Modeling neural coding in the auditory brain with high resolution and accuracy

Fotios Drakopoulos[1]*, Shievanie Sabesan[1], Yiqing Xia[1], Andreas Fragner[2], Nicholas A Lesica[1]

[1]Ear Institute, University College London, United Kingdom, [2]Perceptual Technologies, United Kingdom

*Correspondence: f.drakopoulos@ucl.ac.uk

Abstract

Computational models of auditory processing can be valuable tools for research and technology development. Models of the cochlea are highly accurate and widely used, but models of the auditory brain lag far behind in both performance and penetration. Here, we present ICNet, a model that provides accurate simulation of neural coding in the inferior colliculus across a wide range of sounds, including near-perfect simulation of responses to speech. We developed ICNet using deep learning and large-scale intracranial recordings from gerbils, addressing three key modeling challenges that are common across all sensory systems: capturing the full statistical complexity of neuronal spike patterns; accounting for physiological and experimental non-stationarity; and extracting features of sensory processing that are common across different brains. ICNet can be used to simulate activity from thousands of neural units or to provide a compact representation of central auditory processing through its latent dynamics, facilitating a wide range of hearing and audio applications.

Introduction

Computational models of sensory systems are widely used in fundamental research. They provide a platform for synthesizing existing knowledge and allow for generation and initial testing of hypotheses without the need for in vivo experiments. They can also facilitate the development of technologies that seek to approximate natural sensory function in artificial systems. But the utility of a model for a particular application depends on the degree to which it provides an accurate simulation of the true biological system at the relevant spatial and temporal scales. Numerous high-resolution models of the cochlea have been developed that faithfully capture its mechanical and electrical properties [1], with widespread uptake across academia and industry. Models of auditory processing in the brain, on the other hand, are generally much less accurate, with even the best missing out on a significant fraction of the explainable variance in sub-cortical and cortical neural activity [2, 3].

Models of the cochlea have typically been designed by hand to mimic the transformations performed by the biological system. This approach works well for the cochlea because it is relatively simple, containing only a few stages with homogeneous elements. The brain, however, with its many interconnected and highly non-linear neural networks, is too complex for such an approach. Instead, models of the brain have traditionally relied on “black box” frameworks with parameters estimated from neural recordings. To facilitate model fitting from limited data, the allowable transformations are typically constrained to be either linear or low-order non-linear [2]. These constraints aid interpretability, but such models lack the capacity to capture the full complexity of auditory processing and, thus, are not accurate enough for many applications.

Advances in deep learning have created a new opportunity to develop accurate models of sensory processing in the brain, making it feasible to fit models that capture arbitrary non-linear transformations directly from data. Recent studies have shown that deep neural network (DNN) models of the sensory periphery can be highly accurate: DNNs perform similarly to biophysical models of the cochlea in predicting a wide range of auditory phenomena [4–7], while models of the retina provide accurate predictions of ganglion cell responses to natural scenes [8]. Initial attempts to build DNN models of sensory processing in the brain have also produced impressive results. Models of primary visual cortex (V1) responses to natural images explained 50-80% of the explainable variance in recorded activity [9, 10], while models of V4 explained 90% [11]. DNN models of primary auditory cortex (A1) responses to natural sounds perform similarly well, explaining 60-70% of the explainable variance in recorded activity [12, 13]. But these models reproduce only low-resolution measurements of activity (calcium transients [9], spike counts over large time windows [10, 11], field potentials [13] or fMRI voxels [12]), ignoring potentially important information in the neural code that is present at fine spatial and temporal scales.

We recently developed a DNN model of the inferior colliculus (IC), the midbrain hub of the central auditory pathway, to simulate neural spike patterns with millisecond precision [14]. To bring the performance of this model up to the level required for it to serve as a useful general model of central auditory processing or as a valid in silico surrogate for the IC in auditory neuroscience, there are three key challenges that must be addressed. The first is to capture the full statistical complexity of neural activity. Models of neural coding typically rely on regression frameworks with assumed distributions (e.g., Poisson) that constrain the relationship between response strength and reliability [2]. While these assumed distributions might be valid for some higher-level brain areas, they are a poor match for early sensory pathways where strength and reliability exhibit complex interactions.

The second challenge is to address non-stationarity. The mapping between sensory input and recorded neural activity can change over time for both physiological reasons (e.g., changes in brain state) and experimental reasons (e.g., movement of the recording electrode within the brain). Because training high-capacity models requires the use of large datasets recorded over long periods of time, non-stationarity is inevitable and must be explicitly accounted for. The final challenge is to overcome the idiosyncrasy of individual brains and the recordings made from them. While models must be trained on individual recordings, the final model of a sensory brain area should reflect only the processing that is common across all healthy individuals.

To address these challenges, we developed a novel encoder-decoder framework for modeling neural coding. The framework comprises a shared encoder that maps a sound waveform to a generic latent representation, followed by separate decoders that map the generic representation to the neural activity of individual brains. Each decoder also receives a second input, specific to each brain, that it uses to modulate the generic representation based on external variables. The final model, which we call ICNet, can simulate neural activity across many brains in response to a wide range of sounds with high accuracy and fast execution, and its latent dynamics provide a compact representation of auditory processing that can serve in a variety of sound engineering applications.

Results

To collect the dataset required for model training, we recorded neural activity from the IC of nine normal-hearing anesthetized gerbils, a common animal model for human hearing, using custom-designed electrode arrays with 512 channels [14, 15]. We presented a wide variety of sounds to each animal across a range of intensities, including speech in quiet and in noise, processed speech, music, environmental sounds, ripples and tones that together totalled more than 10 hr in duration. We processed the recordings to extract multi-unit activity (MUA) spike counts for each recording channel using 1.3 ms time bins [14, 16]. We included

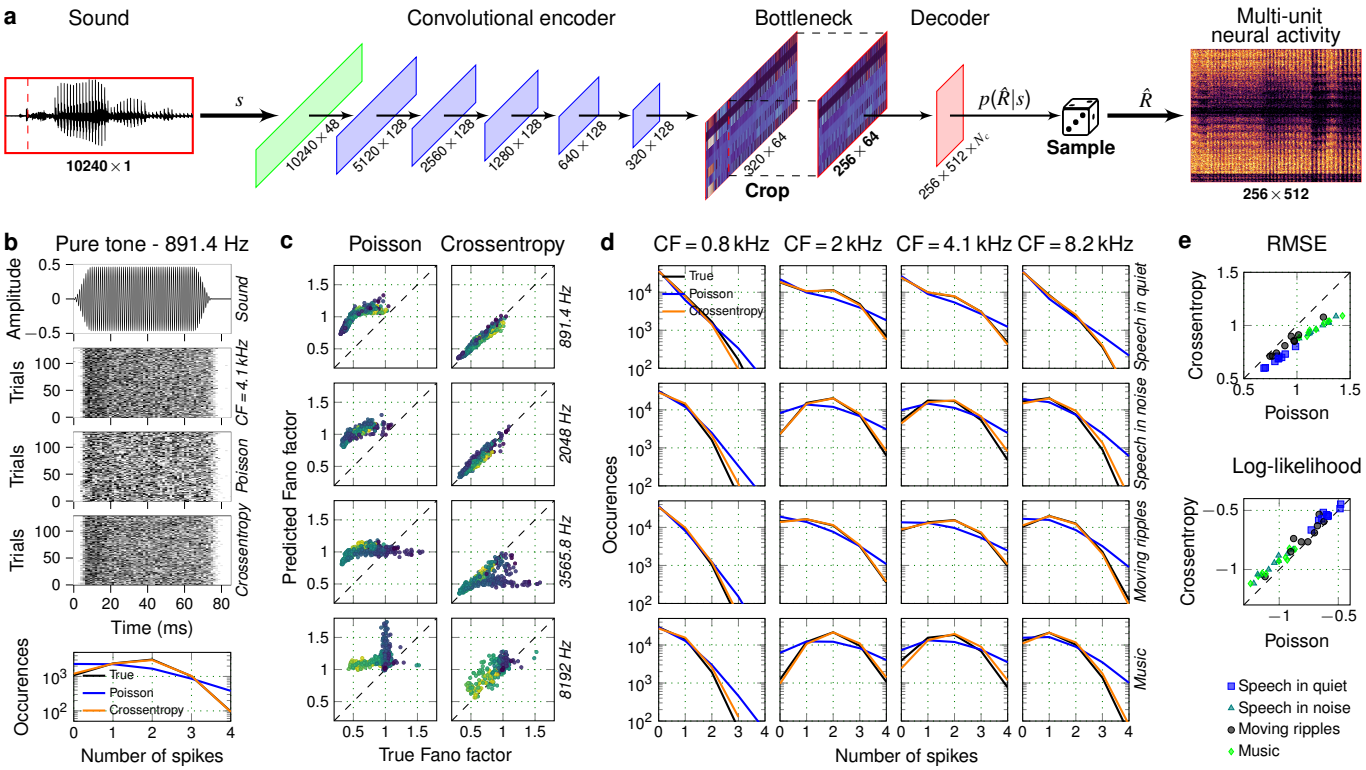

Fig 1. Modeling the full statistical properties of neural activity. a. Schematic diagram of the deep neural network (DNN) architecture trained to map a sound waveform to neural activity for 512 units from the inferior colliculus (IC) of a single animal. N_c is the number of parameters in the probability distribution $p(\hat{R}|s)$, which is 1 for Poisson models and 5 for crossentropy models. **b.** Raster plots of recorded and predicted neural activity for one unit (characteristic frequency (CF) of 4096 Hz) in response to 128 presentations of a 891.44 Hz pure tone at 85 dB SPL. Neural activity was simulated using either a Poisson or crossentropy model. The bottom panel shows the distribution of spike counts across trials for the responses in the raster plots, averaged across time bins during the presentation of the tone. **c.** Comparison of the variance-to-mean ratio (Fano factor) between recorded and predicted neural activity (483 units from one animal) for the Poisson and crossentropy models (columns) in response to 4 tones (rows). All points are colored based on the CF of the corresponding unit, with brighter colors indicating higher CFs. **d.** Comparison of spike count distributions for recorded and predicted neural activity for 4 units (columns) in response to 4 sounds (rows). **e.** Performance comparison across 9 animals and 4 sounds between Poisson and crossentropy models. Each symbol represents the average across all time bins and units for one animal in response to one sound.

all recorded units in model training, but only those units with a signal correlation (across repeated trials of broadband noise) of 0.2 or higher are included in our analyses (494 $\pm$ 16.7 units from each animal, resulting in a total of 4446 units across all animals).

