that emerge: 1) the peak is lower
223 in spectral density than the intended peak, and 2) the peak is lower in power with
224 increasing density.

225 In the case of high-density simulated-electric ripples with low overall spectral
226 modulation power, it is worth examining the *relative* power of spectral modulations that
227 are indeed present, even if they are weaker than those for low-density ripples. To do
228 this, we normalized each log-spectral modulation spectrum to a 0-1 scale such that the
229 peak level is scaled to 1, for each ripple density. The resulting *normalized* log-spectral
230 modulation spectra – shown in Figure 3, right column – are relatively unchanged in
231 appearance for the lower densities, but show a new and complicated pattern for higher
232 densities (above 2.5). Specifically, the output peak is always below the intended peak
233 (see lack of match between peaks indicated as red circles, and the target gray line), and
234 there are relatively strong contributions of other peaks across the spectrum. We marked
235 the spectral density of the maximum peak and any peak whose power was at least half
236 that of the maximum peak.

238

[[Figure 3 here]]

239

240 Fig 3. Log-frequency spectral modulation spectra indicating the presence of octave-

241 spaced spectral modulations for a variety of densities of ripples analyzed through the

242 synthetic CI processor. Vertical gray lines indicate the “target” location of the peak,

243 corresponding to the peak in the idealized acoustic input. The left panel shows raw

244 spectral modulation power, including virtually absent modulation strength for high-

245 density ripples whose modulation depth is neutralized. The right panel shows

246 normalized spectral modulation power, where the peak – at whatever raw power – is

247 scaled to 1. In the right panel, the absolute peak is marked with a closed black point, and

248 any local peaks that are at least half the strength of the maximum peak are marked with

249 open points.

250

251

252 The distribution of those spectral densities as a function of intended input
253 density is shown in Figure 4. The pattern reveals that the relationship between spectral
254 density input and output is clean and orderly for densities up to about 2.1 RPO. Higher
255 densities show three distinct features: 1) aliased maxima, where the peak spectral
256 density is roughly 2.1 minus the difference between the target and 2.1, 2) decreased
257 spectral modulation power at higher RPOs, indicated by smaller points on the graph and
258 3) chaotic mini-maxima whose relation to the target is likely mathematically rule-
259 governed, but not representative of the intended spectral density. Even the predictable
260 pattern of the aliased maxima breaks down at the density two times the roughly
261 estimated critical point of 2.1 RPO. We conclude that the non-monotonicity of output
262 spectral density as a function of input spectral density is a serious flaw for an
263 experimenter who wants to exert control over the stimulus and interpret correlations
264 between spectral density perception and perception of other sounds such as speech.

265

266

[[Figure 4 here]]

267

Fig 4. Peaks in log-frequency spectral modulation density (i.e. presence of different RPOs) as a function of intended RPO input to an idealized CI processor. Black peaks indicate maximum spectral modulation strength, and open circles indicate any local peaks that are at least half the strength of the maximum peak; these points correspond to the points of similar marking in Figure 3). Size of the point scales with raw modulation power.

It has been proposed by Aronoff and Landsberger (2013) that some of the limitations of stimuli with “static” spectral ripples can be addressed by introducing drifting phase changes in the spectrum. We have simulated and measured the log-frequency modulation spectra of each of the previously illustrated RPOs across the entire phase space in increments of ten degrees. However, when phase is shifted in

increments of ten degrees across the entire phase spectrum, results show no appreciable reduction in the aliasing problem. Higher-density ripples still contain log-frequency modulation spectra that are lower in power and which contain disorderly distributions of modulation spectra peaks past the same critical RPO limit around 2.1 RPO for this simulated Cochlear device.

Spectral modulation depth saturation

In addition to the aforementioned issue of spectral aliasing, there remains another equally problematic issue: as spectral ripple density increases, the spectral modulation depth of the output stimulus gets smaller. This is a mathematical inevitability, as any phase change that occurs within a single frequency-analysis channel results in averaging the power of its included components, resulting in output that is equal to neither the maximum peak nor the minimum valley of the input. Therefore, the difference in spectral power between the maxima and minima *across the entire spectrum* will be compromised. Additionally, if the width of the frequency channel becomes narrower than the entire period ($1/\text{ripple density}$) of a spectral modulation (i.e. if the peak and valley complete a whole modulation cycle within one channel), the spectral modulation is neutralized and discrimination becomes unreliable (c.f. O'Neil et al. 2019). Figure 5 illustrates idealized electrical ripple output spectra with phase inversions, across different spectral densities. It can be seen that at higher densities, the depth of spectral modulation change is shallower. Still, as long as a listener can detect a difference in the channel with maximum phase-inverted power change, the task can be completed successfully.

[[Figure 5 here]]

Fig. 5. Phase inversion of simulated electric output of spectral ripple stimuli. Gray space between phase-inverted patterns is the difference in channel-specific output energy, and the darkest shaded region indicates the channel with greatest change in energy across phase inversion.

An advocate of spectral ripple testing could counter the importance of this problem by pointing out that maintaining good auditory discrimination performance in the face of shallower modulation depth is indeed still a sign of success. However, the input-output function of the ripple sent through the CI processor violates the concept of parametrically varying a single feature of the stimulus. As the experimenter varies the RPO, she or he wishes to conclude that change in performance is a change related to a change in RPO, not some other unintended parameter change. It is unfavorable to

wonder whether success was attributable to good spectral resolution or to excellent intensity discrimination within a single channel.

Echoing the message mentioned by Anderson et al. (2011) and O'Neill et al. (2019), we analyzed the relationship between CI frequency-channel bandwidth and resulting change in spectral modulation depth when a ripple's phase is inverted. Figure 6 illustrates the maximum potential power change (normalized to 1 as "full" change from minimum to maximum or vice versa) as a function of RPO density. Each channel, at each phase from 0 to 170 in 10-degree increments, contributes a single data point to this image. Therefore, maximum depth achieved at one RPO might be from a different spectral phase than that maximum achievable depth for a different RPO. It can be observed that modulation depth change is extremely variable across channels even within the same spectral densities, and even at relatively low RPOs. There are three critical points to be highlighted here. First, modulation depth is intended to be *fixed* across all of these stimuli, with only spectral density parametrically changed; there is more than one stimulus change happening at the same time. Second, although the peak spectral modulation change possible regularly occurs on channel 9, other changes are variable across phase, suggesting that spectral density changes across the array in a non-monotonic fashion. Some channels show maximum-depth power changes at 2.5 RPO that are actually weaker than those at 3.5 RPO, violating the apparent linear relationship observed on the figure. Additionally, the *maximum possible* change in modulation depth declines even at relatively low densities, and eventually asymptotes around 5.5 RPO. Supplemental materials reveal further that across the frequency-channel processing schemes of the other CI manufacturers, these problems are even more

glaring (with a particularly interesting pattern in the Advanced Bionics device, whose 0.25-octave-wide filters result necessarily in zero modulation depth for 4 RPO stimuli).

