Title: Envelope reconstruction of speech and music highlights unique tracking of speech at low frequencies

Authors: Nathaniel J Zuk^{1,2,4}, Jeremy W Murphy^{1,*}, Richard B Reilly^{1,2,3}, Edmund C Lalor^{1,4,5}

Affiliations: 1) Department of Electronic & Electrical Engineering, Trinity College Dublin, Dublin, Ireland; 2) Trinity College Institute of Neuroscience, Trinity College Dublin, Dublin, Ireland; 3) Trinity Centre for Bioengineering, Trinity College Dublin, Dublin, Ireland; 4) Department of Biomedical Engineering, University of Rochester, Rochester, NY, USA; 5) Del Monte Institute of Neuroscience, University of Rochester Medical Center, Rochester, NY, USA

* Jeremy Murphy is now located at the Department of Neuroscience, Brown University, Providence, RI, USA.

Abstract:

The human brain tracks amplitude fluctuations of both speech and music, which reflects acoustic processing in addition to the processing of higher-order features and one's cognitive state. Comparing neural tracking of speech and music envelopes can elucidate stimulus-general mechanisms, but direct comparisons are confounded by differences in their envelope spectra. Here, we use a novel method of frequency-constrained reconstruction of stimulus envelopes using EEG recorded during passive listening. We expected to see music reconstruction match speech in a narrow range of frequencies, but instead we found that speech was reconstructed better than music for all frequencies we examined. Additionally, speech envelope tracking at low frequencies, below 1 Hz, was uniquely associated with increased weighting over parietal channels. Our results highlight the importance of low-frequency speech tracking and its origin from speech-specific processing in the brain.

Introduction:

Sound carries information in speech and music across a wide range of time scales, and theoretically the brain must have some parsing mechanism that operates on the sound prior to extracting information. There is considerable debate about how speech and music are parsed by the brain, but it is fairly clear that the sound's rhythmic structure enables speech decoding (Ghitza and Greenberg, 2009), facilitates pattern recognition (Rajendran et al., 2016), and affects discrimination of orthogonal acoustic attributes (Geiser et al., 2012; Chang et al., 2019). Fluctuations in the envelopes of speech and music capture rhythmic structure, so if the brain is responding to this structure, brain activity should also be tracking their envelopes.

Indeed, past research has shown that brain signals track envelopes for both speech (for reviews see: Ding and Simon, 2014; Obleser and Kayser, 2019) and music (Doelling and Poeppel, 2015; Cantisani et al., 2019; Di Liberto et al., 2020b). Furthermore, the strength of envelope tracking can be affected by the listener's locus of attention (Power et al., 2011; Ding and Simon, 2012; Zion Golumbic et al., 2013; O'Sullivan et al., 2015), level of comprehension (Di Liberto et al., 2018; Etard and Reichenbach, 2019), and perceived musical meter (Iversen et al., 2009; Nozaradan et al., 2011; Fujioka et al., 2015). Thus, rather than simply representing a response to the acoustical structure, envelope tracking may include neural mechanisms that respond to more complex, non-acoustic features which correlate with the envelope but don't represent the stimulus envelope per se (Di Liberto et al., 2015; Ding et al., 2015; Baltzell et al., 2019; Oganian and Chang, 2019; Di Liberto et al., 2020b).

The mechanisms that produce this envelope tracking are still hotly debated, however, and the various interpretations typically stem from the method of analysis used in each study. One theory posits that neural oscillations are tuned to particular frequencies relevant for parsing the speech or music into these discrete units (Giraud and Poeppel, 2012; Large et al., 2015). For speech, the phase consistency of the neural responses observed at theta band (4-8 Hz) and delta band (< 4 Hz) correspond roughly to syllabic and phrase rates respectively. In another theory, the neural tracking of the envelope measured with EEG and MEG is represented by its convolution with an evoked response (Ding and Simon, 2009; Lalor and Foxe, 2010), which has been shown to respond particularly to syllable onsets in speech (Oganian and Chang, 2019; Kojima et al., 2020). Yet it is not clear if the two mechanisms are irreconcilably different. Phase consistency can come about if characteristic evoked responses occur at transient, regular times in the sound

(Kojima et al., 2020). Likewise, an evoked response model that captures envelope tracking acts like a filter (de Cheveigné et al., 2018) and, via regularization, focuses on lower frequencies that overlap with the frequencies exhibiting the largest phase consistency. Evoked-response-type models have been used to quantify and compare the strength of neural tracking in theta and delta bands (Di Liberto et al., 2015; Etard and Reichenbach, 2019; Donhauser and Baillet, 2020), which also suggests that, in spite of the differing theories for the mechanisms of neural tracking, the distinct results across studies could be related. This can be examined further by computing models responsible for tracking various frequency bands of the envelope and looking at the corresponding evoked responses captured by those models.

Neural tracking in the theta band appears to be strongest for speech and music compared to other natural sounds (Zuk et al., 2020), yet it remains to be determined if the underlying mechanisms of envelope tracking are distinct between speech and music. A model-based approach can test this directly by examining how stimulus-response models differ across stimulus types. Unfortunately, comparing model-based quantifications of speech and music neural tracking is complicated by the different spectral profiles between the two types of sounds. Additionally, neural tracking is also affected by feedback mechanisms relating to non-acoustic features (Ding et al., 2015; Baltzell et al., 2019; Broderick et al., 2019) or one's cognitive state (Nozaradan et al., 2011; Power et al., 2011; Mesgarani and Chang, 2012; O'Sullivan et al., 2015, 2019). But to compare these "higher-level" effects between speech and music in naturalistic stimuli it is useful to first isolate differences in neural tracking that cannot be explained by envelope spectral differences alone.

In this study, we apply a variant on linear modeling to focus analysis on a narrow range of frequency bands. We then compare reconstruction accuracies for various shifts of the frequency range in order to understand precisely how EEG tracks the signal envelope of speech and music during passive listening. We hypothesized that speech and music would have different frequency bands at which their envelopes are tracked best, and by using linear modeling, we could observe the spatiotemporal patterns of the responses that are most relevant for envelope tracking in each respective frequency region.

Results:

In our experiment, 16 subjects (7 female; ages 18-44, median 22) passively listened to four types of stimuli, interleaved across trials: rock songs in full, excerpts from classical songs, the vocals isolated from the rock songs, and several segments of an audiobook (called “speech” stimuli in the figures). Each subject listened to 6-7 of the 10 possible stimuli available for each stimulus type (see Table S1 for details on the stimuli). We will first step through the model construction and optimization procedure. Then, we will show the differences in neural tracking we observed and examine the spatiotemporal EEG responses that gave rise to these differences in neural tracking.

Envelope reconstruction with PCA and basis spline transformation

Our initial goal was to quantify how well EEG tracked the music envelopes in comparison to speech. Previous research had shown that EEG tracks envelope fluctuations in speech and music particularly above 2 Hz (Luo and Poeppel, 2007; Doelling and Poeppel, 2015; Kaneshiro et al., 2020). Based on this prior work, we expected that music may track speech comparably when reconstructions are focused on a narrower bandwidth of frequencies above 2 Hz. Additionally, speech and vocals might show low frequency tracking of envelopes, below 1 Hz, where the energy in classical and rock envelopes is comparatively low (Figure 1a, b; see also Figure S1). On the other hand, if we assume that the neural signal in the EEG is directly proportional to the dB

Figure 1: (A) 20 seconds of example dB envelopes for each stimulus type. (B) Power spectra of the dB envelopes using Welch’s method with a 16 s Hamming window and half-overlap. Lines indicate the median across stimuli of each type, and shaded regions indicate 95% quantiles of the distribution of 1000 bootstrapped median values.

envelope, we might expect speech tracking to be largest across all modulation frequencies, followed by vocals, then classical, and then rock music (Figure 1b, also Figure S1). With this in mind, we decided to use an approach to envelope reconstruction that focused on narrow bands of modulation frequencies. If envelope tracking for speech is better (based on reconstruction accuracy) than the music stimuli (see Figure 7c), is it because of the relative variance in the modulations in the dB envelope, or is it because of neural tracking at specific frequencies? (see also Figure 7 in the Methods)

We used a method of modeling that allowed us to focus on envelope tracking in narrow frequency bands. This required us to control for the spatial and temporal correlations, or “collinearity”, inherent in EEG data (Figure 2a). First, the moving average of both the stimulus envelope and EEG, using a window size equal to the maximum EEG delay in the model, was removed in order to remove effects of neural tracking outside of the model’s spectral range (see

Optimizing model to
Natural Speech dataset:

Figure 2: (A) Diagram of the stages for fitting the PCA & spline model. Throughout this study, the model is trained on all trials with one left out and tested on the left-out trial. (B) The difference in reconstruction accuracy (Pearson's r) between the PCA & spline method and a standard approach based on regularized linear regression for each of the left-out trials was then examined; negative values mean that the standard approach performs better. Error bars show the interquartile range of the reconstruction accuracy differences. Of all hyperparameter pairs, only 64 PCs and 32 Hz splines exhibited performance that was no different than the standard approach. (C) The combined effect of removing the moving average and using basis splines restricts the frequency content of the PCA-based reconstruction to a three-octave range; for a 500 ms window, this is restricted from 2-16 Hz.

