



(51) **International Patent Classification:**
G01L 19/00 (2006.01) *G10L 25/30* (2013.01)

(21) **International Application Number:**
 PCT/US2019/026824

(22) **International Filing Date:**
 10 April 2019 (10.04.2019)

(25) **Filing Language:** English

(26) **Publication Language:** English

(30) **Priority Data:**

| | | |
|------------|----------------------------|----|
| 62/656,275 | 11 April 2018 (11.04.2018) | US |
| 18173673.7 | 22 May 2018 (22.05.2018) | EP |
| 62/829,552 | 04 April 2019 (04.04.2019) | US |

(74) **Agent:** ANDERSEN, Robert L. et al.; Dolby Laboratories, Inc., Intellectual Property Group, 1275 Market Street, San Francisco, California 94103 (US).

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(71) **Applicant:** DOLBY LABORATORIES LICENSING CORPORATION [US/US]; 1275 Market Street, San Francisco, California 94103 (US).

(84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(72) **Inventors:** FEJGIN, Roy M.; c/o Dolby Laboratories, Inc., 1275 Market Street, San Francisco, California 94103 (US). DAVIDSON, Grant A.; c/o Dolby Laboratories, Inc., 1275 Market Street, San Francisco, California 94103 (US). WU, Chih-Wei; 419 19th Ave Apt. 5, San Francisco, California 94121 (US). KUMAR, Vivek; c/o Dolby Laboratories, Inc., 1275 Market Street, San Francisco, California 94103 (US).

(54) **Title:** PERCEPTUALLY-BASED LOSS FUNCTIONS FOR AUDIO ENCODING AND DECODING BASED ON MACHINE LEARNING

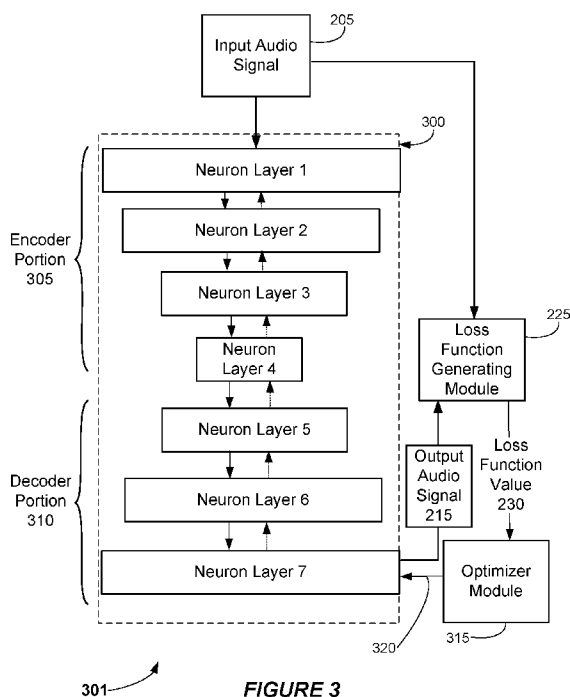


FIGURE 3

(57) **Abstract:** Computer-implemented methods for training a neural network, as well as for implementing audio encoders and decoders via trained neural networks, are provided. The neural network may receive an input audio signal, generate an encoded audio signal and decode the encoded audio signal. A loss function generating module may receive the decoded audio signal and a ground truth audio signal, and may generate a loss function value corresponding to the decoded audio signal. Generating the loss function value may involve applying a psychoacoustic model. The neural network may be trained based on the loss function value. The training may involve updating at least one weight of the neural network.

WO 2019/199995 A1

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*

physical state of at least one non-transitory storage medium location corresponding with at least one weight of the neural network.

[0005] In some implementations, a first portion of the neural network may generate the encoded audio signal and a second portion of the neural network may decode the encoded audio signal. In some such implementations, the first portion of the neural network may include an input neuron layer and a plurality of hidden neuron layers. The input neuron layer may, in some instances, include more neurons than a final hidden neuron layer. At least some neurons of the first portion of the neural network may be configured with rectified linear unit (ReLU) activation functions. In some examples, at least some neurons in hidden layers of the second portion of the neural network may be configured with ReLU activation functions and at least some neurons in an output layer of the second portion may be configured with sigmoidal activation functions.

[0006] According to some examples, the psychoacoustic model may be based, at least in part, on one or more psychoacoustic masking thresholds. In some implementations, the psychoacoustic model may involve modeling an outer ear transfer function, grouping into critical bands, frequency-domain masking (including but not limited to level-dependent spreading), modeling of a frequency-dependent hearing threshold and/or calculation of a noise-to-mask ratio. In some examples, the loss function may involve calculating an average noise-to-masking ratio and the training may involve minimizing the average noise-to-masking ratio.

[0007] Some audio encoding methods and devices are disclosed herein. In some examples, an audio encoding method may involve receiving, by a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, a currently-input audio signal. The control system may be configured to implement an audio encoder comprising a neural network that has been trained according to any of the methods disclosed herein. Such methods may involve encoding, via the audio encoder, the currently-input audio signal in a compressed audio format and outputting an encoded audio signal in the compressed audio format.

[0008] Some audio decoding methods and devices are disclosed herein. In some examples, an audio decoding method may involve receiving, by a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, a currently-input compressed audio signal. The control system may be configured to implement an audio decoder comprising a neural network that has been trained according to any of the methods disclosed herein. Such methods may involve

decoding, via the audio decoder, the currently-input compressed audio signal and outputting a decoded audio signal. Some such methods may involve reproducing the decoded audio signal via one or more transducers.

[0009] Some or all of the methods described herein may be performed by one or more
5 devices according to instructions (e.g., software) stored on one or more non-transitory media. Such non-transitory media may include memory devices such as those described herein, including but not limited to random access memory (RAM) devices, read-only memory (ROM) devices, etc. Accordingly, various innovative aspects of the subject matter described
10 in this disclosure can be implemented in a non-transitory medium having software stored thereon. The software may, for example, include instructions for controlling at least one device to process audio data. The software may, for example, be executable by one or more components of a control system such as those disclosed herein. The software may, for example, include instructions for performing one or more of the methods disclosed herein.

[0010] At least some aspects of the present disclosure may be implemented via
15 apparatus. For example, one or more devices may be configured for performing, at least in part, the methods disclosed herein. In some implementations, an apparatus may include an interface system and a control system. The interface system may include one or more network interfaces, one or more interfaces between the control system and a memory system, one or more interfaces between the control system and another device and/or one or more
20 external device interfaces. The control system may include at least one of a general purpose single- or multi-chip processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, or discrete hardware components. Accordingly, in some implementations the control system may include one or more processors and one or
25 more non-transitory storage media operatively coupled to the one or more processors.

[0011] According to some such examples, the apparatus may include an interface system and a control system. The control system may, for example, be configured for implementing one or more of the methods disclosed herein. For example, the control system may be configured to implement an audio encoder. The audio encoder may include a neural
30 network that has been trained according to one or more of the methods disclosed herein. The control system may be configured for receiving a currently-input audio signal, for encoding the currently-input audio signal in a compressed audio format and for outputting (e.g., via the interface system) an encoded audio signal in the compressed audio format.

[0012] Alternatively, or additionally, the control system may be configured to implement an audio decoder. The audio decoder may include a neural network that has been trained according to a process that involves receiving, by the neural network and via the interface system, an input training audio signal and generating, by the neural network and
5 based on the input training audio signal, an encoded training audio signal. The process may involve decoding, via the control system, the encoded training audio signal to produce a decoded training audio signal and receiving, by a loss function generating module that is implemented via the control system, the decoded training audio signal and a ground truth
10 audio signal. The process may involve generating, by the loss function generating module, a loss function value corresponding to the decoded training audio signal. Generating the loss function value may involve applying a psychoacoustic model. The process may involve training the neural network based on the loss function value.

[0013] The audio encoder may be further configured for receiving a currently-input audio signal, for encoding the currently-input audio signal in a compressed audio format and
15 for outputting an encoded audio signal in the compressed audio format.

[0014] In some implementations, a disclosed system may include an audio decoding apparatus. The audio decoding apparatus may include an interface system and a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors. The control system may be configured to
20 implement an audio decoder.

[0015] The audio decoder may include a neural network that has been trained according to a process that involves receiving, by the neural network and via the interface system, an input training audio signal and generating, by the neural network and based on the input training audio signal, an encoded training audio signal. The process may involve
25 decoding, via the control system, the encoded training audio signal to produce a decoded training audio signal and receiving, by a loss function generating module that is implemented via the control system, the decoded training audio signal and a ground truth audio signal. The process may involve generating, by the loss function generating module, a loss function value corresponding to the decoded training audio signal. Generating the loss function value may
30 involve applying a psychoacoustic model. The process may involve training the neural network based on the loss function value.

[0016] The audio decoder may be further configured for receiving a currently-input encoded audio signal in a compressed audio format, for decoding the currently-input encoded audio signal in a decompressed audio format and for outputting a decoded audio signal in the

decompressed audio format. According to some implementations, the system may include one or more transducers configured for reproducing the decoded audio signal.

[0017] Details of one or more implementations of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages will become apparent from the description, the drawings, and the claims. Note that the relative dimensions of the following figures may not be drawn to scale. Like reference numbers and designations in the various drawings generally indicate like elements.

10 BRIEF DESCRIPTION OF THE DRAWINGS

[0018] Figure 1 is a block diagram that shows examples of components of an apparatus that may be configured to perform at least some of the methods disclosed herein.

[0019] Figure 2 shows blocks for implementing a process of machine learning according to a perceptually-based loss function according to one example.

15 [0020] Figure 3 shows an example of a neural network training process according to some implementations disclosed herein.

[0021] Figures 4A-4D show alternative examples of neural networks that are suitable for implementing some of the methods disclosed herein.

20 [0022] Figure 5A is a flow diagram that outlines blocks of a method of training a neural network for audio encoding and decoding according to one example.

[0023] Figure 5B is a flow diagram that outlines blocks of a method of using a trained neural network for audio encoding according to one example.

[0024] Figure 5C is a flow diagram that outlines blocks of a method of using a trained neural network for audio decoding according to one example.

25 [0025] Figure 6 is a block diagram that illustrates a loss function generating module that is configured to generate a loss function based on mean squared error.

[0026] Figure 7A is a graph of a function that approximates a typical acoustic response of a human ear canal.

30 [0027] Figure 7B illustrates a loss function generating module that is configured to generate a loss function based on the typical acoustic response of a human ear canal.

[0028] Figure 8 illustrates a loss function generating module that is configured to generate a loss function based on banding operations.

[0029] Figure 9A illustrates processes involved in frequency masking according to some examples.

[0030] Figure 9B illustrates an example of a spreading function.

[0031] Figure 10 shows one example of an alternative implementation of a loss
5 function generating module.

[0032] Figure 11 shows an example of objective test results for some disclosed implementations.

[0033] Figure 12 shows examples of subjective test results for audio data
corresponding to male talkers that was produced by a neural network trained using various
10 types of loss functions.

[0034] Figure 13 shows examples of subjective test results for audio data
corresponding to female talkers that was produced by a neural network trained using the same
types of loss functions that are represented in Figure 12.

15

DESCRIPTION OF EXAMPLE EMBODIMENTS

[0035] The following description is directed to certain implementations for the purposes of describing some innovative aspects of this disclosure, as well as examples of contexts in which these innovative aspects may be implemented. However, the teachings herein can be applied in various different ways. Moreover, the described embodiments may
20 be implemented in a variety of hardware, software, firmware, etc. For example, aspects of the present application may be embodied, at least in part, in an apparatus, a system that includes more than one device, a method, a computer program product, etc. Accordingly, aspects of the present application may take the form of a hardware embodiment, a software embodiment (including firmware, resident software, microcodes, etc.) and/or an embodiment
25 combining both software and hardware aspects. Such embodiments may be referred to herein as a “circuit,” a “module” or “engine.” Some aspects of the present application may take the form of a computer program product embodied in one or more non-transitory media having computer readable program code embodied thereon. Such non-transitory media may,
30 (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. Accordingly, the teachings of this disclosure are not intended to be limited to the implementations shown in the figures

and/or described herein, but instead have wide applicability.