[[Figure 6 here]]

Fig 6. Change in potential channel output energy across ripple density. Each small gray/black point reflects the change in energy resulting from phase inversion (illustrated on Fig. 5). For each ripple density between 0.5 and 8 (spaced by 0.5 RPO), data were collected for each channel, across all phase pairs sampled at 10-degree intervals (e.g. 0-180, 10-190, 20-200, etc.). Large blue points reflect the median change across all channel and all sampled phases. Red triangles represent the maximum achievable change in modulation depth for any single channel across all sampled phases. Encircled numbers indicate the channel where that maximum depth emerged.

Predictably, the change in modulation depth power of a single channel is limited by its input frequency bandwidth. In the Cochlear device, which has been the example illustrated in all images thus far, the first nine channels have 125 Hz bandwidth, followed by four channels of 250 Hz, two each at 375, 500 and 625Hz, and one each at 750, 875 and 1000 Hz. When dealing in linear frequency space, this would render the first nine channels equal in their frequency-resolving power. However, when dealing in log-frequency space, as we do for the spectral ripple test, one must convert these linear Hz bandwidths to octave-spaced bandwidths. When doing so, channel 9 emerges as that which has the finest log-frequency-resolving power. See Figure 7 for an illustration of this computation.

368

[[Figure 7 here]]

369

370 Fig. 7. Bandwidth of each CI processor channel in linear space (upper panel) and octave
 371 space (lower panel). Channel 9 has the narrowest octave bandwidth, and is highlighted
 372 because of its role in producing the maximum achievable change in spectral modulation
 373 depth resulting from phase inversion (see Fig. 7).

374

375

376 The spectral modulation period that would saturate channel 9 in the cochlear
 377 device is a ripple whose density is 3.46 RPO. The channel is 0.1443 octaves wide,
 378 allowing a spectral modulation period of $1 / (2 * 0.1443)$. This results in a maximum of

3.46 ripples per octave before the modulation period equals the width of the filter. Interestingly, a recent study by O'Neill et al (2019) verified this limitation using a behavioral study in which the mean threshold for discriminable ripple density processed with no spectral spread (3.7) was almost perfectly matched to the calculated upper limit of discrimination predicted by this formula for the CI processor used in their study (3.73).

DISCUSSION

Spectral ripple stimuli processed through an idealized CI speech processor contain a troubling number of distortions that limit their interpretation in psychophysical studies. Whether through spectral aliasing or the saturation of frequency modulation depth, there is an upper limit in spectral ripple density that can be conveyed through a CI before the output no longer reflects the experimenter's intended stimulus. The changes in spectral density are non-monotonic, rendering high-density CI-processed ripples *different* but not necessarily *ordered* in their density. In any case, they are not meaningfully described by the designation of ripple density (RPO), and therefore, differences in spectral ripple density thresholds above the aliasing limit cannot be reliably interpreted as being better or worse than each other. Below, we discuss some implications of these findings in the context of previous literature, the relationship to speech perception, and motivation for how to avoid these complicating factors.

Why are spectral ripple scores correlated with speech perception measures if all these problems are true?

In light of the aforementioned imitations, there is an obvious counter-argument: “Correlations have been found between the spectral ripple test and other meaningful abilities such as perception of speech and music. Why should I not use the spectral ripple test?” This is a fair question that deserves a response. Here, we address two matters: 1) a reason why correlations have likely emerged, and 2) the potential problems of interpreting the test where the causal mechanism is unknown or misleading.

Correlations might have emerged between ripple discrimination scores and speech recognition because the aliasing and distortion resulting from spectral under-sampling render spectral ripple stimuli accidentally rather sparse in the spectral modulation domain, and hence more like natural vowel stimuli (see earlier discussion of the modulations present in real speech sounds). It is thus feasible to suspect that two ripples of high density (say, 5 and 6 RPO) could differ from each other not in terms of spectral density in any meaningful way, but in terms of the locations of particular peaks and valleys in the spectrum. This description neatly describes vowel perception as well. Those who are good at spectral pattern discrimination for the purpose of the ripple test are likely also empowered to distinguish between spectral patterns that characterize vowel spectra. Furthermore, the locations of peaks and valleys changes non-monotonically (see Figure 4); the listener who is able to monitor the whole spectrum for these changes is likely going to be a better-performing listener. In these cases, good listening skills correspond to better spectral ripple scores. So what is the problem? It is that the ripple score is described by a number that loses meaning beyond a critical point

near 2.1 RPO. In other words, the 5 RPO stimulus is not *more spectrally dense* than the 4 RPO stimulus; it is simply *different* than the 4 RPO stimulus. Past the critical spectral sampling point, the rippled stimuli are no longer monotonically ordered, so a score of 5 is not necessarily better than a score of 4.

Jones et al. (2013) argued that even if there are contaminating factors affecting the description of spectral ripples through cochlear implants, that the history of positively correlated effects preserves the test's utility in clinical auditory science. Based on the analyses put forth in the current article, we contend that there is good reason to think that something other than spectral resolution is driving (or at least majorly contributing to) these correlations, potentially misleading experimenters to make conclusions about the wrong domain of processing. Effort spent continuing in the same line of testing while knowing the limitations is effort that might alternatively be spent 1) developing a test that more purely taps into spectral resolution, or 2) more directly and intentionally taps into the auditory and/or cognitive abilities tapped by the current spectral ripple test. After all, if it's auditory pattern detection or other auditory-cognitive skills that we're interested in, there are likely ways of probing that more directly than hoping a high-density spectrum will accidentally present the opportunity for a listener to showcase those skills. For those experimenters who are invested in a deeper understanding of spectral resolution and its role in outcomes for CIs, or more specifically its role in perceiving speech sounds, the test of ripple discrimination using high densities is at best limited, and at worst misleading.