Figure 7c in Methods). Next, the EEG data were spatially transformed into orthogonal components using principal components analysis (PCA) in order to remove spatial collinearity in the EEG. Lastly, before fitting the model, the time-delayed EEG signal was transformed into a basis of cubic splines. Cubic splines are smooth functions constructed based on a sequence of knots. By collating these splines across the range of delays, the splines can be used to interpolate any delayed EEG with a smoother, lower-frequency representation. The number of basis splines can be equivalently defined by the sampling frequency of the evenly-distributed knots, which also relates to the lowpass cutoff frequency of the splines (see legend in Figure 2b). Overall, for a specific window size, this model contains two hyperparameters: 1) the number of principal components to retain, and 2) the number of basis splines within the model window size.

To optimize these hyperparameters, we fitted the model using the Natural Speech dataset which is freely available online (Broderick et al., 2018). In this dataset, subjects ($N = 19$) listened to 20 trials of an audiobook that were approximately three minutes long. In order to identify the optimal hyperparameters, we compared the PCA & spline model performance to a regularized model using ridge regression, which is commonly used for speech envelope reconstruction (Ding and Simon, 2012; Crosse et al., 2015; Wong et al., 2018; Das et al., 2020). Specifically, we compared the performance of the two models on each trial by training and testing the models using a leave-one-trial-out procedure. The model containing 64 principal components and 19 splines (or equivalently a spline knot sampling rate of 32 Hz) exhibited the best performance and was the only hyperparameter pair whose performance was not significantly different than the model with regularization (Figure 2b) (Wilcoxon rank-sum test, $p > 0.05$). The optimal sampling of spline knots constrained the frequency content of the time-delayed EEG to 16 Hz, so combined with the removal of the moving average, the optimal PCA & spline model constrained to a three-octave range (Figure 2c).

Speech envelope tracking is better than music at all modulation frequencies

Using the number of principal components and splines optimized on the Natural Speech dataset, we computed the reconstruction accuracies for each trial in the current dataset, separately for each stimulus type, using a leave-one-out procedure (Figure 3a). We analyzed neural tracking of the envelope across a wide range of frequencies by varying the window size of the model

Figure 3: (A) For each stimulus type (6-7 trials per subject) the model was iteratively fit to all trials with one left out and tested on the left-out trial. In order to get a null distribution of reconstruction accuracies, we repeated this procedure, leaving out one trial at a time, after randomly circularly shifting the envelopes in each trial by the same amount. This was repeated 50 times for each stimulus type. (B) Schematic of our expectation for how reconstruction accuracy varies with frequency. We expected that EEG may be tracking a particular frequency range of the stimulus envelope. This could be identified by varying the range of the three-octave model bandwidth. The reconstruction accuracy increases from chance as the model bandwidth overlaps the relevant frequency range, and plateaus when the model bandwidth is fully contained within the relevant frequency range. (C) As lower frequencies are introduced into the stimulus envelope, the variance of the null distribution increases. Because of this, we z-scored the true trial-by-trial reconstruction accuracies relative to the null distribution to ease cross-frequency comparisons. (D) Shown are the median reconstruction accuracies across subjects and trials. Shaded regions show the 95% quantiles of the distribution of 1000 median values calculated using bootstrap resampling with replacement. Thicker lines indicate frequency ranges where median z-scored reconstruction accuracies were significantly greater than zero (one-tailed Wilcoxon signed-rank test with Bonferroni correction for 40 comparisons, $p < 0.001$). (E) Throughout the frequency range tracked by speech, speech reconstructions were significantly better than all other musical stimuli, with a difference peaking in the 0.5-4 Hz range. Thick lines indicate differences in reconstruction accuracy that are significantly greater than zero (two-tailed permutation test with Bonferroni correction for 40 comparisons, $p < 0.001$).

between 31.25 ms and 16 s, corresponding to an upper frequency range of 32-256 Hz and a lower range of 0.0625 – 0.5 Hz respectively. By varying the frequency content of the EEG involved in reconstruction while simultaneously controlling the low frequency cutoff of the envelope, we expected to see an increase in reconstruction accuracy as the bandwidth of the model encompassed the relevant frequency range for envelope tracking for the stimulus (Figure 3b). We can quantify the upper and lower cutoff frequencies of envelope tracking based on the skirts of the reconstruction accuracy plots, which should vary over a three-octave range. However, the reconstruction accuracy needed to be corrected by the variance of the distribution corresponding to chance performance; as the model includes more low frequency content, the variance in the null distribution also increases (Figure 3c). To control for this, we normalized the trial-by-trial reconstruction accuracy with respect to a distribution of null values computed by circularly shifting the envelopes relative to the EEG (Figure 3a, c). In this way, reconstruction accuracies were “z-scored” relative to the null distribution.

Overall, we found that all stimuli exhibited above-chance reconstruction accuracies, but the ranges exhibiting this were slightly different. As expected, the trends in the reconstruction accuracy curves increased from chance to near-peak reconstruction accuracy over a three-octave change in the upper cutoff frequency, indicating that EEG tracks the rock envelopes above 2 Hz, vocals and classical above 1 Hz, and speech above 0.5 Hz (Figure 3d). Reconstruction accuracies for all stimuli were above chance up to the highest range we could examine for the 512 Hz recorded sampling rate of the EEG. There is considerable evidence of cortical and brainstem tracking of fluctuations at the fundamental frequency of the stimulus pitch (Forte et al., 2017; Maddox and Lee, 2018), which might explain the above-chance reconstructions at very high modulation frequencies. Additionally, the ranges of peak reconstruction varied across stimulus types: while

the music stimuli had peak reconstructions within 2 – 32 Hz, speech showed peak reconstruction at a lower frequency range. We thought that the music reconstruction accuracies might be comparable to speech above 2 Hz, but contrary to our expectations, frequency tracking for speech was better throughout the range of frequencies at which performance was above chance (Figure 3e). The difference in reconstruction accuracy between speech and the music stimuli increased with lower frequencies, peaking at the 0.5-4 Hz range (for z-scored reconstruction accuracies for all subjects, see Figure S2).

We used the Natural Speech dataset to validate that this trend of reconstruction accuracies as a function of modulation frequency was not specific to the data shown here (Figure S3). The Natural Speech data also showed decreasing reconstruction accuracies for models with frequencies below 0.5 Hz (as in Figure 3d), but z-scored reconstruction accuracies were still above zero at the lowest frequency range we used in this study. This suggests that the limit found in Figure 3d is a consequence of the limited amount of speech data available in this dataset (6-7 trials here compared to 20 trials in Natural Speech), and speech may be weakly tracked at even lower frequencies below our lowest frequency range of 0.0625-0.5 Hz).

Because the model hyperparameters were optimized to speech, we considered that this might explain the high reconstruction accuracies for speech. We repeated this analysis using optimal hyperparameters for the music stimuli, but after doing so speech reconstruction still outperformed music (Figure S4). Another potential reason was that the speech envelopes had lower cross-trial spectral variance than the music envelopes (see Figure 2b); the speech stimuli came from the same audiobook, but the music stimuli came from different songs with different tempos. To control for this possibility, we repeated the envelope reconstruction analysis by fitting and testing models on randomly sampled data within each trial, which maximized the possibility of the testing data containing the same spectrum as the training data (Figure S5a). We found that within-trial reconstruction accuracies were still significantly larger for speech than the other stimuli across all measured frequencies (Figure S5b,c).

Together, these results show that envelope reconstruction better captures the neural representation of speech compared to music. The differences in envelope reconstruction performance cannot be explained by spectral variability across trials or by the model optimization procedure. Our results also support the possibility that speech envelope tracking is better due to the relatively larger variability in the speech envelope (Figure 1b).

Speech envelope tracking at low frequencies is uniquely associated with increased weighting of parietal electrodes

Theoretically, the envelope reconstruction model weights spatiotemporal components of the EEG response that most strongly track the stimulus envelope. To understand how these weights correspond to activity evoked by the envelope fluctuations, we inverted the reconstruction models and transformed them from splines and principal components into delays and EEG channels (see Methods). The EEG response shows temporally consistent peaks and troughs across frequency ranges which vary in magnitude depending on the frequency range represented, indicating an overlap in the spatiotemporal neural activity that is being picked up by each model (Figure 4a).

The temporal variation in the weights are generally consistent across stimulus types, which appears to be particularly true at low frequencies (< 4 Hz) (Figure 4b). Except for increased weighting for speech in channels over posterior scalp, the spatial distribution of the weights was also fairly consistent (Figure 4d). At higher frequencies (4-32 Hz), however, the models appear to be of different magnitudes and potentially phase-shifted relative to each other (Figure 4b, c; see also Figure S6 for individual model weights for each subject). For example, the speech and classical responses are well aligned, but the primary peak for vocals happens about 20 ms earlier and the secondary peak for rock happens 20 ms later. To what extent do these differences affect neural tracking of the envelopes for the different stimuli? Additionally, speech tracking is better at low frequencies (0.5 – 4 Hz) than music tracking, but this could be a consequence of differences in variance in the dB envelope at those frequencies (Figure 1b). If so, then a common mechanism would produce comparatively strong neural tracking for speech than music. Is the speech tracking at low frequencies a result of common processing across stimulus types?

