

Published in final edited form as:

*Nat Mach Intell.* 2021 February ; 3(2): 134–143. doi:10.1038/s42256-020-00286-8.

## A convolutional neural-network model of human cochlear mechanics and filter tuning for real-time applications

**Deepak Baby, Arthur Van Den Broucke, Sarah Verhulst**

Hearing Technology @ WAVES, Dept. of Information Technology, Ghent University, 9000 Ghent, Belgium

### Abstract

Auditory models are commonly used as feature extractors for automatic speech-recognition systems or as front-ends for robotics, machine-hearing and hearing-aid applications. Although auditory models can capture the biophysical and nonlinear properties of human hearing in great detail, these biophysical models are computationally expensive and cannot be used in real-time applications. We present a hybrid approach where convolutional neural networks are combined with computational neuroscience to yield a real-time end-to-end model for human cochlear mechanics, including level-dependent filter tuning (CoNNear). The CoNNear model was trained on acoustic speech material and its performance and applicability were evaluated using (unseen) sound stimuli commonly employed in cochlear mechanics research. The CoNNear model accurately simulates human cochlear frequency selectivity and its dependence on sound intensity, an essential quality for robust speech intelligibility at negative speech-to-background-noise ratios. The CoNNear architecture is based on parallel and differentiable computations and has the power to achieve real-time human performance. These unique CoNNear features will enable the next generation of human-like machine-hearing applications.

### Introduction

The human cochlea is an active, nonlinear system that transforms sound-induced vibrations of the middle-ear bones to cochlear travelling waves of basilar-membrane (BM) motion [1]. Cochlear mechanics and travelling waves are responsible for hallmark features of mammalian hearing, including the level-dependent frequency selectivity [2–5] that results from a cascade of cochlear mechanical filters with centre frequencies (CFs) between 20 kHz and 40 Hz from the human cochlear base to apex [6].

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:[http://www.nature.com/authors/editorial\\_policies/license.html#terms](http://www.nature.com/authors/editorial_policies/license.html#terms)

Correspondence to: Deepak Baby; Sarah Verhulst.

\* s.verhulst@ugent.be; deepakbabycet@gmail.com.

#### Additional Information

DB: conceptualisation, methodology, software, validation, formal analysis, investigation, data curation, writing: original draft, visualisation; AVDB: software, validation, investigation, writing: visualisation; SV: conceptualisation, methodology, resources, writing: original draft, review & editing, supervision, project administration, funding acquisition.

#### Competing interests

A patent application (PCTEP2020065893) was filed by UGent on the basis of the research presented in this manuscript. Inventors on the application are Sarah Verhulst, Deepak Baby, Fotios Drakopoulos and Arthur Van Den Broucke.

Modelling cochlear mechanics has been an active field of research because computational methods can help characterise the mechanisms underlying normal or impaired hearing and thereby improve hearing diagnostics [7,8] and treatment [9,10], or inspire machine-hearing applications [11,12]. One popular approach involves representing the cochlea as a transmission line (TL) by discretising the space along the BM and describing each section using a system of ordinary differential equations that approximate the local biophysical filter characteristics (Fig. 1a; state-of-the-art model) [13–19]. Analytical TL models represent the cochlea as a cascaded system in which the response of each section depends on the responses of all previous sections. This architecture makes them computationally expensive, as the filter operations in the different sections cannot be computed in parallel. The computational complexity is even greater when nonlinearities or feedback pathways are included to faithfully approximate cochlear mechanics [15,20].

Computational complexity is the main reason that real-time applications for hearing-aid [21], robotics [22], and automatic speech recognition applications do not adopt cochlear travelling-wave models in their preprocessing. Instead, they use computationally efficient approximations of auditory filtering that compromise on key auditory features. A common simplification implements the cochlear filters as a parallel, rather than cascaded, filterbank [23,24]. However, this parallel architecture captures neither the longitudinal coupling properties of the BM [25] nor the generation of otoacoustic emissions [26]. Another popular model, the gammatone filterbank model [27], ignores the stimulus-level dependence of cochlear filtering. Lastly, a number of models simulate the level-dependence of cochlear filtering but fail to match the performance of TL models [28]: they either simulate the longitudinal cochlear coupling locally within the individual filters of the uncoupled filterbank [29] or introduce distortion artefacts when combining an automatic-gain-control type of level-dependence with cascaded digital filters [30–32].

Thus, the computational complexity of biophysically realistic cochlear models poses a serious impediment to the development of human-like machine-hearing applications. This complexity motivated our search for an efficient model that matches the performance of state-of-the-art analytical TL models while offering real-time execution. Here, we investigate whether convolutional neural networks (CNNs) can be used for this purpose. Neural networks of this type can deliver end-to-end waveform predictions [33,34] with real-time properties [35] and are based on convolutions akin to the filtering process associated with cochlear processing.

This paper details how CNNs can best be connected and trained to approximate the computations performed by TL cochlear models [19, 36, 37], with a specific emphasis on simultaneously capturing the tuning, level-dependence, and longitudinal coupling characteristics of human cochlear processing. The proposed model (*CoNNear*) converts speech stimuli into corresponding BM displacements across 201 cochlear filters distributed along the length of the BM. Unlike TL models, the CoNNear architecture is based on parallel CPU computations that can be accelerated through GPU computing. Consequently, CoNNear can easily be integrated with real-time auditory applications that use deep learning. The quality of the CoNNear predictions and the generalisability of the method are evaluated on the basis of cochlear mechanical properties such as filter tuning estimates [38],

nonlinear distortion characteristics [39], and spatial excitation patterns [40] obtained using sound stimuli commonly adopted in experimental studies of cochlear mechanics but not included in the training.

## The CoNNear model

The CoNNear model has an auto-encoder CNN architecture and transforms a 20-kHz sampled acoustic waveform (in [Pa]) to  $N_{CF}$  cochlear BM displacement waveforms (in [ $\mu\text{m}$ ]) using several CNN layers and dimension changes (Fig. 1a). The first four layers are *encoder* layers and use strided convolutions to halve the temporal dimension after every CNN layer. The following four, *decoder*, layers map the condensed representation onto  $L \times N_{CF}$  outputs using deconvolution operations.  $L$  corresponds to the initial size of the audio input and  $N_{CF}$  to 201 cochlear filters with centre frequencies (CFs) between 0.1 and 12 kHz. The adopted CFs were spaced according to the Greenwood place-frequency map of the cochlea [6] and span the most sensitive frequency range of human hearing [41]. It is important to preserve the temporal alignment (or, phase) of the inputs across the architecture, because this information is essential for speech perception [42]. We used  $U$ -shaped skip connections for this purpose. Skip connections have earlier been adopted in image-to-image translation [43] and speech-enhancement applications [33,34]; they pass temporal information directly from encoder to decoder layers (Fig. 1a; dashed arrows). Aside from preserving phase information, skip connections may also improve the model's ability to learn how best to combine the nonlinearities of several CNN layers to simulate the level-dependent properties of human cochlear processing.

Every CNN layer is comprised of a set of filterbanks followed by a nonlinear operation [44] and the CNN filter weights were trained using TL-simulated BM displacements from  $N_{CF}$  cochlear channels [37]. While training was conducted using a speech corpus [45] presented at 70 dB SPL, model evaluation was based on the ability to reproduce key cochlear mechanical properties using basic acoustic stimuli (e.g. clicks, pure-tones) unseen during training (Fig. 1c). During training and evaluation, the audio input was segmented into 2048-sample windows ( $\approx 100$  ms), after which the corresponding BM displacements were simulated and concatenated over time. Because CoNNear treats each input independently, and resets its adaptation properties at the start of each simulation, this concatenation procedure could result in discontinuities near the window boundaries. To address this issue, we also evaluated an architecture that had the previous and following (256) input samples available as context (Fig. 1b). Different from the no-context architecture (Fig. 1a), a final cropping layer was added to remove the simulated context and yield the final  $L$ -sized BM displacement waveforms. Additional details on the CoNNear architecture and training procedure are given in Methods and Extended Data Fig. 1. Lastly, training CoNNear using audio inputs of fixed duration does not prevent it from handling inputs of other durations after training, thanks to its convolutional architecture. This flexibility is a clear benefit over matrix-multiplication-based neural network architectures, which can operate only on inputs of fixed-duration.

## CoNNear hyperparameter tuning

Critical to this work is proper determination of the optimal CNN architecture and its hyperparameters that yield a realistic model for cochlear processing. Ideally, CoNNear should both simulate the speech training dataset with a sufficiently low L1 loss [i.e., the average mean-absolute difference between simulated BM displacements of the reference TL and CoNNear model] and also reproduce key cochlear response properties. Details on the hyperparameters that can be adjusted to define the final CoNNear architecture (number of layers, nonlinearity, context) are given in Extended Data Fig. 1. To determine the hyperparameters, we followed an iterative principled approach taking into account: (i) the L1 loss on speech material, (ii) the desired frequency- and level-dependent cochlear filter tuning characteristics, and (iii), the computational load needed for real-time execution. Afterwards, the prediction accuracy of CoNNear was evaluated on a broader set of response features. We considered a total of four metrics derived from cochlear responses evoked using simple stimuli of different frequencies and levels: filter tuning ( $Q_{ERB}$ ), click-evoked dispersion, pure-tone excitation patterns, and the generation of cochlear distortion products. Together, these evaluation metrics form a comprehensive description of cochlear processing. Further information on the tests and their implementation is given in Methods. A benefit of our approach is that the evaluation stimuli were *unseen* to the model to ensure an independent evaluation procedure. Even though any fragment of the speech training material can be seen as a combination of basic acoustic elements such as impulses and pure tones of varying levels and frequencies, the cochlear mechanics stimuli were not explicitly present in the training material.

The first hyperparameter specifies the total number of encoder/decoder layers, which we set to 8 on the basis of  $Q_{ERB}$  simulations. Aside from the reference experimental human  $Q_{ERB}$  curve for low stimulus levels [46], Fig. 2a shows simulated  $Q_{ERB}$  curves from the reference TL-model (red) overlaid with CoNNear model simulations. CoNNear captures the frequency-dependence of the human and reference TL-model  $Q_{ERB}$  function better as the layer depth was increased from 4 to 8. Models with 4 and 6 layers underestimated the overall  $Q_{ERB}$  function and performed worse for CFs below 1 kHz where reference ERBs were narrower, and corresponding BM impulse responses longer, than for higher CFs. Extending the number of layers beyond 8 increased the required computational resources without producing a substantial improvement in the quality of the fits. These resources are listed in Extended Data Fig. 2 for models of different layers, along with L1-loss predictions on the speech training set and a small set of evaluation metrics.

