Inductive biases, pretraining and fine-tuning jointly account for brain responses to speech

Juliette MILLET^{1 2 3 4} Jean-Rémi KING^{1 5}

Abstract

Our ability to comprehend speech remains, to date, unrivaled by deep learning models. This feat could result from the brain’s ability to fine-tune generic sound representations for speech-specific processes. To test this hypothesis, we compare i) five types of deep neural networks to ii) human brain responses elicited by spoken sentences and recorded in 102 Dutch subjects using functional Magnetic Resonance Imaging (fMRI). Each network was either trained on an acoustics scene classification, a speech-to-text task (based on Bengali, English, or Dutch), or not trained. The similarity between each model and the brain is assessed by correlating their respective activations after an optimal linear projection. The differences in brain-similarity across networks revealed three main results. First, speech representations in the brain can be accounted for by random deep networks. Second, learning to classify acoustic scenes leads deep nets to increase their brain similarity. Third, learning to process phonetically-related speech inputs (i.e., Dutch vs English) leads deep nets to reach higher levels of brain-similarity than learning to process phonetically-distant speech inputs (i.e. Dutch vs Bengali). Together, these results suggest that the human brain fine-tunes its heavily-trained auditory hierarchy to learn to process speech.

1. Introduction

To understand speech, the human brain continuously transforms and parses the stream of auditory features generated by the cochleas. To date, this ability remains unrivaled by

deep learning architectures, especially in noisy or accented speech conditions (Zhang et al., 2020; Unni et al., 2020), partly because the exact nature of speech representations remains underdetermined. Indeed, while the brain – and the superior temporal cortex in particular – represents and clusters phonetic features (Mesgarani et al., 2014; Kell et al., 2018), the corresponding representations may not be specific to speech but generic to sounds (Galantucci et al., 2006; Daube et al., 2019). It is thus unclear whether speech models should be trained with speech input only or from a wide variety of sounds.

We propose to address this issue by comparing i) the brain representations elicited during speech comprehension to ii) the representations of artificial neural networks, input with the same sounds, and trained in three distinct regimes: random neural networks (i.e. not trained), neural networks trained to process generic sounds (i.e. acoustic scene classification), and neural networks trained to process speech (i.e. speech-to-text using Dutch, English or Bengali). Previous studies have already shown that deep convolutional neural networks trained to classify words, musical genres (Kell et al., 2018; Kumar et al., 2020) or natural sounds (Koumura et al., 2019), generate brain-like representations, in the sense that one can find a linear correspondence between the activation of the neural networks and the activations of the brain (Figure 1). This similarity can be quantified with a “brain score” (Jung et al., 2019), a correlation between the brain measurements and a linear projection of the model’s activations, under the assumption that representations are defined as linearly exploitable information (Hung et al., 2005; Kamitani & Tong, 2005; Kriegeskorte & Kievit, 2013; King et al., 2018). We thus hypothesize that the nature of speech representations in the brain can be elucidated by comparing them to those of random, sound-generic and speech-specific neural networks (Figure 1).

To this end, we analyze the functional Magnetic Resonance Imaging (fMRI) recordings of 102 Dutch speakers who each underwent a 1 h speech comprehension task, consisting of sequences of unrelated sentences. We then compared, for each subject separately, the resulting blood oxygenation level dependent (BOLD) responses to the activations of 30 deep neural networks, all derived from the same architecture

¹Facebook AI Research, Paris ²Université de Paris, LLF, CNRS, Paris, France ³CoML, ENS/CNRS/EHESS/INRIA/PSL University, Paris, France ⁴CRI, Département Frontières du Vivant et de l’Apprendre, IIFR, Université de Paris ⁵École normale supérieure, PSL University, CNRS, Paris, France. Correspondence to: Juliette MILLET <juliette.millet@cri-paris.org>.

taken from DeepSpeech 2 (Amodei et al., 2016).

We formalize our hypotheses with the following non-mutually exclusive predictions (Figure 1): the brain representations of speech are likely to benefit from an **inductive bias** if a random neural network can generate representations similar to those of the brain, **generic** if a neural network trained on non-speech sounds leads to representations that are more similar to brain activity than a random network, **speech-specific** if a neural network trained on speech sounds leads to representations that are more similar to brain activity than a network trained on non-speech sounds, **phoneme-specific** if a neural network trained on a language that roughly contains the same speech sounds (or phonemes) than the participants’ native language (e.g. English vs Dutch) leads to representations that are more similar to brain activity than a network trained on a language that has a very different phonemes inventory (e.g. Bengali vs Dutch), and **language-specific** if a neural network trained on the same language as the participants’ native language (i.e. Dutch) leads to representations that are more similar to brain activity than a network trained on a distinct, but phonetically-related language (e.g. Dutch vs English).

