

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property  
Organization  
International Bureau

(43) International Publication Date  
08 October 2020 (08.10.2020)



(10) International Publication Number  
**WO 2020/199990 A1**

- (51) **International Patent Classification:**  
*G10L 21/02* (2013.01)      *G10L 15/16* (2006.01)  
*G10L 25/30* (2013.01)      *G10L 17/18* (2013.01)
- (21) **International Application Number:** PCT/CN2020/080959
- (22) **International Filing Date:** 24 March 2020 (24.03.2020)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
19166398.8      29 March 2019 (29.03.2019)      EP
- (71) **Applicant: GOODIX TECHNOLOGY (HK) COMPANY LIMITED** [CN/CN]; Room 2113, Level 21, Landmark North, 39 Lung Sum Avenue, Sheung Shui, Hong Kong (CN).
- (72) **Inventors: STRAKE, Maximilian;** Rebenring 8, 38106 Braunschweig (DE). **TIRRY, Wouter Joos;** Abbey House 25 Clarendon Road, Redhill Surrey RH1 1QZ (GB). **DEFRAENE, Bruno Gabriel Paul G;** Abbey House 25 Clarendon Road, Redhill Surrey RH1 1QZ (GB). **FINGSCHIEDT, Tim;** P.O.Box, Walter-Flex-Str.8, 38104 Braunschweig (DE).
- (74) **Agent: LONGSUN LEAD IP LTD.;** Rm. 101, Building 3, No. 68 Beiqing Road, Haidian District, Beijing 100094 (CN).
- (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, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.
- (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).

(54) **Title:** SPEECH PROCESSING SYSTEM AND METHOD THEREFOR




FIG. 1

(57) **Abstract:** A speech processing system (100, 200) comprises an input that receives an input speech signal (101, 201) that comprises clean speech and acoustic interference; a first neural network circuit (102, 216) configured to suppress the acoustic interference in the input speech signal (101, 201) and output (120) a representation of an estimate of a clean speech signal (209); and a second neural network circuit (103, 217) configured to receive the representation of the estimated clean speech signal (209) and restore at least one speech component of the clean speech in the input speech signal (101, 201) and output a reconstructed speech signal (104, 215).

[Continued on next page]

WO 2020/199990 A1

**Declarations under Rule 4.17:**

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

**Published:**

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

## **SPEECH PROCESSING SYSTEM AND METHOD THEREFOR**

### **PRIORITY CLAIMS AND RELATED PATENT APPLICATIONS**

[01] This patent document claims the benefits and priorities of European Patent Application No. EP19166398.8 entitled “SPEECH PROCESSING SYSTEM AND  
5 METHOD THEREFOR”. The entire contents of the before-mentioned patent applications are incorporated by reference as part of the disclosure of this document.

### **TECHNICAL FIELD**

[02] The field of the invention relates to a speech processing system and a method therefor. In particular, the field of the invention relates to a neural network-based speech  
10 enhancement technique.

### **BACKGROUND**

[03] In some speech enhancement tasks, neural network models are able to achieve a strong attenuation of acoustic interference, such as noise, reverberation or acoustic echo. A neural network is regarded as a computing system that is made up of a number of simple,  
15 highly interconnected processing circuits or elements, which process information based on their dynamic state in response to external inputs. In some instances, the processed information is based on an internal dynamic state that crosses all layers of a neural network. In some instances, for recurrent neural networks, the processed information based on a dynamic state relates to an internal state of the neural network over different time steps. Most  
20 speech-based system designers aim to improve the speech quality whilst attempting to maintain high interference attenuation. Classical interference suppression algorithms often suffer from a strong dependency of interference attenuation on the actual interference instance (e.g. different noise types), whilst state-of-the-art neural-network-based systems

show a similar dependency regarding speech distortions.

[04] These state-of-the art systems typically learn a direct mapping from the degraded speech (e.g. clean speech plus acoustic interference) to the clean speech domain. Alternatively, these state-of-the art systems may estimate time-frequency (T-F) masks, in order to separate the clean speech signal from interferences. Such methods, aiming at a strong attenuation of interferences, suffer from audible distortions of the speech component as well as temporal discontinuities or distortions in the residual interference component.