DNNs can capture the full statistics of neural spike patterns

We first trained a set of baseline models to simulate activity across all units for each individual animal through Poisson regression [14]. These models use a cascade of convolutional layers to encode a sound waveform s to a latent representation (bottleneck), which is then passed to a simple decoder (linear readout) to predict the response $\hat{R}$ for all recorded units (Fig. 1a). The output of the decoder for each unit in each time bin is the rate parameter λ that reflects the expected spike count given the recent sound input (N_c = 1 in the decoder layer). This parameter defines a probability distribution $p(\hat{R}|s)$, from which we can then sample to simulate the activity $\hat{R}$ of each unit across time. Once trained, the DNN can be used to simulate neural activity for any sound, whether it was presented during neural recordings or not, with the activations in the

bottleneck layer providing a compact representation of the underlying neural dynamics [14].

The Poisson assumption is common in models of neural coding. It provides a theoretical basis for the optimal estimation of model parameters via classical statistical methods, and it is a reasonable approximation of the true distribution of activity in some brain areas. But activity in sub-cortical areas like the IC is typically underdispered, i.e., it is much more reliable across trials than is expected for a Poisson process [17, 18]. This phenomenon is illustrated in Fig. 1b, which compares recorded and simulated neural responses for an example IC unit to 128 trials of a pure tone. The Poisson model captured the mean activity across trials well, but overestimated the variability across trials, overpredicting both low and high counts.

Fortunately, DNNs can be trained to predict full activity distributions without any assumptions by framing the problem as classification rather than regression [19]. We trained the same DNN architecture for each animal with a crossentropy loss and categorical distributions with $N_c = 5$ discrete classes (corresponding to the probability of each unit firing 0 to 4 spikes within a given time bin). The trained crossentropy model accurately captured both the mean activity and the variability across trials for the example unit (Fig. 1b). For all time bins within the duration of the tone, the Fano factor (ratio of variance to mean) was 0.55 for the true response, 0.54 for crossentropy model, and 1.16 for the Poisson model. (Note that because the rate of the Poisson model varies across time bins, the overall Fano factor does not need to be exactly 1.) For all units from one example animal in response to four different tone frequencies, the Fano factor for the crossentropy model was much closer to the true value than the Fano factor for the Poisson model (Fig. 1c; the average absolute error was 0.38 for the Poisson model and 0.07 for the crossentropy model).

To compare the overall performance of the Poisson and crossentropy models, we simulated responses for all animals to four different sounds: speech in quiet, speech in noise, dynamic moving ripples, and music. As shown in Fig. 1d for four example units, the crossentropy model accurately predicted the full spike count distributions for each sound. We compared the overall performance of the models for each animal and sound using two metrics (averaged across all time bins and units): the root-mean-square error (RMSE) between recorded neural activity R and predicted activity $\hat{R}$, and the logarithm of the probability (log-likelihood) of observing the recorded activity R from the inferred probability distribution $p(\hat{R})$. We found consistent improvements for the crossentropy models across all animals and sounds (Fig. 1e; improvement of 15.98% [15.88, 16.02] for RMSE (median and 95% confidence intervals over bootstrap samples) and 8.13% [8.04, 8.16] for log-likelihood). These results demonstrate that characterizing the full distributions of spiking activity for each neural unit is feasible and improves the accuracy of simulated neural coding in the IC.

DNNs can account for non-stationarity in neural recordings

Capturing the full non-linearity of auditory processing requires the use of high-capacity models, which in turn requires large datasets recorded over long periods of time. Non-stationarity in such recordings is inevitable, as the physiological state of the animal or the electrode position can vary over the course of the recording, altering the mapping from sound to neural activity expressed by individual units. Figure 2a illustrates this for two of our recordings, where neural activity in response to the same speech was either stable (stationary) or variable (non-stationary) over the 10 hours of recording. While there is an overall trend of decreased activity in the second example, the non-stationarity can vary widely across individual units (Fig. 2b), making it challenging to account for its effects within a model.

A standard encoder-decoder model can only capture a time-invariant encoding that will reflect the average mapping from sound to neural activity over a large dataset. To overcome this limitation, we designed a time-variant version of the crossentropy model that takes an additional input (alongside the sound input) that reflects the time during recording at which the sound was presented (Fig. 2c). The interactions between the two inputs within the encoder enable the model to alter its sound encoding as needed to account for non-stationarity. After the model is trained, neural activity that reflects the sound encoding at specific time

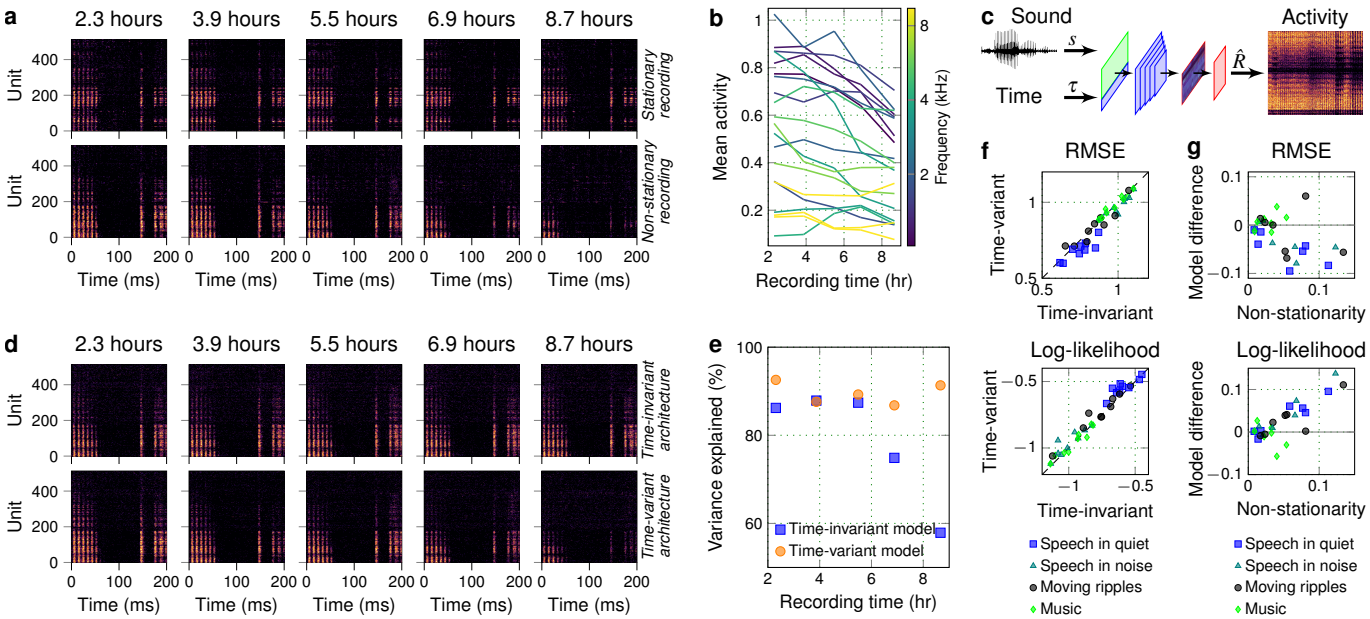

Fig 2. Accounting for non-stationarity in neural recordings. a. Example neural activity in response to the same sound (speech at 60 dB SPL) presented at five different times within two individual recordings. Units were ordered based on their CFs. Brighter colors indicate higher activity. **b.** The mean spike rate of 20 randomly selected units from the non-stationary recording shown in **a**. **c.** Schematic diagram of the time-variant DNN architecture. **d.** Predicted neural activity for the non-stationary recording in **a** using the time-invariant and the time-variant models. **e.** Median predictive power of the time-invariant and time-variant models across all units at five different times during the recording. **f.** Performance comparison across 9 animals and 4 sounds between time-variant and time-invariant DNN models. **g.** Performance difference between models (time-variant - time-invariant) as a function of recording non-stationarity across 7 animals and 4 sounds.

points during the recording can be simulated by providing the appropriate additional inputs (or the simulation can made stationary by keeping the additional input constant across all sounds).

As a first test, we trained the time-invariant and time-variant models to predict the neural activity of the example non-stationary recording, and evaluated their performance in predicting the responses to speech shown in Fig. 2a. The time-invariant model, by construction, predicts the same responses each time the speech is presented. But the time-variant model learned to utilize the additional input to capture the changes over the course of the recording, predicting strong responses at early recording times and weak responses at late recording times (Fig. 2d). We quantified the differences in model performance for this example by computing the fraction of explainable variance that each model explained across all recorded units for each of the five speech presentations. The results demonstrate that the time-variant model can capture the changes in the mapping from sound to neural activity over time, consistently explaining 89.5% ($\pm$ 2.4) of the explainable variance across the entire recording (Fig. 2e), while the performance of the time-invariant model was much more variable and worse overall (78.8 $\pm$ 12.9%).