Our results support several of these predictions. First, the architecture of the network, without training, suffices to significantly account for the hierarchy of brain responses to speech. Second, training on auditory-scene classification increases the brain-similarity of the networks. Third, training on English leads to higher brain scores than training on Bengali. Together, these results suggest that the brain combines inductive biases, generic, and phoneme-specific representations to process speech.

2. Methods

2.1. Stimuli

One-hundred and two subjects performed a simple speech comprehension task while being recorded with functional Magnetic Resonance Imaging (fMRI) by Schoffelen and colleagues (Schoffelen et al., 2019). Each participant listened to 60 isolated sentences, grouped in series of five and selected from a pool of 360 sentences, varying from 9 to 15 words. Each sentence lasted on average 4.28 ± 0.6 seconds (maximal length: 6s, minimal length: 2.8s), for a total of 28 min and 49 seconds. All sentences were recorded by a unique female native Dutch speaker in stereo at 44.1 kHz. We downsample the recordings to 16 kHz and convert them to mono with the Sox software¹. Note that the dataset also contains word lists that are not presently analyzed. However, these sound sequences included silences between each word and thus lead to abnormal speech inputs.

¹<http://sox.sourceforge.net/>

2.2. fMRI preprocessing

Data were provided (in part) by the Donders Institute for Brain, Cognition and Behaviour. Each subject performed two MRI scans performed: a structural T1-weighted magnetization-prepared rapid gradient-echo (MP-RAGE) pulse sequence (TR=2,300 ms, TE=3.03 ms, 8 degree flip-angle, 1 slab, slice-matrix size=256×256, slice thickness=1 mm, field of view=256 mm, isotropic voxel-size=1.0×1.0×1.0 mm) and a functional T2*-weighted echo planar blood oxygenation level dependent (EPI-BOLD) sequence (TR=2.0 seconds, TE=35ms, flip angle=90 degrees, anisotropic voxel size=3.5×3.5×3.0 mm extracted from 29 oblique slices.) Additional acquisition details can be found in Schoffelen et al. (2019).

Structural images were defaced by Schoffelen and colleagues. We then preprocessed them with Freesurfer (Fischl, 2012), using the `recon-all` pipeline with default parameters and a manual inspection of the cortical segmentations. Region-of-interest analyses of regions the primary auditory cortex A1, Belt (LBelt MBelt PBelt) and TPJ (STV, PFm, PGI) were selected from the PALS Brodmann’ area atlas (Van Essen, 2005) and the Destrieux Atlas segmentation (Destrieux et al., 2010).

Functional images are preprocessed with fMRIPrep with default parameters (Esteban et al., 2019). The Blood Oxygenation Level Dependent (BOLD) time series are detrended and de-confounded from 18 variables: the six head motion parameters ($\text{trans}_{x,y,z}$, $\text{rot}_{x,y,z}$) estimated analytically, the first 6 noise components derived from CompCorr (Behzadi et al., 2007) and the 6 DCT-basis regressors derived from Nilearn’s `cleaning` pipeline and otherwise default parameters (Abraham et al., 2014b). The resulting volumetric data are projected onto the surface, after subselecting voxels along a 3mm “line” orthogonal to the mid-thickness. Surface projections are spatially decimated by 10, and are hereafter referred to as voxels, for simplicity.

To mitigate slow drifts in fMRI acquisition, the recordings of each group of 5 sentences are separately and linearly detrended. The 12 cross-validation splits are applied across 5-sentences blocks (i.e., across detrended segments) to ensure no information leakage between the training and testing sets.