[05] A post-processing method, which uses a convolutional neural network (CNN) network topology and is purely for the enhancement of 'transcoded' speech, is proposed by the authors Ziyue Zhao, Huijun Liu, and Tim Fingscheidt, in the papers titled: "Convolutional Neural Networks to Enhance Coded Speech" and "Neural-Network-based Postprocessing for the Enhancement of Transcoded Speech", arXiv preprint arXiv:1806.09411, 2018. The techniques proposed in these papers stipulate a single stage of neural network processing that is used for speech enhancement of transcoded signals, and is specifically designed to reduce quantization noise introduced by media encoding/decoding and transmission of speech signals.

[06] Zhao, Yan, Zhong-Qiu Wang, and DeLiang Wang, in their paper titled "A two-stage algorithm for noisy and reverberant speech enhancement", published in the Proc. of ICASSP, 2017, identify two neural network stages, each targeting specific interferences, i.e. de-noising and de-reverberation. However, the system proposed in this paper is incompatible with tasks requiring low delay, because of the usage of the iterative Griffin-Lim algorithm for time signal reconstruction and the usage of context from subsequent frames for input feature generation for the second stage network.

[07] The inventors of the present invention have identified a need for an improved speech system, for example by acoustic interference suppression with an ability to restore natural sounding speech.

## SUMMARY

[08] The present invention provides a speech processing system and a method therefor,

together with a computer program product, as described in the accompanying claims.

[09] Specific embodiments of the invention are set forth in the dependent claims.

[10] These and other aspects of the invention will be apparent from and elucidated with reference to the embodiments described hereinafter.

5

## DETAILED DESCRIPTION

[11] Further details, aspects and embodiments of the invention will be described, by way of example only, with reference to the drawings. In the drawings, like reference numbers are used to identify like or functionally similar elements. Elements in the figures are illustrated for simplicity and clarity and have not necessarily been drawn to scale.

10 [12] FIG. 1 schematically illustrates a block diagram of a speech processing system that includes a neural network-based two-stage speech enhancement according to an example of the present invention.

[13] FIG. 2 schematically illustrates a simplified diagram of an example speech processing system, together with some example pre-processing steps before acquiring the  
15 input signal, according to some examples of the present invention.

[14] FIG. 3 schematically illustrates a block diagram of separate training schemes for acoustic interference suppression and speech restoration neural networks, according to an example of the present invention.

[15] FIG. 4 schematically illustrates a block diagram of a joint neural network training  
20 scheme, according to some examples of the present invention.

[16] FIG. 5 illustrates a flowchart of an example of a topology for an acoustic background interference suppression first stage neural network, in accordance with some examples of the invention.

[17] FIG. 6 illustrates a flowchart of an example of an encoder-decoder topology for a  
25 speech restoration second stage neural network, in accordance with some examples of the invention.

[18] FIG. 7 illustrates a flowchart of an alternative example of an encoder-decoder topology for a second stage neural network, which can be employed to restore at least one

speech component and that includes a convolutional LSTM bottleneck layer, in accordance with some examples of the invention.

### BRIEF DESCRIPTION OF THE DRAWINGS

**[019]** Examples of the invention aim to address the abovementioned problems with known speech processing systems by provision of a neural-network-based approach that uses two distinct processing stages, namely a first neural network circuit (or stage) that is configured for acoustic interference suppression and a second neural network circuit (or stage) that is configured for restoration of a natural sounding speech signal. Examples of the invention employ specific neural network topologies for these two successive tasks to be able to obtain a strong acoustic interference suppression, advantageously even in highly non-stationary and/or poor speech-to-interference ratio acoustic conditions, whilst still maintaining a high-quality speech signal. For example, non-stationary noises exhibit dynamic spectral characteristics over time, as opposed to stationary noises, which have a more stable/static spectral characteristic over time. An example of the former are impulsive noises (e.g. slamming door, rattling dishes), while an example of the latter are vehicle noises (car, train).

**[020]** Advantageously, examples of the neural-network-based approach for speech enhancement described herein, that use two distinct processing stages, reduce the known design dependency of targeting attenuation of a specific interference. This is achieved by targeting very high acoustic interference attenuation in the first neural network circuit (or stage). Notably, in accordance with examples of the invention, any residual acoustic interference may be alleviated by the second neural-network speech restoration stage.

**[021]** In some examples, this enables the second neural network circuit (or stage) to be configured to generalize well to different interference instances (e.g. acoustic interference), in terms of speech restoration and residual acoustic interference suppression, e.g. any residual acoustic interference exhibited in the signal passed to the second neural network circuit (or stage) following acoustic background interference suppression performed by the first neural network circuit.