The second hyperparameter controls the activation function, or nonlinearity, which is placed between the CNN layers. To mimic the original shape of the outer-hair-cell input/output function [47] responsible for cochlear compression, we required that the activation function cross the x-axis. We therefore considered only the parametric rectified linear unit (PReLU) and hyperbolic-tangent (tanh) nonlinearities rather than standard activation functions such as the sigmoid and ReLU. Figure 2b depicts how the PReLU and tanh activation functions affected the simulated  $Q_{ERB}$ s across CFs. Whereas the PReLU activation function was unable to capture the level-dependence of cochlear filter tuning, the tanh nonlinearity reproduced both the level and frequency-dependence of human  $Q_{ERB}$ s. Additionally,

Extended Data Fig. 3 (no context simulations) shows that the model with tanh nonlinearity reached a lower L1 loss on both training and test sets compared to the PReLU nonlinearity, while requiring a smaller number of parameters. The benefit of the tanh nonlinearity is further illustrated when comparing simulated pure-tone excitation patterns across model architectures. Figure 3 shows that the tanh nonlinearity (c) outperforms the PReLU nonlinearity (b) in capturing the compressive growth of BM displacement with stimulus level, as observed in the excitation pattern maxima (a). Although both activation functions were able to code negative input deflections, only the model with tanh nonlinearity (whose shape best resembles the cochlear input/output function) performed well. These simulations show that it is essential to consider the shape of the activation function when the reference system is nonlinear.

The final hyperparameter relates to adding context information when simulating an input of size  $L$ . As noted earlier, CNN architectures treat each input independently, which can produce discontinuities near the window boundaries. The effect of adding context [cf. architecture (b) vs (a) in Fig. 1] is best observed when simulating speech. Figure 4 shows simulated BM displacements of the reference and CoNNear models to a segment of the speech test set which was not seen during training. Panel (d) shows that providing context prevents discontinuities near the window boundaries. Context was also beneficial when simulating the reference 70-dB  $Q_{\text{ERB}}$  function in Fig. 2c. The L1-loss simulations in Extended Data Fig. 3 (tanh simulations) confirmed the overall performance improvement when providing context, and hence the final CoNNear architecture included a context window of 256 samples on either side of the 2048 sample input (Fig. 1b).

## Generalisability of CoNNear

Overfitting occurs when the trained model merely memorises the training material and fails to generalise to data not present in the training set. To investigate whether the final CoNNear architecture was robust against overfitting, we tested how well the trained CoNNear model performed on unseen stimuli from (i) the same database, (ii) a different database of the same language, (iii) a different database of a different language, and (iv) a music piece.

Figure 5 shows boxplots of the L1-loss distributions of all simulated windows in the following audio material (from left to right): the TIMIT training set, the TIMIT test set using different speakers, the Dutch Matrix sentence database [49], Radioheads' OK Computer album, and 400 randomly-drawn samples from different event categories in the AudioSet database [50]. In addition to calculated loss distributions, Fig. 5 compares instantaneous BM displacement intensities between the reference TL-model (g) and CoNNear (h) for three stimuli unseen during training: an English sentence from the Wall Street Journal Corpus [48] to show how CoNNear adapts to different recording settings (first column). A sentence from the Dutch matrix test [49] to investigate performance on a different language stimulus (middle column), and a music segment to test performance on a non-speech acoustic stimulus (right column). From an application perspective, we also tested how CoNNear handles audio input of arbitrary length. Extended Data Fig. 4 shows that CoNNear generalises well to an unseen speech stimulus of length 0.5 s (10048 samples) and a music

stimulus of length 0.8 s (16384 samples), even though training was performed using shorter 2048-sample windows of fixed duration.

The similar loss distributions across tested speech conditions, together with the small intensity prediction errors seen in Fig. 5(i), show that CoNNear generalises well to unfamiliar stimuli. When comparing L1-loss distributions between speech (b,c) and non-speech (d,e) audio, we observe increased mean L1-losses along with a greater performance variability, which reflects the large variety of spectral and temporal features present in music or the AudioSet database. Even though median performance is still acceptable for non-speech audio, we conclude that CoNNear performs most comparable to the original TL-model simulations for speech. CoNNear's prediction accuracy on non-speech audio may be improved by retraining with a larger and more-diverse stimulus set than the adopted TIMIT training set. However, a clear benefit from our hybrid, neuroscience-inspired training method is that we can yield a generalisable CoNNear architecture even when using a training dataset which is much smaller than those adopted in standard data-driven NN-approaches [44]. Lastly, we note that the non-speech audio were downsampled to 16 kHz to match the frequency content of the TIMIT speech material to allow for a fair comparison between speech and non-speech audio. We can hence not guarantee that CoNNear performs well on non-speech audio with frequency content above 8 kHz.

## CoNNear as a model for human cochlear signal processing

Since the primary goal of this work was to develop a neural-network based model for human cochlear mechanics, we evaluated how well the trained CoNNear model collectively performs on simulating key cochlear mechanics metrics (see Methods for a detailed description). Figure 2c shows that the final CoNNear architecture with context simulates the frequency- and level-dependence of human cochlear tuning. At the same time, CoNNear faithfully captures the shape and compression properties of pure-tone cochlear excitation patterns (Fig. 3d). However, we did observe small excitation pattern fluctuations at CFs below the stimulus frequency for the 1 and 2 kHz simulations (middle and bottom row in Fig. 3). These fluctuations had levels of  $\approx 30$  dB below the peak of the excitation pattern and are hence not expected to degrade the sound-driven CoNNear response to complex stimuli such as speech. This latter statement is supported by (i) the CoNNear speech simulations in Fig. 5i, which show minimally visible noise in the error patterns, and (ii) the root mean-square error (RMSE) observed between simulated TL and CoNNear model excitation patterns for a broad range of frequencies and levels (Extended Data Fig. 6c). To obtain a meaningful error estimate, we normalised the RMSE to the TL excitation pattern maximum to yield a CoNNear error percentage. The error for the final CoNNear architecture remained below 5 % for pure-tone frequencies up to 4 kHz, and stimulation levels below 80 dB. Increased errors were observed for higher stimulus levels and at CFs of 8 and 10 kHz and stemmed from the overall level and frequency content of the speech material used for training; the dominant energy in speech occurs below 5 kHz, and we presented the speech corpus at 70 dB. However, it is noteworthy that the visual difference between simulated high-level excitation patterns was mostly associated with low-level fluctuations at CFs below the stimulation frequency and not at the stimulation frequency itself (cf. the first and last columns of Fig. 3). This means that the stimulus-driven error that occurs when using



broadband stimuli is likely smaller than that observed here for tonal stimuli. In summary, the error appears acceptable for processing speech-like audio. We speculate that CoNNear's application range can be extended by retraining with stimuli at different levels and/or with greater high-frequency content.

The top panel of Fig. 6 shows the dispersion characteristics of the TL model and trained CoNNear models, illustrating the characteristic 12-ms onset delay from basal (high CF, low channel numbers) to apical (low CF, high channel numbers) cochlear sections in both models. In the cochlea, dispersion arises from the biophysical properties of the coupled BM (i.e., its spatial variation of stiffness and damping), and the CoNNear architecture captures this phenomenon. Adding context did not affect the simulation quality.

Finally, we tested CoNNear's ability to simulate a last, key feature of cochlear mechanics: distortion products, which travel in reverse along the BM to generate pressure waveforms in the ear canal (i.e., distortion-product otoacoustic emissions, DPOAEs). The bottom panel of Fig. 6 compares reference TL-model simulations with those of different CoNNear models (b)-(d). DPOAE frequencies are visible as spectral components that are not present in the stimulus (which consists of two pure tones at frequencies  $f_1 = 2.0$  kHz and  $f_2 = 2.4$  kHz). In humans, the strongest DP component occurs at  $2f_1 - f_2 = 1.6$  kHz, and as the simulations show, the level of this DP was best captured using the tanh activation function. As observed before in the excitation pattern simulations, the activation function most resembling the shape of the cochlear nonlinearity performed best when simulating responses that rely on cochlear compression. Adding context removed the high-frequency distortions that were visible in panels (b) and (c).

The quality of the DPOAE simulations across a range of stimulation levels and frequencies was further investigated in Fig. 6e. Up to  $L_2$  levels of 60 dB SPL, CoNNear matched the characteristic nonlinear growth of both TL-simulated and human DPOAE level functions well [52]. At higher stimulation levels, DPOAE levels started fluctuating in both models. Although this may reflect the chaotic character of cochlear DPs interacting with cochlear irregularities in TL models and in human hearing [26], it may simply reflect the limited training material we had available for CoNNear at these higher levels. Given that human DPOAEs are generally recorded for stimulation levels below 60 dB SPL in clinical or research contexts [53], we can conclude that CoNNear was able to capture this important epiphenomenon of hearing. The cochlear mechanics evaluations we performed in Figs. 2, 3, 6 and Extended Data Fig. 6 together demonstrate that the 8-layer, tanh, CoNNear model with context performed best on four crucial aspects of human cochlear mechanics. Despite training on a limited speech corpus presented at 70 dB SPL, CoNNear learned to simulate outputs which matched those of biophysically-realistic analytical models of human cochlear processing across level and frequency.

## CoNNear as a real-time model for audio applications

In addition to its ability to simulate realistic cochlear mechanical responses, CoNNear operates in real-time. In audio applications, real-time is commonly defined as a computation duration less than 10 ms; below this limit no delay is perceived. Table 1 summarises the

necessary time to compute the final CoNNear-context model for a stimulus window of 1048 samples on CPU or GPU architectures. On a CPU, the CoNNear model outperforms the TL-model by a factor of 129 and on a GPU, CoNNear is 2142 times faster. Additionally, the GPU computations show that the trained CoNNear model (8-layers, tanh, with context) has a latency of 7.27 ms, and hence reaches real-time audio processing performance.

## Discussion

This paper details how a hybrid, deep-neural-net and analytical approach can be used to develop a real-time model of human cochlear processing (CoNNear), with performance matching that of human cochlear processing. As we demonstrate here for the first time, neither real-time performance nor biophysically realistic responses need be compromised. CoNNear combines both in a single auditory model, laying the groundwork for a new generation of human-like robotic, speech-recognition and machine hearing applications. Prior work has demonstrated the clear benefits of using biophysically realistic cochlear models as front-ends for auditory applications: e.g. for capturing cochlear compression, [7,11,28], speech enhancement at negative signal-to-noise ratio's [12], realistic sound perception predictions [54] and for simulating the generation of human auditory brainstem responses [37, 55]. Hence, CoNNear can dramatically improve performance in application areas which, to date, have relied on computationally intensive biophysical models. Not only can CoNNear operate on running audio input with a latency below 7.5 ms, the model offers a differentiable solution which can be used in closed-loop systems for auditory feature enhancement or augmented hearing.

With the rise of neural-network (NN) based methods, computational neuroscience has seen an opportunity to map audio or auditory brain signals directly to sound perception [56–58] and to develop computationally efficient methods to compute large-scale differential-equation-based neuronal networks [59]. These developments are transformative, as they can unravel the functional role of hard-to-probe brain areas in perception and yield computationally fast neuromorphic applications. The key to these breakthroughs is the hybrid approach in which knowledge from neuroscience is combined with that of NN-architectures [60]. While the possibilities of NN approaches are numerous when large amounts of training data are available, this is rarely the case for biological systems and human-extracted data. It therefore remains challenging to develop models of biophysical systems which can generalise to a broad range of unseen conditions or stimuli.