To address this question, we reanalyzed the EEG data using envelope reconstruction models trained on all stimulus types and then tested on each trial individually. These stimulus-general models would necessarily capture consistent trends across all stimulus types relevant to envelope tracking. If any trends observed in the stimulus-specific models were unique to tracking the envelope of the respective stimulus, we expected to see higher reconstruction accuracies for the stimulus-specific model than the stimulus-general model. Yet contrary to our expectations for the faster frequencies, speech reconstruction using the stimulus-specific model performed no better

Figure 4: (A) Model weights, median across subjects and averaged across channels. Models are color-coded based on their frequency range (see to the right of the plots). The model of the range 0.0625-0.5 Hz was excluded because none of the stimuli exhibited significant neural tracking in this range, and the large values for the weights obscured the trends in the other models. (B) Mean and standard error across subjects of model weights for two frequency ranges (see Figure S4 for the models for individual subjects). (C) and (D) show the topographies of the model weights averaged over the range of delays corresponding to peaks and troughs in the 4-32 Hz and 0.5-4 Hz models respectively. Note that the range of delays vary across stimulus types in the 4-32 Hz model in order to capture similar peaks and troughs.

than the stimulus-general model, and for all other stimuli the stimulus-specific model performed worse (Wilcoxon signed-rank test with Bonferroni correction for 36 comparisons, $p < 0.001$) (Figure 5a). In contrast, the stimulus-specific model performed better for speech reconstruction only at 1-8 Hz and 0.5-4 Hz (for reconstruction accuracies for individual subjects, see Figure S2). We were concerned that the large frontal negativity at 200-500 ms might be indicative of eyeblinks

(Figure 4d, 5c) (Talsma and Woldorff, 2005). We looked at the individual subject topographies and found 4/16 subjects with strong eyeblink-like topographies in their models, but after removing these subjects, only the differences in the 1-8 Hz range were no longer significant, and the topographies continued to show strong frontal activity (Figure S7).

On average, the stimulus-general models were topographically and temporally similar to the stimulus-specific models (Figure 5b, c, compare to Figure 4). To better understand what differences between the speech-specific models and stimulus-general models might have produced the improved reconstruction accuracies for speech, we chose to examine these differences trial by

Figure 5: (A) The difference between the trial-by-trial reconstruction accuracies using the stimulus-specific and stimulus-general models was then computed. Lines show the median reconstruction accuracy differences across subjects and trials. Shaded regions show the 95-percentile range of bootstrapped resampled median values. Significance values are based on a Wilcoxon signed-rank test with Bonferroni correction for 32 comparisons, *** $p < 0.001$, ** $p < 0.01$). The stimulus-specific speech model outperformed the stimulus-general model at 0.5-4 Hz and 1-8 Hz (blue text), while the other stimulus-specific models performed worse than the stimulus-general model for most frequency ranges (black text, solid lines at top). (B, C) The stimulus-general model was quite similar in its temporal (B) and spatial (C) pattern compared to the stimulus-specific models. (D) We assumed that stimulus-general model is a scaled and phase shifted version of each of the stimulus-specific models, so by circularly shifting and scaling the model we could quantify the difference between the model weights. The scaling was computed separately for each EEG channel, but we assumed the shift would be identical for all EEG channels. (E) R^2 fits of the scaled and shifted stimulus-general model to each stimulus-specific model on a trial-by-trial basis, based on the summed errors across all EEG channels. Solid black lines show the median values across all trials and subjects. The grey lines to the left of each set of black dots designates the 5% and 95% range of the chance R^2 distribution. Red asterisks show the stimulus-types for which the fits were significantly better than chance (Wilcoxon rank-sum, $p < 0.001$). (F) Distribution of circular shifts plotted identically to E. Red asterisks show distributions whose medians are significantly different than zero (Wilcoxon signed-rank, $p < 0.001$). (G) Topography of scaling factors for vocals and speech.

trial. Noting that differences between the stimulus-general and stimulus-specific models were not restricted to specific delays, we instead chose to quantify the scaling and circular shift of the stimulus-general model to best match the stimulus-specific model tested on the same trial (Figure 5d). The circular shift quantifies the change in delay between the two models, and the positively constrained scaling quantifies the change in magnitude.

The ability to fit the stimulus-specific model varied across stimuli: for both 4-32 Hz and 0.5-4 Hz models, the fit was not significantly better than chance for rock music, and for classical the 0.5-4 Hz fit was not above chance (Figure 5e). This indicates that there were either differences in the rock and classical stimulus-specific models compared to the stimulus-general model. The vocals models were well-fit by the scaled and shifted stimulus-general model, and the shifts of the stimulus-general model matched the expected shifts from the averaged models (Figure 5f).

We also looked at the scaling of the stimulus-general model as a function of EEG channel (Figure 5g). For both frequency ranges, the scaling was around two for speech, but interestingly the topographies were different. For 4-32 Hz, the scaling was largest in frontal and central channels, consistent with the spatial patterns of the model weights for all other stimuli (Figure 5c), and consistent with the weights found for models constrained to this frequency range in prior work (Lalor et al., 2009; Di Liberto et al., 2015). For 0.5-4 Hz, the scaling was largest in parietal channels, where model weights were somewhat larger for speech than the other stimuli (Figure 4d).

Taken together, we found that, at higher frequencies, temporal responses seem to differ between stimulus types, but these differences are not consistent enough or robust enough to outperform a model fit to all stimulus types. However, the stimulus-specific speech model

outperformed the stimulus-general model at lower frequencies, which was associated with increased model weighting over parietal electrodes in the EEG.

Drums in rock music tracked at multiples of the musical beat frequency, but tracking was worse than speech

Our previous analyses demonstrated that EEG tracks both speech and music above 1 Hz, but speech is tracked more strongly than the music stimuli (Figure 3d, e). To some extent, this could be explained by the greater magnitude of fluctuations in the dB envelope for speech, which is a consequence of the sparsity of speech (McDermott et al., 2013; Kell and McDermott, 2019). In contrast, the rock and classical stimuli were multi-instrumental, which have reduced envelope power and reduces sparsity. However, it is plausible that the brain might more strongly track one of the instruments in these stimuli, perhaps matching the strength of tracking observed for speech. Furthermore, the rock reconstruction accuracies were not significantly different than zero using models with frequency content strictly below 2 Hz. This contradicts some evidence showing neural tracking of music with note rates below this frequency (Doelling and Poeppel, 2015; Doelling et al., 2019) (note also that three out of 10 of the rock songs in this study had tempos, or musical beat rates, below 2 Hz). Perhaps individual instruments may be tracked better at lower frequencies.

To address these possibilities, we re-examined EEG neural tracking of the rock stimuli by quantifying reconstruction accuracies to the envelopes of individual instruments. This was possible because the rock songs were multi-tracked, so to retrieve the envelope of an individual instrument all of the other instruments in the track were muted. We focused our analysis on the vocals, guitar, bass, and drums because these instruments were present in all songs. Models were then trained and tested on individual instruments using the EEG data collected while subjects were listening to the multi-instrument rock songs.

Compared to all other instruments, the drums were reconstructed best; only this instrument produced z-scored reconstruction accuracies that were not significantly lower than the accuracies for the multi-instrument rock envelope (Figure 6a, b; a Wilcoxon signed-rank test with Bonferroni correction for 40 comparisons showed that the drum z-scored reconstruction accuracy was slightly better than the accuracy for the full rock envelope for the 8-64 Hz model: $z = 3.28$, $p = 0.042$) (see Figure S8 for the reconstruction accuracy curves and differences for individual subjects).

However, the z-scored reconstruction accuracy was still considerably less than what we observed for speech (Figure 3d).

At the start of this study, we assumed that EEG might track the amplitude fluctuations in music as a consequence of tracking rhythmic structure (Nozaradan et al., 2012; Kaneshiro et al., 2020). To better understand the neural tracking of the drums, which strongly suggests tracking of musical rhythm, we looked at the power spectral density of the drums reconstructions for each track averaged across subjects. We focused particularly on reconstructions using the 2-16 Hz, 4-

Figure 6: (A) Using the EEG data recorded while subjects were listening to the rock songs, we trained and tested PCA & spline models on the dB envelopes for the vocals, guitar, bass, and drums individually. Z-scored reconstruction accuracies were quantified as in Figure 3a-d. All instrument envelopes were reconstructed above chance when the model included frequencies above 2 Hz (Wilcoxon signed-rank test: $p < 0.001$ with Bonferroni correction for 40 comparisons). The full rock envelope, shown with a dashed black line, is equivalent to the values shown in Figure 3d. (B) Pairwise differences between the z-scored reconstruction accuracy for the full envelope and the envelope for each individual instrument. The z-scored reconstruction accuracies for drum were not significantly different than the same pairwise reconstruction accuracies for the rock envelope based on the multi-tracked recording with all instruments, except for the 8-64 Hz model where reconstruction accuracies were slightly but significantly better than full rock envelope (Wilcoxon signed-rank test with Bonferroni correction for 40 comparisons: $z = 3.28$, $p = 0.042$). (C) Welch's power spectral density of the reconstructions was computed for each stimulus and averaged across subjects. The noise floor of the power spectra is shown with dashed lines. (D) We then adjusted the power spectral density by subtracting the true spectrum from the average of the null spectra, which made the peaks associated with temporally coherent reconstructions across subjects clearer. The maximum values in the adjusted power spectral density were then identified relative to the expected tempo of the music (1x tempo) as well as 2x to 4x the tempo. (E) Each of the 10 rock stimuli are plotted as a different color, and each dot corresponds to 1 – 4x the music's tempo with increasing frequency. The darkest blue line and dots correspond to the example stimulus shown in C and D.