[022] Thus, a speech processing system is described that comprises: an input that receives an input speech signal that comprises clean speech and acoustic interference. A first neural network circuit is configured to suppress the acoustic interference in the input speech signal and output a representation of an estimate of a clean speech signal. A second neural network circuit is configured to receive the representation of the estimated clean speech signal and restore speech components of the clean speech in the input speech signal, and suppress any residual acoustic interference, and output a reconstructed speech signal.

[023] Further, in some examples of the invention, a joint training of both neural network circuits (or stages) is proposed, in order to refine the performance of the total speech processing system. In some examples, joint training (and potentially additional pre-training) may be employed to configure the first neural network circuit (or stage) in order to generate well-suited inputs for the second neural network stage, which is simultaneously further adapted to estimate natural sounding speech. In these examples, the training of the speech processing system is such that the neural network circuit adaptively learns rules for acoustic interference suppression by using a large specific data set. In examples of the invention, an algorithm employed for training is focused on learning and iteratively adapting weights and biases of the neural networks to perform background acoustic interference suppression and speech restoration respectively.

[024] Thus, examples of the invention describes an approach that decouples processing into two stages, each of which can focus on one task, i.e. a first stage configured to suppress acoustic interference, and a second stage configured to restore speech, improve speech quality/naturalness. Furthermore, the outputs of the first stage are configured to be well suited as inputs to the second stage, because most of the acoustic interference has been suppressed, which allows the second stage to focus on speech quality restoration.

[025] Although example embodiments of the invention are designed for acoustic interference suppression in speech signals, in a first neural network circuit (or stage) of a two-stage neural network system, it is envisaged that the concepts described herein may be equally employed for other types of interference, by adapting the first neural network stage to the specific interference suppression role. This is distinct from the cited art that uses a single neural network that is trained to suppress quantization distortion and noise of transcoded

signals introduced by encoding and decoding operations and/or distortion due to digital signal processing.

[026] Furthermore, although the example embodiments of neural network circuits (or stages) propose to use estimators based on real-valued neural networks in both stages, it is envisaged that in other examples that complex-valued neural networks or a combination of any kind of neural network may be employed in the two-stage approach described herein, including classical interference suppression or speech restoration approaches.

[027] Thus, examples of the present invention may be targeted for both (or either) of real-valued networks and complex-valued networks. In a complex-valued network, all operations carried out as part of the neural network are defined as operations on complex numbers, so that especially the trainable parameters may be complex-valued and the loss function for (separate or joint) network training may be defined using complex-valued operations..

[028] In contrast, in a real-valued neural network, all internal mathematical operations are assumed to be operating on real-valued variables. In examples of the invention, when using a combination of a two-stage neural network design and conventional digital speech processing approaches to interference suppression and/or speech restoration, the conventional approach has to be implementable as a number of operations with well-defined derivatives for backpropagation training of the system combination.

[029] Referring now to FIG. 1, a schematic block diagram of a two-stage neural network-based speech enhancement circuit 100 is illustrated, according to an example of the present invention. In accordance with examples of the invention, the two-stage approach for speech enhancement uses neural network models in both processing stages. A first neural network circuit (or stage) 102 is configured to suppress acoustic interference in a degraded input speech signal 101, which in some instances is configured for suppression of acoustic background interference. In accordance with examples of the invention the first neural network circuit 102 of the two-stage neural network system can be trained to suppress any acoustic background interference. In some examples, the first neural network circuit (or stage) 102 may optionally receive one or more auxiliary input signals 105, e.g. additional microphone signals or acoustic interference estimate signals, which may be utilised by the

first neural network circuit (or stage) 102 to better suppress the targeted acoustic interference.

**[030]** In accordance with examples of the invention, the second neural network circuit (or stage) 103 is configured to restore natural sounding speech. In some examples, the second neural network circuit (or stage) 103 may also be configured to potentially suppress residual  
5 acoustic interferences emanating from the output 120 of the first neural network circuit (or stage) 102.

**[031]** The second neural network circuit 103 is configured to restore one or more speech components of the clean speech signal. Some examples of speech components that are relevant to the speech quality and intelligibility include the following: pitch frequency, pitch  
10 harmonic structure, formant structure, spectral envelope, spectral phase. Pitch is the fundamental frequency of vibration of the vocal folds, and the pitch frequency is a measure of how high or low the frequency of a person's voice sounds. The perception of voiced sounds relies heavily on the presence of the clean speech pitch frequency and pitch harmonics, i.e. integer multiples of the pitch frequency. Formants are concentrations of acoustic energy  
15 around a particular frequency in the speech wave, and occur around frequencies that correspond to the resonances of the vocal tract. The spectral envelope is the envelope curve of the amplitude spectrum. Next to the spectral amplitudes, it should be noted that also the spectral phase is important to the perceived speech quality. Thus, in examples of the invention, the second neural network circuit 103 is configured to restore at least one speech  
20 component, which in the context of some examples of the invention may encompass restoring one or more of the above.