[0036] The present inventors have researched various methods of machine learning that is related to audio data processing, including but not limited to audio data encoding and decoding. In particular, the present inventors have researched various methods of training
5 different types of neural networks using loss functions that are relevant to the manner in which human beings perceive sound. The effectiveness of each of these loss functions was evaluated according to audio data produced by neural network encoding. The audio data was evaluated according to objective and subjective criteria. In some examples, audio data processed by a neural network that had been trained by using a loss function based on mean
10 squared error was used as a basis for evaluating audio data produced according to the methods disclosed herein. In some instances, the process of evaluation via subjective criteria involved having human listeners evaluate the resulting audio data and obtaining the listeners' feedback.

[0037] The techniques disclosed herein are based on the above-referenced research.
15 This disclosure provides various examples of using perceptually-based loss functions for training a neural network for audio data encoding and/or decoding. In some examples, the perceptually-based loss functions are based on a psychoacoustic model. The psychoacoustic model may, for example, be based at least in part on one or more psychoacoustic masking thresholds. In some examples, the psychoacoustic model may involve modeling an outer ear
20 transfer function, grouping audio data into critical bands, frequency-domain masking (including but not limited to level-dependent spreading), modeling of a frequency-dependent hearing threshold and/or calculation of a noise-to-mask ratio. In some implementations, the loss function may involve calculating an average noise-to-masking ratio. In some such examples, the training process may involve minimizing the average noise-to-masking ratio.

[0038] Figure 1 is a block diagram that shows examples of components of an
25 apparatus that may be configured to perform at least some of the methods disclosed herein. In some examples, the apparatus 105 may be, or may include, a personal computer, a desktop computer or other local device that is configured to provide audio processing. In some examples, the apparatus 105 may be, or may include, a server. According to some
30 examples, the apparatus 105 may be a client device that is configured for communication with a server, via a network interface. The components of the apparatus 105 may be implemented via hardware, via software stored on non-transitory media, via firmware and/or by combinations thereof. The types and numbers of components shown in Figure 1, as well as other figures disclosed herein, are merely shown by way of example. Alternative

implementations may include more, fewer and/or different components.

[0039] In this example, the apparatus 105 includes an interface system 110 and a control system 115. The interface system 110 may include one or more network interfaces, one or more interfaces between the control system 115 and a memory system and/or one or more external device interfaces (such as one or more universal serial bus (USB) interfaces).
5 In some implementations, the interface system 110 may include a user interface system. The user interface system may be configured for receiving input from a user. In some implementations, the user interface system may be configured for providing feedback to a user. For example, the user interface system may include one or more displays with
10 corresponding touch and/or gesture detection systems. In some examples, the user interface system may include one or more microphones and/or speakers. According to some examples, the user interface system may include apparatus for providing haptic feedback, such as a motor, a vibrator, etc. The control system 115 may, for example, include a general purpose single- or multi-chip processor, a digital signal processor (DSP), an application
15 specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, and/or discrete hardware components.

[0040] In some examples, the apparatus 105 may be implemented in a single device. However, in some implementations, the apparatus 105 may be implemented in more than
20 one device. In some such implementations, functionality of the control system 115 may be included in more than one device. In some examples, the apparatus 105 may be a component of another device.

[0041] Figure 2 shows blocks for implementing a process of machine learning according to a perceptually-based loss function according to one example. In this example,
25 an input audio signal 205 is provided to a machine learning module 210. The input audio signal 205 may, in some examples, correspond with human speech. However, in other examples the input audio signal 205 may correspond to other sounds, such as music, etc.

[0042] According to some examples, the elements of the system 200, including but not limited to the machine learning module 210, may be implemented via one or more
30 control systems such as the control system 115. The machine learning module 210 may, for example, receive the input audio signal 205 via an interface system such as the interface system 110. In some instances, the machine learning module 210 may be configured to implement one or more neural networks, such as the neural networks disclosed herein. However, in other implementations the machine learning module 210 may be configured to

implement one or more other types of machine learning, such as Non-negative Matrix Factorization, Robust Principal Component Analysis, Sparse Coding, Probabilistic Latent Component Analysis, etc.

[0043] In the examples shown in Figure 2, the machine learning module 210 provides
5 the output audio signal 215 to the loss function generating module 220. The loss function
generating module 225 and the optional ground truth module 220 may, for example, be
implemented via a control system such as the control system 115. In some examples the loss
function generating module 225, the machine learning module 210 and the optional ground
10 truth module 220 may be implemented via the same device, whereas in other examples the
loss function generating module 225, the optional ground truth module 220 and the machine
learning module 210 may be implemented via different devices.

[0044] According to this example, the loss function generating module 225 receives
the input audio signal 205 and uses the input audio signal 205 as the “ground truth” for error
determination. However, in some alternative implementations the loss function generating
15 module 225 may receive ground truth data from the optional ground truth module 220. Such
implementations may, for example, involve tasks such as speech enhancement or speech de-
noising, in which the ground truth is not the original input audio signal. Whether the ground
truth data is the input audio signal 205 or data that is received from the optional ground truth
module, the loss function generating module 225 evaluates the output audio signal according
20 to a loss function algorithm and the ground truth data, and provides a loss function value 230
to the machine learning module 210. In some such implementations, the machine learning
module 210 includes an implementation of the optimizer module 315, which is described
below with reference to Figure 3. In other examples, the system 200 includes an
implementation of the optimizer module 315 that is separate from, but in communication
25 with, the machine learning module 210 and the loss function generating module 225. Various
examples of loss functions are disclosed herein. In this example, the loss function generating
module 225 applies a perceptually-based loss function, which may be based on a
psychoacoustic model. According to this example, the process of machine learning that is
implemented via the machine learning module 210 (for example, the process of training a
30 neural network) is based, in part, on the loss function value 230.

[0045] The employment of the perceptually-based loss function, such as a loss
function that is based on a psychoacoustic model, for machine learning (e.g., for training a
neural network) can increase the perceptual quality of the output audio signal 215, as
compared to the perceptual quality of an output audio signal produced by machine learning

processes using traditional loss functions based on a mean squared error (MSE), an L1-norm, etc. For example, a neural network trained for a given length of time via a loss function based on a psychoacoustic model can increase the perceptual quality of the output audio signal 215, as compared to the perceptual quality of an output audio signal produced by a neural network having the same architecture that has been trained via a loss function based on MSE for the same length of time. Moreover, a neural network that is trained to convergence via a loss function based on a psychoacoustic model will generally produce an output audio signal a higher perceptual quality than the output audio signal of a neural network having the same architecture that has been trained to convergence via a loss function based on MSE.

[0046] Some disclosed loss functions take advantage of psychoacoustic principles to determine which differences in the output audio signal 215 will be audible to an average person and which will not be not be audible to an average person. In some examples, a loss function that is based on a psychoacoustic model may employ psychoacoustic phenomena such as time masking, frequency masking, equal loudness curves, level dependent masking, and/or human hearing thresholds. In some implementations, the perceptual loss function may operate in the time domain, whereas in other implementations the perceptual loss function may operate in the frequency domain. In alternative implementations the perceptual loss function may involve both time-domain and frequency-domain operations. In some examples, the loss function may use one-frame inputs to calculate the loss function, whereas in other examples the loss function may use multiple input frames to calculate the loss function.

[0047] Figure 3 shows an example of a neural network training process according to some implementations disclosed herein. As with other figures provided herein, the numbers and types of elements are merely made by way of example. According to some examples, the elements of the system 301 may be implemented via one or more control systems such as the control system 115. In the example shown in Figure 3, the neural network 300 is an autoencoder. Techniques for designing an autoencoder are described in chapter 14 of Goodfellow, Ian, Yoshua Bengio, and Aaron Courville, *Deep Learning* (MIT Press, 2016), which is hereby incorporated by reference.

[0048] The neural network 300 includes layers of nodes, which are also referred to herein as “neurons.” Each neuron has a real-valued activation function, whose output is commonly referred to as an “activation,” which defines the output of that neuron given an input or set of inputs. According to some examples, neurons of the neural network 300 may employ sigmoidal activation functions, ELU activation functions and/or tanh activation

functions. Alternatively, or additionally, neurons of the voice neural network 300 may employ rectified linear unit (ReLU) activation functions.

[0049] Each connection (also referred to as a “synapse”) between neurons has a modifiable real-valued weight. A neuron may be an input neuron (receiving data from outside the network), an output neuron, or a hidden neuron that modifies the data en route from an input neuron to an output neuron. In the example shown in Figure 3, the neurons in Neuron Layer 1 are input neurons, the neurons in Neuron Layer 7 are output neurons and the neurons in Neuron Layers 2–6 are hidden neurons. Although five hidden layers are shown in Figure 3, some implementations may include more or fewer hidden layers. Some implementations of the neural network 300 may include more or fewer hidden layers, e.g., 10 or more hidden layers. For example, some implementations may include 10, 20, 30, 40, 50, 60, 70, 80, 90 or more hidden layers.

[0050] Here, a first portion of the neural network 300 (the encoder portion 305) is configured to generate an encoded audio signal and a second portion of the neural network 300 (the decoder portion 310) is configured to decode the encoded audio signal. In this example, the encoded audio signal is a compressed audio signal and the decoded audio signal is a decompressed audio signal. Accordingly, the input audio signal 205 is compressed by the encoder portion 305, as suggested by the decreasing sizes of the blocks that are used to illustrate neuron layers 1–4. In some examples, the input neuron layer may include more neurons than at least one of the hidden neuron layers of the encoder portion 305. However, in alternative implementations the neuron layers 1–4 may all have the same number of neurons, or substantially similar numbers of neurons.

[0051] Accordingly, the compressed audio signal provided by the encoder portion 305 is then decoded via the neuron layers of the decoder portion 310 to construct the output signal 215, which is an estimate of the input audio signal 205. A perceptual loss function, such as a psychoacoustically-based loss function, may then be used to determine updates for parameters of the neural network 300 during the training phase. These parameters can subsequently be used to decode (e.g., to decompress) any audio signals that have been encoded (e.g., compressed) using the weights determined by the parameters received from the training algorithm. In other words, encoding and decoding may be done separately from the training process after satisfactory weights are determined for the neural network 300.

[0052] According to this example, the loss function generating module 225 receives the at least a portion of the audio input signal 205 and uses this as ground truth data. Here, the loss function generating module 225 evaluates the output audio signal according to a loss

function algorithm and the ground truth data, and provides a loss function value 230 to the optimizer module 315. In this example, the optimizer module 315 is initialized with information about the neural network and the loss function(s) used by the loss function generating module 225. According to this example, the optimizer module 315 uses that
5 information, along with the loss value that the optimizer module 315 receives from the loss function generating module 225, to compute the gradient of the loss function with respect to the neural network's weights. Once this gradient is known, the optimizer module 315 uses an optimization algorithm to generate updates 320 to the neural network's weights. According to some implementations, the optimizer module 315 may employ an optimization algorithm
10 such as Stochastic Gradient Descent or the Adam optimization algorithm. The Adam optimization algorithm is disclosed in D. P. Kingma and J. L. Ba, "Adam: a Method for Stochastic Optimization," in Proceedings of the International Conference on Learning Representations (ICLR), 2015, pp. 1–15, which is hereby incorporated by reference. In the example shown in Figure 3, the optimizer module 315 is configured to provide the updates
15 320 to the neural network 300. In this example, the loss function generating module 225 applies a perceptually-based loss function, which may be based on a psychoacoustic model. According to this example, the process of training the neural network 300 is based, at least in part, on backward propagation. This backward propagation is indicated in Figure 3 by the stippled arrows between the neuron layers. Backward propagation (also known as
20 "backpropagation") is a method used in neural networks to calculate the error contribution of each neuron after a batch of data is processed. The backpropagation technique is also sometimes called backward propagation of errors, because the error may be calculated at the output and distributed back through the neural network layers.

[0053] The neural network 300 may be implemented by a control system such as the
25 control system 115 that is described above with reference to Figure 1. Accordingly, training the neural network 300 may involve changing the physical state of non-transitory storage medium locations corresponding to the weights in the neural network 300. The storage medium locations may be parts of one or more storage media that are accessible by, or part of, the control system. The weights, as noted above, correspond to connections between
30 neurons. Training the neural network 300 also may involve changing the physical state of non-transitory storage medium locations corresponding to the values of the neurons' activation functions.