32 Hz, and 8-64 Hz models which produced the best reconstructions on average (Figure 6c). We then subtracted the spectra from null spectra generated by randomizing the phases in the reconstructions for each subject and averaging these randomized reconstructions (see Zhu et al., 2013; Keshishzadeh et al., 2020) (Figure 6d). From these adjusted spectral values we identified the peaks occurring at the tempo of the music (see Methods) as well as 2 – 4x the tempo, since the peak energy may occur at multiples of the expected musical beat frequency of the music based on the acoustics (Ding et al., 2017) or neural activity following subcortical processing (Zuk et al., 2018).

Firstly, we found that, most often, the peak energy in the reconstruction often did not occur at the tempo of the music (the left-most dots of Figure 6e). Interestingly, the peak energy for most reconstructions was restricted to a 2-8 Hz range, peaking around 5 Hz. This result supports our earlier observation that reconstruction accuracies for the rock music and drums were best using models fully encompassing this frequency range (Figure 3d and 6a). The peak energy around 5 Hz also supports other observations of frequency tracking to real music (Kaneshiro et al., 2020) as well as rhythmic synthetic stimuli (Nozaradan et al., 2012, 2016). The reduced energy below 2 Hz can be explained by model constraints, which were restricted to a low cutoff of 2 Hz. However, all models contained energy above 8 Hz, suggesting that the upper cutoff reflects neural tracking limits and not the model specifications. This could be a consequence of the evoked responses to the musical events, which may have little energy above 8 Hz.

Overall, EEG tracks drums better than any other instruments at a rate associated with the rhythmic structure of the music, but not necessarily at the musical beat. In particular, this tracking peaked around 5 Hz and tracking was not observed above 8 Hz. However, the neural tracking of the drums alone was still worse than what we observed for speech, implying that the difference in

reconstruction accuracies that we observed between the speech and music could not be explained by attention to, or passive neural tracking of, individual auditory objects in the mixture.

Discussion:

In this study, subjects passively listened to separate trials containing speech, rock music, classical music, or the vocals from the rock music. Using the recorded EEG on each trial we used linear modeling with PCA to control spatial correlations in the EEG and a basis spline transformation to control temporal correlations. These transformations, in addition to highpass filtering, restricted the frequency content in the EEG used for envelope reconstruction. We then compared reconstruction accuracies for a wide range of frequency bands in order to identify frequency ranges that were tracked best by each stimulus type. We found, however, that speech was tracked better than the other stimuli for all frequencies, and speech tracking continued below 1 Hz where all other stimuli showed reconstruction accuracies not significantly better than chance. Closer inspection of EEG tracking for the rock music showed that the drums, or perhaps more generally the rhythmic structure, was tracked most strongly around 5 Hz, but this tracking was still worse than what was observed for speech in this frequency range. For modulation frequencies above 1 Hz, a model trained on all stimuli did just as well at reconstructing speech and better for the music stimuli, suggesting common mechanisms involved in this range. However, unique tracking of speech below 1 Hz appeared to be associated with increased weighting of parietal channels. Together, this suggests that both music and speech are tracked above 1 Hz using mechanisms with largely similar spatiotemporal responses in the EEG, but below 1 Hz mechanisms exist that specifically track speech.

Our aim in this analysis was to understand how the acoustics of speech and music might produce observed differences in envelope reconstruction accuracy. In the average spectrum, speech shows a second, lower frequency peak around 0.3 Hz, within the range where we observed neural tracking that was not present for the other stimuli (Figure 1b, S1). This energy corresponds to phrase-level fluctuations in the envelope. Thus, it is possible that the unique neural tracking we observed for speech is due to the relatively high energy in that frequency range compared to the other stimuli. However, it is surprising that the enhanced neural tracking was not observed for the vocals stimuli, which also showed fluctuations in a slightly lower frequency range (~ 0.2 Hz). Additionally, we found that the models for vocals did not have enhanced weighting over parietal

channels, like we observed for speech (Figure 5g). The importance of low-frequency tracking (within the range generally considered “delta” in most studies) for speech is not new (Bourguignon et al., 2013; Etard and Reichenbach, 2019; Donhauser and Baillet, 2020), but to our knowledge no study has suggested that low-frequency tracking is unique to speech and not present for other naturalistic stimuli. Still, a future study could determine if the low-frequency tracking we observed here is specific to speech by presenting both speech and single-instrument stimuli with phrase-level fluctuations in amplitude. If the low frequency tracking truly is unique to speech, we would expect that other sounds explicitly designed to have high energy in that frequency range will not elicit the same level of neural tracking as observed for speech.

Additionally, we found that speech reconstructions were better than the various musical stimuli across all frequency ranges. The magnitude of the weights for the speech models suggests that the evoked responses to speech are larger, which increases the signal-to-noise ratio in the EEG and could produce greater reconstruction accuracies of the speech envelope. If auditory cortex tracked the dB envelope exactly, this could explain our result; the power spectrum of the speech dB envelope was larger than the music stimuli for all of the frequencies we examined (Figure 1b, S1). However, this reasoning does not match the results for music, because rock music produced reconstruction accuracies on par with the vocals alone even though the average spectrum was smallest of the three stimulus types. More regions of auditory cortex are recruited when listening to speech than music (Norman-Haignere et al., 2015, 2019), which could also explain the increased EEG activation we observed for speech. Alternatively, some studies have argued that temporal modulations are more important for speech processing, while spectral modulations are more important for music (Zatorre et al., 2002; Albouy et al., 2020; but see Norman-Haignere and McDermott, 2018). It is possible that a different acoustic feature might be tracked better by the EEG during music than the temporal envelope. Lastly, many studies of EEG processing of speech have shown the importance of speech features in affecting evoked responses in this frequency range, including phoneme coding (Di Liberto et al., 2015), phoneme probability encoding (Di Liberto et al., 2019), surprise and uncertainty (Etard and Reichenbach, 2019; Donhauser and Baillet, 2020).

Our analysis focused on isolating specific frequency contributions to envelope reconstruction while also identifying the spatiotemporal characteristics of the EEG responses. One intention of the analysis in this study was to bridge the gap between evoked and oscillations-based

interpretations of the neural tracking of speech and music. Here our stimuli are naturalistic sounds, but we use frequency-constrained modeling to isolate the spatiotemporal responses which track amplitude modulations in these sounds, showing increased tracking at frequencies commonly associated with theorized oscillatory tracking for parsing sounds (theta: 4-8 Hz, delta: < 4 Hz). Given our results, it seems reasonable that theories of evoked responses and phase-related tracking can be parsimoniously explained by frequency-tuned evoked responses (see also Kojima et al., 2020). Further modeling work comparing the theories of evoked tracking and oscillations-based tracking will also be beneficial to reconciling their differences and the observations in EEG (for example, Doelling et al., 2019).

When reconstructing the envelopes of individual instruments in the rock music, we found that drums were reconstructed best, and reconstructions showed that peak energy was usually between 2-8 Hz with a maximum around 5 Hz. This analysis was inspired by several studies that have examined the frequency content of EEG or MEG in order to quantify the neural tracking of musical beats (Nozaradan et al., 2012, 2016; Tal et al., 2017; Kaneshiro et al., 2020). We did not find strong or consistent tracking at the frequency of the tempo, but our results suggest a particular importance for neural tracking around 5 Hz, consistent with a recent study of EEG tracking of Indian music (Kaneshiro et al., 2020). The potential importance of neural tracking at a multiple of the beat frequency has been observed in other studies modeling stages of subcortical processing (Ding et al., 2017; Zuk et al., 2018) but they have not identified the 2-8 Hz frequency range as specifically important. Interestingly, this result is also contrary to an analysis of temporal modulations in speech and music, which showed a prominence of 5 Hz for speech and 2 Hz for music (Ding et al., 2017). 5-8 Hz encompasses the limit of auditory-motor synchronization found across various studies (Repp, 2005). However, these fast events might be relevant for representing the smallest temporal unit in a musical piece defining the grid on which the musical beat, rhythmic structure, and meter are based (London, 2012). Understanding the relationship between peak synchronization to musical events, the limits of synchronization, and listener experience could be relevant for biophysical modeling of music perception in the future.