**[032]** In some optional examples, the second neural network circuit (or stage) 103, configured for speech restoration, may be configured to additionally use the degraded input speech signal 101 as an additional input 150 for estimating an enhanced speech signal 104 to  
25 be output from the two-stage neural network-based speech enhancement circuit 100.

**[033]** In some optional examples of the invention, a neural network training system is also described, whereby the two-stage neural networks may be first trained separately and then the neural network training system successively performs a joint training of a combination of the neural network circuits (or stages) 102, 103. In some optional examples,  
30 the joint training may be performed using a single loss function at the output 104 of the

two-stage system or using a combined loss function 125, 135. In some optional examples, the joint training may additionally include an intermediate loss 125 introduced at the output 120 of the acoustic interference suppression first neural network circuit (or stage) 102. In a joint training scheme approach, where a loss function 125 is applied to an output of the first neural network circuit (or stage) 102 and a further loss function 135 applied to an output 104 of the second neural network circuit (or stage) 103, two loss functions 125, 135 that contribute to parameter updates in the neural network can be advantageously achieved. In examples of the invention, the use of an intermediate loss function enables the speech system designer to have control about the goals of the neural network scheme at intermediate points. In this two loss function example, the speech system designer is able to employ a loss function that favours very aggressive acoustic interference suppression in the first stage. It is envisaged that in other examples, the speech system designer may be able to employ a loss function that favours lower acoustic interference suppression in a first stage, if the acoustic interference was determined as not being excessive. Thus, in some optional examples, it is envisaged that joint training of both neural network stages may be performed with or without an intermediate loss function 125 being introduced, and in some instances without performing separate training of each neural network stage beforehand.

**[034]** Referring now to FIG. 2, a simplified diagram of an example processing scheme, together with some example signal pre-processing circuits that may be configured to generate the input signal, is illustrated, according to some examples of the invention. First, FIG. 2 illustrates an example of a pre-processing system 250 that includes some potential pre-processing circuits, each of which typically remove interference, leaving residual acoustic interference in their respective output signals and each of which may be validly selected to provide the input signal to the two neural network circuit (or stages) 220, according to some examples of the invention. For example, the pre-processing system 250 may generate an input signal  $y(n)$  201 for the two neural network circuit (or stages) 220 that contains a mixture of speech and interference, which may be a digitized microphone signal 260 or an output of a pre-processing operation. Thus, the two-stage neural network design example of the invention may be included to suppress unwanted acoustic background interference in the processed speech signal even after a conventional acoustic interference suppression circuit,

such as an acoustic echo canceller.

[035] In this example, a mixture of speech and residual acoustic interference is contained within one or more digitized microphone signal(s) 260, input to an acoustic echo cancellation circuit that comprises an echo estimation circuit 254, as well as a subtraction  
5 junction 264 that is arranged to subtract an acoustic echo estimate signal 262, in response to an input far-end signal 252 and a sampled output 256 of the summing junction(s) 264. The signal resulting from the subtraction of the echo estimate provides the input of the two-neural-network circuit.

[036] In other examples, the pre-processing system 250 may generate a beamformer  
10 output signal 272, output from a beamforming circuit 270 that may also be used as an input of the two-neural-network circuit. Similarly a de-reverberation signal output from a de-reverberation circuit 280, may also be used as an input of the two-neural-network circuit, as such example circuits that typically remove interference, leaving residual acoustic interference in their respective output signals. They are configured to provide a signal with  
15 reduced interference. In this manner, the pre-processing system 250 may provide one or a combination of the aforementioned signals (i.e. speech and unwanted residual acoustic background interference) to the input of the two neural network circuit (or stages) 220.