[0054] Figures 4A–4C show alternative examples of neural networks that are suitable for implementing some of the methods disclosed herein. According to these examples, the

input neurons and the hidden neurons employ rectified linear unit (ReLU) activation functions, whereas the output neurons employ sigmoidal activation functions. However, alternative implementations of the neural network 300 may include other activation functions and/or other combinations of activation functions, including but not limited to Exponential
5 Linear Unit (ELU) and/or tanh activation functions.

[0055] According to these examples, the input audio data is 256-dimensional audio data. In the example shown in Figure 4A, the encoder portion 305 compresses the input audio data to 32-dimensional audio data, providing up to an 8x reduction. According to the example shown in Figure 4B, the encoder portion 305 compresses the input audio data to 16-
10 dimensional audio data, providing up to a 16x reduction. The neural network 300 that is illustrated in Figure 4C includes an encoder portion 305 that compresses the input audio data to 8-dimensional audio data, providing up to a 32x reduction. The present inventors conducted listening tests based on the type of neural network that is illustrated in Figure 4B, some results of which are described below.

[0056] Figure 4D shows an example of blocks of an encoder portion of an autoencoder according to an alternative example. The encoder portion 305 may, for example, be implemented by a control system such as the control system 115 that is described above with reference to Figure 1. The encoder portion 305 may, for example, be implemented by one or more processors of the control system according to software stored in one or more
20 non-transitory storage media. The numbers and types of elements shown in Figure 4D are merely examples. Other implementations of the encoder portion 305 may include more, fewer or different elements.

[0057] In this example, the encoder portion 305 includes three layers of neurons. According to some examples, neurons of the encoder portion 305 may employ ReLU
25 activation functions. However, according to some alternative examples, neurons of the encoder portion 305 may employ sigmoidal activation functions and/or tanh activation functions. The neurons in Neuron Layers 1–3 process the N-dimensional input data while maintaining its N-dimensional state. The layer 450 is configured to receive the output of Neuron Layer 3 and apply a pooling algorithm. Pooling is a form of non-linear down-
30 sampling. According to this example, the layer 450 is configured to apply a max pooling function that partitions the output of Neuron Layer 3 into a set of M non-overlapping portions or “sub-regions” and, for each such sub-region, outputs the maximum value.

[0058] Figure 5A is a flow diagram that outlines blocks of a method of training a neural network for audio encoding and decoding according to one example. The method 500

may, in some instances, be performed by the apparatus of Figure 1 or by another type of apparatus. In some examples, the blocks of method 500 may be implemented via software stored on one or more non-transitory media. The blocks of method 500, like other methods described herein, are not necessarily performed in the order indicated. Moreover, such
5 methods may include more or fewer blocks than shown and/or described.

[0059] Here, block 505 involves receiving, by a neural network implemented via a control system comprising one or more processors and one or more non-transitory storage media, an input audio signal. In some examples, the neural network may be, or may include, an autoencoder. According to some examples, block 505 may involve the control system
10 of Figure 1 receiving the input audio signal via the interface system 110. In some examples, block 505 may involve the neural network 300 receiving the input audio signal 205, as described above with reference to Figures 2–4C. In some implementations, the input audio signal 205 may include at least a portion of a speech dataset, such as the publicly available speech dataset known as TIMIT. TIMIT is a dataset of phonemically and lexically
15 transcribed speech of American English speakers of different sexes and dialects. TIMIT was commissioned by the Defense Advanced Research Projects Agency (DARPA). The corpus design of TIMIT was a joint effort between Texas Instruments (TI), the Massachusetts Institute of Technology (MIT) and SRI International. According to some examples, method 500 may involve transforming the input audio signal 205 from the time domain to the
20 frequency domain, e.g., via a fast Fourier transform (FFT), a discrete cosine transform (DCT) or a short-time Fourier transform (STFT). In some implementations, min/max scaling may be applied to the input audio signal 205 prior to block 510.

[0060] According to this example, block 510 involves generating, by the neural network and based on the input audio signal, an encoded audio signal. The encoded audio
25 signal may be, or may include, a compressed audio signal. Block 510 may, for example, be performed by an encoder portion of the neural network, such as the encoder portion 305 of the neural network 300 described herein. However, in other examples block 510 may involve generating an encoded audio signal via an encoder that is not part of the neural network. In some such examples, a control system that implements the neural network also
30 may include an encoder that is not part of the neural network. For example, the neural network may include a decoding portion but not an encoding portion.

[0061] In this example, block 515 involves decoding, via the control system, the encoded audio signal to produce a decoded audio signal. The decoded audio signal may be, or may include, a decompressed audio signal. In some implementations, block 515 may

involve producing decoded transform coefficients. Block 515 may, for example, be performed by a decoder portion of the neural network, such as the decoder portion 310 of the neural network 300 described herein. However, in other examples block 510 may involve generating a decoded audio signal and/or decoded transform coefficients via a decoder that is not part of the neural network. In some such examples, a control system that implements the neural network also may include a decoder that is not part of the neural network. For example, the neural network may include an encoding portion but not a decoding portion.

[0062] Accordingly, in some implementations a first portion of the neural network may be configured to generate the encoded audio signal and a second portion of the neural network may be configured to decode the encoded audio signal. In some such implementations, the first portion of the neural network may include an input neuron layer and a plurality of hidden neuron layers. In some examples, the input neuron layer may include more neurons than at least one of the hidden neuron layers of the first portion. However, in alternative implementations the input neuron layer may have the same number of neurons, or substantially similar numbers of neurons, as the hidden neuron layers of the first portion.

[0063] According to some examples, at least some neurons of the first portion of the neural network may be configured with rectified linear unit (ReLU) activation functions. In some implementations, at least some neurons in hidden layers of the second portion of the neural network may be configured with rectified linear unit (ReLU) activation functions. According to some such implementations, at least some neurons in an output layer of the second portion may be configured with sigmoidal activation functions.

[0064] In some implementations, block 520 may involve receiving, by a loss function generating module that is implemented via the control system, the decoded audio signal and/or the decoded transform coefficients, and a ground truth signal. The ground truth signal may, for example, include a ground truth audio signal and/or ground truth transform coefficients. In some such examples, the ground truth signal may be received from a ground truth module, such as the ground truth module 220 that is shown in Figure 2 and described above. However, in some implementations the ground truth signal may be (or may include) the input audio signal, or a portion of the input audio signal. The loss function generating module may, for example, be an instance of the loss function generating module 225 that is disclosed herein.

[0065] According to some implementations, block 525 may involve generating, by the loss function generating module, a loss function value corresponding to the decoded audio signal

and/or the decoded transform coefficients. In some such implementations, generating the loss function value may involve applying a psychoacoustic model. In the example shown in Figure 5A, block 530 involves training the neural network based on the loss function value. The training may involve updating at least one weight in the neural network. In some such examples, an optimizer such as the optimizer module 315 that is described above with reference to Figure 3 may have been initialized with information about the neural network and the loss function(s) used by the loss function generating module 225. The optimizer module 315 may be configured to use that information, along with the loss function value that the optimizer module 315 receives from the loss function generating module 225, to compute the gradient of the loss function with respect to the neural network's weights. After computing the gradient, the optimizer module 315 may use an optimization algorithm to generate updates to the neural network's weights and to provide these updates to the neural network. Training the neural network may involve backward propagation based on the updates provided by the optimizer module 315. Techniques for detecting and addressing overfitting during the process of training a neural network are described in chapters 5 and 7 of Goodfellow, Ian, Yoshua Bengio, and Aaron Courville, *Deep Learning* (MIT Press, 2016), which are hereby incorporated by reference. Training the neural network may involve changing a physical state of at least one non-transitory storage medium location corresponding with at least one weight or at least one activation function value of the neural network.

[0066] The psychoacoustic model may vary according to the particular implementation. According to some examples, the psychoacoustic model may be based, at least in part, on one or more psychoacoustic masking thresholds. In some implementations, applying the psychoacoustic model may involve modeling an outer ear transfer function, grouping into critical bands, frequency-domain masking (including but not limited to level-dependent spreading), modeling of a frequency-dependent hearing threshold and/or calculation of a noise-to-mask ratio. Some examples are described below with reference to Figures 6–10B.

[0067] In some implementations, the loss function generating module's determination of a loss function may involve calculating a noise-to-masking ratio (NMR), such as an average NMR. The training process may involve minimizing the average NMR. Some examples are described below.

[0068] According to some examples, training the neural network may continue until the loss function is relatively “flat,” such that the difference between a current loss function

value and a prior loss function value (such as the previous loss function value) is at or below a threshold value. In the example shown in Figure 5, training the neural network may involve repeating at least some of blocks 505 through 535 until a difference between a current loss function value and a prior loss function value is less than or equal to a
5 predetermined value.

[0069] After the neural network has been trained, the neural network (or a portion thereof) may be used for processing audio data, e.g., for encoding or for decoding audio data. Figure 5B is a flow diagram that outlines blocks of a method of using a trained neural network for audio encoding according to one example. The method 540 may, in some
10 instances, be performed by the apparatus of Figure 1 or by another type of apparatus. In some examples, the blocks of method 540 may be implemented via software stored on one or more non-transitory media. The blocks of method 540, like other methods described herein, are not necessarily performed in the order indicated. Moreover, such methods may include more or fewer blocks than shown and/or described.

[0070] In this example, block 545 involves receiving a currently-input audio signal. In this example block 545 involves receiving the currently-input audio signal by a control system that includes one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors. Here, the control system is configured to implement an audio encoder comprising a neural network that has been trained according to
15 one or more of the methods disclosed herein.

[0071] In some examples, the training process may involve: receiving, by the neural network and via the interface system, an input training audio signal; generating, by the neural network and based on the input training audio signal, an encoded training audio signal; decoding, via the control system, the encoded training audio signal to produce a
25 decoded training audio signal; receiving, by a loss function generating module that is implemented via the control system, the decoded training audio signal and a ground truth audio signal; generating, by the loss function generating module, a loss function value corresponding to the decoded training audio signal, wherein generating the loss function value comprises applying a psychoacoustic model; and training the neural network based on
30 the loss function value.

[0072] According to this implementation, block 550 involves encoding, via the audio encoder, the currently-input audio signal in a compressed audio format. Here, block 555 involves outputting an encoded audio signal in the compressed audio format.

[0073] Figure 5C is a flow diagram that outlines blocks of a method of using a

trained neural network for audio decoding according to one example. The method 560 may, in some instances, be performed by the apparatus of Figure 1 or by another type of apparatus. In some examples, the blocks of method 560 may be implemented via software stored on one or more non-transitory media. The blocks of method 560, like other methods
5 described herein, are not necessarily performed in the order indicated. Moreover, such methods may include more or fewer blocks than shown and/or described.

[0074] In this example, block 565 involves receiving a currently-input compressed audio signal. In some such examples, the currently-input compressed audio signal may have been produced according to method 540 or by a similar method. In this example block 565
10 involves receiving the currently-input compressed audio signal by a control system that includes one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors. Here, the control system is configured to implement an audio decoder comprising a neural network that has been trained according to one or more of the methods disclosed herein.

[0075] According to this implementation, block 570 involves decoding, via the audio decoder, the currently-input compressed audio signal. For example, block 570 may involve decompressing the currently-input compressed audio signal. Here, block 575 involves outputting a decoded audio signal. According to some examples, method 540 may involve reproducing the decoded audio signal via one or more transducers.
15

[0076] As noted above, the present inventors have researched various methods of training different types of neural networks using loss functions that are relevant to the manner in which human beings perceive sound. The effectiveness of each of these loss functions was evaluated according to audio data produced by neural network encoding. In some examples, audio data processed by a neural network that had been trained by using a loss function based
20 on mean squared error (MSE) was used as a basis for evaluating audio data produced according to the methods disclosed herein.

[0077] Figure 6 is a block diagram that illustrates a loss function generating module that is configured to generate a loss function based on mean squared error. Here, the estimated magnitude of an audio signal produced by a neural network and magnitude of a
30 ground truth/true audio signal are both provided to a loss function generating module 225. The loss function generating module 225 generates a loss function value 230 based on the MSE value. The loss function value 230 may be provided to an optimizer module that is configured to generate updates to weights of the neural network for training.

[0078] The present inventors have evaluated some implementations of loss functions that are based, at least in part, on models of the acoustic response of one or more portions of the human ear, which also may be referred to as “ear models.” Figure 7A is a graph of a function that approximates a typical acoustic response of a human ear canal.