2.3. Networks

We train and analyze the activations of five distinct types of neural networks (training details are in the next section). All networks share the same architecture derived from DeepSpeech 2 (Amodei et al., 2016), and are all input with spectrograms $U \in \mathbb{R}^{d_\tau \times d_u}$ ($d_u = 160$) of mono-sound waveforms computed with librosa’s `stft` function (McFee et al., 2019) squared using: window size 20 milliseconds, window

Figure 1. A. The functional similarity between an artificial neural network and the human brain was quantified by correlating i) brain responses to speech to ii) the networks’ activations in response to the same input, and optimally projected with a linear operator W onto the subject brain. To this end, a linear ridge regression (W) is fit from the model’s activations convolved by an HRF (X) to the brain response (Y) to the same stimulus sequence. The resulting “brain score” (Jung et al., 2019) is a Pearson correlation between true brain activity and the brain activity predicted by a linear projection of the model’s activation. **B.** Four types of networks, using the same architecture were trained on Dutch, English, and Bengali speech-to-text and on a acoustic scene classification, respectively. The fifth network corresponds to an untrained version, randomly initialised. Each type was trained with six random seeds, leading to a total of 30 networks. **C.** The neural networks will be considered to show no similarity with brain activity if they show brain scores (R) not significantly above zero.

stride of 10 ms, and $n_fft = 320$. These spectrograms are transformed by a hierarchy of three 2D convolutional layers (with Layer 1: kernel $k_1 = 11 \times 41$, stride $s_1 = 2 \times 2$, number of channels $c_1 = 32$; Layer 2: $k_2 = 11 \times 21$, $s_2 = 1 \times 2$, $c_2 = 64$, Layer 3: $k_3 = 11 \times 21$, $s_3 = 1 \times 2$, $c_3 = 96$), followed by one bidirectional gated recurrent unit (GRU, 1 layer with dimension 256) and ending in a fully connected layer, whose dimensionality varies with the task (see below). A rectified linear unit (ReLU) function is used between each layer. Finally, a softmax layer is applied on the linear layer’s outputs to obtain a probability distribution over characters or acoustic-scene classes.

2.4. Tasks

To test our hypotheses, We train each network on one of four different datasets (Figure 1). The Dutch and the English datasets are taken from the aligned part of the Spoken Wikipedia Corpora (Baumann et al., 2018). The Bengali dataset is taken from Driedger et al. (2014). In order to have approximately the same number of characters possible for the three languages we romanize the Bengali transcriptions using URoman (Hermjakob et al., 2018). Finally, the acoustic scene classification dataset is extracted from the curated subset of the FSDKaggle2019 dataset (Fonseca et al., 2019) with manual annotations (Fonseca et al., 2017).

We use 66 hours and 20 minutes (roughly 100,000 sentences) of the first three datasets to match the Dutch dataset, and

ensure that the networks had a similar data exposure. We restrict the scene acoustic dataset to its curated version (10 hours long) to avoid introducing noisy labels. Each dataset is then divided into a training (80%), a validation (10%) and a test set (10%). Input samples are limited to a maximum length of 20 seconds.

Each model m was trained using Connectionist Temporal Classification (CTC) (Graves et al., 2006) parameterized over θ :

$$\arg \min_{\theta} - \log \sum_{a \in a_{U,V}} \prod_t^{d_t} p_{CTC}(a_t | m_{\theta}(U)) \quad (1)$$

where $m_{\theta}(U) \in \mathbb{R}^{d_{\tau} \times d_v}$ are the probabilistic predictions of the networks at each τ time sample given the spectrograms $U \in \mathbb{R}^{d_{\tau} \times d_u}$, $V \in \mathbb{R}^{d_t \times d_v}$ are the true transcriptions of U , and $a_{U,V}$ is the set of all possible alignments between U and V .

For the models trained on speech recognition, V is a textual transcription. Consequently, d_v is the number of possible characters: $d_v = 37$ for Dutch, $d_v = 28$ for English and Bengali. For the models trained on acoustic scene classification, V is a one hot vector indicating the class of U : i.e. $d_v = 80$ classes.

The models are trained with Openseq2seq (Kuchaiev et al.,

2018), with six different initialisation seeds for each dataset, using stochastic gradient descent with the following optimization hyperparameters: momentum of 0.9, an initial learning rate of 0.001, a polynomial decay policy with power 2, a batch size of 8. After each gradient descent, we select the model based on its performance on the validation set, with a maximum number of training steps of 70,000. For each setting, we compute the brain scores for each initialisation and present the average results across initialisations, for simplicity.

On average, the Dutch, English and Bengali models achieved a Word Error Rate (WER) of 0.59, 0.78, and 0.85, on their respective test sets. The acoustic scene classification models achieve a class error rate of 0.64. These modest performances results from a design choice: to avoid semantic representations, we deliberately decided not to incorporate a pretrained language model feeding the architecture.