Across all animals and sounds, we found that the time-variant models improved performance by an average of 2.67% [2.60, 2.69] for RMSE and by 3.22% [3.11, 3.26] for log-likelihood (Fig. 2f). But we did not expect performance improvements across all animals, as the time-variant model should only achieve better predictions for recordings that are non-stationary. We categorized the recordings for which we had repeated trials of the same sounds presented at different times in the recording (seven out of nine animals) based on their non-stationarity. We measured non-stationarity as the absolute difference of the covariance of neural activity between two successive trials and two trials recorded several hours apart. As expected, we found that the benefit of the time-variant model increased with non-stationarity for both the RMSE and the log-likelihood measures (Fig. 2f), approaching 10% improvement on RMSE and 15% improvement on

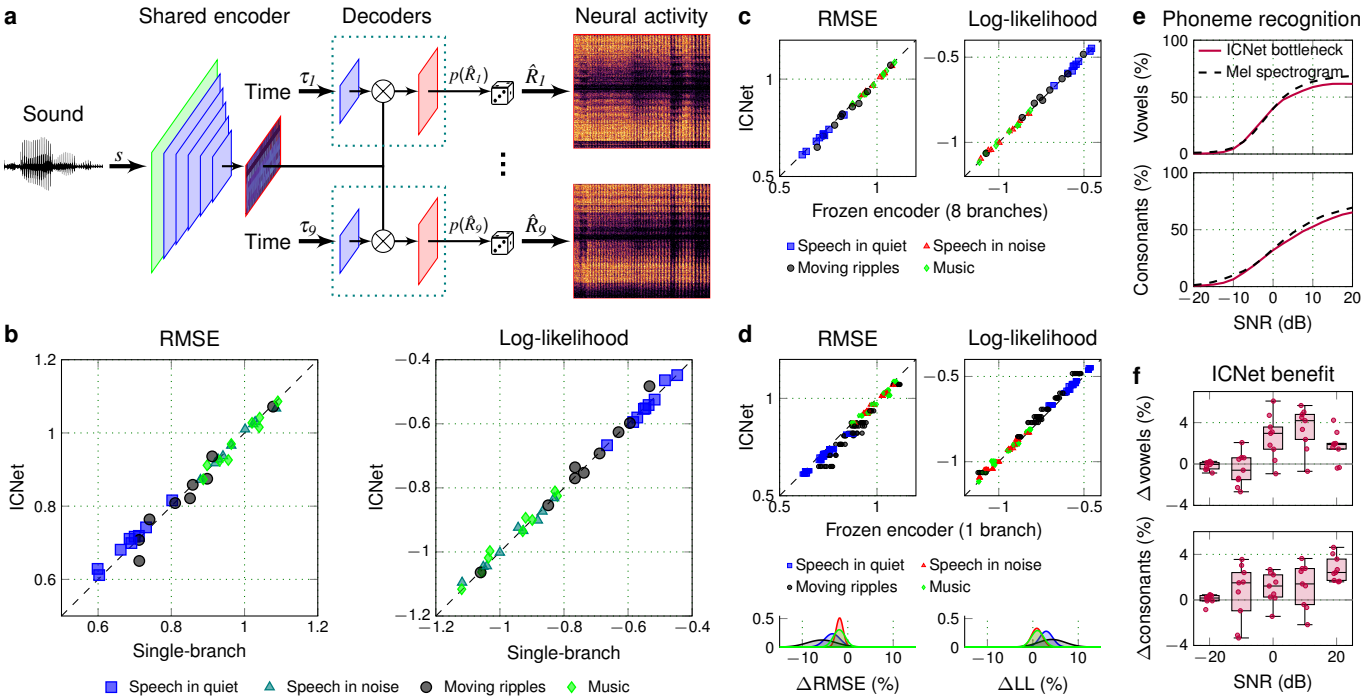

Fig 3. Capturing shared sound encoding across animals. a. Schematic diagram of the ICNet model trained to map a sound waveform to the neural activity of 9 animals. **b.** Performance comparison for 9 animals between ICNet and the corresponding single-branch models. **c.** Cross-validation comparison between ICNet and individual eight-branch models with frozen encoders and decoders trained to predict the ninth animal. The similar performance across animals and sounds demonstrates that multi-branch models learn generic latent dynamics that can be transferred to other animals. **d.** Cross-validation comparison between ICNet and individual single-branch models with frozen encoders and decoders trained to predict the remaining eight animals. Distributions in the bottom show the spread of ICNet benefit for each sound compared to all 72 frozen-encoder models. The decrease in performance across animals and sounds demonstrates that single-branch models cannot fully describe the sound encoding of other animals. **e.** Vowel and consonant recognition performance of an automatic-speech-recognition (ASR) model trained to predict phonemes from either the ICNet bottleneck responses or a Mel spectrogram. Phoneme recognition was computed for sentences presented at 65 dB SPL and mixed with speech-shaped noise at SNRs between -20 and 20 dB. **f.** Benefit in recognition performance for the ASR model trained with the ICNet bottleneck response as input, compared to ASR models trained with the bottleneck responses of the 9 single-branch models.

log-likelihood for the most non-stationary recordings. These results demonstrate that time-variant DNNs can flexibly model non-stationarity, improving performance for recordings that are unstable without compromising performance for recordings that are stable.

Multi-branch DNNs can capture shared latent dynamics across animals

Our goal was not to model any individual animal, but rather to develop a generic model of the IC that describes the features of sound encoding that are common across all animals. Recordings from any individual animal will inevitably contain idiosyncrasies that reflect variations in anatomy or physiology, or in the position of the recording electrode within the IC. These idiosyncrasies may be critical to preserve when developing models for applications such as hearing aid development, which requires sound processing to be unique to each individual hearing loss profile [20]. But for applications that require a generic model of normal auditory processing, architectures that can effectively make use of recordings from different individuals to extract what is common across them are more suitable.

We have previously demonstrated that the acoustic information in IC activity is restricted to a low-

dimensional subspace which is shared across animals with the same hearing status [14]. This result suggests that it should be possible for a DNN with shared latent dynamics to predict neural activity for any normal-hearing animal. We thus designed a multi-branch DNN architecture comprising a shared encoder and separate decoder branches that are used to simulate neural activity for each of nine individual animals (Fig. 3a). To simulate activity for all animals with a shared encoder, we also had to modify the way in which the additional time input was used so that the non-stationarity of each individual recording could be accounted for. We introduced animal-specific time inputs at the level of the decoder and combined each time input with the bottleneck output, allowing for a unique time-variant mapping from the generic latent representation to the neural activity for each animal. After training, the multi-branch model, which we call ICNet, can simulate neural activity for nine individual animals in response to any sound in a manner that refects the state of each animal at a given time.

We first evaluated ICNet by simulating neural activity in response to four sounds (with the time input matched to the time at which each sound was presented to each animal). ICNet matched the performance of the corresponding single-branch models for all nine animals (Fig. 3b; average RMSE increase of 0.04% [0.01, 0.06] and log-likelihood decrease of 0.39% [0.34, 0.41] compared to time-variant models trained to predict the activity of a single animal). To establish that ICNet encoded neural dynamics that were shared across all animals (rather than the idiosyncrasies of each), we performed a cross-validation experiment. We held out each of the nine individual animals in turn while training an eight-branch model to predict the neural activity of the remaining animals. We then froze the encoder of each eight-branch model and trained a new decoder that used the output of the frozen encoder to predict the neural activity of the held-out animal. As shown in Fig. 3c, the models that used a frozen encoder and an optimized decoder performed similarly to the original ICNet (1.19% [1.16, 1.20] decrease of RMSE and 0.41% [0.38, 0.42] increase of log-likelihood for ICNet compared to the frozen-encoder models), which suggests that ICNet performs a generic encoding that can predict the neural activity of any normal-hearing animal.

We then reversed the cross-validation experiment to establish the benefit of our multi-branch DNN architecture, i.e., the degree to which the single-branch models captured the idiosyncracies of each animal without fully capturing the neural dynamics that are shared across animals. We froze the encoder of each of the nine single-branch models and trained eight new decoders to predict the neural activity of the remaining animals. As shown in Fig. 3d, the difference in the performance of ICNet and these frozen-encoder single-branch models was small, but the frozen-encoder models were almost always worse, especially for speech in quiet and ripples (average RMSE decrease of 3.07% [3.05, 3.07] and log-likelihood increase of 2.01% [2.00, 2.02] for ICNet compared to the frozen-encoder models). Taken together, the results of these cross-validation tests establish that ICNet is a generic model of sound encoding.

Because we force the output of the sound encoder through a shared low-dimensional bottleneck and constrain the individual decoder branches to simple linear readouts, the bottleneck can learn a generic latent representation of auditory processing that can be used on its own. As a demonstration, we trained a simple automatic-speech-recognition (ASR) DNN to classify phonemes from the ICNet bottleneck representation across time. We found that the ASR system trained on the ICNet bottleneck achieved similar performance to a baseline system trained with a Mel spectrogram input (Fig. 3e). We also found that training on the ICNet bottleneck representation resulted in better performance than training the same ASR system with the bottleneck representations from the nine single-branch models, especially at relatively high signal-to-noise ratios (SNRs) (Fig. 3f; average improvement of 1.24% [1.02, 1.76] for consonant recognition and 1.60% [1.25, 2.22] for vowel recognition across all SNRs and animals). Together, these results suggest that the latent representation of ICNet can serve in the same manner as widely used acoustic front-ends such as auditory filterbanks or cochlear models for applications where simulation of central (rather than peripheral) auditory processing is desired.

ICNet is a highly accurate model of neural coding

Figure 4a shows example segments of recorded activity and predictions from ICNet for four sounds, with each image showing the response of 1000 units (the units were selected to achieve a uniform distribution of characteristic frequencies (CFs) on a logarithmic scale, reflecting the tonotopic organization of the IC). ICNet's simulations were nearly indistinguishable from the recordings across all units and all sounds. The full non-linearity of the mapping from sound to neural activity that is captured by ICNet is difficult to appreciate from visual inspection, but Fig. 4a illustrates some of its remarkable features. For example, the speech in quiet elicited the strongest activity even though it was presented at the lowest intensity of any of the sounds (60 dB SPL with energy concentrated at low frequencies), while the speech in noise, which was presented at a much higher intensity (85 dB SPL with a much broader range of frequencies), elicited significantly less activity. The response to the moving ripples was particularly complex: the sound was presented at high intensity (85 dB SPL) with all energy concentrated at high frequencies, yet it elicited relatively weak activity in units with high CFs and relatively strong activity in units with low CFs. This frequency transposition indicates a high degree of non-linearity that was captured well by ICNet.

To provide an overall measure of the predictive power of ICNet, we used seven of the nine animals for which we had recordings of responses to the same sounds on successive trials. We computed the fractions of the explainable variance and correlation (covariance and correlation of responses on successive trials) that ICNet explained for sounds that were not included during training (measured after flattening the responses by collapsing across time and units for each animal; see Methods). We found that ICNet performed best for speech in quiet, explaining $\sim$92% of the explainable variance and correlation across all animals (Fig. 4b; 91.5% [91.4, 91.8] for variance and 92.9% [92.9, 93.1] for correlation). Performance was also excellent for speech in noise (83.6% [83.3 84.3] for variance, 91.2% [91.1 91.6] for correlation) and for music (83.6% [83.3 84.3] for variance, 91.3% [91.2 91.7] for correlation), and decreased only for the band-limited moving ripples (64.4% [64.0, 65.2] for variance, 73.5% [73.3, 74.0] for correlation). To further test the generality of ICNet, we also measured performance for four more sounds that were presented on successive trials for four of the nine animals: speech in quiet at a higher intensity, speech in noise at a higher signal-to-noise ratio (SNR), broadband moving ripples, and a solo violin recording. We found that ICNet again explained most of the explainable variance and correlation for all sounds (Fig. 4c; median values for variance and correlation between 88% and 93%).