Our work presents a solution to this problem for cochlear processing by constraining the CoNNear architecture and its hyperparameters on the basis of a state-of-the-art TL cochlear model. Our general approach consists of four steps: First, derive an analytical description of the biophysical system on the basis of available experimental data. Second, use the analytical model to generate a training data set consisting of responses to a representative range of relevant stimuli. Third, use this training data set to determine the NN-model architecture and constrain its hyperparameters. Finally, verify the ability of the model to generalise to unseen inputs. Here, we demonstrated that CoNNear predictions generalise to a diverse set of audio stimuli by faithfully predicting key cochlear mechanics features to sounds which were excluded from the training. We note that CoNNear performed best on



unseen speech and cochlear mechanics stimuli with levels below 80 dB. We expect that improved generalisability to complex non-speech audio such as environmental sounds and music can be obtained when retraining CoNNear with a more diverse audio set.

Our proposed method is by no means limited to NN-based models of cochlear processing. Indeed, the method can be widely applied to other nonlinear and/or coupled biophysical models of sensory and biophysical systems. Over the years, analytical descriptions of cochlear processing have evolved based on the available experimental data from human and animal cochleae, and they will continue to improve. It is straightforward to train CoNNear to an updated/improved analytical model in step (i), as well as to include different or additional training data in (ii) to further optimise its performance.

## Conclusion

We present a hybrid method which uniquely combines expert knowledge from the fields of computational auditory neuroscience and machine-learning-based audio processing to develop a CoNNear model of human cochlear processing. CoNNear presents an architecture with differentiable equations and operates in real time (<7.5 ms delay) at speeds 2000 times faster than state-of-the-art biophysically realistic models. We have high hopes that the CoNNear framework will inspire a new generation of human-like machine hearing, augmented hearing and automatic speech-recognition systems.

## Methods

Extended Data Fig. 1 describes the parameters which define the CNN-based CoNNear auto-encoder architecture. The encoder layers transform audio input of size  $L$  into a condensed representation of size  $L/2^M \times k_M$ , where  $k_M$  equals the number of filters in the  $M^{\text{th}}$  CNN layer. We used 128 filters per layer with a filter length of 64 samples. The encoder layers use strided convolutions, i.e. the filters were shifted by a time-step of two to halve the temporal dimension after every CNN layer. The decoder contains  $M$  deconvolution layers to re-obtain the original temporal dimension of the audio input ( $L$ ). The first three decoder layers used 128 filters per layer and the final decoder layer had ( $N_{\text{CF}}$ ) filters, equalling the number of cochlear sections simulated in the reference cochlear TL model.

## Training CoNNear

CoNNear was trained using TL-model simulations of  $N_{\text{CF}}$  BM displacement waveforms in response to audio input windows of  $L=2048$  samples. For the context architecture (Fig. 1b), the previous and following  $L_{\text{F}}=L_{\text{B}}=256$  input samples were also available to CoNNear. This architecture included a final cropping layer to remove the context after the last CNN decoder and yield an output size  $L \times N_{\text{CF}}$ . Note that the CoNNear model output units are BM displacement  $y_{\text{BM}}$  in [ $\mu\text{m}$ ], whereas the TL-model outputs are in [ $\text{m}$ ]. This scaling was necessary to enforce a training procedure with sufficiently high digital numbers. For training purposes, and visual comparison between the TL and CoNNear outputs, the  $y_{\text{BM}}$  values of the TL model were multiplied by a factor of  $10^6$  in all figures and analyses.

The TIMIT speech corpus [45] was used for training and contains phonetically balanced sentences with sufficient acoustic diversity. To generate the training data, 2310 sentences from the TIMIT corpus were used. They were upsampled from 16 kHz to 100 kHz to solve the TL-model accurately [36] and the root-mean square (RMS) energy of every utterance was adjusted to 70 dB sound pressure level (SPL). 550 sentences of the TIMIT dataset were omitted from the training and considered as the test set. BM displacements were simulated for 1000 cochlear sections with centre frequencies between 25 Hz and 20 kHz using a nonlinear time-domain TL model of the cochlea [37]. From the TL-model output representation (i.e., 1000  $y_{BM}$  waveforms sampled at 20 kHz), outputs from 201 CFs between 100 Hz and 12 kHz, spaced according to the Greenwood map [6], were chosen to train CoNNear. Above 12 kHz, human hearing sensitivity becomes very poor [41], motivating our choice for the upper limit of considered CFs. CoNNear model parameters were optimised to minimise the mean absolute error (dubbed L1 loss) between the predicted model outputs and the reference TL model outputs. A learning rate of 0.0001 was used with Adam optimiser [61] and the entire framework was developed using the Keras machine learning library [62] with a Tensorflow [63] back-end.

### Cochlear mechanics evaluation metrics

The evaluation stimuli were sampled at 20 kHz and had a duration of 102.4 ms (2048 samples) and 128 ms (2560 samples) for the CoNNear and CoNNear-context model, respectively. Stimulus levels were adjusted using the reference pressure of  $p_0 = 2 \cdot 10^{-5}$  Pa. Only when evaluating how CoNNear generalised to longer duration continuous stimuli (Extended Data Fig. 4), or when investigating its real-time capabilities, did we deviate from this procedure. The following sections describe the four cochlear mechanics evaluation metrics we considered to evaluate the CoNNear predictions. Together, these metrics form a comprehensive description of cochlear processing.

### Cochlear filter tuning

A common approach to characterise auditory or cochlear filters is by means of the equivalent-rectangular bandwidth (ERB) or  $Q_{ERB}$ . The ERB describes the bandwidth of a rectangular filter which passes the same total power than the filter shape estimated from behavioural or cochlear tuning curve experiments [64], and presents a standardised way to characterise the tuning of the asymmetric auditory/cochlear filter shapes. The ERB is commonly used to describe the frequency and level-dependence of human cochlear filtering [4,46,65], and  $Q_{ERB}$  to describe level-dependent cochlear filter characteristics from BM impulse response data [19,66]. We calculated  $Q_{ERB}$  as:

$$Q_{ERB} = \frac{CF}{ERB} \cdot \quad (1)$$

The ERB was determined from the power spectrum of a simulated BM time-domain response to an acoustic click stimulus using the following steps [66]: (i) compute the Fast Fourier Transform of the BM displacement at the considered CF, (ii) compute the area underneath the power spectrum, and (iii) divide the area by the CF. The frequency- and

level-dependence of CoNNear predicted cochlear filters were compared against TL-model predictions and experimental  $Q_{ERB}$  values reported for humans [4].

The acoustic stimuli were condensation clicks of 100- $\mu$ s duration and were scaled to the desired peak-equivalent sound pressure level (dB peSPL), to yield a peak-to-peak click amplitude that matched that of a pure-tone with the same dB SPL level ( $L$ ):

$$\text{click}(t) = 2\sqrt{2} \cdot p_0 \cdot 10^{L/20} \cdot x(t) \text{ with } x(t) = \begin{cases} 1 & \text{for } t \leq 100\mu\text{s} \\ 0 & \text{for } t > 100\mu\text{s} \end{cases} \quad (2)$$

### Cochlear Excitation Patterns

Cochlear excitation patterns can be constructed from the RMS energy of the BM displacement or velocity at each measured CF in response to tonal stimuli of different levels. Cochlear excitation patterns show a characteristic half-octave basal-ward shift of their maxima as stimulus level increases [40]. Cochlear excitation patterns also reflect the level-dependent nonlinear compressive growth of BM-responses when stimulating the cochlea with a pure-tone of the same frequency as the CF of the cochlear measurement site [3]. Cochlear pure-tone transfer-functions and excitation patterns have in several studies been used to describe the level-dependence and tuning properties of cochlear mechanics [2,3,40]. We calculated excitation patterns for all 201 simulated BM displacement waveforms in response to pure tones of 0.5, 1 and 2 kHz frequencies and levels between 10 and 90 dB SPL using:

$$\text{tone}(t) = p_0 \cdot \sqrt{2} \cdot 10^{L/20} \cdot \sin(2\pi f_{\text{tone}} t), \quad (3)$$

where  $t$  corresponds to a time vector of 2048 samples,  $L$  to the desired RMS level in dB SPL, and  $f_{\text{tone}}$  to the stimulus frequencies. The pure-tones were multiplied with a Hanning-shaped 10-ms on- and offset ramp to ensure a gradual onset.

### Cochlear Dispersion

Click stimuli can also be used to characterise the cochlear dispersion properties, as their short duration allows for an easy separation of the cochlear response from the evoking stimulus. At the same time, the broad frequency spectrum of the click excites a large portion of the BM. Cochlear dispersion stems from the longitudinal-coupling and tuning properties of BM mechanics [67] and is observed through later click response onsets for BM responses associated with more apical CFs. In humans, the cochlear dispersion delay mounts up to 10-12 ms for stimulus frequencies associated with apical processing [68]. Here, we use clicks of various sound intensities to evaluate whether CoNNear produced cochlear dispersion and BM click responses in line with predictions from the TL-model.

### Distortion-product otoacoustic emissions (DPOAEs)

DPOAEs can be recorded in the ear-canal using a sensitive microphone and are evoked by two pure-tones with frequencies  $f_1$  and  $f_2$  and SPLs of  $L_1$  and  $L_2$ , respectively. For pure tones with frequency ratios between 1.1 and 1.3 [69], local nonlinear cochlear interactions generate distortion products, which can be seen in the ear-canal recordings as frequency

components which were not originally present in the stimulus. Their strength and shape depends on the properties of the compressive cochlear nonlinearity associated with the electro-mechanical properties of cochlear outer-hair-cells [39], and the most prominent DPOAEs appear at frequencies of  $2f_2 - f_1$  and  $2f_1 - f_2$ . Even though CoNNear was not designed or trained to simulate DPs, they form an excellent evaluation metric, as realistically simulated DPOAE properties would demonstrate that CoNNear was able to capture even the epiphenomena associated with cochlear processing. As a proxy measure for ear-canal recorded DPOAEs, we considered the BM displacement at the highest simulated CF which, in the real ear, would drive the middle-ear and eardrum to yield the ear-canal pressure waveform in an OAE recording. We compared simulated DPs extracted from the fast Fourier transform of the BM displacement response to simultaneously presented pure tones of  $f_1$  and  $f_2 = 1.2 f_1$  with levels according to the commonly adopted experimental scissors paradigm:  $L_1 = 0.4 L_2 + 39$  [51]. We considered  $f_1$  frequencies between 1 and 6 kHz and  $L_2$  levels between 10 and 100 dB SPL.