2.5. Preprocessing of model activations

The activations of the models are sampled at a different frequency ($1/t = 50$ Hz) than the fMRI BOLD signals ($1/\hat{t} = 0.5$ Hz). Furthermore, BOLD signals are known to be slow and delayed responses spanning over several seconds. To align the activations of the models and the fMRI BOLD signals, we thus convolve the model activations with a standard hemodynamic response function (HRF). Specifically, activations of the network $m(U) \in \mathbb{R}^{d_t \times d_x}$ are normalized between $[0, 1]$. We then use nistats (Abraham et al., 2014a) `compute_regressor` function with the Glover model to temporally convolve ($h \in \mathbb{R}^{d_t}$) and temporally downsample ($g : \mathbb{R}^{d_t} \rightarrow \mathbb{R}^{d_t}$) each artificial neuron j :

$$x^{(j)} = g\left(m^{(j)}(U) * h\right) \quad (2)$$

2.6. Brain score

The similarity between the preprocessed model’s activations $X \in \mathbb{R}^{d_t \times d_x}$ and the BOLD response $Y \in \mathbb{R}^{d_t \times d_y}$ is assessed with a ”brain score” (Jung et al., 2019). Specifically, we use a cross-validation within each subject (12 train-test splits across 5 sentences series) to fit an L_2 -regularized linear model $W \in \mathbb{R}^{d_x \times d_y}$ trained to predict the BOLD time series from the model activations for each voxel independently:

$$\arg \min_w \sum_i (w^T X_i - y_i)^2 + \lambda \|w\|^2 \quad (3)$$

with X and y normalized across samples to have a mean of 0 and a variance of 1, and the training samples $i \in I$.

Optimization is done using RidgeCV from scikit-learn (Pedregosa et al., 2011), with the penalization hyper-parameter

λ varying between 10 and 10^8 (20 values scaled logarithmically) chosen within nested cross-validation of the training set and independently selected for each dimension.

The performance of the ridge regression is summarized with a Pearson’s correlation score between the predicted and the fMRI BOLD responses on the test samples I^* , a.k.a ”brain score”:

$$R = \text{Corr}(y_{I^*}, w^T \cdot X_{I^*}) \quad (4)$$

Note that we also compute a baseline brain-score with X being the preprocessed mel filterbanks coefficients (mel) of the auditory stimuli, using librosa (McFee et al., 2019), 25 milliseconds windowing and strides of 10 milliseconds. These spectro-temporal decompositions were originally engineered to mimic the cochlear transformations of sounds. They thus provide a reasonable approximation of the features input to the primary auditory cortex (Stevens et al., 1937; Stevens & Volkmann, 1940).

2.7. Statistics

The reliability of brain scores is assessed with second-level analyses across subjects, by applying a Wilcoxon signed-rank test across subjects. Statistical correction for multiple comparisons is performed with Benjamini–Hochberg False Discovery Rate (FDR) across voxels (Moore & Jorgenson, 1993). Figures only plot brain scores that are significant across subjects after this correction.

3. Results

3.1. Random networks reveal the inductive bias of hierarchical convolutions

We first aim to test whether random (i.e. not trained) convolutional architectures generate representations similar to those of the brain. To this end, we assess, for each subject independently, whether their fMRI responses Y can be predicted with a linear regression W from the concatenated activations X of the random intialisation model input with the same speech sounds, as quantified with a Pearson correlation, hereafter referred to as ”brain scores” (R , Figure 2.A). Sound onset and volume are known to trigger attentional capture in large brain networks (Daube et al., 2019). Consequently, significant brain scores result from trivial low-level features. To mitigate this issue, we compare the brain-scores obtained with all layers $L \in [1, 5]$ to those obtained with the mel filterbanks (denoted as L_0 , Figure 2.B):

$$\Delta R = R(X_{L \in [0,5]}) - R(X_0) \quad (5)$$

We perform a similar comparison for each pair of successive layers to identify the layer-specific representations that