5 **[0079]** Figure 7B illustrates a loss function generating module that is configured to generate a loss function based on the typical acoustic response of a human ear canal. In this example, a function W is applied to both an audio signal produced by a neural network and a ground truth/true audio signal.

[0080] In some examples, the function W may be as follows:

$$10 \quad W(f) = -2.184 \left(\frac{f}{1000} \right)^{-0.8} + 6.5 \exp \left(-0.6 \left(\frac{f}{1000} - 3.3 \right)^2 \right) - 0.001 \left(\frac{f}{1000} \right)^{3.6}$$

(Equation 1)

[0081] Equation 1 has been used in an implementation of the Perceptual Evaluation of Audio Quality (PEAQ) algorithm for the purpose of modeling the acoustic response of a human ear canal. In Equation 1, f represents a frequency of the audio signal. In this example, 15 the loss function generating module 225 generates a loss function value 230 based on the difference between the two resulting values. The loss function value 230 may be provided to an optimizer module that is configured to generate updates to weights of the neural network for training.

[0082] When compared to audio signals produced by a neural network trained 20 according to a loss function based on MSE, the audio signals produced by training a neural network using a loss function such as that shown in Figure 7B provided only slight improvements. For example, using an objective standard based on the Perceptual Objective Listening Quality Analysis (POLQA), the MSE-based audio data achieved a score of 3.41, whereas the audio data produced by training a neural network using a loss function such as 25 that shown in Figure 7B achieved a score of 3.48.

[0083] In some experiments, the inventors tested audio signals produced by a neural network trained according to a loss function based on banding operations. Figure 8 illustrates a loss function generating module that is configured to generate a loss function based on banding operations. In this example, the loss function generating module 225 is configured 30 to perform banding operations on an audio signal produced by a neural network and a ground truth/true audio signal, and to calculate a difference between the results.

[0084] In some implementations, the banding operations were based on “Zwicker” bands, which are critical bands defined according to chapter 6 (Critical Bands and Excitation) of

Fastl, H., & Zwicker, E. (2007), *Psychoacoustics: Facts and Models* (3rd ed., Springer), which is hereby incorporated by reference. In alternative implementations, the banding operations were based on “Moore” bands, which are critical bands defined according to chapter 3 (Frequency Selectivity, Masking, and the Critical Band) of Moore, B. C. J. (2012),
 5 *An Introduction to the Psychology of Hearing* (Emerald Group Publishing), which is hereby incorporated by reference. However, other examples may involve other types of banding operations known by those of skill in the art.

[0085] Based on their experiments, the inventors concluded that banding operations alone were unlikely to provide satisfactory results. For example, using an objective standard
 10 based on POLQA, the MSE-based audio data achieved a score of 3.41, whereas in one example the audio data produced by training a neural network using banding operations achieved a score of only 1.62.

[0086] In some experiments, the inventors tested audio signals produced by a neural network trained according to loss functions based, at least in part, on frequency masking.
 15 Figure 9A illustrates processes involved in frequency masking according to some examples. In this example, a spreading function is calculated in the frequency domain. This spreading function may, for example, be a level- and frequency-dependent function that can be estimated from the input audio data, e.g., from each input audio frame. A convolution with the frequency spectrum of the input audio may then be performed, which yields an excitation
 20 pattern. The result of the convolution between the input audio data and the spreading function is an approximation of how human auditory filters react to the excitations of the incoming sounds. Accordingly, the process is a simulation of the human hearing mechanism. In some implementations, the audio data is grouped into frequency bins and the convolution process involves convolving a spreading function for each frequency bin with
 25 the corresponding audio data for that frequency bin.

[0087] The excitation pattern may be adjusted to produce a masking pattern. In some examples, the excitation pattern may be adjusted downward, e.g., by 20 dB, to produce a masking pattern.

[0088] Figure 9B illustrates an example of a spreading function. According to this
 30 example, the spreading function is a simplified asymmetric triangular function that can be pre-computed for efficient implementation. In this simplified example, the vertical axis represents decibels and the horizontal axis represents Bark sub-bands. According to one such example, the spreading function is calculated as follows:

$$S_l = 27 \quad (\text{Equation 2})$$

$$S_u = -24 - \frac{230}{f_c} + 0.2 * L \quad (\text{Equation 3})$$

[0089] In Equations 2 and 3, S_l represents the slope of the portion of the spreading function of Figure 9B that is to the left of the peak frequency, whereas S_u represents the slope of the portion of the spreading function that is to the right of the peak frequency. The slope units are dB/Bark. In Equation 3, f_c represents the center or peak frequency of the spreading function and L represents the level or amplitude of the audio data. In some examples, L may be assumed to be constant, in order to simplify the calculation of the spreading function. According to some such examples, L may be 70 dB.

[0090] In some such implementations, the excitation patterns may be calculated as follows:

$$E = (SF^{0.4} \cdot BP^{0.4})^{\frac{1}{0.4}} \quad (\text{Equation 4})$$

[0091] In Equation 4, E represents the excitation function (also referred to herein as the excitation pattern), SF represents the spreading function and BP represents the banded pattern of frequency-binned audio data. In some implementations, the excitation pattern may be adjusted to produce a masking pattern. In some examples, the excitation pattern may be adjusted downward, e.g., by 20 dB, by 24 dB, by 27 dB, etc., to produce a masking pattern.

[0092] Figure 10 shows one example of an alternative implementation of a loss function generating module. The elements of the loss function generating module 225 may, for example, be implemented by a control system such as the control system 115 that is described above with reference to Figure 1.

[0093] In this example, a reference audio signal x_{ref} , which is an instance of a ground truth signal referenced elsewhere herein, is provided to a fast Fourier transform (FFT) block 1005a of the loss function generating module 225. A test audio signal x , which has been produced by a neural network such as one of those disclosed herein, is provided to FFT block 1005b of the loss function generating module 225.

[0094] According to this example, the output of the FFT block 1005a is provided to the ear model block 1010a and the output of the FFT block 1005b is provided to the ear model block 1010b. The ear model blocks 1010a and 1010b may, for example, be configured to apply a function that is based on the typical acoustic response of one or more portions of a human ear. In one such example, the ear model blocks 1010a and 1010b may be configured to apply the function that is shown above in Equation 1.

[0095] According to this implementation, the output of the ear model blocks 1010a

and 1010b is provided to the difference calculating block 1015, which is configured to compute the difference between the output of the ear model block 1010a and the output of the ear model block 1010b. The output of the difference calculating block 1015 may be regarded as an approximation of the noise that is in the test signal x .

5 **[0096]** In this example, the output of the ear model block 1010a is provided to the banding block 1020a and the output of the difference calculating block 1015 is provided to the banding block 1020b. The banding blocks 1020a and 1020b are configured to apply the same type of banding process, which may be one of the banding processes disclosed above (e.g., Zwicker or Moore banding processes). However, in alternative implementations the
10 banding blocks 1020a and 1020b may be configured to apply any suitable banding process known by those of skill in the art.

[0097] The output of the banding block 1020a is provided to the frequency masking block 1025, which is configured to apply frequency masking operations. The masking block 1025 may, for example, be configured to apply one or more of the frequency masking
15 operations disclosed herein. As noted above with reference to Figure 9B, using a simplified frequency masking process can provide potential advantages. However, in alternative implementations the masking block 1025 may be configured to apply one or more other frequency masking operations known by those of skill in the art.

[0098] According to this example, the output of the masking block 1025 and the
20 output of the banding block 1020b are both provided to the noise-to-mask ratio (NMR) calculating block 1030. As noted above, the output of the difference calculating block 1015 may be regarded as an approximation of the noise that is in the test signal x . Therefore, the output of the banding block 1020b may be regarded as a frequency-banded version of the noise that is in the test signal x . According to one example, the NMR calculating block 1030
25 may calculate the NMR as follows:

$$\text{NMR} = 10 \times \log_{10}\left(\frac{BP_{noise}}{MP}\right) \text{ (Equation 5)}$$

[0099] In Equation 5, BP_{noise} represents the output of the banding block 1020b and MP represents the output of the masking block 1025. According to some examples, the NMR calculated by the NMR calculating block 1030 may be an average NMR across all frequency
30 bands output by the banding blocks 1020a and 1020b. The NMR calculated by the NMR calculating block 1030 may be used as a loss function value 230 for training a neural network, e.g. as described above. For example, the loss function value 230 may be provided to an optimizer module that is configured to generate updated weights for the neural network.

[00100] Figure 11 shows an example of objective test results for some disclosed implementations. Figure 11 shows a comparison between PESQ scores of audio data produced by a neural network trained using loss functions based on MSE, power-law, NMR-Zwicker (NMR based on a banding process like that of the Zwicker banding process, but with bands that were fractionally narrower than those defined by Zwicker) and NMR-Moore (NMR based on a Moore banding process). These results, which are based on the output of the neural network described above with reference to Figure 4B, show that the NMR-Zwicker and NMR-Moore results are both somewhat better than the MSE and power-law results.

10 **[00101]** Figure 12 shows examples of subjective test results for audio data corresponding to male talkers that was produced by a neural network trained using various types of loss functions. In this example, the subjective test results are Multiple Stimulus test with Hidden Reference and Anchor (MUSHRA) ratings. MUSHRA, which is described in ITU-R BS.1534, is a well-known methodology for conducting a codec listening test to evaluate the perceived quality of the output from lossy audio compression algorithms. The MUSHRA method has the advantage of displaying many stimuli simultaneously so that subjects can carry out any comparison between them directly. The time taken to perform the test using the MUSHRA method can be significantly reduced relative to other methods. This is true in part because the results from all codecs are presented at the same time, on the same samples, so that a paired t-test or a repeated measures analysis of variance can be used for statistical analysis. The numbers along the x axis in Figure 12 are the identification numbers of different audio files.

[00102] More specifically, Figure 12 shows a comparison between MUSHRA ratings of audio data produced by the same neural network trained using a loss function based on MSE, using a loss function based on power-law, using a loss function based on NMR-Zwicker and using a loss function based on NMR-Moore, audio data produced by applying a 3.5 kHz low-pass filter (one of the standard “anchors” of the MUSHRA methodology), and reference audio data. In this example, the MUSHRA ratings were obtained from 11 different listeners. As shown in Figure 12, the mean MUSHRA ratings for audio data produced by the neural network trained using a loss function based on NMR-Moore were significantly higher than any of the others. The difference, which was roughly 30 MUSHRA points, is a large effect that is rarely seen. The second-highest mean MUSHRA ratings were for audio data produced by the neural network trained using a loss function based on NMR-Zwicker.

[00103] Figure 13 shows examples of subjective test results for audio data

corresponding to female talkers that was produced by a neural network trained using the same types of loss functions that are represented in Figure 12. As in Figure 12, the numbers along the x axis in Figure 13 are the identification numbers of different audio files. In this example, the highest mean MUSHRA ratings were once again assigned to audio data
5 produced by the neural network after being trained using loss functions based on NMR. Although the perceived differences between the NMR-Moore and NMR-Zwicker audio data and the other audio data were not as pronounced in this example as the perceived differences shown in Figure 12, the results shown in Figure 13 nonetheless indicate significant improvements.

10 **[00104]** The general principles defined herein may be applied to other implementations without departing from the scope of this disclosure. Thus, the claims are not intended to be limited to the implementations shown herein, but are to be accorded the widest scope consistent with this disclosure, the principles and the novel features disclosed herein.

15 Various aspects of the present invention may be appreciated from the following enumerated example embodiments (EEEs):

1. A computer-implemented audio processing method, comprising:
 - receiving, by a neural network implemented via a control system comprising one or more processors and one or more non-transitory storage media, an input audio signal;
 - 20 generating, by the neural network and based on the input audio signal, an encoded audio signal;
 - decoding, via the control system, the encoded audio signal to produce a decoded audio signal;
 - receiving, by a loss function generating module that is implemented via the control
 - 25 system, the decoded audio signal and a ground truth audio signal;
 - generating, by the loss function generating module, a loss function value corresponding to the decoded audio signal, wherein generating the loss function value comprises applying a psychoacoustic model; and
 - training the neural network based on the loss function value, wherein the training
 - 30 involves updating at least one weight of the neural network.