Because the spectral ripple test likely taps into a listener's ability to distinguish irregular spectral profiles (which might accidentally resemble vowel-like spectra), and

because the test is affected by experience and cognitive abilities, we argue that the spectral ripple test at even moderate ripple densities (≥ 2.1 RPO) is not a parametric test of spectral resolution. It is instead a test of general auditory change-detection skills whose parameters are not neatly described in any single dimension. Put simply, someone who is good at discriminating changes in spectral peak locations, and someone who has higher cognitive function, is an individual who is likely to both do well on the spectral ripple test and do well in speech recognition. But the quantification of spectral ripple performance is affected by too many potentially contaminating factors to know *why* speech recognition was better or poorer, particularly in light of the non-monotonicity of the output stimulus. In a case where very tight experimental control was used on a per-frequency basis, Anderson et al. (2011) revealed no correlations between spectral ripple discrimination and speech recognition. We therefore assert that previous positive correlations were not driven by the mechanism intended by the experimenter, but instead by a host of mechanisms that were not under experimental control.

Using 2.56 RPO as a generous upper limit of possible spectral density transmitted through the Cochlear CI processor, we explored two different ways of handling conspicuously high ripple-discrimination scores. The first was to use a subset of scores preserving only those below 2.56. The second method explored for handling conspicuously high ripple-discrimination scores was Winsorizing the data. In this procedure, scores above the critical limit are not dismissed, but instead reassigned to be the maximum possible value (e.g. scores of 3.9, 4.1, or 5.6 would all be coded as 2.56). Each method was used for either A) scores for individual runs, or B) final averages of

three runs. Using these methods on ripple discrimination data collected by Winn et al. (2016) in which three runs for each participant were collected, correlation between ripple scores and word recognition become *stronger*. The published correlation between RPO and word recognition using the full data set was notable ($r^2 = 0.452$), but considerably stronger when the outlying RPO scores were dismissed from the dataset ($r^2 = 0.565$), stronger still when that action was confined to individual runs ($r^2 = 0.59$), even stronger when Winsorizing average scores ($r^2 = 0.619$) and strongest yet when Winsorizing individual runs ($r^2 = 0.631$). These results suggest that individual runs in which CI listeners achieve high RPO discrimination scores ultimately weaken the predictive power of the test.

The dataset from the study by Winn et al. (2016) also includes scores for categorization of spectral phonetic cues in the form of formant transitions that distinguish /b/ from /d/. Full methods are described in the articles by Winn and Litovsky (2015) and Winn et al. (2016); briefly, higher empirical logit coefficients indicate steeper categorization functions as a function of changing formant transition and thus suggest more optimal (or at least more NH-like) perception of the formant cue. After applying the four correction methods described above, a similar trend emerged as for the word recognition scores. Spectral ripple thresholds from the full dataset were correlated with performance scores for formant perception ($r^2 = 0.331$), but this correlation became *stronger* when individual runs above 2.56 were omitted from the dataset ($r^2 = 0.497$), when average scores above 2.56 were omitted ($r^2 = 0.593$), when individual runs were Winsorized ($r^2 = 0.568$) and when average scores were Winsorized ($r^2 = 0.540$). These results corroborate those reported in the previous paragraph; ripple

discrimination scores beyond the limit of spectral aliasing weaken the predictive power of the spectral ripple test in CI listeners. Figure 8 illustrates each of the correlations described here for the original (full) data set as well as all four of the correction methods.

[[Figure 8 here]]

Fig 8. Scatterplots showing relationship between spectral ripple discrimination and word recognition (top row) or logistic cue coefficient (i.e. psychometric function slope) for consonant formant transition perception (bottom row). Within each row, different panels reflect different approaches to handling RPO data above 2.56, including maintaining it, dropping it on a per-session basis, dropping listeners who average above

2.56, reassigning any score higher than 2.56 to be 2.56 on a session basis or on a participant basis (i.e. Winsorizing).

Evidence for an upper limit of spectral density (RPO) in previous literature

There are signs in the literature that ripple stimuli sent through CI processors may behave unusually above a critical ripple density. Based on the frequency allocation tables of modern devices, the highest spectral density that can be faithfully sampled should fall between 2 and 2.5 ripples per octave (RPO). Curiously, in most published spectral ripple tests, the majority of CI participants have thresholds below 2 ripples per octave, while there are occasional outliers who score considerably higher (3-7 ripples per octave). In some data sets, correlations with speech tasks are heavily influenced or leveraged by outliers with unusually high thresholds (Won et al. 2007, for example). Similar predominance of CI listener RPO thresholds limited by a density around 2 was found by the first author of this article in a previous study with other more experienced users of the ripple test (Winn Won & Moon 2016).

The apparent upper limit of 2 RPO extends also to the modified version of the ripple task where listeners are probed for the minimum spectral modulation depth needed for auditory discrimination. Smith et al. (2013) tested low-density ripples of 0.5, 1 or 2 RPO. It was found that with monopolar stimulation, one listener could not discriminate any ripples with 1 RPO, and three (out of nine) could not perform the task at all with 2 RPO stimuli.

Neural signatures of auditory discrimination of ripples measured by Lopez Valdes et al. (2014) corroborate an upper limit near 2 RPO for CI listeners. In their study, neural responses to ripple phase inversion were statistically separable up to 2 RPO, where they were statistically indistinguishable. Insofar as their measurement reflected encoding of peripheral features, this evidence further challenges the idea that high-density ripples convey the auditory features intended by the experimenter.

When experimenters limit their stimuli to spectral densities below 2, correlations with speech perception are *stronger* than when such a limit is not imposed. Henry et al. (2005) concluded that good vowel and consonant recognition could be achieved with spectral discrimination between 1 and 2 RPO. Litvak et al. (2007) showed that identification of consonants and vowels was correlated with spectral modulation detection at low RPOs of 0.25 and 0.5. Saoji et al. (2009) examined this result further, finding that vowel identification was most strongly associated with ripple discrimination thresholds for 0.25 RPO, and consonant perception was most strongly associated with discrimination at 0.5 RPO (which happened to be the *sole* predictor in a multiple regression analysis for consonants).