## Extended Data

Fixed parameters	Summary
Window length	The model input size ( $L$ ) was 2048 samples ( $\approx 100$ ms) for the CoNNear model and $2048 + 2 \cdot 256 = 2560$ samples for the CoNNear-context model. These input sizes were adopted for all training procedures and evaluation metrics, but do not preclude the use of other input durations for the trained CoNNear model. The generalisability of CoNNear towards longer duration stimuli was evaluated using inputs of 10048 and 16384 samples, while its real-time capabilities were evaluated using inputs of 1048 samples.
CNN parameters	The filters in the encoder layers use a stride of 2 for dimensionality reduction, and were upsampled by a factor of 2 in the decoder layers. All CNN layers had a fixed filter length of 64 samples, using 128 filters per layer except for the last decoder layer which had $N_{CF}$ filters to match the BM vibration outputs of the reference TL-model.
Number of cochlear output channels	CoNNear was trained to generate BM displacements of 201 cochlear filters with CFs between 100 Hz and 12 kHz according to the Greenwood map [6].
Hyperparameters	Summary
Number of layers	Increasing the number of layers increases the representational capacity of the model, but evidently also increases the computational complexity (see also Extended Data Fig.2). We investigated CoNNear architectures with layer depths of 4, 6 and 8 to determine the minimum required model size that reliably captured the desired cochlear mechanical properties.
Activation function	The shape of the activation function, or nonlinearity, between the CNN layers is crucial to enable CoNNear to learn the cochlear compression properties associated with the role of the cochlear outer hair cells (OHCs). To mimic the original shape of the OHC input/output function [47], the activation function should cross the x-axis. Standard activation functions (e.g., sigmoid, ReLUs) do not cross the x-axis, and may end up limiting the CoNNear model from learning realistic level-dependent BM impulse response properties. We hence compare the performance of architectures with activation functions that crossed the x-axis, i.e., the parametric rectified linear unit (PReLU) and hyperbolic-tangent (tanh) nonlinearities.
Context	Adding context information to CoNNear can help to avoid discontinuities near the boundaries of the simulated $L$ -sized windows. We examined whether adding a context window of $L_l = L_r = 256$ samples would increase the quality of the CoNNear model predictions.

**Extended Data Fig. 1. Overview of the CoNNear architecture parameters.**

Model	#Param	Time/Epoch [h]	Average L1 loss			
			Training set	Click	1kHz	Test set
tanh/4 lay./context	5,398,528	<b>2.13</b>	0.0187	0.0035	0.0171	0.0190
tanh/6 lay./context	8,544,256	2.15	0.0102	0.0022	0.0123	0.0105
tanh/8 lay./context	11,689,984	2.18	<b>0.0087</b>	<b>0.0017</b>	<b>0.0107</b>	<b>0.0081</b>

**Extended Data Fig. 2. CNN layer depth comparison.**

The first column details the CoNNear architecture. The next columns describe the total number of required model parameters, the required training time per epoch of 2310 TIMIT training sentences and average L1 loss across all windows of the TIMIT training set.

Average L1 losses were also computed for BM displacement predictions to a number of unseen acoustic stimuli (click and 1-kHz pure tones) with levels between 0 and 90 dB SPL. Lastly, average L1 loss was also computed for the 550 sentences of the TIMIT test set. For each evaluated category, the best performing architecture is highlighted in bold font.

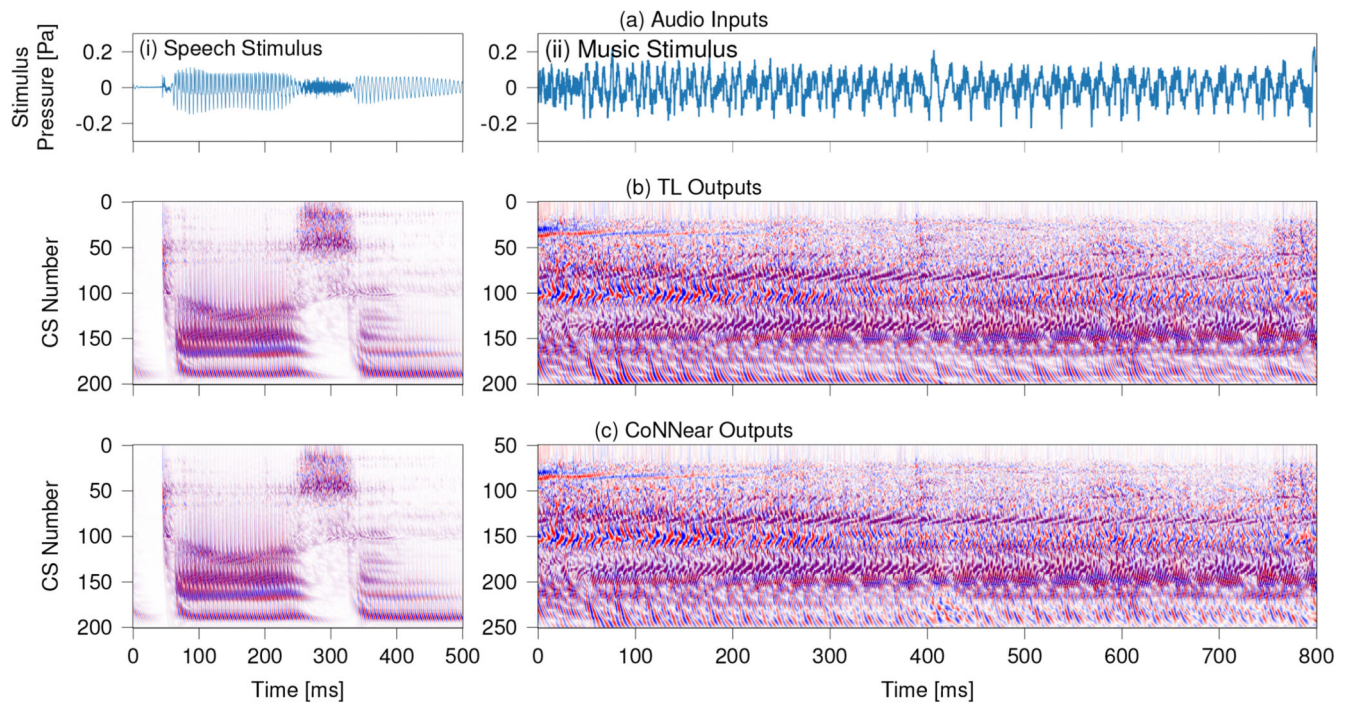
Model	#Param	Time/Epoch [h]	Average L1 loss			
			Training set	Click	1kHz	Test set
PReLU/8 lay./no context	11,982,464	1.89	0.0267	0.0126	0.0302	0.0312
tanh/8 lay./no context	11,507,328	<b>1.84</b>	0.0123	0.0023	0.0116	0.0124
tanh/8 lay./context	11,689,984	2.18	<b>0.0087</b>	<b>0.0017</b>	<b>0.0107</b>	<b>0.0081</b>

**Extended Data Fig. 3. Activation function comparison.**

The first column details the CoNNear architecture. The next columns describe the total number of required model parameters, the required training time per epoch of 2310 TIMIT training sentences and average L1 loss across all windows of the TIMIT training set.

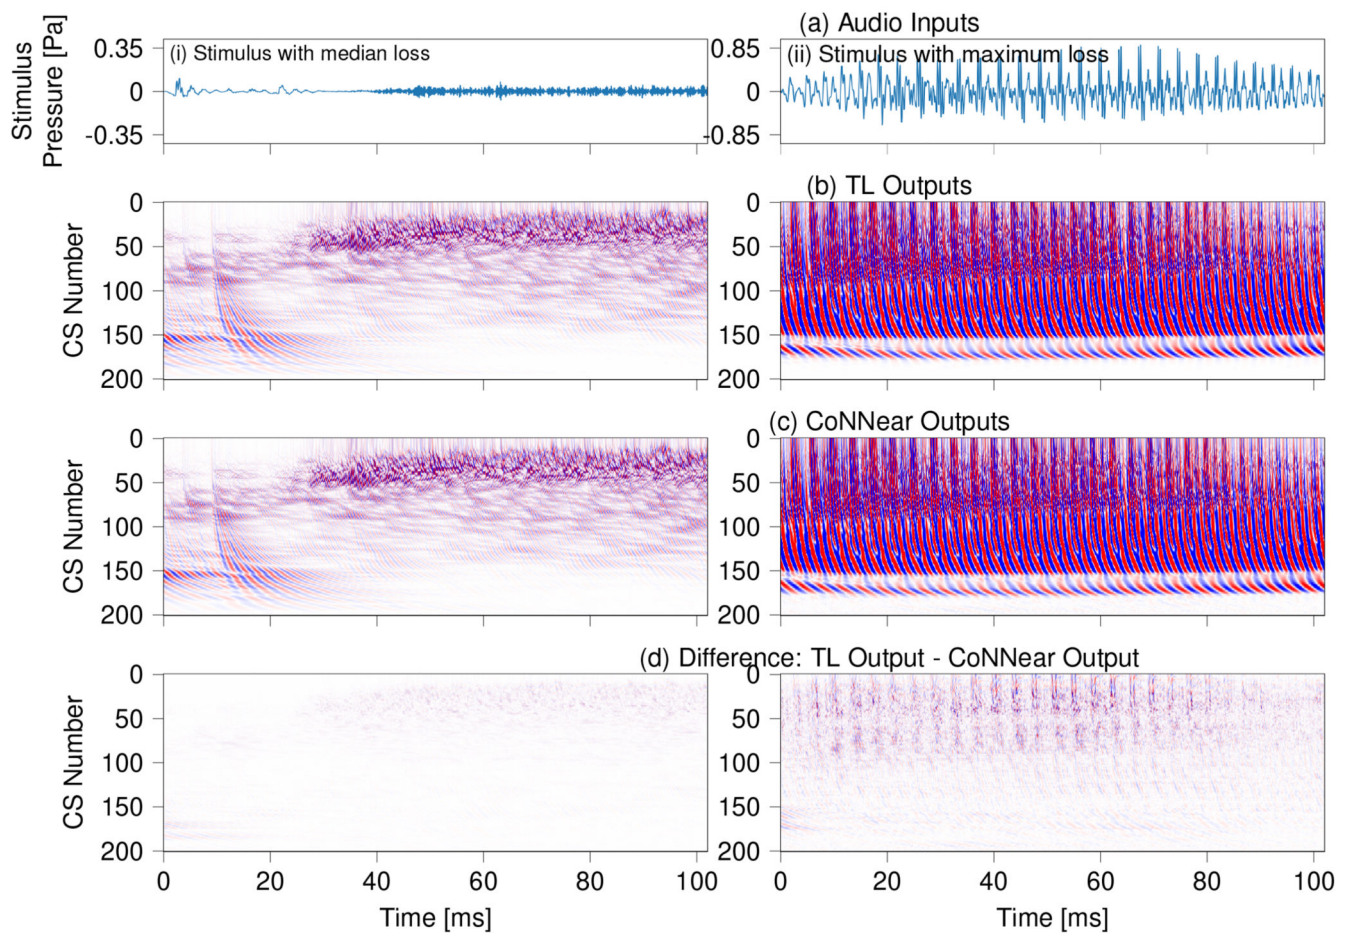
Average L1 losses were also computed for BM displacement predictions to a number of unseen acoustic stimuli (click and 1-kHz pure tones) with levels between 0 and 90 dB SPL. Lastly, average L1 loss was also computed for the 550 sentences of the TIMIT test set. For each evaluated category, the best performing architecture is highlighted in bold font.