The flattened metrics computed over all units are appropriate for assessing the overall predictive power of the model, but are dominated by the units with the most explainable variance or correlation (which are, presumably, the most important units to model accurately because their activity carries the most information about sound). When we instead examined ICNet's performance across individual units, we found lower values overall and substantial variability across units and sounds (Fig. 4d). The lower values are a result of the difference in the way the flattened and unit-by-unit metrics account for the variation in mean spike rate across units for a given sound. This variation is fully accounted for by the flattened metrics but is either partly (for variance) or fully (for correlation) ignored when the metrics are computed unit by unit. This variation makes up a substantial portion of the explainable variance and correlation in the full neural activity and is well captured by the model, so the resulting performance values for the flattened metrics are higher than those for the individual units.

The variability in performance across units could have many sources. From the responses in Fig. 4a, it is clear that units can vary in their activity depending on the match between the acoustic properties of a sound and the features to which a unit is responsive. Both overall activity (measured as mean spike rate) and reliability (measured as the correlation of responses on successive trials) of individual units spanned a wide range that varied across sounds (Fig. 4e), with a clear dependence on CF (color code of each unit). To determine whether any of these properties (overall activity, reliability, preferred frequency) could explain the variation in model performance across units, we measured their correlation with the fraction of explainable

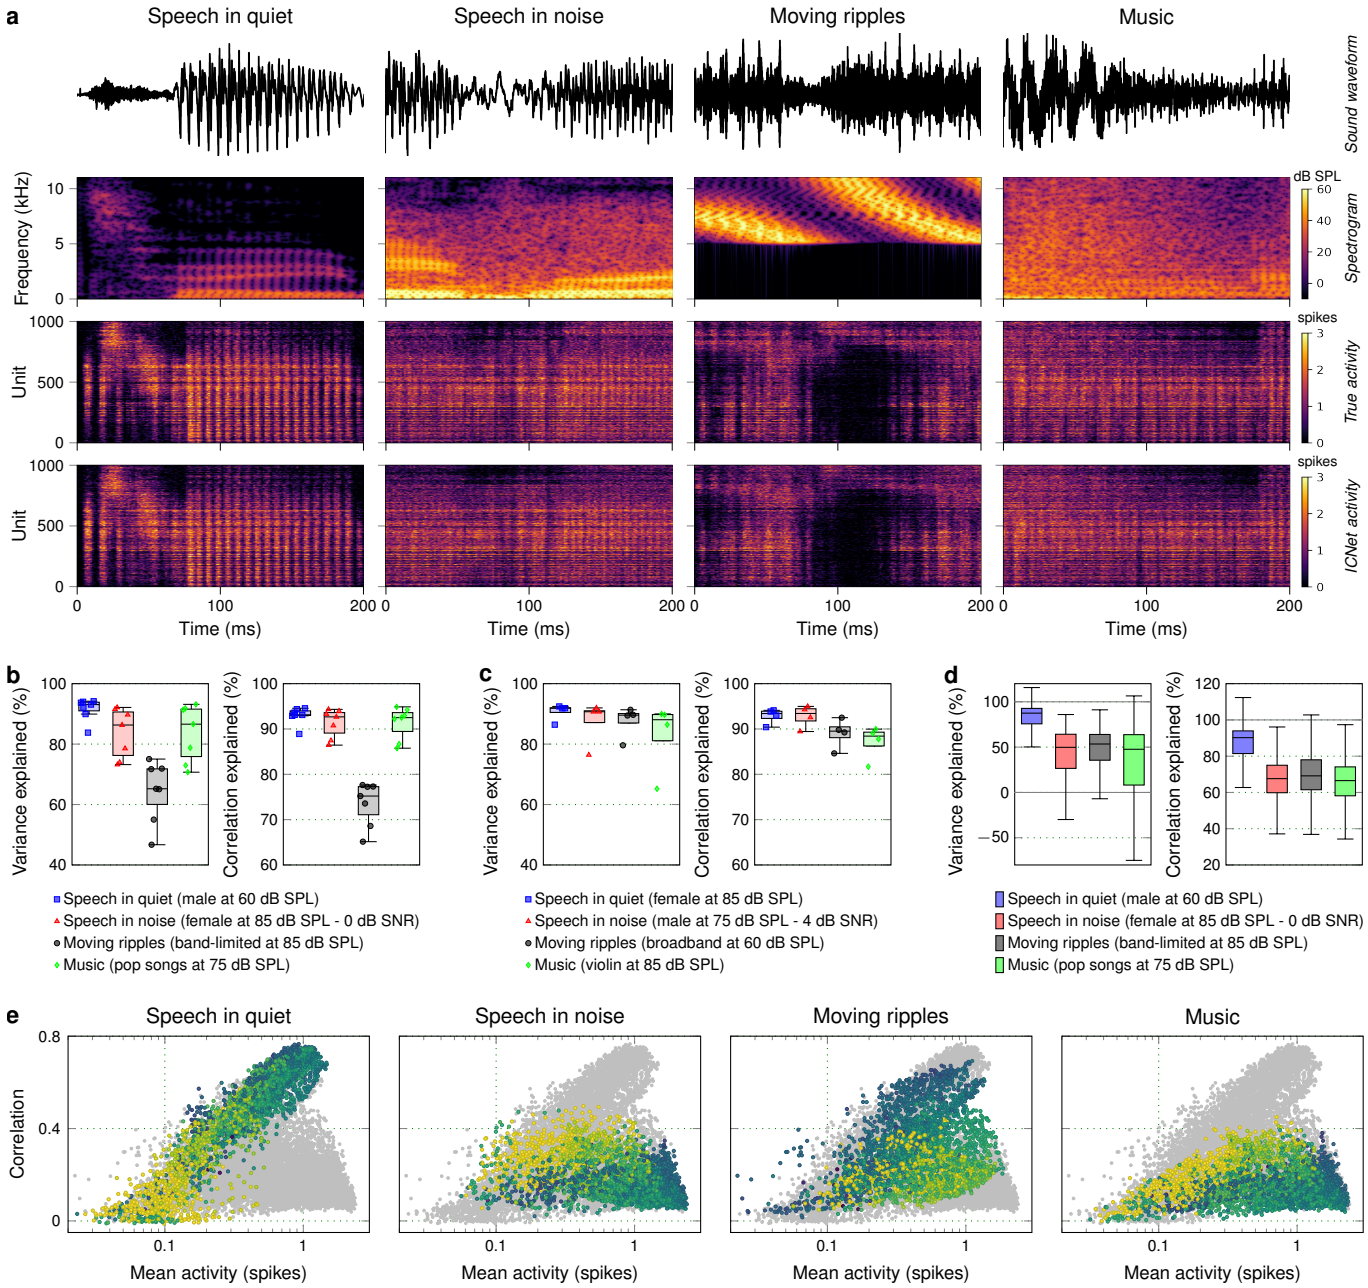

Fig 4. ICNet performance across sounds. a. Segments of recorded and ICNet predicted neural activity in response to 4 sounds. Each column shows the sound stimulus and corresponding spectrogram, as well as the neural responses of 1000 units. Units were selected from all animals to achieve a uniform logarithmic spacing of CFs between 300 and 12000 Hz. Brighter colors indicate higher stimulus intensity for the spectrograms and higher activity for the neural responses. **b.** Predictive power of ICNet across 7 animals and 4 sounds, computed as the fraction of explainable variance and correlation explained by the model across all time bins and units of each animal. **c.** Predictive power of ICNet across 4 animals and 4 sounds, computed as the fraction of explainable variance and correlation explained by the model across all time bins and units of each animal. **d.** Predictive power of ICNet across 4 sounds and 3476 recorded units (from 7 animals), computed as the fraction of explainable variance and correlation explained by the model across all time bins of each unit. **e.** Overall activity (mean spiking activity) is compared to reliability (correlation between two successive trials) of all 3476 units from 7 animals. The gray points correspond to the results across four sounds and are overlaid with the results of each individual sound, colored based on the CF of the respective unit. Brighter colors indicate higher CFs.

correlation explained across all sounds. We found only a weak correlation between overall activity and model performance (-0.12 [-0.14, -0.10] for 3406 units from seven animals, after excluding units with the top and bottom 1% of explainable correlation explained) and no correlation between CF and model performance (-0.02 [-0.03, 0]). But the correlation between reliability and model performance was much stronger (0.54 [0.53, 0.55]). This suggests that the reliability of neural responses to a given sound is a strong determinant of how well ICNet can learn to simulate those responses.

Discussion

In this study, we combined large-scale neural recordings and deep learning to develop ICNet, a comprehensive model of neural coding in the early central auditory pathway. We presented a novel multi-branch encoder-decoder framework to describe the common features of sound encoding across a set of normal-hearing animals, while also accounting for the idiosyncrasy and non-stationarity of each individual recording. ICNet captures the full statistical properties of neural activity with high resolution, simulating spiking across thousands of units with millisecond precision. The simulated responses capture nearly all of the explainable variance in neural activity across a wide range of sounds. And because ICNet is based on convolutional operations, it is computationally fast: neural activity can be simulated more than 10x faster than real time, even on a CPU (for a 420 ms sound input, simulation of 512 units takes ∼49 ms on an Intel Xeon w7-2495X CPU and ∼39 ms on an NVidia RTX 4090 GPU).

Model performance

An ideal model of neural coding would produce activity that is indistinguishable from that recorded in a real brain for any sensory stimulus. ICNet produced near-perfect simulations of responses to speech in quiet, capturing ∼92% of the explainable variance and correlation in recorded activity across all units and animals. This is a substantial improvement over our previous model [14], which captured ∼85% of the explainable variance for speech in quiet. The superior performance of ICNet is particularly remarkable because the previous model was trained only on speech and evaluated only on speech from the same source, whereas ICNet was trained on a wide range of sounds and evaluated only on sounds from sources that were not included in the training set. ICNet also performed well for a range of other sounds, explaining at least 86% of the explainable variance and correlation for speech in noise, music, and broadband ripples. The only sound for which performance was substanially lower was band-limited ripples with frequencies of ∼5-10 kHz, which provide only indirect stimulation to units that prefer low frequencies.