Ripple discrimination thresholds above 2.5 found in the literature are sometimes accompanied by dubious red flags. In the study of ripple discrimination by Winn et al. (2016), five out of 19 CI participants achieved average RPO scores above 2.5. Of these, three showed that their scores were carried by abnormal variation on at least one run, including one listener whose single-run score of 7.450 was complemented by second- and third-run scores of 2.56 and 1.81. The single-run high score of 7.45 was clearly not reliable, and likely reflected a known consequence of getting “lost” in an adaptive

tracking algorithm. Yet inclusion of this score inflated the listener's overall performance score. With only limited trial-level data available to the current authors, it remains unknown how often this issue might have arisen in earlier literature. Other studies corroborate an interpretation similar to what is offered here. For example, in addition to identifying a suggestive correlation between better spatial resolution (as measured by eCAP) and higher-density ripple discrimination, Anderson et al. (2011) found that two of their CI listeners with especially *poor* spatial resolution attained the *highest* spectral ripple discrimination scores. O'Neill et al. (2019) observed that CI users who had especially poor speech perception scores did not have noticeably poorer spectral ripple thresholds compared to other participants in the study. We contend that such surprisingly high ripple scores are enabled by perception of spectral and temporal cues other than spectral density; when spectral density perception is severely limited by poor resolution, alternative strategies might be even *more likely* to arise more strongly.

Lack of correspondence to speech acoustics

Spectrally- and spectro-temporally modulated ripples are attractive partly because they contain broadband spectral modulations, as does speech. The degree of similarity between commonly-used ripple stimuli and actual speech sounds is quantifiable, and using such an approach we will next consider the extent to which typical ripple stimuli are representative of speech. Liu and Eddins (2008) calculated log-frequency spectral modulation spectra of vowels and illustrated such a spectrum for the vowel /æ/. Energy was concentrated at very low modulation rates, primarily below 0.5 peaks per octave, with a secondary peak at 1.5 peaks per octave. They found that

selectively transmitting higher-density spectral modulations resulted in more *errors* in vowel recognition, especially for back vowels.

Extending the goals of Liu and Eddins (2008), we have analyzed the log-frequency spectral modulation spectra of English phonemes to discover the primary modulations contained in them. Corroborating the conclusions of Liu and Eddins, we found that spectral modulations are confined primarily to the region below 1 component per octave. Figure 9 illustrates these spectral modulation spectra for selected phonemes, including the most frequently produced vowels and consonants in English (according to frequency data of the Hoosier Mental Lexicon; Nusbaum et al. 1984). Perhaps unsurprisingly, independent studies by Saoji et al (2009), Winn and Litovsky (2015) and Feng and Oxenham (2018) all suggest that speech recognition in CI listeners is driven more by perception of global spectral-shape properties rather than dense spectral components.

586

[[Figure 9 here]]

587

588 Fig. 9. Log-frequency modulation spectra for the most frequently-spoken vowels and
 589 consonants in English (excluding schwa, whose spectral characteristics are not stable
 590 across contexts). Greater strength is observed at very low frequency modulation rates.

591

592

593 Other CI processors

594 The illustrations in this article have focused on the speech processor frequency
 595 channels in the device manufactured by the Cochlear corporation, for the simple reasons
 596 that 1) it is the most familiar to the first author, 2) it would appear to offer the most
 597 generous input frequency sampling, as it contains 22 processing channels (setting aside
 598 the concept of virtual-channel selection in the Advanced Bionics device). This analysis
 599 was chosen as a simple first step rather than a statement of priority. The spectral

aliasing and modulation-depth saturation points, as well as the log-frequency modulation spectra, can be calculated for other processors as well. Code used to generate the figures is available in Supplemental Digital Content 1, which includes options to toggle between frequency-channel allocation for the Cochlear, Advanced Bionics and Med-El devices (SDC Figures 2 through 15). We note however, that the analysis for the Med-El device should be interpreted critically, as the analysis used here was designed to replicate FFT-based processing, but the Med-El device uses bell-shaped filterbanks. Further research could be done to determine the impact of the filter shape on the calculations offered here. With these limitations set aside for the moment, the analysis suggests that the spectral aliasing values for the Advanced Bionics device is about 2 RPO, and about 1.1 RPO for the Med-El device. The threshold of modulation depth saturation (i.e. smallest ripple whose period fits into the narrowest filter) is 2.02 for the Advanced Bionics device and 1.26 for the Med-El device.

[The role of experience / learning / cognition in the spectral ripple test](#)

One of the primary motivations of the spectral ripple test is to avoid the long adaptation process that characterizes a newly implanted patient's performance with real speech. It has been argued that spectral ripple discrimination should give stable performance metric independent of experience with the implant. Support for this mission comes from Drennan et al. (2010), who found no effects of learning when the ripple task was repeated on the same day, and from Won et al. (2007) who found no learning effects when the test was repeated on separate days. However, a study by de Jong et al. (2018) explored repeated measures over a longer period of time challenge the

stability of the spectral ripple measure; they identified significant learning effects after a 6-week experience in all ten CI participants tested. The authors suggested that earlier studies were affected by experience as well, calling attention to participants in a study by Drennan et al. (2015) of whom 20% showed improvement and 20% showed decline in scores.

Performance on ripple discrimination is potentially impacted by listener age, but results thus far are not clear. Using the temporally-modulated ripple test developed by Aronoff and Landsberger (2013), Kirby et al. (2015) found systematically improved performance across age for children with hearing loss aged 6 to 12. Horn et al. (2017) found no evidence for maturation of spectral resolution beyond early school-age in prelingually deaf implanted CI users, so it is possible that those age effects reflected changes in test-taking maturity or increased cognitive capacity. Extending this approach to children with CIs, Landsberger et al. (2017) did not find evidence of age effects in children (including non-significant correlations with age, age of implantation or years of experience), but did find systematically lower ripple perception for older adults. Ideally the spectral ripple test should probe peripheral auditory processing as purely as possible, and should thus be independent of factors like executive function, cognitive status, etc., so long as the task can at least be completed. However, Kirby et al. (2018) found that greater nonverbal intelligence and visual working memory were associated with better spectro-temporal ripple discrimination thresholds. In their study, verbal working memory, executive function, and language ability were not significantly correlated with ripple discrimination performance. It should be noted that there are other occasions where non-verbal intelligence has been correlated with psychophysical

tasks (Watson 1991; Hartley et al. 2000; Rammseyer & Brandler 2004; Hulslander et al. 2004), so to the extent that this is a concern, it is not limited just to spectral ripple tasks. We draw attention to this association here simply to address previous assertions that the spectral ripple test should in theory be immune to such correlations; it is not.