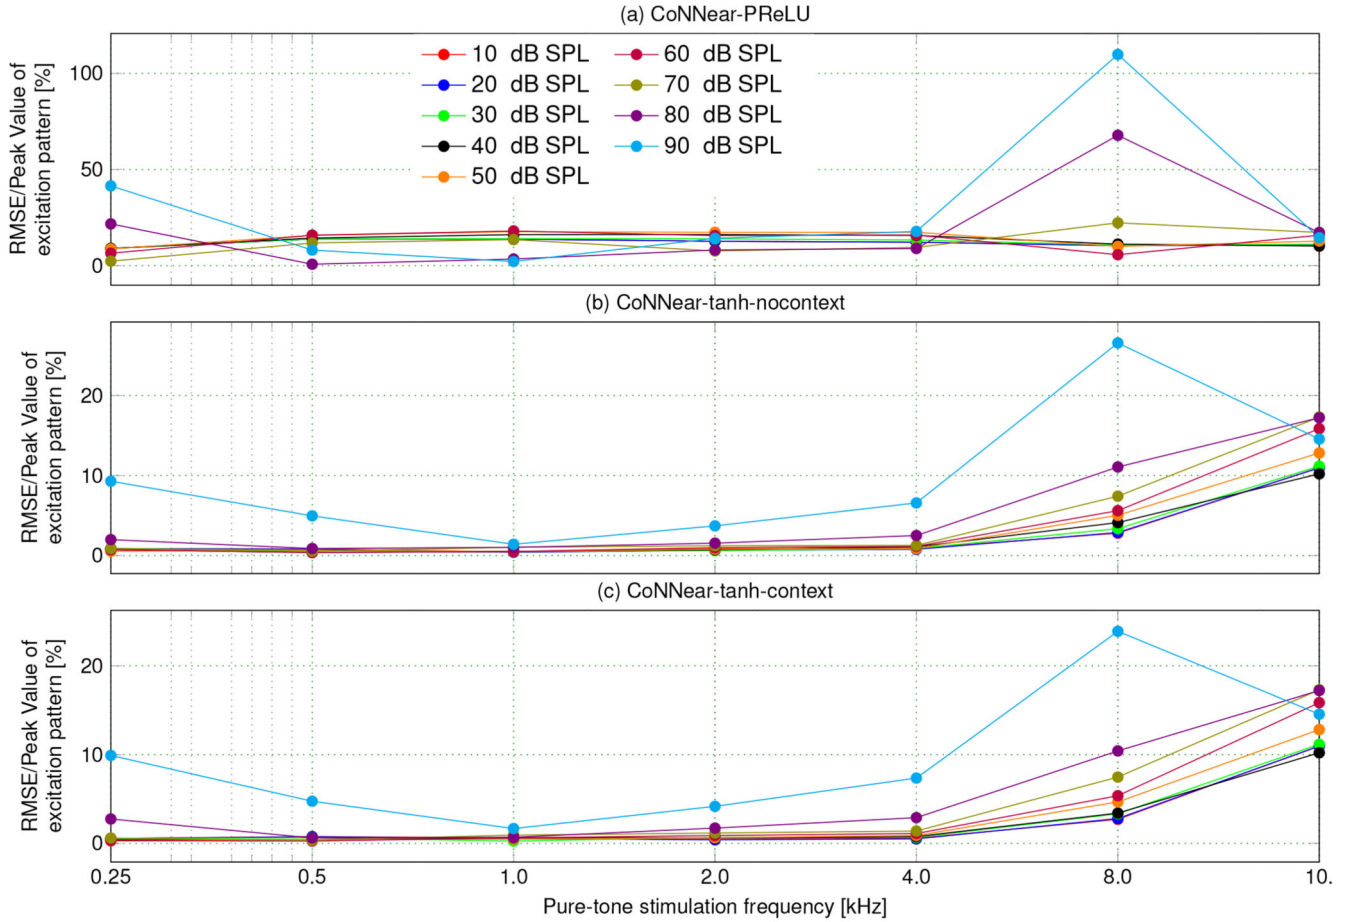
**Extended Data Fig. 4. Simulated BM displacements for a 10048-sample speech stimulus and a 16384-sample music stimulus.**

The stimulus waveform is depicted in panel (a) and panels (b)-(c) depict instantaneous BM displacement intensities (darker colours = higher intensities) of the simulated TL-model (b) and CoNNear (c) outputs. The  $N_{CF}=201$  considered output channels are labelled per channel number: channel 1 corresponds to a CF of 12 kHz and channel 201 to a CF of 100 Hz. The same colour scale was used for both simulations and ranged between  $-0.5 \mu\text{m}$  (blue) and  $0.5 \mu\text{m}$  (red). The left panels show simulations to a speech stimulus from the Dutch matrix test [49] and the right panels shows simulations to a music fragment (Radiohead - No Surprises).



**Extended Data Fig. 5. Comparing TL and CoNNear model predictions at the median and maximum L1 prediction error.**

The figure visually compares BM displacement intensities of the BM (b) and CoNNear (c) model to audio fragments which resulted in the median and maximum L1 errors of 0.008 and 0.038 simulated for the TIMIT test set (Fig. 5). The  $N_{CF}=201$  considered output channels are labelled per channel number: channel 1 corresponds to a CF of 12 kHz and channel 201 to a CF of 100 Hz. The same colour scale was used for both simulations and ranged between  $-0.5 \mu\text{m}$  (blue) and  $0.5 \mu\text{m}$  (red).



**Extended Data Fig. 6. Root mean-square error (RMSE) between simulated excitation patterns of the TL and CoNNear models reported as fraction of the TL excitation pattern maximum (cf. Fig. 3).**

Using the PReLU activation function (a) leads to an overall high RMSE as this architecture failed to learn the level-dependent cochlear compression characteristics and filter shapes. The models using the tanh nonlinearity (b),(c) did learn to capture the level-dependent properties of cochlear excitation patterns, and performed with errors below 5% for the frequency ranges and stimulus levels captured by the speech training data (for CFs below 5 kHz, and stimulation levels below 90 dB SPL) The RMSE increased above 5% for all architectures when evaluating its performance on 8- and 10-kHz excitation patterns. This decreased performance results from the limited frequency content of the TIMIT training material.

### Acknowledgments

This work was supported by the European Research Council (ERC) under the Horizon 2020 Research and Innovation Programme (grant agreement No 678120 RobSpear). The authors would like to thank Christopher and Sarita Shera for their help with the final edits.



## Data availability

The source code of the TL-model used for training is available via [10.5281/zenodo.3717431](https://zenodo.org/record/3717431) or [github/HearingTechnology/Verhulstetal2018Model](https://github.com/HearingTechnology/Verhulstetal2018Model), the TIMIT speech corpus used for training can be found online [45]. Most figures in this paper can be reproduced using the CoNNear model repository.

## Code availability

The code for the trained CoNNear model, including instructions of how to execute it is available from [github.com/HearingTechnology/CoNNear cochlea](https://github.com/HearingTechnology/CoNNear_cochlea) or [10.5281/zenodo.4056552](https://zenodo.org/record/4056552). A non-commercial, academic UGent license applies.

## References

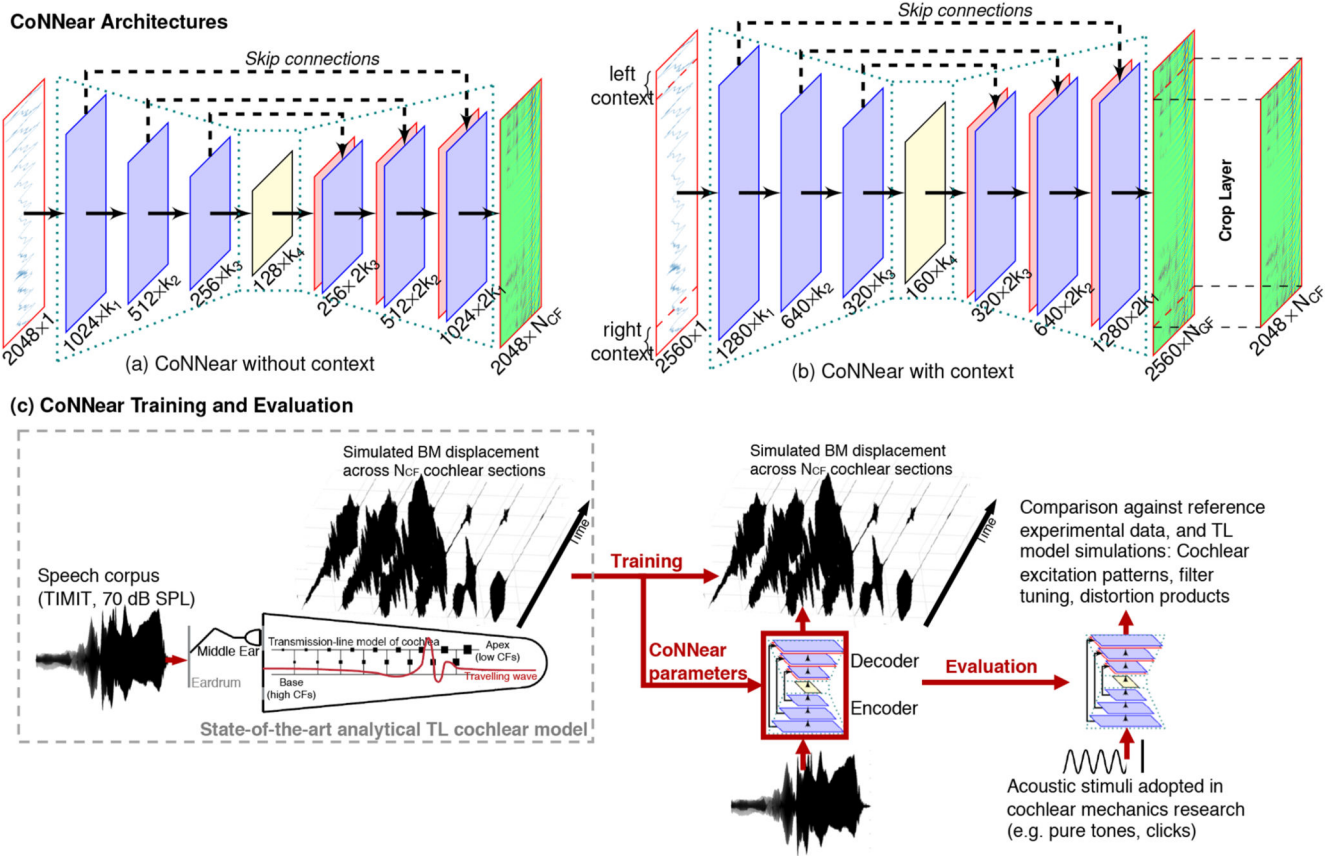
1. von Békésy G. Travelling Waves as Frequency Analysers in the Cochlea. *Nature*. 1970; 225:1207–1209. [PubMed: 5435349]
2. Narayan SS, Temchin AN, Recio A, Ruggero MA. Frequency tuning of basilar membrane and auditory nerve fibers in the same cochleae. *Science*. 1998; 282:1882–1884. [PubMed: 9836636]
3. Robles L, Ruggero MA. Mechanics of the mammalian cochlea. *Phys Rev*. 2001; 81:1305–1352.
4. Shera CA, Guinan JJ, Oxenham AJ. Revised estimates of human cochlear tuning from otoacoustic and behavioral measurements. *Proc Nat Acad Sci*. 2002; 99:3318–3323. [PubMed: 11867706]
5. Oxenham AJ, Shera CA. Estimates of human cochlear tuning at low levels using forward and simultaneous masking. *J Assoc Res Otolaryngol*. 2003; 4:541–554. [PubMed: 14716510]
6. Greenwood DD. A cochlear frequency-position function for several species—29 years later. *J Acoust Soc Am*. 1990; 87:2592–2605. [PubMed: 2373794]
7. Jepsen ML, Dau T. Characterizing auditory processing and perception in individual listeners with sensorineural hearing loss. *J Acoust Soc Am*. 2011; 129:262–281. [PubMed: 21303008]
8. Bondy J, Becker S, Bruce I, Trainor L, Haykin S. A novel signal-processing strategy for hearing-aid design: neurocompensation. *Sig Process*. 2004; 84:1239–1253.
9. Ewert SD, Kortlang S, Hohmann V. A Model-based hearing aid: Psychoacoustics, models and algorithms. *Proc Meetings on Acoust*. 2013; 19
10. Mondol S, Lee S. A machine learning approach to fitting prescription for hearing aids. *Electronics*. 2019; 8:736.
11. Lyon, RF. *Human and Machine Hearing: Extracting Meaning from Sound*. Cambridge University Press; 2017.
12. Baby, D, Van hamme, H. Investigating modulation spectrogram features for deep neural network-based automatic speech recognition *Proc Insterspeech*. Dresden; Germany: 2015. 2479–2483.
13. de Boer E. Auditory physics. physical principles in hearing theory. i. *Phys Rep*. 1980; 62:87–174.
14. Diependaal RJ, Duifhuis H, Hoogstraten HW, Viergever MA. Numerical methods for solving one-dimensional cochlear models in the time domain. *J Acoust Soc Am*. 1987; 82:1655–1666. [PubMed: 3693707]
15. Zweig G. Finding the impedance of the organ of corti. *J Acoust Soc Am*. 1991; 89:1229–1254. [PubMed: 2030212]
16. Talmadge CL, Tubis A, Wit HP, Long GR. Are spontaneous otoacoustic emissions generated by self-sustained cochlear oscillators? *J Acoust Soc Am*. 1991; 89:2391–2399. [PubMed: 1860998]
17. Moleti A, et al. Transient evoked otoacoustic emission latency and estimates of cochlear tuning in preterm neonates. *J Acoust Soc Am*. 2008; 124:2984–2994. [PubMed: 19045786]
18. Epp B, Verhey JL, Mauermann M. Modeling cochlear dynamics: Interrelation between cochlea mechanics and psychoacoustics. *J Acoust Soc Am*. 2010; 128:1870–1883. [PubMed: 20968359]