Quantifying the information carried by EEG for decoding speech is an active area of research in brain-computer interfaces and auditory attention decoding, which is now focused primarily on using reconstruction accuracy to identify a subject's locus of attention (O'Sullivan et al., 2015; Miran et al., 2018; Ciccarelli et al., 2019; Vandecappelle et al., 2020). Our interest here

was also to use reconstruction accuracy as a means of quantifying how sufficiently the spatiotemporal responses represent neural tracking of the stimulus type. But we think that our observation of unique low-frequency tracking (< 1 Hz) for speech is notable for auditory attention decoding work. Low-frequency tracking may relate to several cognitive aspects of speech processing such as semantics, surprise, attention, comprehension, and language proficiency (Broderick et al., 2018; Donhauser and Baillet, 2019; Etard and Reichenbach, 2019; Chien and Honey, 2020; Di Liberto et al., 2020). If other naturalistic sounds are not sensitive to this frequency range, then there could be considerable benefit to focusing on this frequency range to identify the locus of attention of a talker and isolate the most relevant speech feature to which a user is engaged.

Materials & Methods:

Experiment and EEG recording:

Stimuli consisted of seven approximately three-minute segments from an audiobook in English (“speech” stimuli), ten rock songs including vocals each 3.5-5 minutes long, ten segments of classical pieces that were 3-4.5 minutes long, and ten tracks of the vocals from the rock songs (2.5-4.5 minutes) (see Table S1 for a list of the stimuli and more detailed information). All of the rock songs were originally multitracked, so the other instruments were muted in order to isolate the vocal track for the vocals stimuli, and silences were manually shortened to reduce the overall length of the track. Each stimulus was preceded by a 10 ms voltage click occurring 0.5 s before the start of the stimulus; the click triggered an Arduino to provide a code to the EEG data collection system signifying the start of the audio in the recorded EEG data to microsecond precision.

17 subjects took part in the experiment. One subject’s data had issues with trigger timing that produced shorter EEG recordings than the actual stimulus durations, and their data was excluded from further analysis. In total, data from 16 subjects (7 female; ages 18-44, median 22) were included. Subjects listened to the stimuli at a comfortable sound level via Sennheiser HD650 headphones. Stimuli were presented to the subject using Presentation software (Neurobehavioral Systems). 128-channel EEG and two mastoid channels were recorded (Biosemi ActiveTwo) at 512 Hz as subjects passively listened to each of the sounds in the following order within each block: rock, classical, vocals, speech. An additional track consisting of tone pips was also included and

was presented every 4-5 trials, but the data was not analyzed for this study. Each subject listened to 6-7 trials of each stimulus type.

After the session, the EEG channels were referenced to the average of the mastoid channels. No other preprocessing was applied to the EEG data prior to modeling.

Extracting the stimulus envelope:

To get the stimulus envelopes, the stimulus waveform was filtered with a bank of 32 gammachirp filters (Irino and Patterson, 2006) logarithmically spaced between 100 Hz and 8 kHz. The amplitude of the Hilbert transform of each channel was then averaged across frequency and normalized to have a peak amplitude of 1 V. Prior to converting the envelope to dB, in order to prevent discontinuities at zero values in the envelope, all voltages below 10^{-5} V (equivalent to -100 dB V) were set to 10^{-5} . The dB envelope was then resampled to 512 Hz for the speech and music data. We chose to work with the dB envelope because it is more linearly related to the perception of loudness and the EEG response to sound level than the raw voltage values (Rodriguez et al., 1986; Picton et al., 2003; Van Eeckhoutte et al., 2016). Additionally, in contrast with the original voltages that are strictly positive, we thought assumptions of Gaussian distributed errors for linear modeling might be more appropriate for the dB envelope.

Effects of frequency content on regularization-based modeling of envelope tracking:

Our initial goal to compare speech and music neural tracking raised several issues associated with stimulus differences. Firstly, the shapes and spectra of the envelopes are very different across stimulus types (Figure 1), so by comparing envelope reconstructions to the broadband envelope, certain frequency ranges may contribute more to the overall error; if a stimulus type isn't being tracked at low frequencies, it might produce a worse reconstruction accuracy despite reasonable tracking at higher frequencies. Secondly, the speech stimuli had much less cross-trial variability than the other stimulus types, so a trial-by-trial approach to modeling the envelope would be poorer for the music stimuli than for the speech stimuli. Because the rhythms of music are known to be more varied than speech (Ding et al., 2017), we think this issue is unlikely to have changed using a different stimulus set.

Typically, envelope reconstruction model weights are constrained using ridge regression or some other form of regularization which minimizes the variance of the weights (Lalor and Foxe, 2010; Crosse et al., 2016; Wong et al., 2018). Ridge-type regularization addresses issues of multicollinearity that are present both spatially (between neighboring EEG channels) and temporally by down-weighting the contribution of low-variance principal components in the input, effectively acting like a low-pass filter (Girosi et al., 1995; Hastie et al., 2009). For the purposes of comparing neural tracking for stimulus types with different spectral characteristics, this is problematic because of the lack of control over the frequency effects of regularization; while we could show that reconstruction accuracies are different between different stimulus types, we would be unable to claim that these differences are not due to spectral differences in the envelopes.

Additionally, depending upon the range of delays used, model weights will affect the spectral content at higher frequencies, but leave lower frequencies untouched. To simulate this, a

Figure 7: (A) Ridge regression was used to reconstruct the envelope from the EEG. Shown in the middle is the averaged reconstruction model across subjects and EEG channels for speech, where ridge regression was used to fit the model. Normally, the reconstruction model takes EEG as an input, but to simulate the spectral effects of the reconstruction model on the input, we have used broadband noise as the input in this example, which has a flat frequency spectrum. (B) When looking at the spectrum of the reconstruction with respect to a broadband noise input (black), it is clear that the reconstruction model accentuates certain frequencies and reduces higher frequencies (blue). However, it has no effect on the magnitude for frequencies corresponding to less than 2x the model width (1 Hz in this example), although it does add a delay that produces a phase shift at these frequencies (not shown). (C) The presence of low-frequency tracking has an effect on reconstruction accuracies. When the envelope and the EEG are highpass filtered by removing the moving average of the model width (500 ms), speech envelope reconstruction significantly drops, showing neural tracking at low frequencies untouched by the model. In contrast, reconstruction accuracies for the music stimuli significantly improves without these lower frequencies. Each color represents the testing reconstruction accuracies for one of the subjects.

regularized reconstruction model fitted to the speech stimuli using EEG delays between 0 and 500 ms (a “model width” of 500 ms, Figure 7a) was applied to a broadband noise input with a flat spectrum. The output of the model shows variation in the spectrum above 1 Hz, but little effect on the magnitudes of the spectrum below 1 Hz (Figure 7b). As a consequence, low-frequency neural tracking may be present but not accounted for. To demonstrate this in Figure 7c, we reconstructed the different stimuli using models trained on each stimulus type separately, first using EEG and stimulus envelopes that were highpass filtered above 0.1 Hz by subtracting the moving average of a 10 s window, and then using a highpass filtered envelope and EEG after similarly removing energy below 2 Hz. Speech reconstruction accuracies significantly decrease when frequencies below 2 Hz are removed (Wilcoxon signed-rank test with Bonferroni correction for 4 comparisons: $z = -7.2$, $p < 0.001$), while the reconstruction accuracies for the music stimuli significantly increase when the low-frequency content is removed (rock: $z = 6.6$, classical: $z = 4.0$, vocals: $z = 6.2$; $p < 0.001$ for all stimuli).

Envelope reconstruction with PCA & spline transformation:

All code used to create and test these models is available on gitlab (https://gitlab.com/eegetracking/speech_music_envelope_tracking).

The model we used for this study was a linear model that reconstructed the stimulus dB envelope using the principal components of the EEG and a spline basis to focus on lower frequencies. The PCA and spline transformations control multicollinearity in the EEG data. Transforming the EEG channels into principal components reduces spatial correlations between channels; this was done using Matlab’s built-in *pca* function. The basis spline transformation reduces the delayed EEG values to a lower-frequency representation using cubic basis splines. Specifically, cubic splines are piecewise polynomial functions defined by a sequence of knots such that the first and second derivatives of the function are continuous at each knot. For a sequence of knots, cubic splines can be collated across the dependent dimension of a function, and any cubic spline function defined by those knots can be represented by linear combinations of these basis splines. The collation matrix was constructed using functions from Matlab’s Curve Fitting Toolbox, *augknt* to construct the sequence of evenly-spaced knots along the range of delays, and *spscol* to create the basis spline matrix. The number of spline knots also define the number of basis

splines used, so the spline basis can also be defined by the sampling rate of the knots (this designation is used in Figure 2b).