Apart from associations between ripple discrimination and non-auditory tasks, there are other internal correlations that are potentially problematic. de Jong et al. (2018) found significant, though fairly low, correlation between SMRT discrimination and scores for basic amplitude modulation detection, consistent with the interpretation that the SMRT does not purely measure spectral resolution but may also influenced by temporal cues, or a contribution cognitive skill to both measures. This is predictable based on the very nature of the SMRT, since the temporal fluctuations within any particular channel can alone be the cue sufficient for stimulus discrimination. de Jong et al. call attention to other earlier studies (Donaldson et al. 2011; Vermeire et al. 2010) that highlight the need to account for performance improvements that arise simply because of exposure to test stimuli.

What should we do instead?

The full development and validation of a new sensitive and robust test of spectral resolution is beyond the scope of this article, but it is worth highlighting some behavioral tests that hold promise to stand in for the conventional task of detecting phase inversion of increasingly-dense acoustic ripples sent through the CI processor. If the experimenter is committed to using non-linguistic spectral ripple stimuli, it could be favorable to maintain a low-density ripple and measure the threshold for detecting

spectral modulations versus a flat-spectrum signal (where detection of shallower spectral modulations is a sign of success), as done by Litvak et al. (2007). Using ripple stimuli with low density (0.5 or 1 RPO), Gifford et al. (2018) found significant correlations between spectral modulation depth detection and speech recognition (words, sentences, and sentences in noise) in a large group of adult CI listeners. Zhang et al. (2013) observed a very strong correlation ($r^2 = 0.81$) between low-density (1 RPO) ripple modulation depth and benefit from complementary acoustic information contralateral to the CI (i.e. bimodal hearing). Bernstein et al. (2016) extended the utility of modulation depth detection to users of hearing aids, with the clever extra step of controlling the phase-drift (i.e. temporal modulation) rate to equal 4 Hz, which approximately equals the rate of syllables in speech. Alternatively, if the experimenter views spectral modulation-depth discrimination as a mere substitution of potential single-channel intensity resolution instead of spectral resolution, one could maintain a low-density ripple and adaptively find the minimum *phase change* that is detectable by a listener. Figure 10 shows idealized stimulus spectra that illustrate these two options.

685

[[Figure 10 here]]

686

687 Fig. 10. Idealized stimuli showing progressive changes in spectral ripple phase shift
 688 (upper panel) or modulation depth (lower panel) while ripple density is held to 0.5 or 1
 689 RPO, respectively.

690

691 It has been claimed that temporally modulated ripples in particular could be a
 692 good probe to model the perception of acoustic properties that reflect speech (Chi et al.

1999; Smith et al. 2013). The spectro-temporally modulated ripple test developed by Aronoff and Landberger (2013) is relevant for this goal, as it has been used in numerous recent studies involving adults and children. It is not clear, however, whether the tested temporal modulations reflect the temporal modulation spectrum of speech. At the time of this writing, there has been no explicit attempt to represent speech acoustics in any direct way in static or temporally modulated ripples.

Another option an experimenter could pursue is to embrace the reality of aliased / distorted spectra, and deliberately present stimuli whose spectral modulations resemble those in speech, but which are still non-speech sounds. For example, one would take the spectral modulation spectra of speech but scramble the phase, and then reverse the process so that a sound is produced which has spectral modulations that are representative of speech, but which are likely not identifiable as a particular vowel or consonant. This is akin to the process of creating speech-shaped noise in the amplitude-spectrum domain, where the spectrum of speech could be phase-randomized to produce a new sound whose amplitude spectrum is equivalent to the original speech but without any intelligible features.

Previous work by the first author and colleagues has explored behavioral testing in CI listeners using parametrically manipulated cues in actual speech. As described in detail by Winn and Litovsky (2015), Winn et al. (2016) and Winn and Stilp (2019), speech is replete with acoustic properties that are eligible for experimental control. One would assume that perception of particular speech properties like formant transitions, spectral slope, peak frequencies, etc. would enable good speech recognition since they underly phonetic contrasts. Manipulation of these properties had been done in earlier

work using synthetic materials (Francis et al. 2000), but more recently also with modified natural speech (Winn et al. 2012; Winn & Litovsky 2015). Indeed, perception of specific acoustic-phonetic cue properties are more successful at predicting word recognition than non-linguistic counterparts such as spectral ripples in CI listeners (Winn et al. 2016; DiNino et al. submitted). While it is tempting to dismiss this trend as “speech scores predicting speech scores”, it is crucial to note that the predictive score is not speech recognition, but perception of change in a tightly-controlled component of a syllable, such as the second formant transition that helps a listener distinguish /b/ from /d/ for example. Using such a cue forces the experimenter to address speech acoustics within a realistic and phonetically plausible framework. It thus should ensure a more direct and interpretable relationship to the target signal (real speech).

The recommendation to consider speech-cue-based tasks instead of spectral ripples has advantages and disadvantages. On the positive side, tests contrasting /b/ and /d/ have been popular in many related scientific fields, and they are mentioned here, not by accident, but for at least two good reasons. First, this is a distinction of consonant place of articulation, which is by far the speech feature most often misperceived by CI listeners (Munson et al. 2003), because it is cued spectrally. Second, it is a phonetic contrast found in nearly all the world’s languages, including the top 20 most widely-used languages and every language with at least 50 million first-language speakers (Lewis et al. 2013), largely freeing the test from constraints of language-specificity. It is certainly feasible to imagine other speech-cue-based tests that could target specific frequency regions, such as low-frequency contrasts (e.g. F1 difference between “hope” and “hoop”), mid-frequency contrasts (e.g. /b/-/d/ or the F2 contrast

between “hid” and “hood”), or high frequency contrasts (e.g. spectral peaks distinguishing /k/ and /t/). Vowel contrasts are often language- and dialect-specific, but consonant contrasts are likely extensible across a very wide range of individuals. Phonetic perception can be measured electrophysiologically as well (Moberly et al. 2016), for individuals who cannot give behavioral responses (such as infants).