19. Verhulst S, Dau T, Shera CA. Nonlinear time-domain cochlear model for transient stimulation and human otoacoustic emission. *J Acoust Soc Am.* 2012; 132:3842–3848. [PubMed: 23231114]
20. Zweig G. Nonlinear cochlear mechanics. *J Acoust Soc Am.* 2016; 139:2561–2578. [PubMed: 27250151]
21. Hohmann, V. *Signal Processing in Hearing Aids.* Springer New York; New York, NY: 2008. 205–212.
22. Rascon C, Meza I. Localization of sound sources in robotics: A review. *Robot Auton Syst.* 2017; 96:184–210.
23. Mogran, N, Boulard, H, Hermansky, H. *Automatic Speech Recognition: An Auditory Perspective.* Springer New York; New York, NY: 2004. 309–338.
24. Patterson RD, Allerhand MH, Giguère C. Time-domain modeling of peripheral auditory processing: A modular architecture and a software platform. *J Acoust Soc Am.* 1995; 98:1890–1894. [PubMed: 7593913]
25. Shera CA. Frequency glides in click responses of the basilar membrane and auditory nerve: Their scaling behavior and origin in traveling-wave dispersion. *J Acoust Soc Am.* 2001; 109:2023–2034. [PubMed: 11386555]
26. Shera, CA, Guinan, JJ. *Mechanisms of mammalian otoacoustic emission* Active Processes and Otoacoustic Emissions in Hearing. Springer; 2008. 305–342.
27. Hohmann V. Frequency analysis and synthesis using a Gammatone filterbank. *Acta Acust United Ac.* 2002; 88:433–442.
28. Saremi A, et al. A comparative study of seven human cochlear filter models. *J Acoust Soc Am.* 2016; 140:1618–1634. [PubMed: 27914400]
29. Lopez-Poveda EA, Meddis R. A human nonlinear cochlear filterbank. *J Acoust Soc Am.* 2001; 110:3107–3118. [PubMed: 11785812]
30. Lyon RF. Cascades of two-pole-two-zero asymmetric resonators are good models of peripheral auditory function. *J Acoust Soc Am.* 2011; 130:3893–3904. [PubMed: 22225045]
31. Saremi A, Lyon RF. Quadratic distortion in a nonlinear cascade model of the human cochlea. *J Acoust Soc Am.* 2018; 143:EL418–EL424. [PubMed: 29857771]
32. Altoè A, Charaziak KK, Shera CA. Dynamics of cochlear nonlinearity: Automatic gain control or instantaneous damping? *J Acoust Soc Am.* 2017; 142:3510–3519. [PubMed: 29289066]
33. Baby, D; Verhulst, S. SERGAN: Speech enhancement using relativistic generative adversarial networks with gradient penalty. *Acoustics, Speech and Signal Processing (ICASSP), IEEE Int Conf; Brighton, UK: 2019. 106–110.*
34. Pascual, S, Bonafonte, A, Serrà, J. SEGAN: speech enhancement generative adversarial network *Interspeech. ISCA; 2017. 3642–3646.*
35. Drakopoulos, F, Baby, D, Verhulst, S. Real-Time Audio Processing on a Raspberry Pi using Deep Neural Networks *23rd Int Congress on Acoustics (ICA). Aachen; Germany: 2019.*
36. Altoè A, Pulkki V, Verhulst S. Transmission line cochlear models: Improved accuracy and efficiency. *J Acoust Soc Am.* 2014; 136:EL302–EL308. [PubMed: 25324114]
37. Verhulst S, Altoè A, Vasilkov V. Computational modeling of the human auditory periphery: Auditory-nerve responses, evoked potentials and hearing loss. *Hear Res.* 2018; 360:55–75. [PubMed: 29472062]
38. Oxenham, AJ, Wojtczak, M. *Frequency selectivity and masking: Perception.* Oxford University Press; 2010.
39. Robles L, Ruggero MA, Rich NC. Two-tone distortion in the basilar membrane of the cochlea. *Nature.* 1991; 349:413. [PubMed: 1992342]
40. Ren T. Longitudinal pattern of basilar membrane vibration in the sensitive cochlea. *Proc Nat Acad Sci.* 2002; 99:17101–17106. [PubMed: 12461165]
41. *Precise and full-range determination of two-dimensional equal loudness contours.* Standard, International Organization for Standardization; Geneva, CH: 2003.
42. Lorenzi C, Gilbert G, Carn H, Garnier S, Moore BC. Speech perception problems of the hearing impaired reflect inability to use temporal fine structure. *Proc Nat Acad Sci.* 2006; 103:18866–18869. [PubMed: 17116863]

43. Isola P, Zhu JY, Zhou T, Efros AA. Image-to-image translation with conditional adversarial networks. *IEEE-CVPR*. 2017:5967–5976.
44. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015; 521:436–444. [PubMed: 26017442]
45. Garofolo JS, et al. Darpa timit acoustic phonetic continuous speech corpus cdrom. 1993
46. Shera CA, Guinan JJ, Oxenham AJ. Otoacoustic estimation of cochlear tuning: validation in the chinchilla. *J Assoc Res Otolaryngol*. 2010; 11:343–365. [PubMed: 20440634]
47. Russell I, Cody A, Richardson G. The responses of inner and outer hair cells in the basal turn of the guinea-pig cochlea and in the mouse cochlea grown in vitro. *Hear Res*. 1986; 22:199–216. [PubMed: 3733540]
48. Paul, DB; Baker, JM. The Second International Conference on Spoken Language Processing, ICSLP; ISCA; 1992.
49. Houben R, et al. Development of a dutch matrix sentence test to assess speech intelligibility in noise. *Int J Audiol*. 2014; 53:760–763. [PubMed: 24959915]
50. Gemmeke, JF; , et al. Audio set: An ontology and human-labeled dataset for audio events. 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP); IEEE; 2017. 776–780.
51. Kummer P, Janssen T, Hulin P, Arnold W. Optimal L1—L2 primary tone level separation remains independent of test frequency in humans. *Hear Res*. 2000; 146:47–56. [PubMed: 10913883]
52. Dorn PA, et al. Distortion product otoacoustic emission input/output functions in normal-hearing and hearing-impaired human ears. *J Acoust Soc Am*. 2001; 110:3119–3131. [PubMed: 11785813]
53. Janssen, T, Müller, J. Otoacoustic emissions as a diagnostic tool in a clinical context *Active Processes and Otoacoustic Emissions in Hearing*. Springer; 2008. 421–460.
54. Verhulst S, Ernst F, Garrett M, Vasilkov V. Suprathreshold psychoacoustics and envelope-following response relations: Normal-hearing, synaptopathy and cochlear gain loss. *Acta Acus united Ac*. 2018; 104:800–803.
55. Verhulst S, Bharadwaj HM, Mehraei G, Shera CA, Shinn-Cunningham BG. Functional modeling of the human auditory brainstem response to broadband stimulation. *J Acoust Soc Am*. 2015; 138:1637–1659. [PubMed: 26428802]
56. Kell AJ, Yamins DL, Shook EN, Norman-Haignere SV, McDermott JH. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron*. 2018; 98:630–644. [PubMed: 29681533]
57. Akbari H, Khalighinejad B, Herrero JL, Mehta AD, Mesgarani N. Towards reconstructing intelligible speech from the human auditory cortex. *Sci Rep*. 2019; 9:1–12. [PubMed: 30626917]
58. Kell AJ, McDermott JH. Deep neural network models of sensory systems: windows onto the role of task constraints. *Curr Opin Neurobiol*. 2019; 55:121–132.
59. Amsalem O, et al. An efficient analytical reduction of detailed nonlinear neuron models. *Nat Comm*. 2020; 11:1–13.
60. Richards BA, et al. A deep learning framework for neuroscience. *Nature Neurosci*. 2019; 22:1761–1770. [PubMed: 31659335]
61. Kingma DP, Ba J. Adam: A method for stochastic optimization. *CoRR*. 2014
62. Chollet, F, , et al. Keras; 2015. <https://keras.io>
63. Abadi M, et al. TensorFlow: Large-scale machine learning on heterogeneous systems. 2015
64. Moore BC, Glasberg BR. Suggested formulae for calculating auditory-filter bandwidths and excitation patterns. *J Acoust Soc Am*. 1983; 74:750–753. [PubMed: 6630731]
65. Glasberg BR, Moore BC. Derivation of auditory filter shapes from notched-noise data. *Hear Res*. 1990; 47:103–138. [PubMed: 2228789]
66. Rauber S, Verhulst S. Otoacoustic emission estimates of human basilar membrane impulse response duration and cochlear filter tuning. *Hear Res*. 2016; 342:150–160. [PubMed: 27989947]
67. Ramamoorthy S, Zha DJ, Nuttall AL. The biophysical origin of traveling-wave dispersion in the cochlea. *Biophys J*. 2010; 99:1687–1695. [PubMed: 20858412]
68. Dau T, Wegner O, Mellert V, Kollmeier B. Auditory brainstem responses with optimized chirp signals compensating basilar-membrane dispersion. *J Acoust Soc Am*. 2000; 107:1530–1540. [PubMed: 10738807]