While the performance of ICNet is excellent overall, there is substantial variation in performance across units. Our analysis suggests that the reliability of neural responses, rather than their overall strength, is a major explanatory factor in this variation (even when performance is measured relative to neural reliability as with the explained-explainable metrics that we used). This is unsurprising given that the reliability of a neural response is essentially its signal-to-noise ratio, and additional neural noise makes it harder for the model to learn the relationship between sound and the neural signal. This effect is related to broader concepts in statistics (e.g., risk bounds) that relate the error in the estimate of a distribution from a fixed number of samples to its entropy. It can also be understood intuitively: the distribution of a perfectly reliable response will have a single value with a probability of one, and a single sample will yield a perfect estimate of this distribution; conversely, for a noisy response in which all spike counts are equally probable, many samples will be required before the estimated distribution converges to uniform.

Responses with high reliability are, presumably, the most important for the model to capture, since they carry robust acoustic information. But for applications that require improved model performance for low-reliablity responses, there are many approaches that could be taken. One possibility is to use data pre-

processing to remove neural noise. In IC responses, this could be achieved through standard dimensionality reduction because only the signal is shared across neurons [16]. Representing IC responses using only the leading components identified by PCA would serve to retain most of the signal while discarding most of the noise [14], and it is possible that non-linear methods (e.g., DNN-based autoencoders) could be even more effective. Another approach is to alter the training framework to focus on low-reliability responses. This could be achieved through modifications to the loss function or to the composition of the training dataset. But if an application requires simulation of responses to a particular sound class for which the model is thought to be underperforming, then the most effective approach might be to fine tune the model on a new dataset containing responses to sounds from that class.

In order to produce a model that would generalize well, we assembled a broad training dataset that contained representative samples of many complex sounds, and we demonstrated that ICNet produced accurate predictions of the neural responses to a range of natural and artificial evaluation sounds (Fig. 4a). The validity of ICNet as an accurate description of the auditory system could be further evaluated with respect to specific physiological phenomena that are traditionally used to characterize auditory processing (e.g., otoacoustic emissions as in [7] or phase-locking as in [5]). This evaluation could be performed on the current model with a small additional dataset (to provide a reference for the manifestation of these phenomena in the IC) and the results could be used to inform the design of an updated training dataset for a new iteration of the model if needed.

Prediction of neural activity as a classification problem

Accurate simulation of any system requires an appropriately specified model. The Poisson assumption that is typical for models of neural coding is often a poor match for the true distribution of activity, particularly in sub-cortical areas where activity can be strongly underdispersed [17, 18]. There have been many attempts to capture this underdispersion by explicitly modeling its physiological sources. For example, much of the underdispersion results from refractoriness, and models can be made refractory either by incorporating negative feedback on the Poisson rate parameter or by incorporating an integrate-and-fire spike generation mechanism [21–24]. Other approaches include using a different distribution that is a better match to the true distribution (for example, a Bernoulli distribution contains an additional degree of freedom that allows it to capture underdispersion [19]) or injecting noise that can be used to reshape the distribution [8].

ICNet demonstrates that is it possible to model the full distribution of neural activity without any assumptions. We trained ICNet to infer non-parametric conditional probability distributions using a crossentropy loss which should, in principle, result in network weights that maximize the likelihood of the observed activity given the model. But since the prevalence of the different classes is imbalanced (58.4%, 23.2%, 13.7%, 4.2%, and 0.5% for spike counts from 0 to 4, respectively, across our dataset), there is a risk that model performance will be biased against the under-represented classes. This risk can be addressed by training with a balanced loss, which weights each class according to its prevalence in the dataset, or a focal loss, which weights each class based on its predicted probability. The performance of ICNet did not improve significantly after training with a balanced or a focal crossentropy loss (results not shown), but explicit consideration of class imbalance may still prove benefical when training models of neural coding for other brain areas.

Modeling non-stationary sensory processing

Non-stationarity is a prominent feature of sensory coding. In our dataset, the non-stationarity is due primarily to variations in brain state over the long duration of the recordings. Because our goal was to develop a generic model, we designed ICNet so that the effects of non-stationarity could be separated from the shared sound encoding and individualized for each animal. After training, ICNet can be used to provide a stationary

simulation of the auditory processing that is common across all animals by setting the time input to a fixed value that is shared across animals. This value should correspond to a time near the beginning of recording, e.g., 1 hour, on the assumption that the state of all animals is comparable at the beginning of recording and then diverges over time due to variations in the cumulative effects of the anesthesia.

The specifics of our design also reflect the nature of our modeling problem. We provided the model with a simple measure of time rather than a relevant biomarker (e.g., an EEG signature of anesthetic depth) because our reason for accounting for non-stationarity was ultimately to remove it. If the goal was instead to simulate the modulation of sensory coding by an external variable of interest, a measurement of this variable could be provided to the model in place of the time input. For example, to model the dependence of sensory processing on arousal, a measurement of pupil diameter could be used. Once trained, the resulting model could then, in principle, be used to simulate the response to any stimulus at any level of arousal and to probe the effects of arousal on the neural response.

Our desire to remove non-stationarity also motivated our choice to provide the time input to the model at the level of the decoder. We found that ICNet performed well with a simple architecture that allowed the time input to scale each channel of the bottleneck output independently (through a convolutional layer with a kernel size of 1, which also allowed the model to map the time input to the appropriate dynamic range). For datasets with more pronounced effects of non-stationarity, more complex convolutional layers could be used to process the external input. Or, if the goal is to model the effects of non-stationarity without the need to remove it, the external input can be provided to the encoder along with the sensory input (as we did for our single-branch time-variant models) to capture arbitrarily complex interactions.

Capturing shared neural dynamics

Most applications of auditory models require a generic simulation of the fundamental features of auditory processing that are common across all normal-hearing individuals. Phenomenological or mechanistic models that are designed by hand will naturally have this property, but models that are learned directly from experimental data will inevitablty reflect the idiosyncracies of the individuals from which the data were collected. We designed ICNet with a shared encoder to capture the features of auditory processing that were common across the individual animals on which it was trained, and the results of our cross-validation experiments (Fig. 3c) suggest that the model is fully generic, i.e., that the inclusion of additional animals during training would not improve its generality.

We limited the number of channels in the bottleneck of the shared encoder to discourage the model from capturing the idiosyncrasies of the individual animals, and also to allow for fair comparisons between multi-branch and single-branch models. Despite this limitation, ICNet performed as well as single-branch models in predicting the responses of individual animals (Fig. 3b), which suggests that the differences in recorded activity across individuals reflect only differences in the sub-sampling of the full neuronal populations rather then fundamental differences in underlying neural dynamics. The performance of ICNet was likely enabled by the size of our electrode arrays (both the span and the high number of channels), which have the capacity to sample from neurons across the entire IC. When training models on datasets with a sparser sampling of neuronal populations, the differences between single- and multi-branch models might be more pronounced.

With its limited dimensionality, ICNet's bottleneck layer provides a compact representation of sound encoding that can be used on its own for applications that require a general auditory model rather than explicit neural simulation. We demonstrated that this representation can be effective as a front-end for speech recognition (Fig. 3e), and we expect that even better performance could be achieved in combination with ASR models with higher capacity. Depending on the needs of the application, ICNet's bottleneck could be modified to have specific properties (for example, smoothness or orthogonality could be imposed by making the encoder variational) or the model could be retrained with additional loss terms that optimize the bottleneck for specific tasks. Aside from the architectural considerations related to the time input, all DNN

models that we trained followed the hyperparameter choices that we report in Fig. 1a (and in Methods). More principled approaches such as Bayesian optimization might also be helpful to increase predictive power without significantly increasing model complexity.

Species differences

For a computational model of auditory processing to be useful in most applications, it needs to provide accurate simulation of neural coding in humans, not other animals. Assessing the degree to which ICNet achieves this is challenging, as there are no intracranial recordings from human IC that can be used for direct validation. There are recordings of field potentials from the human auditory nerve, and these have been used together with other measures to compare the functional properties of the cochlea across mammalian species. The results indicate that humans have sharper spectral tuning than other mammals [25], but other properties such as the upper limit of phase locking [26] and the timescales of adaptation [27] appear to be similar.

Overall, the direct evidence that exists suggests that the differences in early auditory processing between humans and other mammals with low-frequency hearing are quantitative rather than qualitative. This suggestion is reinforced by the ability of cochlear models to predict human auditory perception without accounting for species differences. Models designed to simulate the cat cochlea have been successful in predicting perceptual outcomes [28–30], auditory evoked potentials [31] and hearing aid benefit [32] in humans. While we cannot rule out the existence of important differences between humans and other mammals at the level of IC, there is currently no specific evidence to such differences.

Explicit validation of ICNet as a model of neural coding in the human IC will not be possible in the near future, but confidence in its validity could be increased in other ways. For example, the validity of ICNet as a general model of neural coding in non-human mammals could be tested by retraining only the decoder layer on IC recordings from other species. This approach should work well for any species with human-like frequency sensitivity; for species with ultrasonic sensitivity, such as mice, the encoder would also need to be retrained. (There are neurons with sensitivity to ultrasonic frequencies in the most ventromedial region of the gerbil IC, but we did not record from these neurons, nor did we present any sounds that contain ultrasonic frequencies.) It should also be possible to link the output of ICNet to neural coding in humans indirectly. For example, new models (phenomenological, mechanistic, or DNN) could be coupled to ICNet to predict auditory evoked potentials (e.g., EEG or ABR recordings) or perceptual outcomes (e.g., speech recognition or quality). Successful prediction would be only a weak indication of ICNet's validity but failure would suggest that significant modifications are needed.

Possible extensions

We developed ICNet to describe MUA rather than the activity of single units, which should be sufficient for most applications. While this has not been explicitly tested in the IC, studies in other brain areas have shown that the latent neural dynamics inferred from MUA are nearly identical to those inferred from the activity of single-units [33], suggesting that single-unit spiking does not carry any additional information about the underlying neural computations. This result also suggests that if simulations of single-unit activity are required, it should be possible to extend ICNet to produce them simply by retraining the decoder with a relatively small dataset of single-unit recordings. If the goal is to maximize the predictive power of the model for particular neural populations (rather than to infer a general model, as we have done), it may be beneficial to use a larger bottleneck or a more complex decoder to capture the unique intrinsic dynamics of each unit.