Despite potential advantages of speech-based categorization experiments, there are at some limitations as well. Categorization experiments are typically done with adults who can perform basic laboratory tasks; they can also be done in school-aged children (Nittrouer & Studdert-Kennedy 1987; Hazan & Barrett 2000), but typically need to be converted to a limited set that might lack the stimulus repetitions or task complexity of an adult task. However, recent developments in psychometric function estimation algorithms enable experimenters to obtain high-quality parameter fits with relatively few stimulus presentations (Shen & Richards 2012). More importantly, the outcome measure of psychometric function slope is a complicated one, driven not only by auditory perception but also by individual differences in perceptual gradiency (Kapnoula et al. 2017), attentional lapse rate (Wichmann & Hill 2001), and a host of other related cognitive factors. Although there have been correlations found between psychometric function slopes and word recognition in CI listeners (Moberly et al. 2014; Winn & Litovsky 2015; Winn et al. 2016), it is difficult to rule out non-auditory factors in auditory categorization. In short, using a psychometric function is *sensitive* to auditory acuity, but not *specific* to it.

CONCLUSIONS

Spectral resolution is an important aspect of cochlear implants worthy of investigation because of its role in speech perception. Spectral ripples are attractive stimuli to test spectral resolution because performance in ripple discrimination has been found to correlate with word recognition, music perception, and various aspects of CI processing. It is also thought to be free from non-spectral contaminating factors such as linguistic experience/acclimatization and cognitive influence. The analysis in the current article shows that there are at least two major flaws in the transmission of spectral ripples through cochlear implants: 1) spectral aliasing and 2) modulation depth neutralization, both of which render the stimulus different from what the experimenter intends. Most importantly, these unintended distortions are nonlinear in nature, rendering parametric ordering of ripples above a certain threshold mathematically unsound. For the cochlear device, the critical limit for spectral aliasing is around 2.1 RPO, and the critical limit for modulation saturation is 3.46 RPO. Phase changes do not solve either problem. Additionally, the spectral modulations at these densities or higher fail to represent the modulations found in actual speech sounds. When we consider only data that fall in the interpretable range of ripple densities, correlations with speech recognition improve. Experience, age and cognitive factors arguably play a role in spectral ripple discrimination, as suggested by some recent work. Various alternative testing approaches exist, some of which have already been used successfully. We conclude that the practice of identifying the highest spectral density at which ripple phase inversion can be detected be abandoned because (1) it lacks the well-behaved mathematical properties a psychophysical experimenter requires, and (2) high spectral

densities are a poor representation of ecologically-relevant cues. We recommend that alternative strategies be taken with the specific consideration of real speech acoustics in their development.

ACKNOWLEDGMENTS

We are grateful Andrew Oxenham, David Landsberger, and Justin Aronoff for thoughtful and critical discussions of this project. Kate Teece contributed to the gathering of information of this manuscript. Funding was received by a grant from the NIH (to M.B.W.) NIH-NIDCD R0301439 and from NIH Auditory Neuroscience Training Grant (to G.E.O) 2T32DC005361-16. Portions of this work were presented at the Conference on Implantable Auditory Prostheses, Lake Tahoe, CA, July 2017, and the Acoustical Society of America, Louisville, KY, May 2019.

REFERENCES

796

797 Anderson, E., Nelson, D., Kreft, H., Nelson, P., Oxenham, A. (2011) Comparing spatial
 798 tuning curves, spectral ripple resolution, and speech perception in cochlear
 799 implant users. *J Acoust Soc Am*, 130, 364–375

800 Aronoff, J., & Landsberger, D. (2013). The development of a modified spectral ripple
 801 test. *J Acoust Soc Am*, 134, EL217–EL222.

802 Azadpour, M., and McKay, C. M. (2012). A psychophysical method for measuring spatial
 803 resolution in cochlear implants. *J Assoc Res Otolaryngol*, 13, 145–157.

804 Bernstein, J., Danielson, H., Hällgren, M., Stenfelt, S., Rönnberg, J., Lunner, T. (2016).
 805 Spectrotemporal modulation sensitivity as a predictor of speech-reception
 806 performance in noise with hearing aids. *Trends Hear*, 20, 1-17.

807 Blamey, P., Artieres, F., Başkent, D., et al. (2013). Factors affecting auditory
 808 performance of postlinguistically deaf adults using cochlear implants: An update
 809 with 2251 patients. *Audiol Neurotol*, 18(1), 36–47.

810 Chi, T. Gao, Y., Guyton, M., Ru, P., Shamma, S. (1999). Spectrotemporal modulation
 811 transfer functions and speech intelligibility. *J Acoust Soc Am*, 106, 2719–32.

812 de Jong, M., Briare, J., Frijns, J. (2018). Learning Effects in Psychophysical Tests of
 813 Spectral and Temporal Resolution. *Ear Hear*, 39, 475–481

814 DiNino M., Arenberg, J., Duchon, A., Winn, M. (2019). Effects of age and cochlear
 815 implantation on spectrally cued speech categorization. *Submitted to J Acoust Soc*
 816 *Am*.

- 817 Donaldson, G., Dawson, P., Borden, L. (2011). Within-subjects comparison of the HiRes
818 and Fidelity120 speech processing strategies: Speech perception and its relation to
819 place-pitch sensitivity. *Ear Hear*, 32, 238–250.
- 820 Drennan, W. R., Won, J. H., Nie, K., Jameyson, E., Rubinstein J. (2010). Sensitivity of
821 psychophysical measures to signal processor modifications in cochlear implant
822 users. *Hear Res*, 262, 1–8.
- 823 Drennan, W., Anderson, E., Won, J. H., Rubinstein, J. (2014). Validation of a clinical
824 assessment of spectral-ripple resolution for cochlear implant users. *Ear Hear*, 35,
825 e92–e98.
- 826 Drennan, W., Won, J. H., Timme, A., Rubinstein, J. (2015). Nonlinguistic outcome
827 measures in adult cochlear implant users over the first year of implantation. *Ear*
828 *Hear*, 37, 354–364.
- 829 Feng, L. & Oxenham, A. (2018). Effects of spectral resolution on spectral contrast effects
830 in cochlear-implant users. *J Acoust Soc Am*, 143, EL468–EL473.
- 831 Francis, A., Baldwin, K., Nusbaum, H. (2000). Effects of training on attention to
832 acoustic cues. *Percept Psychophys*, 62, 1668–1680.
- 833 Friesen, L., Shannon, R., Başkent, D., et al. (2001). Speech recognition in noise as a
834 function of the number of spectral channels: Comparison of acoustic hearing and
835 cochlear implants. *J Acoust Soc Am*, 110, 1150–1163.
- 836 Hartley, D., Wright, B., Hogan, S., Moore, D. (2000). Age-related improvements in
837 auditory backward and simultaneous masking in 6-to 10-year-old children. *J*
838 *Speech Lang Hear Res*, 43, 1402–1415.