When fitting the model, firstly the moving average of the stimulus envelope and the EEG, using a window size equal to the maximum lag in the model, was subtracted. One trial was left out for testing, and the rest of the trials were used for model training. Next, the EEG data in the training trials were converted into principal components, and the principal components were z-scored for each trial. To create the design matrix for model fitting, each principal component was delayed from 0 ms delay up to the maximum delay in the model. For a lag matrix for a principal component X_d , the lags were converted into a matrix of basis splines X_s :

$$X_s = X_d S (S^T S)^{-1}$$

where S is the collocation matrix of basis splines with delays along rows and basis cubic splines along columns. Both the spline-transformed design matrix and the stimulus envelope were z-scored for each trial, and then concatenated across trials for model fitting. The model was then computed using linear regression:

$$w_s = (X_s^T X_s)^{-1} X_s^T s(t)$$

where $s(t)$ is the stimulus envelope. The first and last 16 seconds of each trial were left-out of model fitting and testing in order to avoid potential edge artifacts produced by the removal of the moving average using the lowest-frequency model (the model with the largest window) in subsequent modeling (see Figure 3). This included the click and 0.5 s of silence before the stimulus started, so overall the first 15.5 s and the last 16 s were left-out of each trial.

For testing, the EEG data in the left-out trial were transformed into principal components using the same transformation matrix calculated for the training data. The delayed principal components were transformed into a spline basis as described for the training data and the stimulus envelope for the testing trial was reconstructed by $s'_{test}(t) = X_{s,test} w_s$. Throughout, reconstruction accuracy was quantified based on the Pearson's correlation between the original stimulus envelope and the reconstructed stimulus envelope.

For later analysis of the model, model weights were converted from splines to delays for each principal component by $w_d = S w_s$.

PCA & spline model optimization:

The envelope reconstruction model contained three model hyperparameters: the number of principal components to use, the window size of the model, and the number of basis cubic splines (or, equivalently, the sampling frequency of the spline knots). The model window size and the number of splines constrain the lower and upper cutoff frequencies contained within the model, respectively. As such, we chose to optimize the number of splines with respect to the window size, with the intention of manipulating both in tandem when examining neural tracking across frequency bands.

We used the Natural Speech dataset to identify the optimal model parameters (Broderick et al., 2018). In this dataset, subjects listened to an audiobook for 20 trials approximately 180 s long (a subset of these segments was used for the speech stimuli in the current study). The dataset contains the raw EEG and the stimulus envelope at 128 Hz. We first converted the stimulus envelope into dB; we did not threshold at -100 dB V like we do later for the envelopes in the current study (see “Extracting the stimulus envelope” above) and instead used the real value of the $\log_{10}$ of the envelope, since thresholding failed to reduce the effect of low-amplitude artifacts in the envelope and the real component of the log value seemed to be more robust to these issues. The EEG was referenced to the average of the mastoid channels.

For modeling, we used a 500 ms window because of its common use in envelope reconstruction (Ding and Simon, 2012; Crosse et al., 2015; Wong et al., 2018; Das et al., 2020). We first removed the moving average of a 500 ms window from both the dB stimulus envelope and the EEG. Then we transformed the EEG into principal components and fitted a model using basis splines. We compared the performance of this model to one based on ridge regression of the original EEG using regularization parameters between 0 (no regularization) and 10^8 ; most often the optimal regularization parameter was around 10^4 to 10^5 . PCA & spline models were optimized relative to the performance of the regularized model using a grid search, using 8, 16, 32, 64, or 128 principal components, and 7, 11, 19, 35 splines (corresponding to a spline knot sampling frequency of 4, 8, 16, or 32 Hz respectively) or all 64 lag parameters of the model without a spline transform. Performance was evaluated on a left-out trial and trained on the other 19 trials.

To estimate the frequency content of the resulting optimum model (Figure 2e), we used broadband noise as the input to the model. First, the moving average of the noise using a 500 ms window was removed. Next, the design matrix was created from the noise using delays -250 to

250 ms, and the matrix was converted into 19 basis splines (32 Hz sampling frequency of the knots); we set 0 ms as the center delay because when 0 ms was the first or last delay (for example, 0 – 500 ms), the edge spline was most heavily weighted and produced edge effects in the resulting spectrum, which inappropriately represented the low-pass effects of the basis splines generally. We then used linear regression to estimate the best fit between the spline-transformed design matrix and the original noise with the moving average removed, which we used to reconstruct the noise. The filter resulting from the combined moving average removal and the spline transformation was computed by getting the ratio of the reconstructed noise to the original broadband noise. Figure 2c shows this ratio in dB V.

Quantifying cross-frequency model performance:

We varied the frequency content of the model by changing the window size while retaining the same number of splines (19 splines), which maintained a 3-octave bandwidth. The low frequency cutoff of each range corresponds to the reciprocal of the maximum delay in the model, so by examining models with maximum delays from 31.25 ms to 16 s, we varied the frequency content of each model from a 32-256 Hz range to a 0.0625-0.5 Hz range respectively. Reconstruction accuracy was quantified based on the Pearson’s correlation between the envelope reconstruction of the left-out trial and the original envelope.

Because chance performance varies with the frequency content of the stimulus envelope and EEG (Figure 3c), we calculated a null distribution for each frequency range separately. For each iteration, the stimulus envelopes were randomly circularly shifted and the same shift was applied to all trials. Then, one of the trials was randomly chosen and left out for testing, and the rest of the model fitting and testing procedures were applied as described above. This was repeated 50 times to get a distribution of null accuracies for each stimulus type and for each subject. The true reconstruction accuracies were normalized relative to the null distribution and reported in standard deviations relative to the null distribution of accuracies to get the “z-scored” reconstruction accuracy (see Figure 3c).

Getting a spatiotemporal EEG response from reconstruction model weights:

Next, we wanted to understand how neural responses were represented by the reconstruction models. However, reconstruction models, also known as “backwards” models, are

not interpretable without accounting for autocorrelations in the data, since the resulting reconstruction model could produce weights which cancel out irrelevant autocorrelations and do not necessarily represent a neural response (Haufe et al., 2014). To account for this possibility, we first inverted the model, using principal components and a spline basis, into a “forward” model that is a better representation of the spatiotemporal evoked neural response (Lalor et al., 2009; Crosse et al., 2016). The transformation was a modification of the approach used by Haufe et al, 2014. The forward model a_S was calculated from the backward model w_S by:

$$a_S = \frac{1}{N} (X_S^T X_S) w_S$$

where N is the total number of sample points in the training data and X_S is the design matrix for the training data. In contrast with the original approach by Haufe et al, 2014, this scales the weights assuming that the input is a representation of the stimulus envelope with unit variance. We chose to use this approach to make the forward models more comparable between stimulus types. Then, the “forward” model weights were converted into delay weights (see Methods: “Envelope reconstruction with PCA & spline transformation”).

The model was then converted from principal components into EEG channels. Because the principal components were z-scored prior to fitting the model, each principal component was then multiplied by the ratio of the variance of the principal component to the summed variance across 128 principal components. Firstly, this ensures that the scaling of each principal component matches its original contribution to the EEG signal. Secondly, the EEG variance can increase considerably at lower frequencies, so the variance in the principal components tends to be larger for the lower frequency EEG range. Dividing by the total variance across components normalizes this effect of frequency on EEG variance and allows for cross-frequency model comparison.

Fitting the stimulus-general model to the stimulus-specific model:

We wanted to understand how differences between the stimulus-specific models might correspond to differences in reconstruction accuracy. We addressed this by looking at the difference in performance of a “stimulus-general” model, fit to all stimulus types, relative to the stimulus-specific model. For the stimulus-general model, trials from all stimulus types (23-27 trials) were included in training data, and the series of steps for computing the principal components, normalizing the components, and creating the basis spline design matrix were

identical to those described in the Methods: “Envelope reconstruction with PCA & spline transformation”. Importantly, both the stimulus-general and stimulus-specific models were tested trial-by-trial, so we used the difference in reconstruction accuracies (Pearson’s r) for the two models on each trial to quantify the change in model performance.

We then wanted to systematically quantify the changes in weights between the two types of models. In a linear model, the effects of model weights are combined to optimize the reconstruction, which can complicate interpretation (Kriegeskorte and Douglas, 2019). Thus, we thought it unlikely that differences at specific delays would be sufficient to explain the change in reconstruction accuracy. Subjective examinations of the difference in weights showed that the models differed primarily in the amplitude (“scaling factor”) and delay (“shift”) relative to each other, even though the shapes were similar. Thus, we thought it would be most appropriate to examine the relative circular shift and scaling of the stimulus-general model that would match the stimulus-specific model at each channel.

To compute the circular shifts and scalings, both the stimulus-specific and stimulus-general models were centered to have a mean weight of zero for each channel. Next, for each circular shift, the best scaling factor for each EEG channel was computed using linear regression. Then the R^2 model fit across all channels was computed. The shift producing the maximum R^2 was identified as the optimum model for that trial. All analyses were then based on the model fits (R^2), circular shifts, and scaling factors computed on a trial-by-trial basis.

Additionally, to evaluate the goodness of fit, we computed a null distribution of R^2 values by randomizing the phases of the stimulus-specific model and refitting the stimulus-general model as described above. This procedure ensured that the spectral amplitudes of the stimulus-specific model would remain the same, since the high R^2 values could be due to relatively high signal power at low frequencies, while ensuring that the temporal relationship between frequencies in the model was destroyed. 20 null R^2 values were computed for each trial, and the distribution of all null values was combined across trials and subjects for statistical testing relative to the true distribution of R^2 values (see Figure 5e).