Figure 2. **A** Average (across subjects and random seeds) brain scores obtained by the random initialisation model. **B**. Average (across subjects and random seeds) gains in brain scores (ΔR) between the random initialisation model and the mel filterbanks, using a concatenation of all sound features and activations. Gray zones are not statistically significant across subjects after FDR correction for multiple comparison. **C**. Average (across subjects and random seeds) gains in brain scores (ΔR) obtained between four representative layers of the random initialisation models. **D**. Average (across voxels and random seeds) brain scores obtained for each level of concatenation for the random initialisation models. Error bars indicate SEM across subjects. Vertical lines and stars indicate significant gains across subjects. Redundant statistical comparisons are omitted for clarity purposes (e.g. $C > A$ is omitted when $C > B$ and $B > A$). **E**. Same as (B) for the networks trained on Dutch. **F**. Same as (C) for the networks trained on Dutch.

correlate with brain responses (Figure 2.C):

$$\Delta R_L = R(X_{[0,\dots,L]}) - R(X_{[0,\dots,L-1]}) \quad (6)$$

The results show that the first layers of the model lead to a significant gain in brain prediction around the primary and secondary auditory areas as well as in the superior temporal gyrus, whereas the deepest layers led to significant gains in a distributed set of brain areas. The average brain score obtained by the random initialisation model for each level of concatenation can be seen in Figure 2.D.

Overall, these results suggest that 1) the representations generated by random convolutional networks suffice to predict a significant proportion of the brain representations of speech *above-and-beyond* mel filterbanks and 2) that the hierarchy of the random neural networks automatically map onto the hierarchy of auditory processing in the cortex.

3.2. Training on Dutch increases the brain-similarity of neural networks

Does training make networks’ representation more brain-like? To tackle this issue, we follow the same brain-scores analysis with networks trained on a Dutch speech-to-text task (Figure 2.E and 2.F). Overall, Dutch models better predict brain responses than random networks (Figure 3.A). Layer-wise comparison of the brain scores between the random networks and the trained networks shows that training leads to a significant gain for each layer but the first one ($p < 10^{-6}$ across all subsequent layers, Figure 3.B).

3.3. Impact of the phonetic distance between training and processed speech on brain-similarity

Training with Dutch increases brain-similarity. Is this effect due to the generation of language-specific representations? To address this issue, we compare networks trained with either Bengali or English, while matching the architecture, training procedure and data exposure to the Dutch networks (see Section 2.4).

Training on English and Bengali induces a significant increase of brain-scores compared to random networks (Figure 3.C). Training on English does not produce significantly higher brain scores than training on Dutch ($p = 0.22$). This lack of effect is reassuring, as it confirms that Dutch models do not implicitly capture lexico-semantic features. On the other hand, the fact that English and Dutch networks achieve similar brain-scores suggests that the sublexical representations of speech in the brain are *not* language specific.

Finally, training on Bengali produces a significantly lower gain in brain scores than training on either Dutch ($p < 0.01$) or English ($p < 0.01$). Given that Dutch and English are closer phonetically than Dutch and Bengali (Schepens et al., 2020), these results suggests that the brain generates

phoneme-specific representations.

3.4. The first levels of the cortical hierarchy are most similar to sound-generic networks

Does training on a non-speech task also increase brain-similarity?

To tackle this issue, we now focus on the networks trained on a curated acoustic scene classification task. On average across all voxels (and across initialization seeds), the brain scores of the acoustic scene networks are not significantly different from those of the Dutch ($p = 0.2$) and English networks ($p = 0.2$).

However, the predictions of the brain activity in the first (A1) and secondary (Belt) auditory cortices (see Figure 3.D) are significantly higher when using acoustic scene classification networks than when using English (Mean A1: $p < 10^{-9}$, Mean Belt: $p < 10^{-6}$) and Dutch networks (Mean A1: $p < 10^{-7}$, Belt: $p < 10^{-5}$). On the contrary, for the temporoparietal junction, where parts of the higher level processing of speech take place (Hickok & Poeppel, 2007), the training on Dutch leads to significantly higher gain in brain score (Mean TPJ: $p < 0.015$) than the acoustic scene classification training.

This result suggests that the brain uses sound-generic representations for the first processing stage of its hierarchy, and then builds speech-specific representations in higher-level processing stages.

4. Discussion

By focusing on speech representations, the present study provides four empirical contributions to the investigation of auditory representations in brains and deep learning models (Yamins & DiCarlo, 2016; Keshishian et al., 2020; Huang et al., 2019; Liu et al., 2021; Kell et al., 2018; Kumar et al., 2020; Koumura et al., 2019).