2. The method of EEE 1, wherein training the neural network comprises backward propagation based on the loss function value.

3. The method of EEE 1 or EEE 2, wherein the neural network comprises an autoencoder.

4. The method of any one of EEES 1–3, wherein training the neural network comprises changing a physical state of at least one non-transitory storage medium location corresponding with at least one weight of the neural network.
5. The method of any one of EEES 1–4, wherein a first portion of the neural network generates the encoded audio signal and a second portion of the neural network decodes the encoded audio signal.
6. The method of EEES 5, wherein the first portion of the neural network comprises an input neuron layer and a plurality of hidden neuron layers, wherein the input neuron layer includes more neurons than a final hidden neuron layer.
- 10 7. The method of EEES 5, wherein at least some neurons of the first portion of the neural network are configured with rectified linear unit (ReLU) activation functions.
8. The method of EEES 5, wherein at least some neurons in hidden layers of the second portion of the neural network are configured with rectified linear unit (ReLU) activation functions and wherein at least some neurons in an output layer of the second portion are configured with sigmoidal activation functions.
- 15 9. The method of any one of EEES 1–8, wherein the psychoacoustic model is based, at least in part, on one or more psychoacoustic masking thresholds.
10. The method of any one of EEES 1–9, wherein the psychoacoustic model involves on one or more of the following: modeling an outer ear transfer function; grouping into critical bands; frequency-domain masking, including but not limited to level-dependent spreading; modeling of a frequency-dependent hearing threshold; or calculation of a noise-to-mask ratio.
- 20 11. The method of any one of EEES 1–10, wherein the loss function involves calculating an average noise-to-masking ratio and wherein the training involves minimizing the average noise-to-masking ratio.
- 25 12. An audio encoding method, comprising:
 - receiving, by a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, a currently-input audio signal, the control system being configured to implement an audio encoder comprising a neural network that has been trained according to any one of the methods recited in EEES 1–11;
 - 30 encoding, via the audio encoder, the currently-input audio signal in a compressed audio format; and
 - outputting an encoded audio signal in the compressed audio format.
13. An audio decoding method, comprising:

- receiving, by a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, a currently-input compressed audio signal, the control system being configured to implement an audio decoder comprising a neural network that has been trained according to any one of the
- 5 methods recited in EEEs 1–11;
- decoding, via the audio decoder, the currently-input compressed audio signal; and outputting a decoded audio signal.
14. The method of EEE 13, further comprising reproducing the decoded audio signal via one or more transducers.
- 10 15. An apparatus, comprising:
- an interface system; and
- a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, the control system configured to implement the method of any one of EEEs 1–14.
- 15 16. One or more non-transitory media having software stored thereon, the software including instructions for controlling one or more devices to perform the method of any one of EEEs 1–14.
17. An audio encoding apparatus, comprising:
- an interface system; and
- 20 a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, the control system configured to implement an audio encoder, the audio encoder comprising a neural network that has been trained according to any one of the methods recited in EEEs 1–11, wherein the control system is configured for:
- 25 receiving a currently-input audio signal;
- encoding the currently-input audio signal in a compressed audio format; and
- outputting an encoded audio signal in the compressed audio format.
18. An audio encoding apparatus, comprising:
- an interface system; and
- 30 a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, the control system configured to implement an audio encoder, the audio encoder comprising a neural network that has been trained according to a process comprising:

receiving, by the neural network and via the interface system, an input training audio signal;

generating, by the neural network and based on the input training audio signal, an encoded training audio signal;

5 decoding, via the control system, the encoded training audio signal to produce a decoded training audio signal;

receiving, by a loss function generating module that is implemented via the control system, the decoded training audio signal and a ground truth audio signal;

10 generating, by the loss function generating module, a loss function value corresponding to the decoded training audio signal, wherein generating the loss function value comprises applying a psychoacoustic model; and

training the neural network based on the loss function value;

wherein the audio encoder is further configured for:

receiving a currently-input audio signal;

15 encoding the currently-input audio signal in a compressed audio format; and outputting an encoded audio signal in the compressed audio format.

19. A system that includes an audio decoding apparatus, comprising:

an interface system;

20 a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, the control system configured to implement an audio decoder, the audio decoder comprising a neural network that has been trained according to a process comprising:

receiving, by the neural network and via the interface system, an input training audio signal;

25 generating, by the neural network and based on the input training audio signal, an encoded training audio signal;

decoding, via the control system, the encoded training audio signal to produce a decoded training audio signal;

30 receiving, by a loss function generating module that is implemented via the control system, the decoded training audio signal and a ground truth audio signal;

generating, by the loss function generating module, a loss function value corresponding to the decoded training audio signal, wherein generating the loss function value comprises applying a psychoacoustic model; and

training the neural network based on the loss function value;

wherein the audio decoder is further configured for:

receiving a currently-input encoded audio signal in a compressed audio format;

5 decoding the currently-input encoded audio signal in a decompressed audio format; and

outputting a decoded audio signal in the decompressed audio format.

20. The system of EEE 19, wherein the system further comprises one or more transducers configured for reproducing the decoded audio signal.

CLAIMS

1. A computer-implemented method for training an autoencoder neural network implemented via a control system comprising one or more processors and one or more non-transitory storage media, comprising:
 - 5 receiving, by the autoencoder neural network, an input audio signal comprising audio data;
 - generating, by an encoder portion of the autoencoder neural network and based on the input audio signal, an encoded audio signal;
 - decoding, by a decoder portion of the autoencoder neural network, the encoded audio
10 signal to produce at least one of a decoded audio signal and decoded transform coefficients;
 - receiving, by a loss function generating module that is implemented via the control system, said at least one of the decoded audio signal and the decoded transform coefficients and a ground truth audio signal;
 - generating, by the loss function generating module, a loss function value
15 corresponding to said at least one of the decoded audio signal and the decoded transform coefficients; and
 - training the autoencoder neural network based on the loss function value, wherein the training involves updating at least one weight of the autoencoder neural network,
wherein generating the loss function value comprises applying a psychoacoustic
20 model, which includes calculation of a noise-to-mask ratio.
2. The method of claim 1, wherein training the autoencoder neural network comprises backward propagation based on the loss function value.
3. The method of claim 1 or claim 2, wherein training the autoencoder neural network
25 comprises changing a physical state of at least one non-transitory storage medium location corresponding with at least one weight of the autoencoder neural network.
4. The method of any of the preceding claims, wherein the encoder portion of the autoencoder neural network comprises an input neuron layer and a plurality of hidden neuron layers, wherein the input neuron layer includes more neurons than a final hidden neuron
30 layer.

5. The method of any of the preceding claims, wherein at least some neurons of the encoder portion of the autoencoder neural network are configured with rectified linear unit (ReLU) activation functions.
6. The method of any of the preceding claims, wherein at least some neurons in hidden
5 layers of the decoder portion of the autoencoder neural network are configured with rectified linear unit (ReLU) activation functions and wherein at least some neurons in an output layer of the decoder portion are configured with sigmoidal activation functions.
7. The method of any one of claims 1-6, wherein the psychoacoustic model is based, at least in part, on one or more psychoacoustic masking thresholds.
- 10 8. The method of any one of claims 1-7, wherein the psychoacoustic model involves on one or more of the following:
modeling an outer ear transfer function;
grouping into critical bands;
frequency-domain masking, including but not limited to level-dependent spreading; and
15 modeling of a frequency-dependent hearing threshold.
9. The method of claim 8, wherein generating the loss function value comprises:
transforming the decoded audio signal to a frequency domain to generate a frequency transformed decoded audio signal;
20 transforming the ground truth audio signal to the frequency domain to generate a frequency transformed ground truth audio signal;
applying an outer ear transfer function to the frequency transformed decoded audio signal to generate a first intermediate signal;
applying the outer ear transfer function to the frequency transformed ground truth
25 audio signal to generate a second intermediate signal;
determining the difference between the first intermediate signal and the second intermediate signal to generate a difference signal;
applying a banding operation to the first intermediate signal to generate a plurality of frequency bands of the first intermediate signal;
30 applying the banding operation to the second intermediate signal to generate a plurality of frequency bands of the second intermediate signal;

applying a frequency masking operation to the plurality of frequency bands of the second intermediate signal to generate a plurality of masked bands;

determining a noise-to-mask ratio using the plurality of frequency bands of the first intermediate signal and the masked band; and

5 using the noise-to-mask ratio as the loss function value.

10. The method of any one of claims 1-9, wherein the loss function involves calculating an average noise-to-masking ratio and wherein the training involves minimizing the average noise-to-masking ratio.

11. The method of any of claims 1-10, wherein the calculation of the noise-to-mask ratio
10 involves a banding process.

12. The method of claim 11, wherein the banding process is a Zwicker banding process or a Moore banding process.

13. Audio encoder, comprising the encoder portion of a neural network that has been trained according to the method of any of the claims 1-12, wherein the audio encoder is further
15 configured to:

receive a currently-input audio signal;

encode the currently-input audio signal in a compressed audio format; and

output an encoded signal in the compressed audio format.

14. An audio encoding apparatus, comprising:

20 an interface system; and

a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, the control system configured to implement an audio encoder according to claim 13.

25 15. An audio decoder, comprising the decoder portion of a neural network that has been trained according to the method of any of the claims 1-12, wherein the audio decoder is further configured to:

receive a currently-input encoded audio signal in a compressed audio format;

decode the currently-input encoded audio signal in a decompressed audio format; and

30 output a decoded audio signal in the decompressed audio format.

16. An audio decoding apparatus, comprising:
an interface system;
a control system comprising one or more processors and one or more non-transitory storage media operatively coupled to the one or more processors, the control system
5 configured to implement an audio decoder, the audio decoder comprising the decoder portion of a neural network that has been trained according to the method of any of the claims 1-12, wherein the audio decoder is further configured for:
receiving a currently-input encoded audio signal in a compressed audio
format;
10 decoding the currently-input encoded audio signal in a decompressed audio format; and
outputting a decoded audio signal in the decompressed audio format.
17. A system including the audio decoding apparatus of claim 16, wherein the system further comprises one or more transducers configured for reproducing the decoded audio
15 signal.
18. One or more non-transitory media having software stored thereon, the software including instructions for controlling one or more devices to perform the method of any one of claims 1-12.
19. A computer program product having instructions which, when executed by a
20 computing device or system, cause said computing device or system to perform the method according to any of the claims 1-12.