It should also be straightforward to extend ICNet to simulate awake activity. We used anesthesia to faciliate the collection of the large-scale recordings required for initial model training. But with the trained ICNet as a starting point, it should be possible to learn the mapping from ICNet's latent representation to

awake activity with a much smaller dataset. The impact of anesthesia in the IC is relatively modest compared to cortex [34, 35], so it is likely that retraining only the decoder (possibly with a more complex architecture) will be sufficient. This applies, of course, only to passive listening; modeling the modulation of auditory processing by attention, arousal, or other contextual factors will require also making use of the external input to ICNet as discussed above.

One critical area for further development is binaural processing. We designed ICNet to capture the mapping between a single sensory input (a diotic stimulus in which the same sound was presented in both ears) and the neural activity recorded from both hemispheres of the IC. But IC activity is also strongly modulated by binaural disparity, and these modulations encode information that is critical for spatial hearing. It is likely that a simple modification to ICNet that allows for two channels of sound input would allow it to capture the full complexity of binaural processing, but the new model would need to be retrained from scratch on a large dataset. However, because the modulations of IC activity that encode binaural disparity operate on much slower timescales than the modulations that encode the monaural acoustic features that ICNet is already trained to capture [36], it may be possible to build binaural models with relatively small extensions to ICNet's trained encoder (especially if the focus is on capturing only a specific binaural property such as interaural masking or echo suppression).

Applications of ICNet

We designed ICNet so that it can be used either as a faithful simulator of neural coding in the IC or as a general model of auditory processing for sound engineering applications. It can be used to simulate either spike patterns from a large population of neural units or the latent neural dynamics that underlie these spike patterns. For hearing research, ICNet opens new possibilities with potential benefits that are both scientific (studies are no longer data limited) and ethical (animal experiments can be limited to confirmatory studies). Much of hearing research is focused on understanding the cortical computations that allow for stream segregation and recognition of objects in complex auditory scenes. However, understanding the cortical mechanisms that facilitate scene analysis is impossible without first specifying precisely the representation of the auditory scene that is made available to the cortex for further analysis. ICNet has the potential to be a valuable tool in this regard, allowing downstream computations to be referenced directly to early neural representations. Cochlear models are already employed widely and successfully in this way [31]; using ICNet as a second reference point midway along the central auditory pathway would enable early and late central processing to be decoupled.

ICNet also provides a comprehensive normal-hearing reference for technology development. It can be used as a front-end in the design of systems that achieve natural (rather than objectively optimal) performance, as its outputs reflect the physiological constraints that are faced by the brain. The use of ICNet for optimization of signal processing (e.g., noise suppression) could unlock additional perceptual benefit beyond that achieved through optimization based directly on acoustics. The same may also be true for the development of assistive listening technologies, with the output of ICNet serving as a target reference against which the efficacy of different hearing restoration strategies can be compared.

Methods

Experimental protocol

Experiments were performed on nine young-adult gerbils of both sexes that were born and raised in standard laboratory conditions. IC recordings were made at an age of 20-24 weeks. All experimental protocols were approved by the UK Home Office (PPL P56840C21).

Preparation for IC recordings

Recordings were made using the same procedures as in previous studies [14, 15]. Animals were placed in a sound-attenuated chamber and anesthetized for surgery with an initial injection of 1 ml per 100 g body weight of ketamine (100 mg per ml), xylazine (20 mg per ml), and saline in a ratio of 5:1:19. The same solution was infused continuously during recording at a rate of approximately 2.2 µl per min. Internal temperature was monitored and maintained at 38.7° C. A small metal rod was mounted on the skull and used to secure the head of the animal in a stereotaxic device. The pinnae were removed and speakers (Etymotic ER-10X) coupled to tubes were inserted into both ear canals. Two craniotomies were made along with incisions in the dura mater, and a 256-channel multi-electrode array was inserted into the central nucleus of the IC in each hemisphere.

Multi-unit activity

MUA was measured from recordings on each channel of the electrode array as follows: (1) a bandpass filter was applied with cutoff frequencies of 700 and 5000 Hz; (2) the standard deviation of the background noise in the bandpass-filtered signal was estimated as the median absolute deviation / 0.6745 (this estimate is more robust to outlier values, e.g., neural spikes, than direct calculation); (3) times at which the bandpass filtered signal made a positive crossing of a threshold of 3.5 standard deviations were identified and grouped into 1.3 ms bins.

Characteristic frequency analysis

To determine the preferred frequency of the recorded units, 50 ms tones were presented with frequencies ranging from 300 Hz to 16000 Hz in 0.2 octave steps and intensities ranging from 4 dB SPL to 85 dB SPL in 9 dB steps with 10 ms cosine on and off ramps and 75 ms pause between tones. Tones were presented 8 times each in random order with a sampling rate of 44.1 kHz. The characteristic frequency (CF) of each unit was defined as the frequency that elicited a significant response at the lowest intensity. Significance was defined as $p < 0.0001$ for the total spike count elicited by a tone given an assumed Poisson distribution with rate equal to the higher of either the mean spike count observed in time windows of the same duration in the absence of sound or 1/8.

Deep neural network models

Time-invariant single-branch architecture

The time-invariant single-branch DNN that was used to predict the neural activity of individual recordings is shown in Fig. 1. The model was trained to map a sound waveform s to neural activity $\hat{R}$ via an inferred conditional distribution of spike counts $p(\hat{R}|s)$ using the following architecture: (1) a SincNet layer [37] with 48 bandpass filters of size 64 and stride of 1, followed by a symmetric logarithmic activation $y = \mathrm{sgn}(x) * \log(|x| + 1)$; (2) a stack of five 1-D convolutional layers with 128 filters of size 64 and stride of 2, each followed by a PReLU activation; (3) a 1-D bottleneck convolutional layer with 64 filters of size 64 and a stride of 1, followed by a PReLU activation; and (4) a decoder convolutional layer without bias and 512 $\times$ N_c filters of size 1, where N_c is the number of parameters in the spike count distribution $p(\hat{R}|s)$ as described below. All convolutional layers in the encoder included a bias term and used a causal kernel. A cropping layer was added after the bottleneck to remove the left context from the output, which eliminated any convolutional edge effects.

For the Poisson model, $N_c = 1$ and the decoder output was followed by a softplus activation to produce the rate function λ that defined the conditional spike count distribution for each unit in each time bin. The

loss function was $\sum_{m=1}^{M}\sum_{t=1}^{T}(\lambda[m,t]-R[m,t]\log\lambda[m,t])$, where M and T are the number of units and time bins in the response, respectively. For the crossentropy model, $N_c = 5$ and the decoder output was followed by a softmax activation to produce the probability of each possible spike count $n \in \{0,...,N_c-1\}$ for each unit in each time bin. The loss function was $-\sum_{m=1}^{M}\sum_{t=1}^{T}\sum_{n=1}^{N_c} R[m,t]\log(p(\hat{R}[m,t]=n-1|s))$. The choice of 4 as the maximum possible spike count was made based on the spike count distributions across all datasets (the percentage of spikes above 4 was less than 0.02%). All neural activity was clipped during training and inference at a maximum value of 4.

Time-variant single-branch architecture

The time-variant single-branch architecture (Figs. 2-3) included an additional time input that matched the dimensionality and sampling rate of the sound input. Each sample of the time input defined the relative time at which the corresponding sound sample was presented during a given recording (range of approximately 0-12 hr). To match the dimensionality and dynamic range of the sound input after the SincNet layer, the time input was given to a convolutional layer with 48 filters of size 1 (without bias), followed by a symmetric logarithmic activation. The resulting output was then concatenated to the output of the SincNet layer to form the input of the first convolutional layer of the encoder.

Time-variant multi-branch architecture

The final ICNet model (Figs. 3-4) comprises a time-variant DNN with a shared encoder and separate decoder branches for each of the individual animals. The architecture of the encoder and decoders were identical to those in the single-branch models, except for the additional time input. Instead of a single time input to the encoder, a separate time input was used for each animal with the same dimensionality and sampling rate as the bottleneck output. Each sample of the time input thus defined the relative time at which the corresponding MUA sample was recorded, and was given to a convolutional layer with 64 filters of size 1 (without bias), followed by a PReLU activation. The output of each of these layers was then multiplied by the bottleneck output and was given to the corresponding decoder.

Training

Models were trained to transform 24414.0625 Hz sound input frames of 10240 samples into 762.9395 Hz neural activity frames of 256 samples (Fig. 1a). Context of 2048 samples was added on the left side of the sound input frame and was cropped after the bottleneck layer (2048 divided by a decimation factor of 2^5 resulted in 64 cropped samples at the level of the bottleneck). Sound inputs were scaled such that an RMS of 0.04 corresponded to a level of 94 dB SPL. For the time-variant models, time inputs were expressed in seconds and were scaled by 1/36000 to match the dynamic range to that of the sound inputs.

The DNN models were trained on NVidia RTX 4090 GPUs using Python and Tensorflow [38]. A batch size of 50 was used with the Adam optimizer and a starting learning rate of 0.0004. All trainings were performed so that the learning rate was halved if the loss in the validation set did not decrease for 2 consecutive epochs. Early stopping was used and determined the total number of training epochs if the validation loss did not decrease for 5 consecutive epochs. This resulted in an average of 38.1 $\pm$ 3.4 epochs for the nine single-branch models, and a total of 42 epochs for the nine-branch ICNet model. To speed up training, the weights from one of the single-branch models were used to initialize the encoder of the nine-branch model.

Training dataset

All sounds that were used for training the DNN models totalled to 7.83 hr and are described below. 10% of the sounds were randomly chosen and formed the validation set during training. The order of the sound presentation was randomized for each animal.

Speech Sentences were taken from the TIMIT corpus [39] that contains speech read by a wide range of American English talkers. The entire corpus excluding "SA" sentences was used and was presented either in quiet or with background noise. The intensity for each sentence was chosen at random from 45, 55, 65, 75, 85, or 95 dB SPL. The speech-to-noise ratio (SNR) was chosen at random from either 0 or 10 when the speech intensity was 55 or 65 dB SPL (as is typical of a quiet setting such as a home or an office) or –10 or 0 when the speech intensity was 75 or 85 dB SPL (as is typical of a noisy setting such as a pub). The total duration of speech in quiet was 1.25 hr and the total duration of speech in noise was 1.58 hr.