Quantifying power spectra of drum reconstructions relative to the music’s tempo:

While we found higher reconstruction accuracies for speech than the other music stimuli, we considered the possibility that EEG may track individual instruments in the rock or classical

music more so than what we observed using the envelopes calculated from the multi-instrument mixture. To examine this further, we computed instrument-specific models using the EEG data recorded during the presentations of rock stimuli by training and testing on the envelopes for the vocals, guitar, bass, and drums in the rock tracks. Because the rock songs were multi-tracked, we muted all other instruments in order to get the waveform for the individual instrument and computed the envelope as described previously (Methods: “Extracting the stimulus envelope”).

After we found that drums exhibited the best reconstruction accuracy of all of the rock instruments (Figure 6a, b), we then examined which frequencies were most strongly tracked in the drum reconstructions. We focused on the 2-16 Hz, 4-32 Hz, and 8-64 Hz models because they produced the best reconstruction accuracies for drums and for the full rock envelope on average (Figure 6a). First, for each rock stimulus and model frequency range, the drum reconstructions were averaged across subjects. Then Welch’s power spectral density (*pwelch* in Matlab) was computed using a Hamming window of 10 s with half-overlap. Then a null distribution of power spectral densities was created by shuffling the phases in each of the reconstructions, averaging the randomized reconstructions, and computing the power spectral density. This technique is based on methods to quantify magnitudes of peaks in frequency-following responses (Zhu et al., 2013; Keshishzadeh et al., 2020), where randomizing the phases of each signal and then averaging captures the spectra associated with the noise floor. 100 null spectra were computed for each model. The average of these null spectra was subtracted from the true spectra to get the adjusted power spectral density for each model.

Next, we focused on the peak values in the adjusted power spectral density with respect to the expected rate of musical beats, since several studies have demonstrated that EEG and MEG tracks frequencies at multiples or fractions of this rate (Nozaradan et al., 2012; Tal et al., 2017; Kaneshiro et al., 2020). The musical beat rate (which we call the “tempo” here) was quantified using a beat-tracking algorithm (Ellis, 2007). This algorithm dynamically computes the timing of beats in a musical recording. The beat timings produced by the model were validated by the lead author of this study. The tempo was computed as the inverse of the median inter-beat interval.

Then, for each adjusted spectrum, peaks were identified with respect to 1 – 4x the music’s tempo. For each scaling of the tempo, we identified the maximum value in the spectrum across all three models, using a frequency range $\pm 8\%$ around the multiple of the tempo (see McKinney et al., 2007).

Acknowledgments: This work was supported by a Science Foundation Ireland Career Development Award (CDA/15/3316) and an Irish Research Council Starter Research Project Grant (RPG2013-1). Additional support was provided by the Del Monte Institute for Neuroscience at the University of Rochester. We would also like to thank Giovanni Di Liberto, Aaron Nidiffer, Andrew Anderson, Michael Broderick, Shyanthony Synigal, and the others in the Lalor Lab for their comments and corrections on this manuscript.

References:

- Albouy P, Benjamin L, Morillon B, Zatorre RJ (2020) Distinct sensitivity to spectrotemporal modulation supports brain asymmetry for speech and melody. *Science* 367:1043–1047.
- Baltzell LS, Srinivasan R, Richards V (2019) Hierarchical organization of melodic sequences is encoded by cortical entrainment. *Neuroimage* 200:490–500.
- Bourguignon M, De Tiège X, de Beeck MO, Ligot N, Paquier P, Van Bogaert P, Goldman S, Hari R, Jousmäki V (2013) The pace of prosodic phrasing couples the listener’s cortex to the reader’s voice. *Hum Brain Mapp* 34:314–326.
- Broderick MP, Anderson AJ, Di Liberto GM, Crosse MJ, Lalor EC (2018) Electrophysiological Correlates of Semantic Dissimilarity Reflect the Comprehension of Natural, Narrative Speech. *Curr Biol* 28:803-809.e3.
- Broderick MP, Anderson AJ, Lalor EC (2019) Semantic Context Enhances the Early Auditory Encoding of Natural Speech. *J Neurosci* 39:7564–7575.
- Cantisani G, Essid S, Richard G (2019) EEG-Based decoding of auditory attention to a target instrument in polyphonic music. In: *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, pp 80–84. Institute of Electrical and Electronics Engineers Inc.
- Chang A, Bosnyak DJ, Trainor LJ (2019) Rhythmicity facilitates pitch discrimination: Differential roles of low and high frequency neural oscillations. *Neuroimage* 198:31–43.
- Chien HYS, Honey CJ (2020) Constructing and Forgetting Temporal Context in the Human Cerebral Cortex. *Neuron* 106:675-686.e11.
- Ciccarelli G, Nolan M, Perricone J, Calamia PT, Haro S, O’Sullivan J, Mesgarani N, Quatieri TF, Smalt CJ (2019) Comparison of Two-Talker Attention Decoding from EEG with Nonlinear Neural Networks and Linear Methods. *Sci Rep* 9:11538.
- Crosse MJ, Butler JS, Lalor EC (2015) Congruent visual speech enhances cortical entrainment to continuous auditory speech in noise-free conditions. *J Neurosci* 35:14195–14204.
- Crosse MJ, Di Liberto GM, Bednar A, Lalor EC (2016) The multivariate temporal response function (mTRF) toolbox: a MATLAB toolbox for relating neural signals to continuous stimuli. *Front Hum Neurosci* 10:604.
- Das N, Vanthornhout J, Francart T, Bertrand A (2020) Stimulus-aware spatial filtering for single-trial neural response and temporal response function estimation in high-density EEG with applications in auditory research. *Neuroimage* 204:116211.
- de Cheveigné A, Wong DE, Di Liberto GM, Hjortkjær J, Slaney M, Lalor E (2018) Decoding the auditory brain with canonical component analysis. *Neuroimage* 172:206–216.
- Di Liberto GM, O’Sullivan JA, Lalor EC (2015) Low-Frequency Cortical Entrainment to Speech

Reflects Phoneme-Level Processing. *Curr Biol* 25:2457–2465.

Di Liberto G, Nie J, Yeaton J, Khalighinejad B, Shamma S, Mesgarani N (2020a) Neural representation of linguistic feature hierarchy reflects second-language proficiency. *bioRxiv*:2020.06.15.142554.

Di Liberto GM, Lalor EC, Millman RE (2018) Causal cortical dynamics of a predictive enhancement of speech intelligibility. *Neuroimage* 166:247–258.

Di Liberto GM, Pelofi C, Bianco R, Patel P, Mehta AD, Herrero JL, de Cheveigné A, Shamma S, Mesgarani N (2020b) Cortical encoding of melodic expectations in human temporal cortex. *Elife* 9.

Di Liberto GM, Wong D, Melnik GA, de Cheveigné A (2019) Low-frequency cortical responses to natural speech reflect probabilistic phonotactics. *Neuroimage* 196:237–247.

Ding N, Melloni L, Zhang H, Tian X, Poeppel D (2015) Cortical tracking of hierarchical linguistic structures in connected speech. *Nat Neurosci* 19:158–164.

Ding N, Patel AD, Chen L, Butler H, Luo C, Poeppel D (2017) Temporal modulations in speech and music. *Neurosci Biobehav Rev* 81:181–187.

Ding N, Simon JZ (2009) Neural Representations of Complex Temporal Modulations in the Human Auditory Cortex. *J Neurophysiol* 102:2731–2743.

Ding N, Simon JZ (2012) Emergence of neural encoding of auditory objects while listening to competing speakers. *Proc Natl Acad Sci* 109:11854–11859.

Ding N, Simon JZ (2014) Cortical entrainment to continuous speech: functional roles and interpretations. *Front Hum Neurosci* 8:311.

Doelling KB, Assaneo MF, Bevilacqua D, Pesaran B, Poeppel D (2019) An oscillator model better predicts cortical entrainment to music. *Proc Natl Acad Sci U S A* 116:10113–10121.

Doelling KB, Poeppel D (2015) Cortical entrainment to music and its modulation by expertise. *Proc Natl Acad Sci* 112:E6233–E6242.

Donhauser PW, Baillet S (2020) Two Distinct Neural Timescales for Predictive Speech Processing. *Neuron* 105:385–395.e9.

Ellis DPW (2007) Beat tracking by dynamic programming. *J New Music Res* 36:51–60.

Etard O, Reichenbach T (2019) Neural speech tracking in the theta and in the delta frequency band differentially encode clarity and comprehension of speech in noise. *J Neurosci* 39:5750–5759.

Forte AE, Etard O, Reichenbach T (2017) The human auditory brainstem response to running speech reveals a subcortical mechanism for selective attention. *Elife* 6:e27203.

Fujioka T, Ross B, Trainor LJ (2015) Beta-Band Oscillations Represent Auditory Beat and Its Metrical Hierarchy in Perception and Imagery. *J Neurosci* 35:15187–15198.