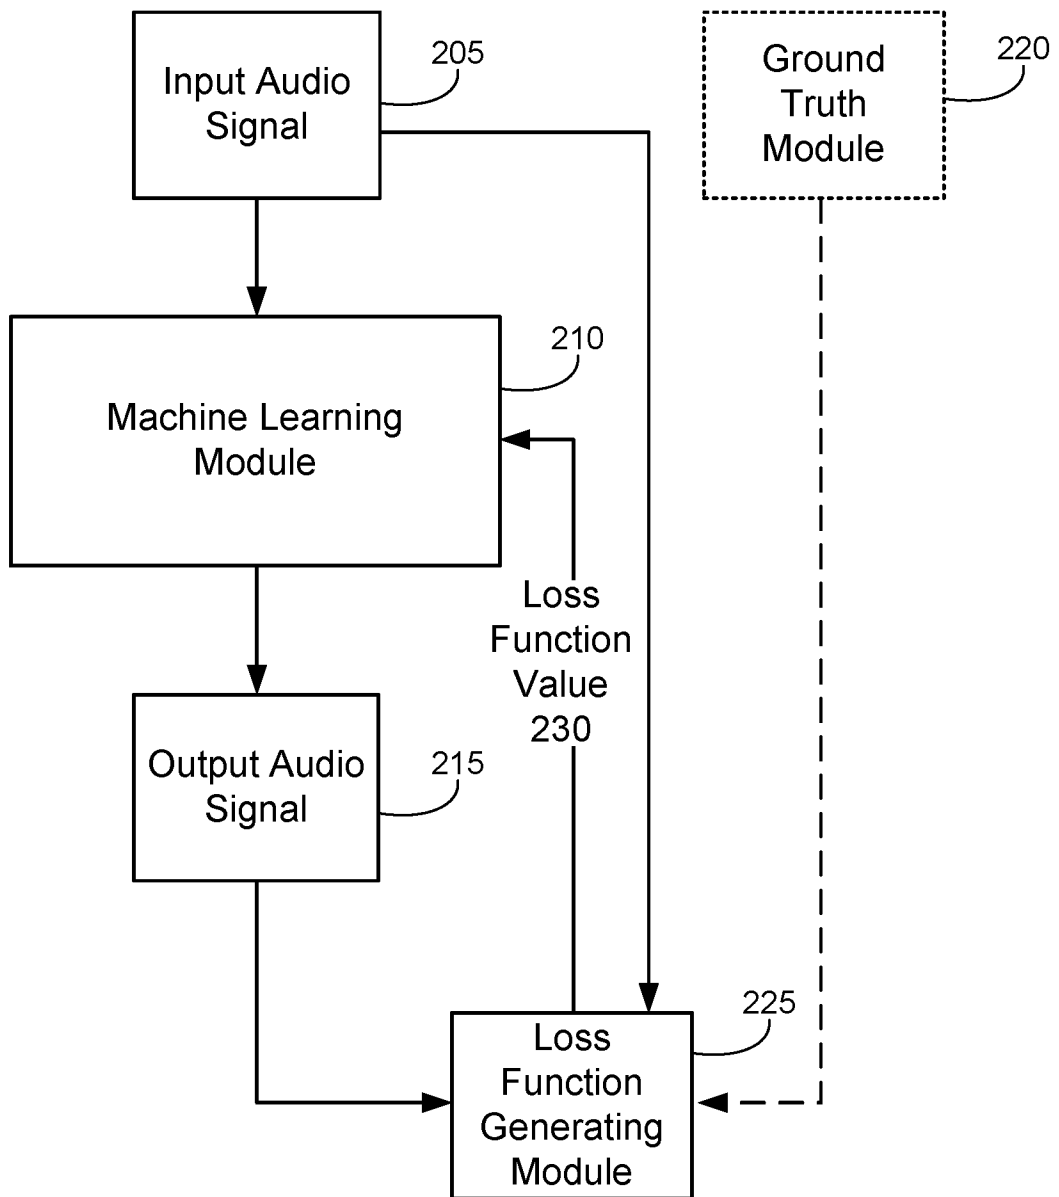
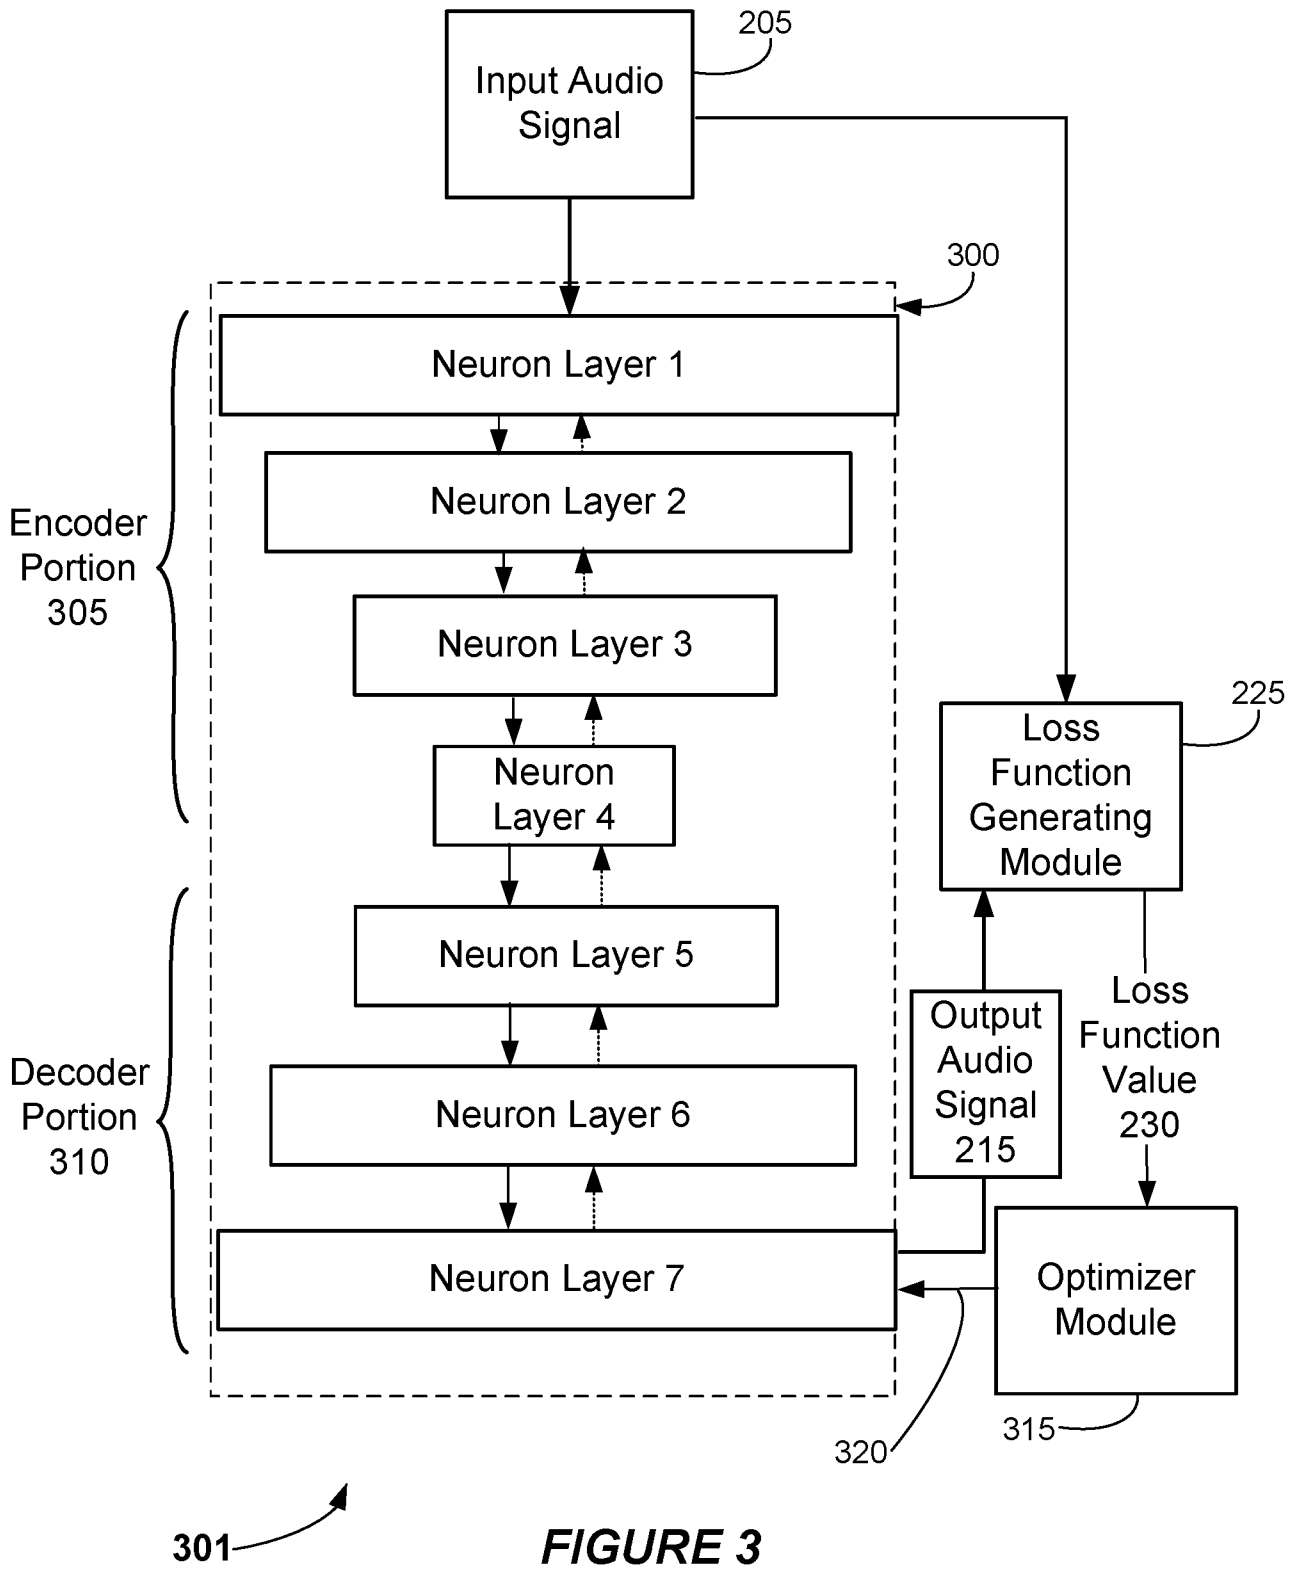