Noise Background noise sounds were taken from the Microsoft Scalable Noisy Speech Dataset [40], which includes recordings of environmental sounds from a large number of different settings (e.g., café, office, roadside) and specific noises (e.g., washer-dryer, copy machine, public address announcements). The intensity of the noise presented with each sentence was determined by the intensity of the speech and the SNR as described above.

Processed speech Speech from the TIMIT corpus was also processed in several ways: (1) the speed was increased by a factor of 2 (via simple resampling without pitch correction of any other additional processing); (2) linear multi-channel amplification was applied, with channels centered at 0.5, 1, 2, 4, and 8 kHz and gains of 3, 10, 17, 22, and 25 dB SPL, respectively; or (3) the speed was increased and linear amplification was applied. The total duration of processed speech was 2.66 hr.

Music Pop music was taken from the musdb18 dataset [41], which contains music in full mixed form as well as in stem form with isolated tracks for drums, bass, vocals and other (e.g., guitar, keyboard). The total duration of the music presented from this dataset was 1.28 hr. Classical music was taken from the musopen dataset (`https://musopen.org`; including piano, violin and orchestral pieces) and was presented either in its original form; after its speed was increased by a factor of 2 or 3; after it was high-pass filtered with a cutoff frequency of 6 kHz; or after its speed was increased and it was high-pass filtered. The total duration of the music presented from this dataset was 0.45 hr.

Moving ripples Dynamic moving ripple sounds were created by modulating a series of sustained sinusoids to achieve a desired distribution of instantaneous amplitude and frequency modulations. The lowest frequency sinusoid was either 300 Hz, 4.7 kHz, or 6.2 kHz. The highest frequency sinusoid was always 10.8 kHz. The series contained sinusoids at frequencies between the lowest and the highest in steps of 0.02 octaves, with the phase of each sinusoid chosen randomly from between 0 and 2π. The modulation envelope was designed so that the instantaneous frequency modulations ranged from 0 to 4 cycles/octave, the instantaneous amplitude modulations ranged from 0 to 10 Hz, and the modulation depth was 50 dB. The total duration of the ripples was 0.22 hr.

Evaluation

To predict neural activity using the trained DNN models, the decoder parameters were used to define either a Poisson distribution or a categorical probability distribution with N_c classes using the Tensorflow Probability

toolbox [42]. The distribution was then sampled from to yield simulated neural activity across time bins and units, as shown in Fig. 1a.

Evaluation dataset

All model evaluations (Figs. 1-4) used only sounds that were not part of the training dataset. Each sound segments was 30 s in duration. The timing of the presentation of the evaluation sounds varied across animals. For two animals, each sound was presented only twice and the presentation times were random. For three animals, each sound was presented four times, twice in successive trials during the first half of the recording and twice in successive trials during the second half of the recording. For the other four animals, each sound was presented ten times, twice in successive trials at times that were approximately 20, 35, 50, 65, and 80% of the total recording time, respectively.

Speech in quiet For all animals, a speech segment from the UCL SCRIBE dataset (`http://www.phon.ucl.ac.uk/resource/scribe`) consisting of sentences spoken by a male talker was presented at 60 dB SPL. For four of the animals (Fig. 4b), another speech segment from the same dataset was chosen and was presented at 85 dB SPL, consisting of sentences spoken by a female talker.

Speech in noise For all animals, a speech segment from the UCL SCRIBE dataset consisting of sentences spoken by a female talker was presented at 85 dB SPL in hallway noise from the Microsoft Scalable Noisy Speech dataset at 0 dB SNR. For four of the animals (Fig. 4b), another speech segment from the same dataset was chosen and was presented at 75 dB SPL, consisting of sentences spoken by a male talker mixed with café noise from the Microsoft Scalable Noisy Speech dataset at 4 dB SNR.

Moving ripples For all animals, dynamic moving ripples with the lowest frequency sinusoid of 4.7 kHz were presented at 85 dB SPL. For four of the animals (Fig. 4b), dynamic moving ripples with the lowest frequency sinusoid of 300 Hz were also used and were presented at 60 dB SPL.

Music For all animals, three seconds from each of 10 mixed pop songs from the musdb18 dataset were presented at 75 dB SPL. For four of the animals (Fig. 4b), a solo violin recording from the musopen dataset was presented at 85 dB SPL.

Pure tones For measuring Fano factors (Fig. 1), 75 ms tones at 4 frequencies (891.44, 2048, 3565.78 and 8192 Hz) were presented at 85 dB SPL with 10 ms cosine on and off ramps. Each tone was repeated 128 times, and the resulting responses were used to compute the average, the variance and the full distribution of spike counts (using time bins from 2.5 ms to 80 ms).

Non-stationarity

For Fig. 2g, we measured non-stationarity in each recording as the absolute difference of the covariance of neural activity between two successive trials and two trials recorded several hours apart. The covariance between two successive trials was the average computed over the covariance values of two pairs of neural activity after flattening (collapsing both neural responses across time and units into one dimension and computing their covariance). The covariance between two trials recorded several hours apart was computed in the same manner as the average over the covariance values of two pairs of neural responses.

Overall performance metrics

RMSE and log-likelihood were used as overall performance metrics (Figs. 1 to 3). The reported results for both metrics were the averages computed over two trials of each sound within each individual recording. One trial was chosen from the first half of the recording and the other from the second, resulting in times of 3.51 $\pm$ 1.83 hr and 8.92 $\pm$ 1.16 hr across all nine animals. For each trial, the RMSE and log-likelihood were computed as follows:

$$\text{RMSE} = \sqrt{\frac{1}{MT}\sum_{m=1}^{M}\sum_{t=1}^{T}(R[m,t]-\hat{R}[m,t])^2} \tag{1}$$

$$\text{Log-likelihood} = \frac{1}{MT}\sum_{m=1}^{M}\sum_{t=1}^{T}\log(p(\hat{R}[m,t]=R[m,t]|s,\tau)), \tag{2}$$

where M is the number of units, T is the number of time bins, s is the sound stimulus, τ is the time input, R is the recorded neural activity, and $\hat{R}$ is the predicted neural activity obtained by sampling from the inferred distribution $p(\hat{R}|s,\tau)$ (or $p(\hat{R}|s)$ for a time-invariant model).

Predictive power metrics

To assess the predictive power of ICNet in Fig. 4, we selected two metrics that were computed using successive trials for the seven animals for which this was possible (chosen from the beginning of the recording, resulting in times of 1.94 $\pm$ 0.46 hr across animals). The two metrics are formulated as the fraction of explainable variance or correlation that the model explained across all units, and were computed after flattening the neural activity responses (collapsing across time and units into one dimension). The fraction of explainable correlation explained across all units was computed as follows:

$$\text{Correlation explained (\%)} = \frac{100}{2}\cdot\frac{\rho(R_1,\hat{R}_1)+\rho(R_2,\hat{R}_2)}{\rho(R_1,R_2)}, \tag{3}$$

where $R_1, R_2 \in \mathbb{N}^{M\times T}$ are the recorded neural activity for each trial, $\hat{R}_1, \hat{R}_2 \in \mathbb{N}^{M\times T}$ are the predicted neural activity, and ρ is the correlation coefficient. The fraction of explainable variance explained was computed as follows: [10]:

$$\text{Variance explained (\%)} = 100\cdot\left(1-\frac{\frac{1}{2MT}\sum_{m=1}^{M}\sum_{t=1}^{T}(R_1-E_c[\hat{R}_1])^2+\frac{1}{2MT}\sum_{m=1}^{M}\sum_{t=1}^{T}(R_2-E_c[\hat{R}_2])^2-\sigma^2_{noise}}{\frac{1}{2}(Var[R_1]+Var[R_2])-\sigma^2_{noise}}\right), \tag{4}$$

$$\sigma^2_{noise} = \frac{1}{2}Var[R_1-R_2], \tag{5}$$

where E_c denotes expectation over spike counts of $p(\hat{R}|s,\tau)$ (or $p(\hat{R}|s)$ for a time-invariant model). To compute the predictive power of ICNet across each unit (Fig. 4d), we computed the formulas of Eqs. *3*-*5* across time only (without collapsing across the unit dimension or averaging across units).

Phoneme recognition

A DNN architecture based on Conv-TasNet [43] was used to train an automatic-speech-recognition (ASR) back-end that predicted phonemes from the ICNet bottleneck response across time. The DNN architecture comprised: (1) a convolutional layer with 64 filters of size 3 and no bias, followed by a PReLU activation, (2) a normalization layer, followed by a convolutional layer with 128 filters of size 1, (3) a block of 8 dilated convolutional layers (dilation from 1 to 2^7) with 128 filters of size 3, including PReLU activations and residual

skip connections in between, (4) a convolutional layer with 256 filters of size 1, followed by a sigmoid activation, (5) an output convolutional layer with 40 filters of size 3, followed by a softmax activation. All convolutional layers were 1-D and used a causal kernel with a stride of 1.

The ASR model was trained using the train subset of the TIMIT corpus [39]. The TIMIT sentences were resampled to 24414.0625 Hz and were segmented into sound input frames of 81920 samples (windows of 65536 samples with 50% overlap and left context of 16384 samples). In each training step, the sound input frames were given to the frozen ICNet model to simulate the bottleneck response at 762.9395 Hz (output frames of 2048 samples with left context of 256 samples). The simulated bottleneck response was then given to the ASR back-end to predict the probabilities of the 40 phoneme classes across time (output frames of 2048 samples at 762.9395 Hz after cropping the context).

The TIMIT sentences were calibrated to randomly selected levels between 40 and 90 dB SPL with a step of 5 dB, and were randomly mixed with the 18 noise types from the DEMAND dataset [44] at SNRs of -30, -20, -10, 0, 10, 20, 30, and 100 dB. The phonetic transcriptions of the TIMIT dataset were downsampled to the sampling frequency of the neural activity (762.9395 Hz) and were grouped into 40 classes (15 vowels and 24 consonants plus the glottal stop). To account for the phonetic class imbalance (high prevalence of silence in the dataset), a focal crossentropy loss function was used for training with the focusing parameter γ set to 5.