Geiser E, Notter M, Gabrieli JDE (2012) A Corticostriatal Neural System Enhances Auditory Perception through Temporal Context Processing. *J Neurosci* 32:6177–6182.

Ghitza O, Greenberg S (2009) On the Possible Role of Brain Rhythms in Speech Perception: Intelligibility of Time-Compressed Speech with Periodic and Aperiodic Insertions of Silence. *Phonetica* 66:113–126.

Giraud A-L, Poeppel D (2012) Cortical oscillations and speech processing: emerging computational principles and operations. *Nat Neurosci* 15:511–517.

Girosi F, Jones M, Poggio T (1995) Regularization Theory and Neural Networks Architectures. *Neural Comput* 7:219–269.

Hastie T, Tibshirani R, Friedman JH (2009) The elements of statistical learning: data mining, inference, and prediction. Springer-Verlag New York.

- Haufe S, Meinecke F, Görgen K, Dähne S, Haynes J-D, Blankertz B, Bießmann F (2014) On the interpretation of weight vectors of linear models in multivariate neuroimaging. *Neuroimage* 87:91–110.
- Irino T, Patterson RD (2006) A dynamic compressive gammachirp auditory filterbank. *IEEE Trans Audio, Speech Lang Process* 14:2222–2232.
- Iversen JR, Repp BH, Patel AD (2009) Top-down control of rhythm perception modulates early auditory responses. *Ann N Y Acad Sci* 1169:58–73.
- Kaneshiro B, Nguyen DT, Norcia AM, Dmochowski JP, Berger J (2020) Natural music evokes correlated EEG responses reflecting temporal structure and beat. *Neuroimage*:116559.
- Kell AJE, McDermott JH (2019) Invariance to background noise as a signature of non-primary auditory cortex. *Nat Commun* 10:3958.
- Keshishzadeh S, Garrett M, Vasilkov V, Verhulst S (2020) The derived-band envelope following response and its sensitivity to sensorineural hearing deficits. *Hear Res*:107979.
- Kojima K, Oganian Y, Cai C, Findlay A, Chang EF, Nagarajan S (2020) Low-frequency neural tracking of speech envelope reflects evoked responses to acoustic edges. *bioRxiv*:2020.04.02.022616.
- Kriegeskorte N, Douglas PK (2019) Interpreting encoding and decoding models. *Curr Opin Neurobiol* 55:167–179.
- Lalor EC, Foxe JJ (2010) Neural responses to uninterrupted natural speech can be extracted with precise temporal resolution. *Eur J Neurosci* 31:189–193.
- Lalor EC, Power AJ, Reilly RB, Foxe JJ (2009) Resolving precise temporal processing properties of the auditory system using continuous stimuli. *J Neurophysiol* 102:349–359.
- Large EW, Herrera JA, Velasco MJ (2015) Neural networks for beat perception in musical rhythm. *Front Syst Neurosci* 9:159.
- London J (2012) *Hearing in time: psychological aspects of musical meter*, 2nd ed. Oxford; New York: Oxford University Press.
- Luo H, Poeppel D (2007) Phase patterns of neuronal responses reliably discriminate speech in human auditory cortex. *Neuron* 54:1001–1010.
- Maddox RK, Lee AKC (2018) Auditory brainstem responses to continuous natural speech in human listeners. *eNeuro* 5.
- McDermott JH, Schemitsch M, Simoncelli EP (2013) Summary statistics in auditory perception. *Nat Neurosci* 16:493–498.
- McKinney MF, Moelants D, Davies MEP, Klapuri A (2007) Evaluation of Audio Beat Tracking and Music Tempo Extraction Algorithms. *J New Music Res* 36:1–16.
- Mesgarani N, Chang EF (2012) Selective cortical representation of attended speaker in multi-talker speech perception. *Nature* 485:233–236.
- Miran S, Akram S, Sheikhattar A, Simon JZ, Zhang T, Babadi B (2018) Real-Time Tracking of Selective Auditory Attention From M/EEG: A Bayesian Filtering Approach. *Front Neurosci* 12:262.
- Norman-Haignere S, Feather J, Brunner P, Ritaccio A, McDermott JH, Schalk G, Kanwisher N (2019) Intracranial recordings from human auditory cortex reveal a neural population selective for musical song. *bioRxiv*:696161.
- Norman-Haignere S, Kanwisher NG, McDermott JH (2015) Distinct Cortical Pathways for Music and Speech Revealed by Hypothesis-Free Voxel Decomposition. *Neuron* 88:1281–1296.
- Norman-Haignere S V., McDermott JH (2018) Neural responses to natural and model-matched

stimuli reveal distinct computations in primary and nonprimary auditory cortex Davis M, ed. PLOS Biol 16:e2005127.

Nozaradan S, Peretz I, Keller PE (2016) Individual Differences in Rhythmic Cortical Entrainment Correlate with Predictive Behavior in Sensorimotor Synchronization. Sci Rep 6:20612.

Nozaradan S, Peretz I, Missal M, Mouraux A (2011) Tagging the neuronal entrainment to beat and meter. J Neurosci 31:10234–10240.

Nozaradan S, Peretz I, Mouraux A (2012) Selective neuronal entrainment to the beat and meter embedded in a musical rhythm. J Neurosci 32:17572–17581.

O’Sullivan J, Herrero J, Smith E, Schevon C, McKhann GM, Sheth SA, Mehta AD, Mesgarani N (2019) Hierarchical Encoding of Attended Auditory Objects in Multi-talker Speech Perception. Neuron 104:1195–1209.e3.

O’Sullivan JA, Power AJ, Mesgarani N, Rajaram S, Foxe JJ, Shinn-Cunningham BG, Slaney M, Shamma SA, Lalor EC (2015) Attentional Selection in a Cocktail Party Environment Can Be Decoded from Single-Trial EEG. Cereb Cortex 25:1697–1706.

Obleser J, Kayser C (2019) Neural Entrainment and Attentional Selection in the Listening Brain. Trends Cogn Sci 23:913–926.

Oganian Y, Chang EF (2019) A speech envelope landmark for syllable encoding in human superior temporal gyrus. Sci Adv 5:eaay6279.

Picton TW, John MS, Dimitrijevic A, Purcell D (2003) Human auditory steady-state responses. Int J Audiol 42:177–219.

Power AJ, Lalor EC, Reilly RB (2011) Endogenous Auditory Spatial Attention Modulates Obligatory Sensory Activity in Auditory Cortex. Cereb Cortex 21:1223–1230.

Rajendran VG, Harper NS, Abdel-Latif KHA, Schnupp JWH (2016) Rhythm facilitates the detection of repeating sound patterns. Front Neurosci 10.

Repp BH (2005) Sensorimotor synchronization: A review of the tapping literature. Psychon Bull Rev 12:969–992.

Rodriguez R, Picton T, Linden D, Hamel G, Laframboise G (1986) Human Auditory Steady State Responses. Ear Hear 7:300–313.

Tal I, Large EW, Rabinovitch E, Wei Y, Schroeder CE, Poeppel D, Golumbic EZ (2017) Neural entrainment to the beat: The “missing-pulse” phenomenon. J Neurosci 37:6331–6341.

Talsma D, Woldorff MG (2005) Methods for the estimation and removal of artifacts and overlap in ERP waveforms. In: Event-related potentials: A methods handbook (Handy TC, ed), pp 115–148. Cambridge, MA: MIT Press.

Van Eeckhoutte M, Wouters J, Francart T (2016) Auditory steady-state responses as neural correlates of loudness growth. Hear Res 342:58–68.

Vandecappelle S, Deckers L, Das N, Ansari AH, Bertrand A, Francart T (2020) EEG-based detection of the attended speaker and the locus of auditory attention with convolutional neural networks. bioRxiv:475673.

Wong DDE, Fuglsang SA, Hjortkjær J, Ceolini E, Slaney M, de Cheveigné A (2018) A Comparison of Regularization Methods in Forward and Backward Models for Auditory Attention Decoding. Front Neurosci 12.

Zatorre RJ, Belin P, Penhune VB (2002) Structure and function of auditory cortex: music and speech. Trends Cogn Sci 6:37–46.

Zhu L, Bharadwaj H, Xia J, Shinn-Cunningham B (2013) A comparison of spectral magnitude and phase-locking value analyses of the frequency-following response to complex tones. J

885 Acoust Soc Am 134:384–395.
886 Zion Golumbic EM, Ding N, Bickel S, Lakatos P, Schevon CA, McKhann GM, Goodman RR,
887 Emerson R, Mehta AD, Simon JZ, Poeppel D, Schroeder CE (2013) Mechanisms
888 Underlying Selective Neuronal Tracking of Attended Speech at a “Cocktail Party.” Neuron
889 77:980–991.
890 Zuk NJ, Carney LH, Lalor EC (2018) Preferred tempo and low-audio-frequency bias emerge
891 from simulated sub-cortical processing of sounds with a musical beat. Front Neurosci 12.
892 Zuk NJ, Teoh ES, Lalor EC (2020) EEG-based classification of natural sounds reveals
893 specialized responses to speech and music. Neuroimage 210:116558.
894