FIGURE 1



200 ↗

FIGURE 2






FIGURE 4A




FIGURE 4B




FIGURE 4C

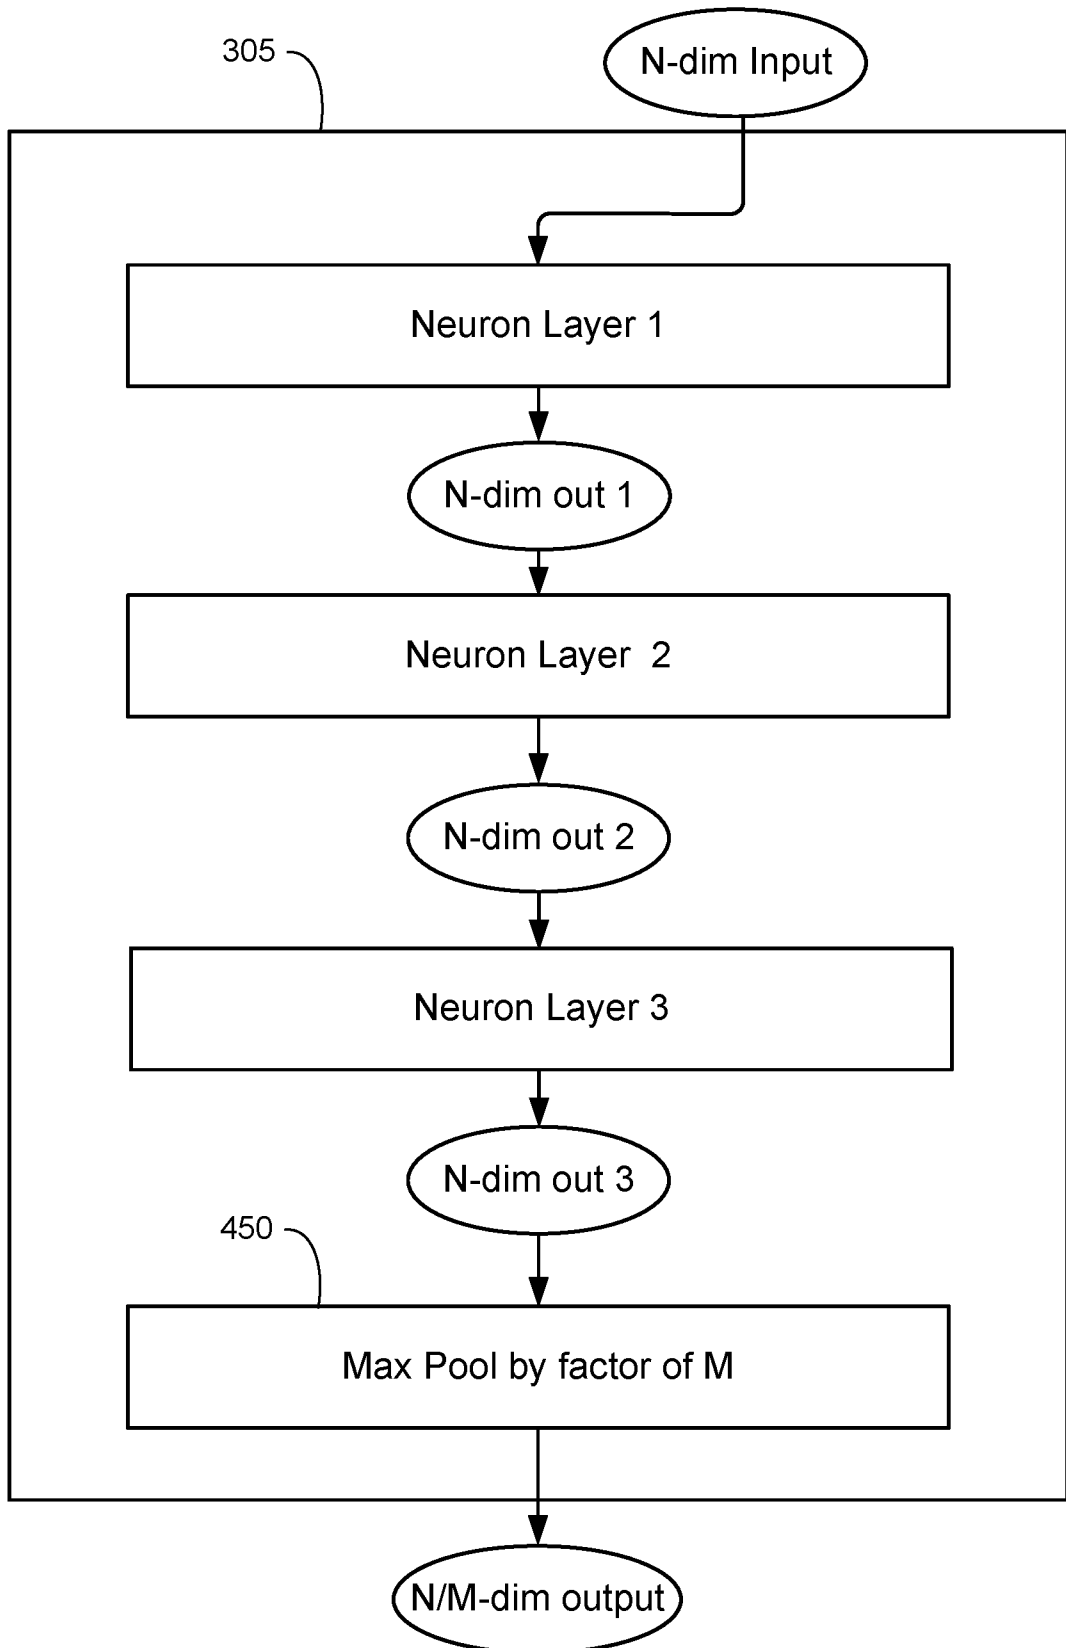


FIGURE 4D

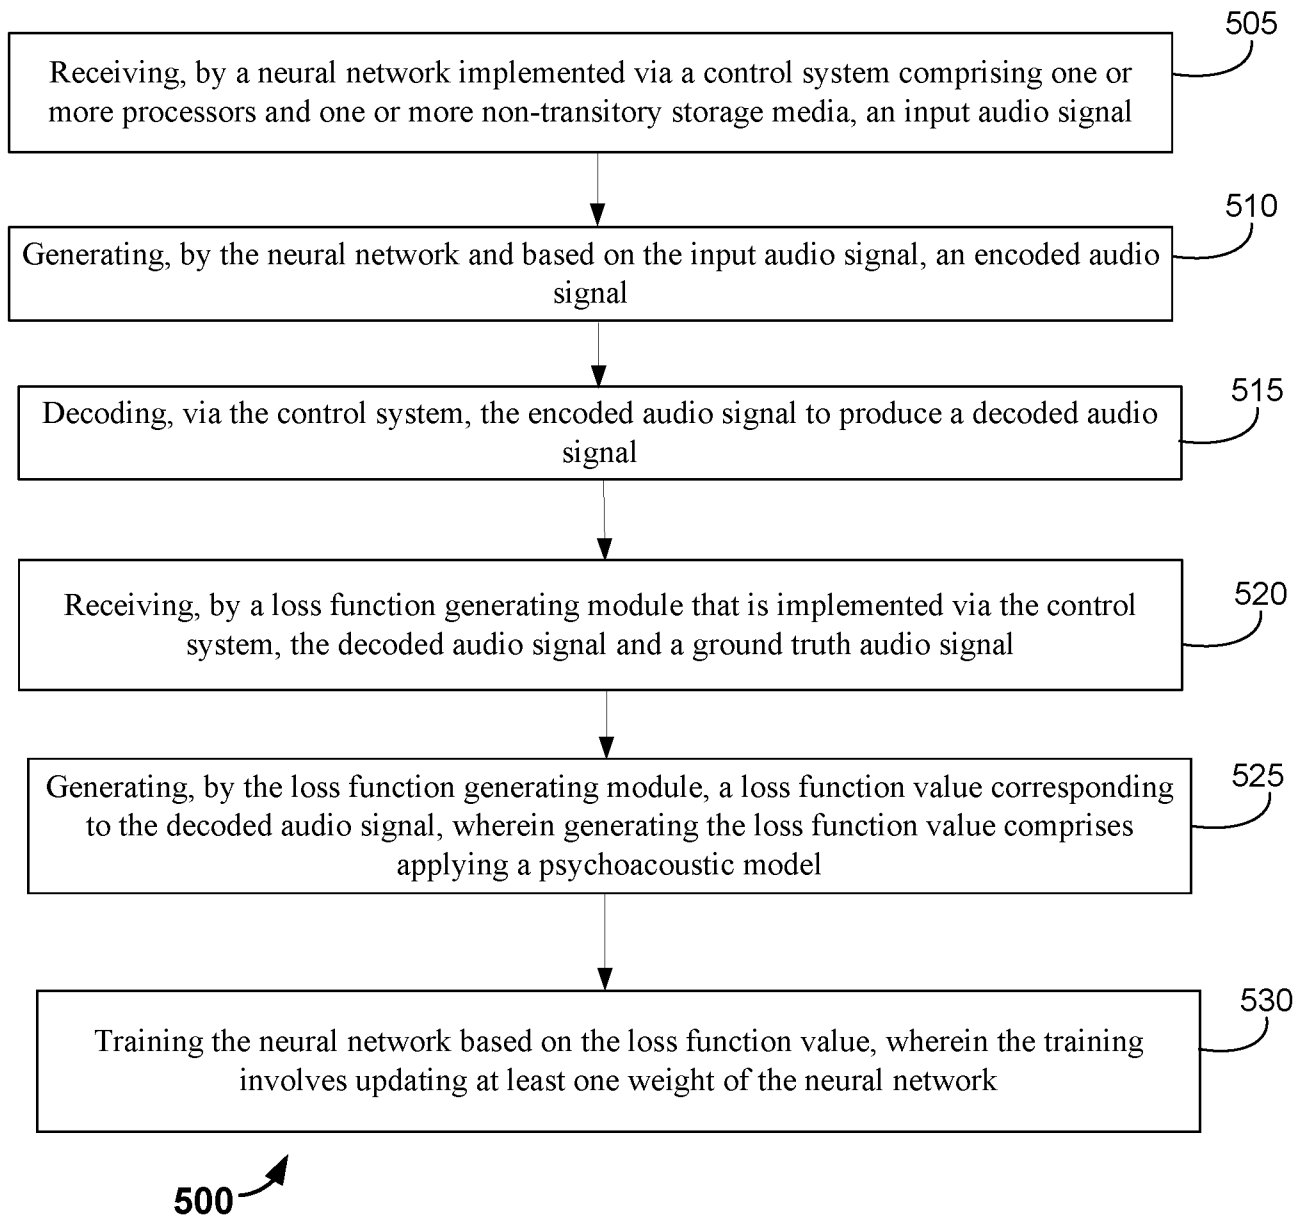


FIGURE 5A

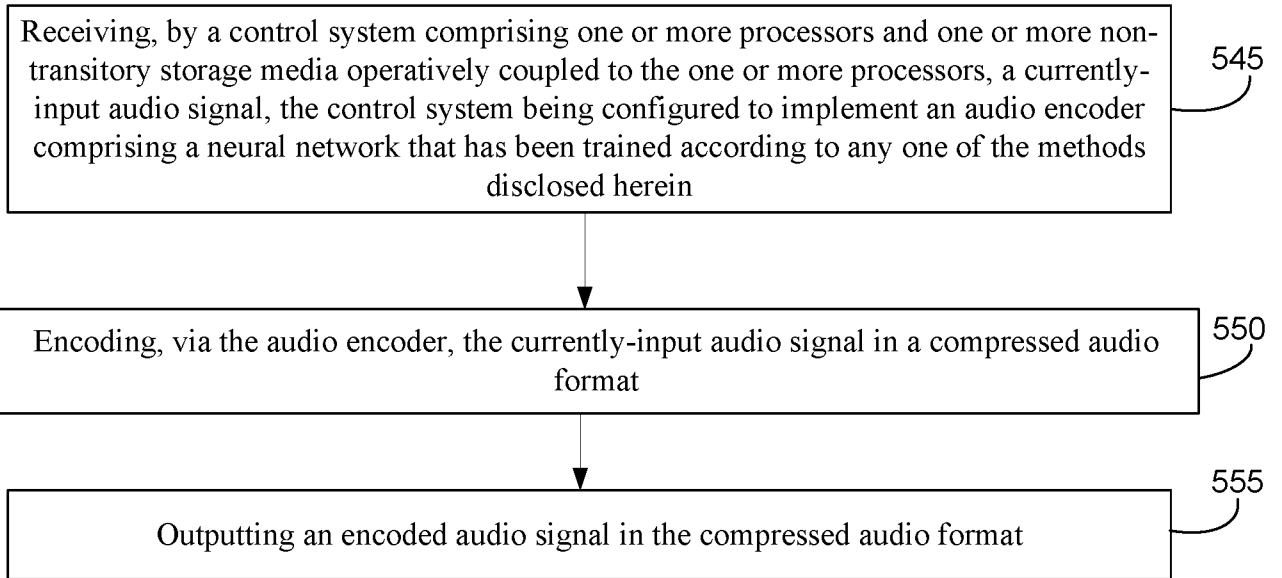


FIGURE 5B

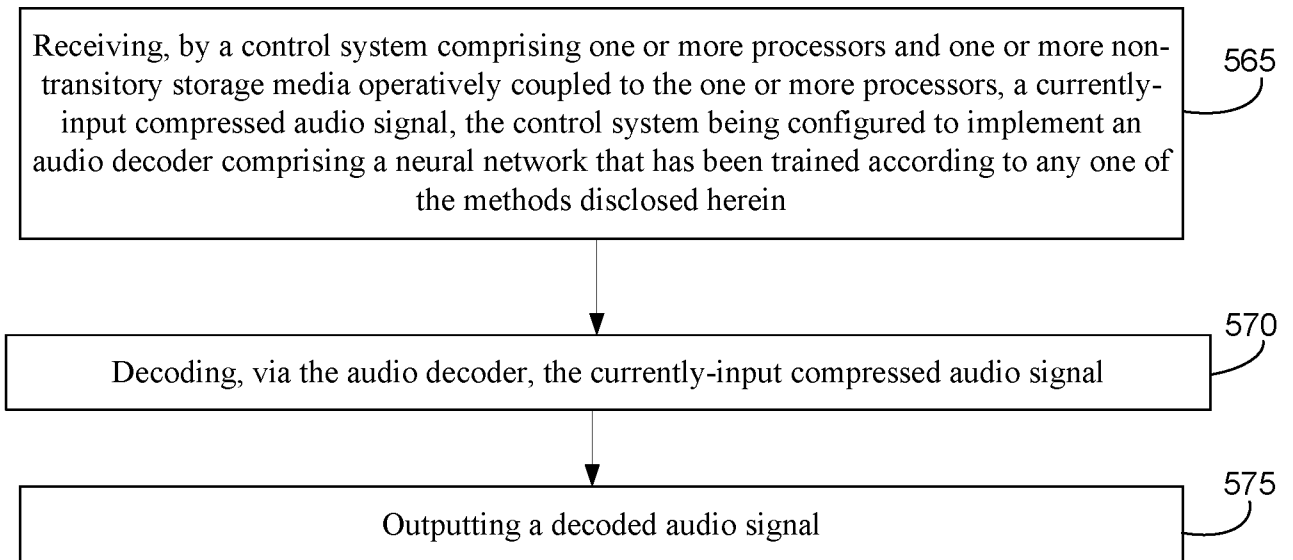


FIGURE 5C

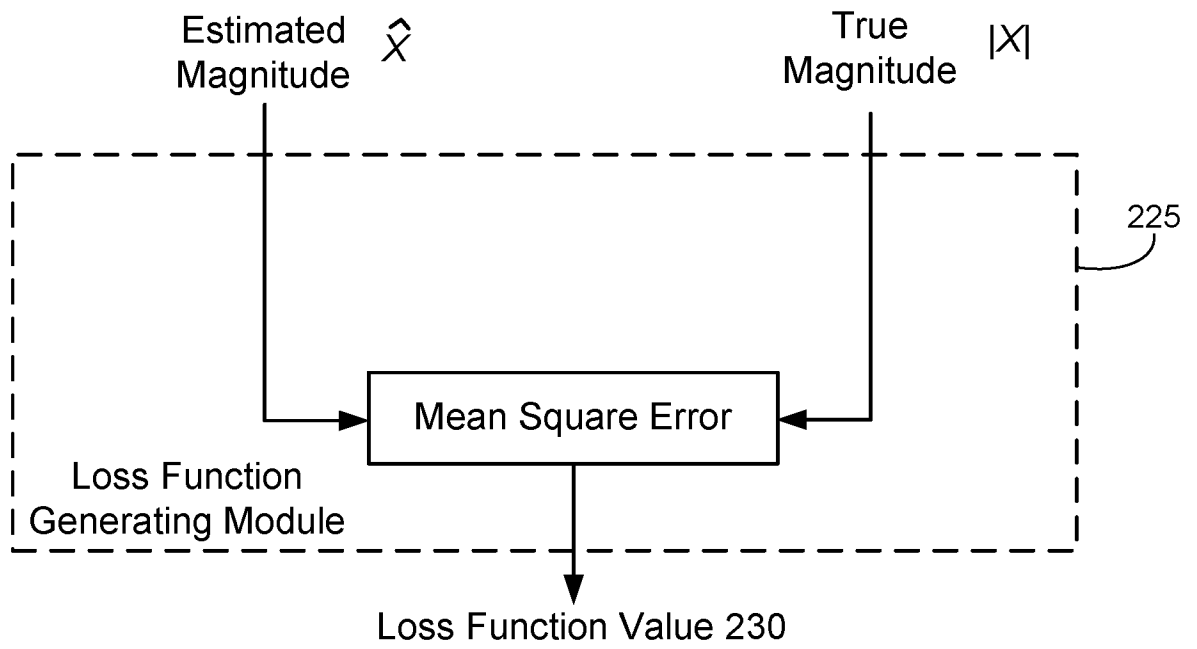


FIGURE 6




FIGURE 7A

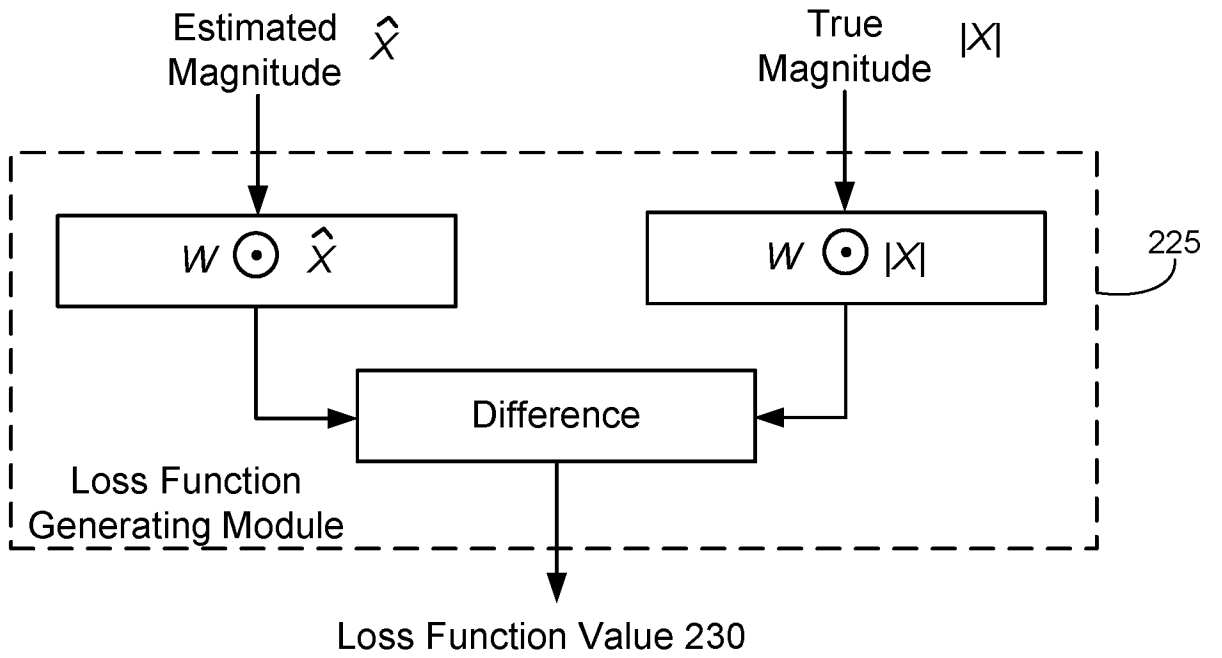


FIGURE 7B




FIGURE 8

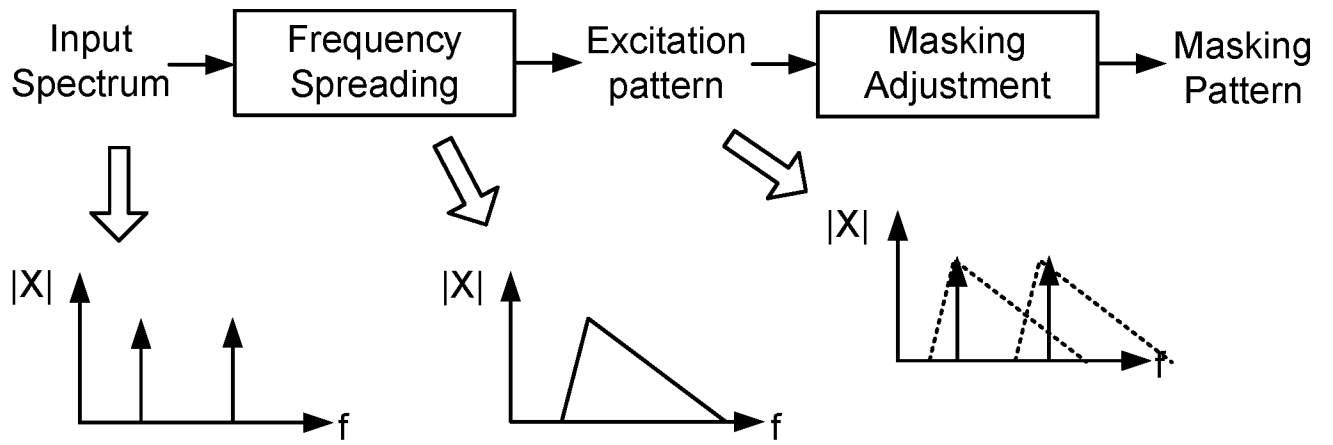


FIGURE 9A




FIGURE 9B

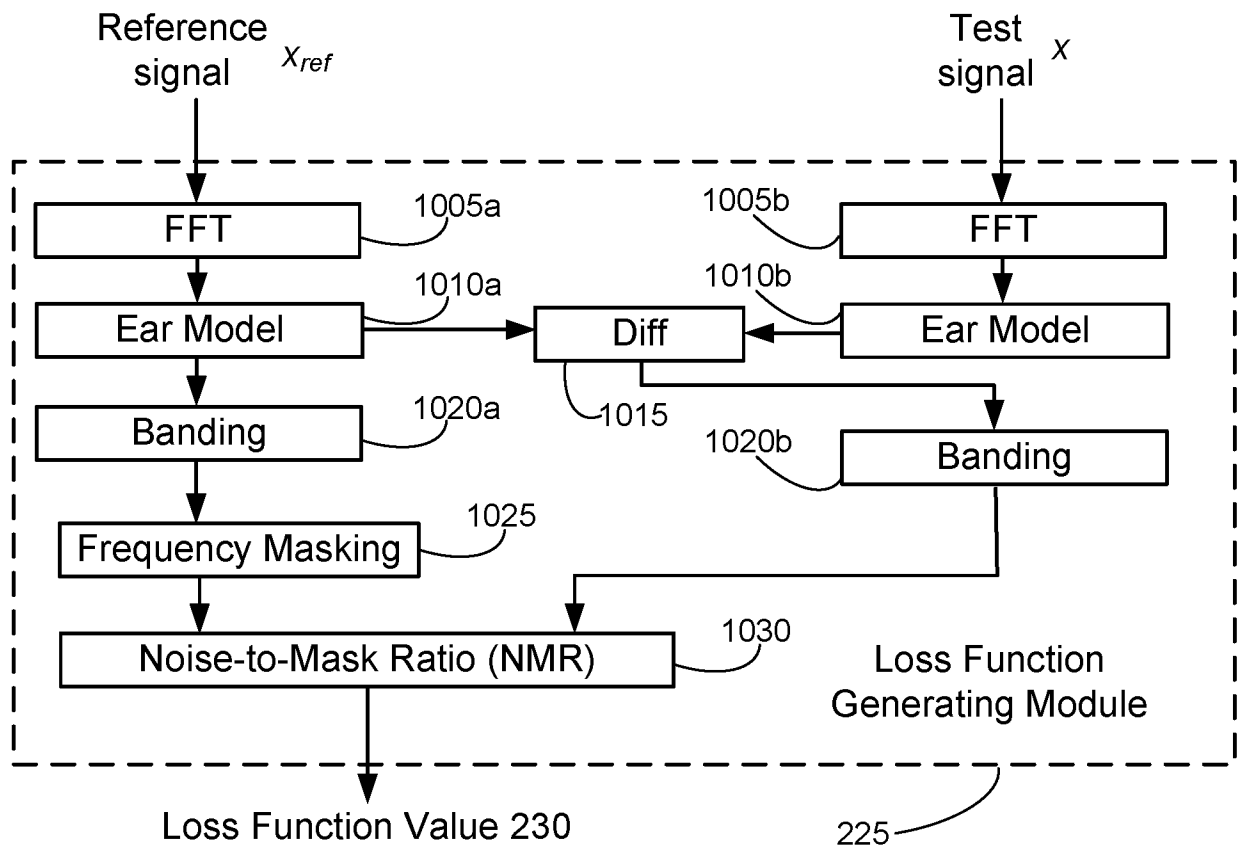


FIGURE 10




FIGURE 11




FIGURE 12




FIGURE 13

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/026824

A. CLASSIFICATION OF SUBJECT MATTER
INV. G01L19/00 G10L25/30
ADD.
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
G10L G06N G01L
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| A | ANAND ATREYA ET AL: "Novel Lossy Compression Algorithms with Stacked Autoencoders", INTERNET CITATION, 11 December 2009 (2009-12-11), XP002768425, Retrieved from the Internet: URL:http://cs229.stanford.edu/proj2009/AtreyaOshea.pdf [retrieved on 2017-03-21] Sections 1.1, 2.1, and 4.2 ----- -/-- | 1-19 |

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

| | |
|---|--|
| Date of the actual completion of the international search 12 July 2019 | Date of mailing of the international search report 19/07/2019 |
|---|--|

| | |
|--|---|
| Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016 | Authorized officer Müller, Achim |
|--|---|

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/026824

| C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT | | |
|--|--|-----------------------|
| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
| A,P | <p>KANKANAHALLI SRIHARI: "End-To-End Optimized Speech Coding with Deep Neural Networks", 2018 IEEE INTERNATIONAL CONFERENCE ON ACOUSTICS, SPEECH AND SIGNAL PROCESSING (ICASSP), IEEE, 15 April 2018 (2018-04-15), pages 2521-2525, XP033404005, DOI: 10.1109/ICASSP.2018.8461487 Section 2</p> <p style="text-align: center;">-----</p> | 1-19 |
| A | <p>MAJUMDAR ANGSUL ET AL: "Asymmetric stacked autoencoder", 2017 INTERNATIONAL JOINT CONFERENCE ON NEURAL NETWORKS (IJCNN), IEEE, 14 May 2017 (2017-05-14), pages 911-918, XP033112151, DOI: 10.1109/IJCNN.2017.7965949 Section 1</p> <p style="text-align: center;">-----</p> | 1-19 |