- 839 Hazan, V. & Barrett, S. (2000). The development of phonemic categorization in children
840 aged 6–12. *J Phonetics*, 28(4), 377-396.
- 841 Henry, B., Turner, C., Behrens, A. (2005). Spectral peak resolution and speech
842 recognition in quiet: Normal hearing, hearing impaired, and cochlear implant
843 listeners. *J Acoust Soc Am*, 118, 1111–1121.
- 844 Horn, D., Won, J.H., Rubinstein, J., Werner, L. (2016). Spectral ripple discrimination in
845 normal-hearing infants. *Ear Hear*, 38, 212-222.
- 846 Horn, D., Dudley, D., Dedhia, K., et al. (2017). Effects of age and hearing mechanism on
847 spectral resolution in normal hearing and cochlear-implanted listeners. *J Acoust*
848 *Soc Am*, 141, 613.
- 849 Hulslander, J., Talcott, J., Witton, C., et al. (2004). Sensory processing, reading, IQ, and
850 attention. *J Exper Child Psych*, 88, 274-295.
- 851 Jones, G., Won, J. H., Drennan, W., Rubinstein, J. (2013). Relationship between
852 channel interaction and spectral-ripple discrimination in cochlear implant users. *J*
853 *Acoust Soc Am*, 133, 425–433.
- 854 Kapnoula, E., Winn, E., Kong, E., Edwards, J., & McMurray, B. (2017). Evaluating the
855 sources and functions of gradiency in phoneme categorization: An individual
856 differences approach. *J Exper Psych: Human Percept Perf*, 43, 1594-1611
- 857 Kirby, B., Browning, J., Brennan, M., et al. (2015). Spectro-temporal modulation
858 detection in children. *J Acoust Soc Am*, 138, EL465–EL468.

- 859 Landsberger, D., Padilla, M., Martinez, A., Eisenberg, L. (2017). Spectral-temporal
860 modulated ripple discrimination by children with cochlear implants. *Ear Hear*, 39,
861 60-68.
- 862 Lawler, M., Yu, J., Aronoff, J. (2017). Comparison of the spectral-temporally modulated
863 ripple test with the Arizona Biomedical Institute sentence test in cochlear implant
864 users. *Ear Hear*, 38, 760-766.
- 865 Lewis, M., Simons, G., Fennig, C. (2013). *Ethnologue: Languages of the World*, 17th ed.
866 (SIL International, Dallas, TX). <http://www.ethnologue.com> (retrieved on August
867 13, 2014).
- 868 Litvak, L., Spahr, A., Saoji, A., Fridman, G. (2007). Relationship between perception of
869 spectral ripple and speech recognition in cochlear implant and vocoder listeners. *J*
870 *Acoust Soc Am*, 122, 982-991.
- 871 Liu, C. & Eddins, D. (2008) Effects of spectral modulation filtering on vowel
872 identification. *J Acoust Soc Am*, 124, 1704-1715.
- 873 Lopez Valdes, A., McLaughlin, M., Viani, L., Walshe, P., Smith, J., Zeng, F.-G., Reilly, R.
874 (2014). Objective assessment of spectral ripple discrimination in cochlear implant
875 listeners using cortical evoked responses to an oddball paradigm. *PLoS One*, 9,
876 e90044.
- 877 Moberly, A., Lowenstein, J., Tarr, E., et al. (2014). Do adults with cochlear implants rely
878 on different acoustic cues for phoneme perception than adults with normal
879 hearing? *J Speech Lang Hear Res*, 57, 566-582.

- 880 Moberly, A., Bhat, J., Shahin, A. (2016). Acoustic cue weighting by adults with cochlear
881 implants: A mismatch negativity study. *Ear Hear*, 37, 465-472.
- 882 Munson, B., Donaldson, G., Allen, S., Collison, E., and Nelson, D. (2003). "Patterns of
883 phoneme misperceptions by individual with cochlear implants. *J Acoust Soc Am*,
884 113, 925-935.
- 885 Nittrouer, S., & Studdert-Kennedy, M. (1987). The role of coarticulatory effects in the
886 perception of fricatives by children and adults. *J Speech Lang Hear Res*, 30, 319-
887 329.
- 888 Nusbaum, H., Pisoni, D., and Davis, C. (1984) Sizing up the Hoosier mental lexicon:
889 Measuring the familiarity of 20,000 words. *Research on Speech Perception*,
890 *Progress Report 10*, (pp. 357-376), Bloomington, IN: Speech Research Laboratory,
891 Indiana University.
- 892 O'Neill, E., Kreft, H., Oxenham, A. (2019). Speech perception with spectrally non-
893 overlapping m maskers as measure of spectral resolution in cochlear implant users.
894 *J Assoc Res Otolaryngol*, 20, 151-167.
- 895 Rammsayer, T. & Brandler, S. (2002). On the relationship between general fluid
896 intelligence and psychophysical indicators of temporal resolution in the brain. *J*
897 *Res Personality*, 36, 507-530.
- 898 Saoji, A., Litvak, L., Spahr, A., Eddins, D. (2009). Spectral modulation detection and
899 vowel and consonant identifications in cochlear implant listeners. *J Acoust Soc*
900 *Am*, 126, 955-958