First, deep convolutional architectures appear to already account for the hierarchy of brain responses to speech, in that their first and deepest layers linearly map onto the primary auditory and associative cortices, respectively. This result suggests that high-level auditory representations can benefit from an inductive bias imputable to a simple and biologically plausible architecture. Indeed, unlike visual conv nets, which require neurons with different visual receptive fields to share weights during training (Yamins & DiCarlo, 2016), audio conv nets apply their convolution over time and frequency – a ubiquitous operation in sensory and associative cortices (David & Shamma, 2013).

Second, training networks with a variety of auditory input appear to systematically improve brain scores. This result suggests that the brain can, in principle, use speech non-

Figure 3. **A.** Average (across subjects and random seeds) gains in brain scores (ΔR) between the models trained on Dutch and the random networks (untrained), using a concatenation of all sound features and activations. Gray zones are not statistically significant across subjects after FDR correction for multiple comparison. **B.** Average (across voxels and random seeds) gains in brain scores between the trained model and the random model (untrained) for each layer independently. Error bars indicate SEM across subjects. Stars indicate significant gains across subjects. **C.** Average improvement in brain score (ΔR) obtained by the models trained –from left to right– on Bengali, acoustic scene classification (ac. sc.), English, and Dutch, compared to their random initialisation (untrained model). Vertical lines and stars indicate significant gains across subjects. **D.** Comparison of the effect of training on Dutch and acoustic scenes for three regions of interest: the primary auditory cortex (A1), the secondary auditory cortex (Belt) and the temporo-parietal junction (TPJ). Position of the three regions on the left hemisphere of the brain is visible under each graph.

specific representations to process sounds. This finding encourages research to test the utility of pretraining speech recognition systems with non-speech data, in order to extract more (and potentially better) features.

Third, training with Dutch or on English – i.e., a language phonetically-related to Dutch – achieved the highest brain scores. This results provides empirical evidence against the hypothesis that phonetic features are reducible to generic auditory representations (Galantucci et al., 2006; Daube et al., 2019). On the contrary, it suggests that learning speech leads the brain to generate specific phonetic representations.

Finally, training on acoustic scene classification led to higher brain scores in the primary and secondary auditory cortex than speech models. This finding supports the idea that the brain builds speech representations “on top” of generic representations, and fine-tunes of “recycles” (Dehaene & Cohen, 2011) higher-level processing stages.

The brain scores reported throughout the study are low. This limitation can be accounted for by two main factors. First, we do not report the absolute brain scores but the difference of brain scores between different layers or different models (Figure 2.A for an example of full brain score, reaching up to $R=0.42$ in the superior temporal lobe). Second, fMRI responses are notoriously noisy, especially at the single-sample and single-voxel level. Fortunately, the large number

of subjects and stimuli allows these weak effects to achieve reasonable statistical significance.

To our surprise, we observe significant brain scores in visual regions (e.g. Figure 3.A). We can speculate that this phenomenon may either relate to i) feedback projection to visual areas specialized in orthographic representations (Dehaene & Cohen, 2011) and/or ii) an auditory-dependent processing of the fixation cross (Petro et al., 2017; Majka et al., 2019). However, this unexpected result requires temporally-resolved brain recordings to confirm and clarify the nature of such auditory-induced visual responses.

Overall, our study strengthens the mutual relevance of machine learning models and neuroscience. While the present investigation remains to be generalized to more efficient audio processing models (Baevski et al., 2020; Collobert et al., 2016), it suggests, once again, that some of the computational solutions designed by machine learning researchers, may actually converge to or re-discover those implemented in the human brain (Marblestone et al., 2016; Caucheteux & King, 2020).