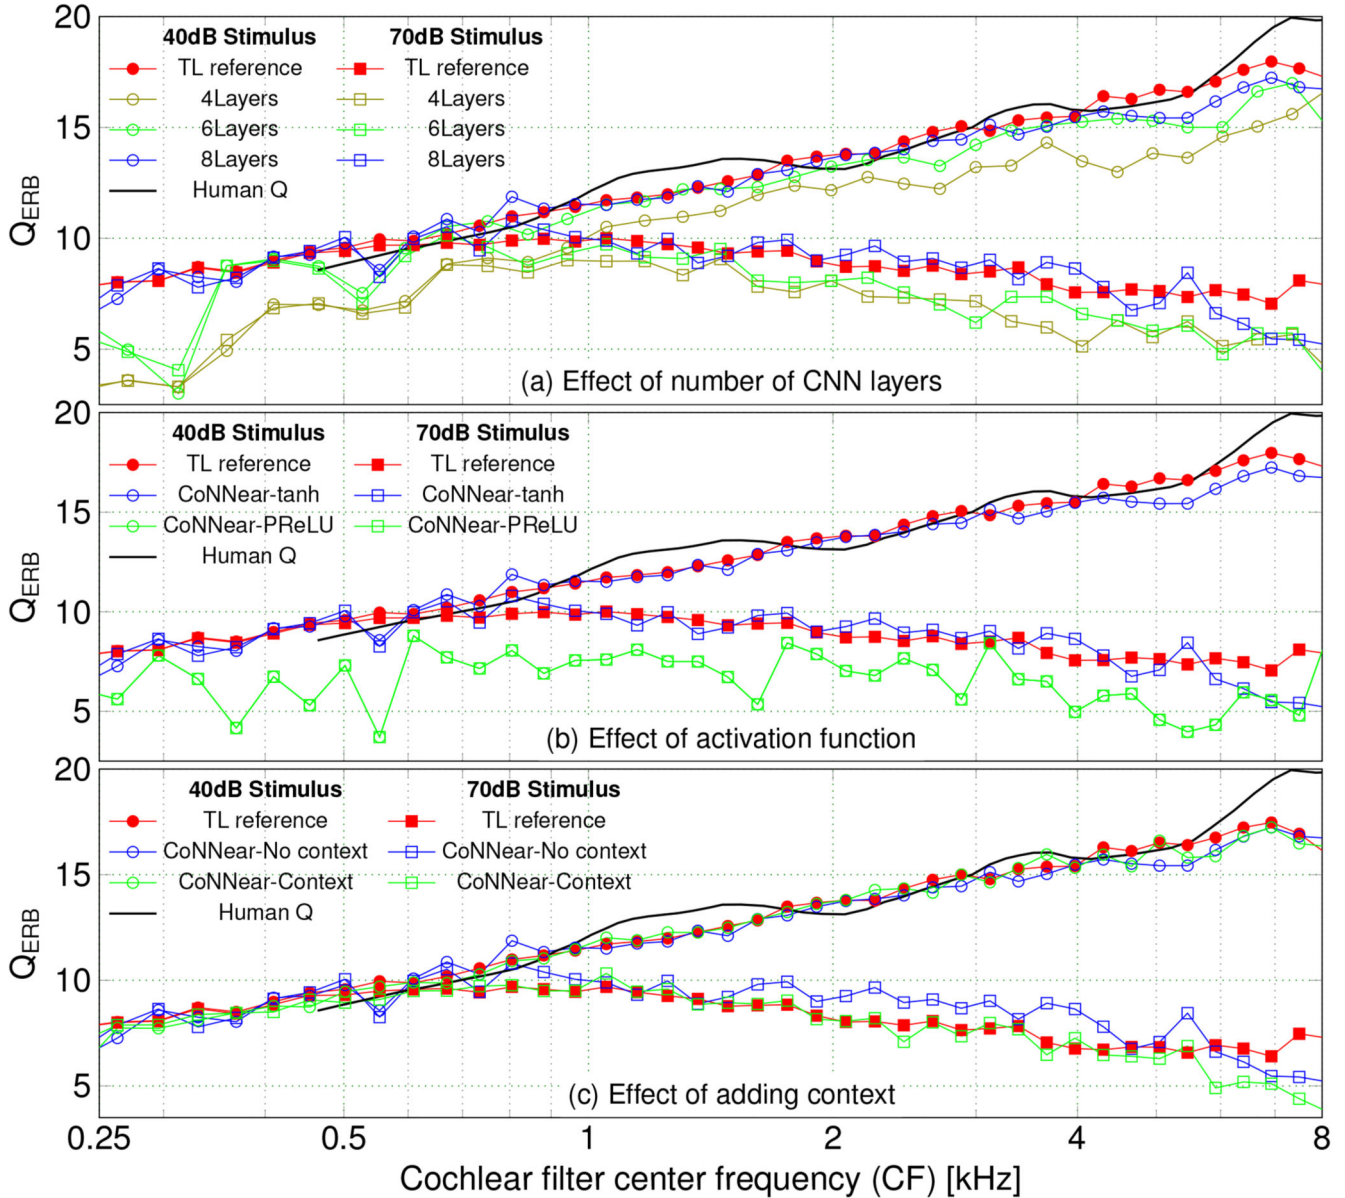


69. Neely ST, Johnson TA, Kopun J, Dierking DM, Gorga MP. Distortion-product otoacoustic emission input/output characteristics in normal-hearing and hearing-impaired human ears. *J Acoust Soc Am*. 2009; 126:728–738. [PubMed: 19640039]



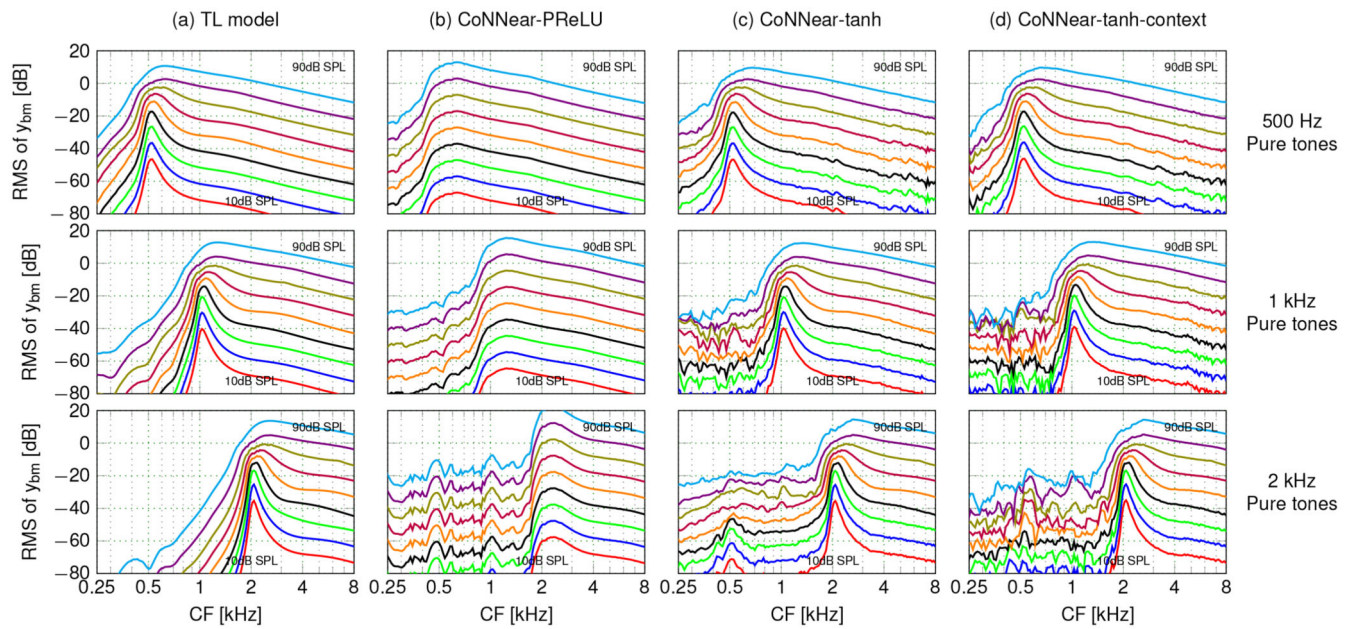
**Fig 1. CoNNear Overview.**

CoNNear is a fully convolutional encoder-decoder neural network with strided convolutions and skip-connections to map audio input to 201 basilar-membrane vibration outputs of different cochlear sections ( $N_{CF}$ ) in the time-domain. CoNNear architectures with (a) and without (b) context are shown. The final CoNNear model has four encoder and decoder layers, uses context and includes a tanh activation function between the CNN layers. (c) provides an overview of the model training and evaluation procedure. Whereas reference, analytical TL-model simulations to a speech corpus were used to train the CoNNear parameters, evaluation of the model was performed using simple acoustic stimuli commonly adopted in cochlear mechanics studies.



**Fig 2. CoNNear hyperparameter tuning.**


Cochlear filter tuning ( $Q_{ERB}$ ) was simulated to 100- $\mu$ s clicks of 40 and 70 dB peSPL for different cochlear filter centre frequencies (CF). Reference human  $Q_{ERB}$  estimates for low-stimulus levels [46] are shown and compared to simulations of the reference TL model and CoNNear models with (c) and without (a,b) context. (a) Increasing the number of CoNNear layers from 4 to 8, improved the  $Q_{ERB}$ -across-CF simulations. (b) Comparing architectures with the PReLU or tanh activation function shows that the PReLU nonlinearity failed to capture level-dependent cochlear filter tuning because the  $Q_{ERB}$  functions remained invariant to stimulus level changes. 8-layer CNN architectures were used for this simulation. (c) Adding context to the 8-layer, tanh model further improved the CoNNear predictions by showing less  $Q_{ERB}$  fluctuations across CF.