- 901 Scheperle, R. & Abbas, P. (2015). Relationships among peripheral and central
902 electrophysiological measures of spatial and spectral selectivity and speech
903 perception in cochlear implant users. *Ear Hear*, 36, 441-453.
- 904 Shen, Y. & Richards, V. (2012). A maximum-likelihood procedure for estimating
905 psychometric functions: Thresholds, slopes, and lapses of attention. *J Acoust Soc*
906 *Am*, 132, 957-967.
- 907 Smith, Z., Parkinson, W., Long, C. (2013). Multipolar current focusing increases spectral
908 resolution in cochlear implants. *Proc Ann Intl Conf IEEE Eng Med Biol Soc*, Osaka,
909 Japan, 2796-2799.
- 910 Vermeire, K., Punte, A., Van de Heyning, P. (2010). Better speech recognition in noise
911 with the fine structure processing coding strategy. *ORL J Otorhinolaryngol Relat*
912 *Spec*, 72, 305–311.
- 913 Watson, B. (1991). Some relationships between intelligence and auditory
914 discrimination. *J Speech Lang Hear Res*, 34, 621-627.
- 915 Wichmann, F. & Hill, N. (2001). The psychometric function: I. Fitting, sampling, and
916 goodness of fit. *Percept Psychophys*, 63, 1293-1313.
- 917 Winn, M., & Litovsky, R. (2015). Using speech sounds to test functional spectral
918 resolution in listeners with cochlear implants. *J Acoust Soc Am*, 137, 1430–1442.
- 919 Winn, M. & Stilp, C. (2019) Phonetics and the auditory system. In W. Katz & P.
920 Assmann (Eds.) *The Routledge Handbook of Phonetics* (pp. 164-192). New York,
921 NY: Routledge.

- 922 Winn, M., Chatterjee, M., Idsardi, W. (2012). The use of acoustic cues for phonetic
923 identification: Effects of spectral degradation and electric hearing. *J Acoust Soc*
924 *Am*, 131, 1465–1479.
- 925 Won, J. H., Drennan, W., Rubinstein, J. (2007). Spectral-ripple resolution correlates
926 with speech reception in noise in cochlear implant users. *J Assoc Res Otolaryngol*,
927 8, 384–392.
- 928 Won, J. H., Drennan, W., Kang, R., et al. (2010). Psychoacoustic abilities associated
929 with music perception in cochlear implant users. *Ear Hear*, 31, 796–805.
- 930 Won, J. H., Clinard, C., Kwon, S., et al. (2011a). Relationship between behavioral and
931 physiological spectral-ripple discrimination. *J Assoc Res Otolaryngol*, 12, 375–
932 393.
- 933 Won, J. H., Jones, G., Drennan, W., et al. (2011b). Evidence of across-channel
934 processing for spectral-ripple discrimination in cochlear implant listeners. *J*
935 *Acoust Soc Am*, 130, 2088–2097.
- 936 Won, J. H., Humphrey, E., Yeager, K., et al. (2014). Relationship among the physiologic
937 channel interactions, spectral-ripple discrimination, and vowel identification in
938 cochlear implant users. *J Acoust Soc Am*, 136, 2714–2725.
- 939 Zhang, T., Spahr, A., Dorman, M., Saoji, A. (2013). Relationship between auditory
940 function of nonimplanted ears and bimodal benefit. *Ear Hear*, 34, 133-141.

Figure Captions

Fig 1. Spectral ripples analyzed through a synthetic CI processor. The gray sinusoidal spectra are idealized acoustic ripples. Banded rectangular shapes are the simulated electric channel output, consisting of the average energy of the acoustic energy within each channel.

Fig 2. Extended sampling of spectral ripple densities following the same pattern as Figure 1. Around 2.1 ripples per octave (RPO), the simulated electric output no longer faithfully encodes the intended spectral modulations.

Fig 3. Log-frequency spectral modulation spectra indicating the presence of octave-spaced spectral modulations for a variety of densities of ripples analyzed through the synthetic CI processor. Vertical gray lines indicate the “target” location of the peak, corresponding to the peak in the idealized acoustic input. The left panel shows raw spectral modulation power, including virtually absent modulation strength for high-density ripples whose modulation depth is neutralized. The right panel shows normalized spectral modulation power, where the peak – at whatever raw power – is scaled to 1. In the right panel, the absolute peak is marked with a closed black point, and any local peaks that are at least half the strength of the maximum peak are marked with open points.

Fig 4. Peaks in log-frequency spectral modulation density (i.e. presence of different RPOs) as a function of intended RPO input to an idealized CI processor. Black peaks indicate maximum spectral modulation strength, and open circles indicate any local peaks that are at least half the strength of the maximum peak; these points correspond to the points of similar marking in Figure 3). Size of the point scales with raw modulation power.

Fig. 5. Phase inversion of simulated electric output of spectral ripple stimuli. Gray space between phase-inverted patterns is the difference in channel-specific output energy, and the darkest shaded region indicates the channel with greatest change in energy across phase inversion.

Fig 6. Change in potential channel output energy across ripple density. Each small gray/black point reflects the change in energy resulting from phase inversion (illustrated on Fig. 5). For each ripple density between 0.5 and 8 (spaced by 0.5 RPO), data were collected for each channel, across all phase pairs sampled at 10-degree intervals (e.g. 0-180, 10-190, 20-200, etc.). Large blue points reflect the median change across all channel and all sampled phases. Red triangles represent the maximum achievable change in modulation depth for any single channel across all sampled phases. Encircled numbers indicate the channel where that maximum depth emerged.

Fig. 7. Bandwidth of each CI processor channel in linear space (upper panel) and octave space (lower panel). Channel 9 has the narrowest octave bandwidth, and is highlighted because of its role in producing the maximum achievable change in spectral modulation depth resulting from phase inversion (see Fig. 7).

Fig 8. Scatterplots showing relationship between spectral ripple discrimination and word recognition (top row) or logistic cue coefficient (i.e. psychometric function slope) for consonant formant transition perception (bottom row). Within each row, different panels reflect different approaches to handling RPO data above 2.56, including maintaining it, dropping it on a per-session basis, dropping listeners who average above 2.56, reassigning any score higher than 2.56 to be 2.56 on a session basis or on a participant basis (i.e. Winsorizing).

Fig. 9. Log-frequency modulation spectra for the most frequently-spoken vowels and consonants in English (excluding schwa, whose spectral characteristics are not stable across contexts). Greater strength is observed at very low frequency modulation rates.

Fig. 10. Idealized stimuli showing progressive changes in spectral ripple phase shift (upper panel) or modulation depth (lower panel) while ripple density is held to 0.5 or 1 RPO, respectively.

1004

Supplemental Digital Content

1005 1: R code used to generate figures in this manuscript, and to extend the analysis to other

1006 cochlear implant manufacturers

1007 2: Re-drawn figures from the manuscript using other cochlear implant frequency-

1008 electrode filters

1009

1010