References

Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., and

- Varoquaux, G. Machine learning for neuroimaging with scikit-learn. *Frontiers in neuroinformatics*, 8:14, 2014a.
- Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., and Varoquaux, G. Machine learning for neuroimaging with scikit-learn. *Frontiers in neuroinformatics*, 8:14, 2014b.
- Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case, C., Casper, J., Catanzaro, B., Cheng, Q., Chen, G., et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pp. 173–182, 2016.
- Baevski, A., Zhou, H., Mohamed, A., and Auli, M. wav2vec 2.0: A framework for self-supervised learning of speech representations. *arXiv preprint arXiv:2006.11477*, 2020.
- Baumann, T., Köhn, A., and Hennig, F. The spoken wikipedia corpus collection: Harvesting, alignment and an application to hyperlistening. *Language Resources and Evaluation*, Jan 2018. ISSN 1574-0218. doi: 10.1007/s10579-017-9410-y. URL <https://doi.org/10.1007/s10579-017-9410-y>.
- Behzadi, Y., Restom, K., Liau, J., and Liu, T. T. A component based noise correction method (compcor) for bold and perfusion based fmri. *Neuroimage*, 37(1):90–101, 2007.
- Caucheteux, C. and King, J.-R. Language processing in brains and deep neural networks: computational convergence and its limits. *BioRxiv*, 2020.
- Collobert, R., Puhersch, C., and Synnaeve, G. Wav2letter: an end-to-end convnet-based speech recognition system. *arXiv preprint arXiv:1609.03193*, 2016.
- Daube, C., Ince, R. A., and Gross, J. Simple acoustic features can explain phoneme-based predictions of cortical responses to speech. *Current Biology*, 29(12):1924–1937, 2019.
- David, S. V. and Shamma, S. A. Integration over multiple timescales in primary auditory cortex. *Journal of Neuroscience*, 33(49):19154–19166, 2013.
- Dehaene, S. and Cohen, L. The unique role of the visual word form area in reading. *Trends in cognitive sciences*, 15(6):254–262, 2011.
- Destrieux, C., Fischl, B., Dale, A., and Halgren, E. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *Neuroimage*, 53(1): 1–15, 2010.
- Driedger, J., Müller, M., and Disch, S. Extending harmonic-percussive separation of audio signals. In *Proceedings of the 15th International Society for Music Information Retrieval Conference*, pp. 611–616, 2014.
- Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., Kent, J. D., Goncalves, M., DuPre, E., Snyder, M., et al. fmripip: a robust preprocessing pipeline for functional mri. *Nature methods*, 16 (1):111–116, 2019.
- Fischl, B. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.
- Fonseca, E., Pons Puig, J., Favory, X., Font Corbera, F., Bogdanov, D., Ferraro, A., Oramas, S., Porter, A., and Serra, X. Freesound datasets: a platform for the creation of open audio datasets. In *Hu X, Cunningham SJ, Turnbull D, Duan Z, editors. Proceedings of the 18th ISMIR Conference; 2017 oct 23-27; Suzhou, China.[Canada]: International Society for Music Information Retrieval; 2017. p. 486-93. International Society for Music Information Retrieval (ISMIR), 2017.*
- Fonseca, E., Plakal, M., Font, F., Ellis, D. P., and Serra, X. Audio tagging with noisy labels and minimal supervision. *arXiv preprint arXiv:1906.02975*, 2019.
- Galantucci, B., Fowler, C. A., and Turvey, M. T. The motor theory of speech perception reviewed. *Psychonomic bulletin & review*, 13(3):361–377, 2006.
- Graves, A., Fernández, S., Gomez, F., and Schmidhuber, J. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pp. 369–376, 2006.
- Hermjakob, U., May, J., and Knight, K. Out-of-the-box universal Romanization tool uroman. In *Proceedings of ACL 2018, System Demonstrations*, pp. 13–18, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-4003. URL <https://www.aclweb.org/anthology/P18-4003>.
- Hickok, G. and Poeppel, D. The cortical organization of speech processing. *Nature reviews neuroscience*, 8(5): 393–402, 2007.
- Huang, J., Lu, H., et al. Acoustic scene classification using deep learning-based ensemble averaging. pp. 94–98, October 2019.
- Hung, C. P., Kreiman, G., Poggio, T., and DiCarlo, J. J. Fast readout of object identity from macaque inferior temporal cortex. *Science*, 310(5749):863–866, 2005.
- Jung, J.-w., Heo, H.-S., Shim, H.-j., and Yu, H.-J. Knowledge distillation with specialist models in acoustic scene classification. June 2019.