**Fig 3. Comparing cochlear excitation patterns across model architectures.**

Simulated root-mean-square (RMS) levels of BM displacement across CF for pure-tone stimuli between 10 and 90 dB SPL. From top to bottom, the stimulus frequencies were 500 Hz, 1 kHz and 2 kHz, respectively. Both the reference, TL-model, level-dependent excitation pattern shape changes and compressive growth of pattern maxima (a), were captured by the tanh architectures (c,d), but not by the PReLU architecture (b).






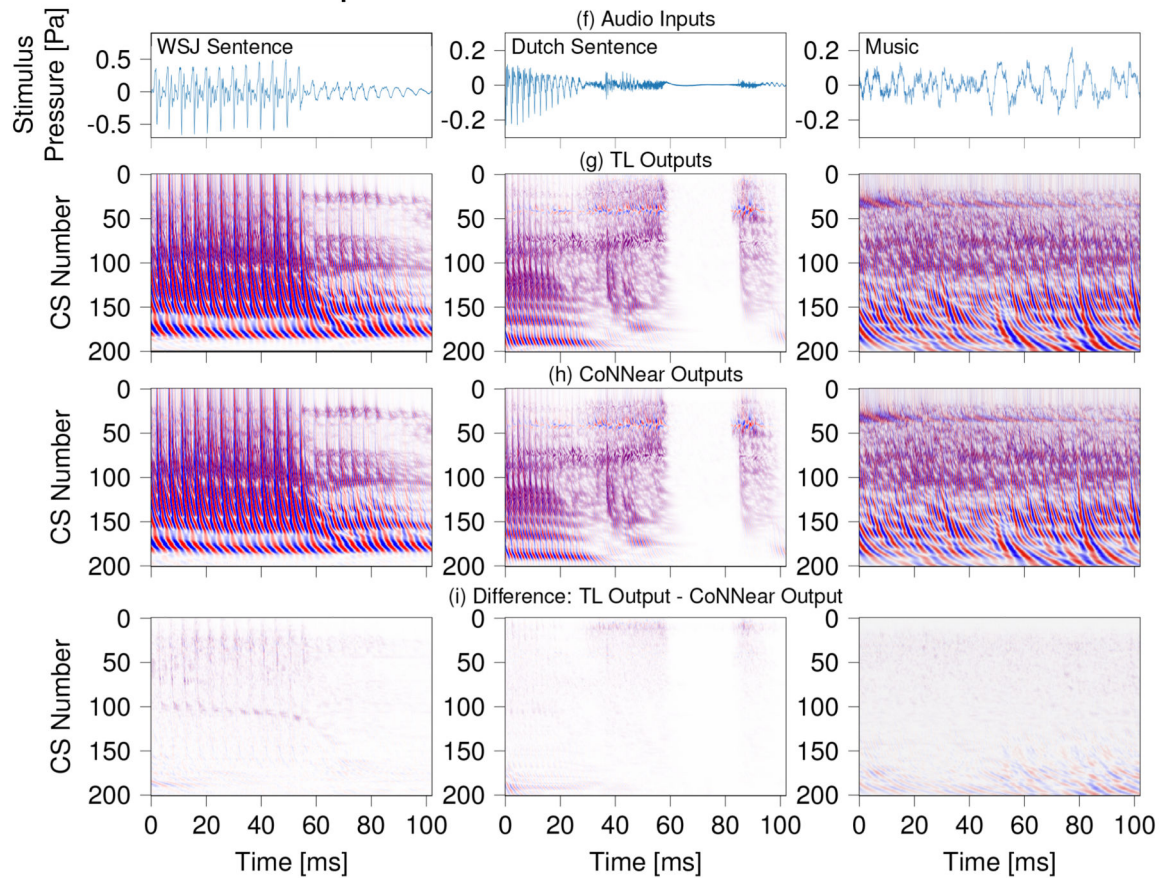
**Fig 4. Effect of adding context to the CoNNear simulations.**

Simulated BM displacements for a 2048-sample-long speech fragment of the TIMIT test set (i.e., unseen during training). The stimulus waveform is shown in panel (a) and panels (b)-(d) depict instantaneous BM displacement intensities (darker colours = higher intensities) of simulated TL-model outputs (b) and two CoNNear architecture outputs: without (c) and with (d) context. The left and right segments show the output of the neighbouring windows to demonstrate the discontinuity over the boundaries. The  $N_{CF}=201$  output channels are labelled per channel number: channel 1 corresponds to a CF of 12 kHz and channel 201 to a CF of 100 Hz. Intensities varied between  $-0.5 \mu\text{m}$  (blue) and  $0.5 \mu\text{m}$  (red) for all panels.

### Boxplots of L1 errors between TL and CoNNear



### Simulation on unseen audio examples




**Fig 5. Generalisability of CoNNear to unseen input.**

Top: boxplots comparing the distribution of L1 losses between TL-model and CoNNear model simulations for 2048-long windows within (a) 2310 sentences from the TIMIT training set (131760 calculated L1 losses, 16-kHz sampled audio) (b) the 550 sentences in the TIMIT test set (32834 L1 losses, 16-kHz sampled audio), (c) 100 sentences from the Dutch Matrix test (7903 L1 losses, 44.1-kHz sampled audio), (d) All songs of Radioheads OK Computer (48022 L1 losses, 16-kHz downsampled audio), (e) 400 random audio-fragments from the natural-sounds AudioSet database (53655 L1 losses, 16-kHz downsampled audio). The red line in the boxplot shows the median of the L1 errors, the black box denotes the values between the 25% (Q1) and 75% (Q3) quartiles. The blue

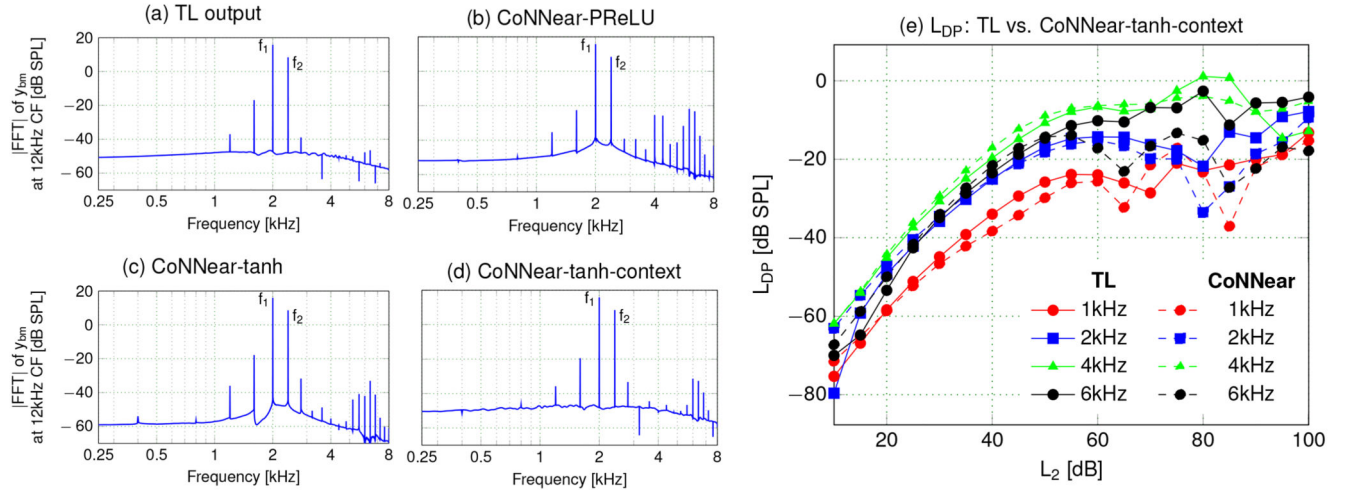


whiskers denote the error values that fall within the  $1.5 * (Q3 - Q1)$  range. Outliers are shown as circles. Cochlear model predictions and prediction errors for the windows associated with median and maximum L1 loss in the TIMIT test set are shown in Extended Data Fig. 5. **Bottom:** Along the columns, simulated BM displacements for three different 2048-long audio stimuli are shown: a recording of the English Wall Street Journal speech corpus [48], a sentence from the Dutch matrix test [49] and a music fragment taken from Radiohead - No surprises. Stimulus waveforms are depicted in panel (f) and instantaneous BM displacement intensities (darker colours = higher intensities) of the simulated TL-model and CoNNear outputs are depicted in (g-h). Panel (i) shows the intensity difference between the TL and CoNNear outputs. The  $N_{CF}=201$  considered output channels are labelled per channel number: channel 1 corresponds to a CF of 12 kHz and channel 201 to a CF of 100 Hz. The same colour map was used for all figures and ranged between  $-0.5 \mu\text{m}$  (blue) and  $0.5 \mu\text{m}$  (red).

**Cochlear dispersion in click response**



**DPOAE simulations**



**Fig 6. Cochlear dispersion and DPOAEs.**

Top: Comparing cochlear dispersion properties across model architectures. Panel (a) shows the stimulus pressure, while panels (b)-(d) show instantaneous BM displacement intensities for CFs (channel numbers, CS) between 100 Hz (channel 201) and 12 kHz (channel 1). The colour scale is the same in all figure panels, and ranges between  $-15 \mu\text{m}$  (blue) and  $15 \mu\text{m}$  (red). **Bottom:** Comparing simulated DPOAEs across model architectures. The frequency response of the 12-kHz CF channel (i.e., fast Fourier transform of the BM displacement waveform) was considered as a proxy for otoacoustic emissions recorded in the ear-canal. Panels (a) - (d) show simulations in response to two pure tones of  $f_1$  of 2.0 and  $f_2$  of 2.4 kHz for different model architectures. In humans, the most pronounced distortion product occurs at  $2f_1 - f_2$  (1.6 kHz). Panel (e) depicts simulated distortion-product levels ( $L_{DP}$ ) compared between TL and the CoNNear-tanh-context model.  $L_{DP}$  was extracted from the frequency response of the 12-kHz CF channel at  $2f_1 - f_2$ . Simulations were conducted for  $L_2$  levels between 10 and 100 dB SPL and  $L_1 = 0.4 L_2 + 39$  [51].  $f_1$  ranged between 1 and 6 kHz, following a  $f_2/f_1$  ratio of 1.2.

**Table 1****Model calculation speed.**

Comparison of the time required to calculate a TL and CoNNear model window of 1048 samples on a CPU (Apple MacBook Air, 1.8 GHz Dual-Core processor) and a GPU (NVIDIA GTX1080). The calculation time for the first window is considered separately for the GPU computations since this window also includes the weight initialisation. For each evaluated category, the best performing architecture is highlighted with a bold font.

<b>Model</b>	<b>#Param</b>	<b>CPU (s/window)</b>	<b>GPU 1st window (s)</b>	<b>GPU (ms/window)</b>
PReLU/8 lay./no context	11,982,464	0.222	1.432	7.70
tanh/8 lay./no context	11,507,328	<b>0.195</b>	1.390	7.59
tanh/8 lay./context	11,689,984	0.236	<b>1.257</b>	<b>7.27</b>
Transmission Line	N/A	25.16	N/A	16918