ASR models were trained to predict speech recognition from the ICNet bottleneck response, as well as from the bottleneck responses of the nine single-branch time-variant models. The trained ASR models were evaluated using 6 random sentences of the test subset of the TIMIT corpus (Fig. 4). The sentences were calibrated at 65 dB SPL, and speech-shaped noise was generated to match the long-term average spectrum of the sentences. A baseline system was also trained with the same DNN architecture, but using a standard Mel spectrogram as the input. The Mel spectrogram was computed from the sound input (24414.0625 Hz) using frames of 512 samples, a hop size of 32 samples and 64 Mel channels with frequencies logarithmically spaced from 50 to 12000 Hz. These parameters were chosen to maximize performance while matching the resolution of the ICNet bottleneck responses (762.9395 Hz output frames with 64 channels).

Statistics

Confidence intervals for all reported values were computed using 1000 bootstrap samples. For the metrics that we used to assess model performance (RMSE, log-likelihood, variance and correlation explained), bootstrapping was performed across units for each animal. For the ASR results of Fig. 3f, bootstrapping was performed across animals.

Code availability

The code of the ICNet model is available via `https://doi.org/10.5281/zenodo.11943090` or `https://github.com/fotisdr/ICNet`. A Jupyter notebook is included with a simple usage example for ICNet. A non-commercial, academic UCL license applies.

Acknowledgments

This work was supported by EPSRC EP/W004275/1 and MRC MR/W019787/1.

Author contributions

F.D.: Conceptualization, Data Curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing: Original Draft Preparation; **S.S.**: Investigation; **Y.X.**: Investigation; **A.F.**: Conceptualization, Project administration, Software, Supervision; **N.A.L.**: Conceptualization, Data Curation, Funding acquisition, Investigation, Methodology, Project Administration, Resources, Supervision, Writing - Review & Editing

Competing interests

N. A. L. is a co-founder of Perceptual Technologies.

References

1. Osses Vecchi, A. *et al.* A comparative study of eight human auditory models of monaural processing. *Acta Acustica* **6**, 17 (2022).

2. Meyer, A. F., Williamson, R. S., Linden, J. F. & Sahani, M. Models of neuronal stimulus-response functions: elaboration, estimation, and evaluation. *Frontiers in Systems Neuroscience* **10**, 109 (2017).

3. Rahman, M., Willmore, B. D., King, A. J. & Harper, N. S. Simple transformations capture auditory input to cortex. *Proceedings of the National Academy of Sciences* **117**, 28442–28451 (2020).

4. Baby, D., Van Den Broucke, A. & Verhulst, S. A convolutional neural-network model of human cochlear mechanics and filter tuning for real-time applications. *Nature Machine Intelligence* **3**, 134–143 (2021).

5. Drakopoulos, F., Baby, D. & Verhulst, S. A convolutional neural-network framework for modelling auditory sensory cells and synapses. *Communications Biology* **4**, 827 (2021).

6. Nagathil, A. & Bruce, I. C. WaveNet-based approximation of a cochlear filtering and hair cell transduction model. *The Journal of the Acoustical Society of America* **154**, 191–202 (2023).

7. Van Den Broucke, A., Drakopoulos, F., Baby, D. & Verhulst, S. Otoacoustic emissions in a deep-neural-network model of cochlear mechanics. In *The 14th International Mechanics of Hearing Workshop (MoH 2022)*, vol. 3062 (2024).

8. McIntosh, L., Maheswaranathan, N., Nayebi, A., Ganguli, S. & Baccus, S. Deep learning models of the retinal response to natural scenes. *Advances in Neural Information Processing Systems* **29** (2016).

9. Walker, E. Y. *et al.* Inception loops discover what excites neurons most using deep predictive models. *Nature Neuroscience* **22**, 2060–2065 (2019).

10. Cadena, S. A. *et al.* Deep convolutional models improve predictions of macaque V1 responses to natural images. *PLoS Computational Biology* **15**, e1006897 (2019).

11. Bashivan, P., Kar, K. & DiCarlo, J. J. Neural population control via deep image synthesis. *Science* **364**, eaav9436 (2019).

12. Kell, A. J., Yamins, D. L., Shook, E. N., Norman-Haignere, S. V. & McDermott, J. H. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron* **98**, 630–644 (2018).

13. Keshishian, M. *et al.* Estimating and interpreting nonlinear receptive field of sensory neural responses with deep neural network models. *eLife* **9**, e53445 (2020).

14. Sabesan, S., Fragner, A., Bench, C., Drakopoulos, F. & Lesica, N. A. Large-scale electrophysiology and deep learning reveal distorted neural signal dynamics after hearing loss. *eLife* **12**, e85108 (2023).

15. Armstrong, A. G., Lam, C. C., Sabesan, S. & Lesica, N. A. Compression and amplification algorithms in hearing aids impair the selectivity of neural responses to speech. *Nature Biomedical Engineering* **6**, 717–730 (2022).

16. Garcia-Lazaro, J. A., Belliveau, L. A. & Lesica, N. A. Independent population coding of speech with sub-millisecond precision. *Journal of Neuroscience* **33**, 19362–19372 (2013).

17. Kara, P., Reinagel, P. & Reid, R. C. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* **27**, 635–646 (2000).

18. Pecka, M., Siveke, I., Grothe, B. & Lesica, N. A. Enhancement of ITD coding within the initial stages of the auditory pathway. *Journal of Neurophysiology* **103**, 38–46 (2010).

19. Williamson, R. S., Sahani, M. & Pillow, J. W. The equivalence of information-theoretic and likelihood-based methods for neural dimensionality reduction. *PLoS Computational Biology* **11**, e1004141 (2015).

20. Drakopoulos, F. & Verhulst, S. A neural-network framework for the design of individualised hearing-loss compensation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* (2023).

21. Berry, M. & Meister, M. Refractoriness and neural precision. *Advances in Neural Information Processing Systems* **10** (1997).

22. Keat, J., Reinagel, P., Reid, R. C. & Meister, M. Predicting every spike: a model for the responses of visual neurons. *Neuron* **30**, 803–817 (2001).

23. Simoncelli, E. P., Paninski, L., Pillow, J., Schwartz, O. *et al.* Characterization of neural responses with stochastic stimuli. *The Cognitive Neurosciences* **3**, 1 (2004).

24. Escabí, M. A., Nassiri, R., Miller, L. M., Schreiner, C. E. & Read, H. L. The contribution of spike threshold to acoustic feature selectivity, spike information content, and information throughput. *Journal of Neuroscience* **25**, 9524–9534 (2005).

25. Sumner, C. J. *et al.* Mammalian behavior and physiology converge to confirm sharper cochlear tuning in humans. *Proceedings of the National Academy of Sciences* **115**, 11322–11326 (2018).

26. Verschooten, E., Desloovere, C. & Joris, P. X. High-resolution frequency tuning but not temporal coding in the human cochlea. *PLoS Biology* **16**, e2005164 (2018).

27. Huet, A. *et al.* Peristimulus time responses predict adaptation and spontaneous firing of auditory-nerve fibers: from rodents data to humans. *Journal of Neuroscience* **42**, 2253–2267 (2022).

28. Wirtzfeld, M. R., Ibrahim, R. A. & Bruce, I. C. Predictions of speech chimaera intelligibility using auditory nerve mean-rate and spike-timing neural cues. *Journal of the Association for Research in Otolaryngology* **18**, 687–710 (2017).

29. Zaar, J. & Carney, L. H. Predicting speech intelligibility in hearing-impaired listeners using a physiologically inspired auditory model. *Hearing Research* **426**, 108553 (2022).

30. Swaminathan, J. & Heinz, M. G. Psychophysiological analyses demonstrate the importance of neural envelope coding for speech perception in noise. *Journal of Neuroscience* **32**, 1747–1756 (2012).

31. Shan, T., Cappelloni, M. S. & Maddox, R. K. Subcortical responses to music and speech are alike while cortical responses diverge. *Scientific Reports* **14**, 789 (2024).

32. Leer, P. *et al.* Neural networks hear you loud and clear: Hearing loss compensation using deep neural networks. *arXiv preprint arXiv:2403.10420* (2024).

33. Trautmann, E. M. *et al.* Accurate estimation of neural population dynamics without spike sorting. *Neuron* **103**, 292–308 (2019).

34. Ter-Mikaelian, M., Sanes, D. H. & Semple, M. N. Transformation of temporal properties between auditory midbrain and cortex in the awake mongolian gerbil. *Journal of Neuroscience* **27**, 6091–6102 (2007).

35. Shaheen, L. A. & Liberman, M. C. Cochlear synaptopathy changes sound-evoked activity without changing spontaneous discharge in the mouse inferior colliculus. *Frontiers in Systems Neuroscience* **12**, 59 (2018).

36. Belliveau, L. A., Lyamzin, D. R. & Lesica, N. A. The neural representation of interaural time differences in gerbils is transformed from midbrain to cortex. *Journal of Neuroscience* **34**, 16796–16808 (2014).

37. Ravanelli, M. & Bengio, Y. Speaker recognition from raw waveform with SincNet. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, 1021–1028 (2018).

38. Abadi, M. *et al.* Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 265–283 (2016).

39. Garofolo, J. S., Lamel, L. F., Fisher, W. M., Fiscus, J. G. & Pallett, D. S. DARPA TIMIT acoustic-phonetic continous speech corpus CD-ROM. NIST speech disc 1-1.1. *NASA STI/Recon technical report n* **93**, 27403 (1993).

40. Reddy, C. K. *et al.* A scalable noisy speech dataset and online subjective test framework. *arXiv preprint arXiv:1909.08050* (2019).

41. Rafii, Z., Liutkus, A., Stöter, F.-R., Mimilakis, S. I. & Bittner, R. The MUSDB18 corpus for music separation (2017).

42. Dillon, J. V. *et al.* Tensorflow distributions. *arXiv preprint arXiv:1711.10604* (2017).

43. Luo, Y. & Mesgarani, N. Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **27**, 1256–1266 (2019).

44. Thiemann, J., Ito, N. & Vincent, E. The diverse environments multi-channel acoustic noise database (DEMAND): A database of multichannel environmental noise recordings. In *Proceedings of Meetings on Acoustics*, vol. 19 (2013).