- Kamitani, Y. and Tong, F. Decoding the visual and subjective contents of the human brain. *Nature neuroscience*, 8 (5):679–685, 2005.
- Kell, A. J., Yamins, D. L., Shook, E. N., Norman-Haignere, S. V., and McDermott, J. H. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron*, 98(3):630–644, 2018.
- Keshishian, M., Akbari, H., Khalighinejad, B., Herrero, J. L., Mehta, A. D., and Mesgarani, N. Estimating and interpreting nonlinear receptive field of sensory neural responses with deep neural network models. *Elife*, 9: e53445, 2020.
- King, J.-R., Gwilliams, L., Holdgraf, C., Sassenhagen, J., Barachant, A., Engemann, D., Larson, E., and Gramfort, A. Encoding and decoding neuronal dynamics: Methodological framework to uncover the algorithms of cognition. 2018.
- Koumura, T., Terashima, H., and Furukawa, S. Cascaded tuning to amplitude modulation for natural sound recognition. *Journal of Neuroscience*, 39(28):5517–5533, 2019.
- Kriegeskorte, N. and Kievit, R. A. Representational geometry: integrating cognition, computation, and the brain. *Trends in cognitive sciences*, 17(8):401–412, 2013.
- Kuchaiev, O., Ginsburg, B., Gitman, I., Lavrukhin, V., Li, J., Nguyen, H., Case, C., and Mickevicus, P. Mixed-precision training for nlp and speech recognition with openseq2seq, 2018.
- Kumar, S., Ellis, C., O’Connell, T. P., Chun, M. M., and Turk-Browne, N. B. Searching through functional space reveals distributed visual, auditory, and semantic coding in the human brain. *bioRxiv*, 2020.
- Liu, Y., Neophytou, A., Sengupta, S., and Sommerlade, E. Cross-modal spectrum transformation network for acoustic scene classification. pp. 830–834, 2021.
- Majka, P., Rosa, M. G., Bai, S., Chan, J. M., Huo, B.-X., Jermakow, N., Lin, M. K., Takahashi, Y. S., Wolkowicz, I. H., Worthy, K. H., et al. Unidirectional monosynaptic connections from auditory areas to the primary visual cortex in the marmoset monkey. *Brain Structure and Function*, 224(1):111–131, 2019.
- Marblestone, A. H., Wayne, G., and Kording, K. P. Toward an integration of deep learning and neuroscience. *Frontiers in computational neuroscience*, 10:94, 2016.
- McFee, B., McVicar, M., Balke, S., Lostanlen, V., Thomé, C., Raffel, C., Lee, D., Kyungyun Lee, Nieto, O., Zalkow, F., Ellis, D., Battenberg, E., Yamamoto, R., Moore, J., Ziyao Wei, Bittner, R., Keunwoo Choi, Nullmightybofo, Friesch, P., Fabian-Robert Stöter, , Thassilo, Vollrath, M., Siddhartha Kumar Golu, Nehz, Waloschek, S., , Seth, Naktinis, R., Repetto, D., Hawthorne, C., and CJ Carr. librosa/librosa: 0.6.3, 2019. URL <https://zenodo.org/record/2564164>.
- Mesgarani, N., Cheung, C., Johnson, K., and Chang, E. F. Phonetic feature encoding in human superior temporal gyrus. *Science*, 343(6174):1006–1010, 2014.
- Moore, A. W. and Jorgenson, J. W. Median filtering for removal of low-frequency background drift. *Analytical Chemistry*, 65(2):188–191, 1993. doi: 10.1021/ac00050a018.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- Petro, L., Paton, A., and Muckli, L. Contextual modulation of primary visual cortex by auditory signals. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1714):20160104, 2017.
- Schepens, J., van Hout, R., and Jaeger, T. F. Big data suggest strong constraints of linguistic similarity on adult language learning. *Cognition*, 194:104056, 2020.
- Schoffelen, J.-M., Oostenveld, R., Lam, N. H., Uddén, J., Hultén, A., and Hagoort, P. A 204-subject multimodal neuroimaging dataset to study language processing. *Scientific data*, 6(1):1–13, 2019.
- Stevens, S. S. and Volkman, J. The relation of pitch to frequency: A revised scale. *The American Journal of Psychology*, 53(3):329–353, 1940.
- Stevens, S. S., Volkman, J., and Newman, E. B. A scale for the measurement of the psychological magnitude pitch. *The journal of the acoustical society of america*, 8(3): 185–190, 1937.
- Unni, V., Joshi, N., and Jyothi, P. Coupled training of sequence-to-sequence models for accented speech recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8254–8258. IEEE, 2020.
- Van Essen, D. C. A population-average, landmark-and surface-based (pals) atlas of human cerebral cortex. *Neuroimage*, 28(3):635–662, 2005.
- Yamins, D. L. and DiCarlo, J. J. Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience*, 19(3):356–365, 2016.

Zhang, S., Do, C.-T., Doddipatla, R., and Renals, S. Learning noise invariant features through transfer learning for robust end-to-end speech recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7024–7028. IEEE, 2020.