[037] In accordance with examples of the invention, the two neural network stage system 220 is described with reference to a first acoustic interference suppression circuit (e.g.  
20 a first stage) 216, followed by a second speech restoration circuit (e.g. a second stage) 217. The two neural network circuits are described with reference to an input signal  $y(n)$  201 that is input to a framing circuit 202 of the first acoustic interference suppression circuit (e.g. stage) 216 and processed in frames of 32msec., for example, with a frame shift of 50%, i.e. 16msec. It is envisaged that in other example implementations that other frame lengths may  
25 be employed, as well as other frame shifts, such as 25%. In this example, each frame with frame index  $i$  is subjected to a square root Hann window and subsequently transformed into the frequency domain via Fast Fourier Transform (FFT) 203 of size  $N$ . It is envisaged that in other example implementations that other windowing techniques may be employed, such as a non-square-root or hamming window.

30 [038] The resulting signal 204 in the first acoustic interference suppression circuit (e.g.

stage) 216, output from FFT 203 forms a basis for feature extraction. In this example, the feature extraction operations are performed by a feature extraction circuit 205 and a mean and variance normalization circuit 206.

**[039]** It is envisaged that some features that may be extracted by feature extraction circuit 205, may include spectral magnitudes, real and imaginary parts of the audio spectrum, mel-frequency-based spectral magnitudes, etc. In some examples, for the computation and extraction of spectral features, a frequency domain transformation (e.g. by FFT 203) may be applied, which decomposes a time domain signal into its constituent frequency components. Example frequency domain transformations include a Discrete Fourier Transform, which results in complex spectral components (consisting of both real and imaginary parts), and the related Discrete Cosine Transform, which results in purely real spectral components. In some examples, these spectral features can be further transformed into spectral magnitudes by taking the Euclidean norm of the spectral components. In speech processing applications, a commonly derived feature from the spectral magnitudes is the mel-frequency cepstrum, where the spectral magnitudes are mapped to the mel-scale using triangular overlapping windows. The mel-scale is a logarithmic scale, which was found to match well the human perception of sounds.

**[040]** It is envisaged that, in other example implementations, a single mean normalisation may be employed. Each of these optional circuits is employed before the first neural network circuit (or stage) 207 configured for acoustic background interference suppression. In some examples, the feature extraction circuit 205 may include a concatenation of neighbouring frames, using features from previous and/or future frames. Such a concatenation is not obligatory and in some examples the feature extraction circuit 205 may alternatively be configured to operate on features based only on the current frame.

**[041]** In some examples, the first neural network circuit (or stage) 207 may be trained to estimate T-F masks, which in some examples may subsequently be multiplied with a matching representation of a frequency domain mixture signal 204 obtained from an output of the FFT 203. Such a frequency domain mixture signal 204 may be applied by a masking circuit 208, to an output of the first neural network circuit (or stage) 207 in order to obtain a frequency domain vector representation of the estimated clean speech signal 209. The

frequency domain vector representation of the estimated clean speech signal 209 may consist of spectral magnitudes, or a concatenation of real and imaginary parts of the complex spectral bins, or mel-frequency-based spectral magnitudes. In some examples, it is envisaged that it may be possible to directly estimate the clean speech signal 209, in contrast to applying the output of the first neural network circuit (or stage) 207 to a masking circuit 208. However, the use of a masking circuit 208, in some applications, provides an advantage of preventing the first neural network circuit 207 from being forced to estimate absolute signal levels of the speech.

**[042]** The frequency domain vector representation of the estimated clean speech signal 209 is input to the second speech restoration stage 217. The input frequency domain vector representation for the second speech restoration stage 217 may be subject to an optional second feature extraction in a second feature extraction circuit 210, which generates features in the processing domain of the second neural network circuit (or stage) 212, e.g. frequency, time or cepstral domain. In some preferred examples, the second neural network stage 217 may not use a second feature extraction circuit 210 and instead enable a second neural network circuit (or stage) 212 to operate directly on the frequency domain vector representation of the estimated clean (acoustic interference suppressed) speech 209, which is first applied to a second mean and variance normalization circuit 211. In this manner, it is possible to deliver features that are more suited for learning the second stage neural network. It is envisaged, one example may be to extract spectral features with a higher resolution for the second stage, in order to focus on bringing back spectral details in the estimated clean speech. Furthermore, it is envisaged that in another scenario, where computational complexity is a concern, it may be possible to use a second feature extraction reducing the feature dimension and therefore enabling the use of a simpler signal reconstruction network 214. In this manner, the output of the second neural network circuit (or stage) 212 is a vector representation of the second-stage clean signal estimate 213 for the current frame . The vector representation of the second-stage clean signal estimate 213 is then applied to a signal reconstruction circuit 214, which is configured to generate a time domain second-stage estimate 215 of the clean signal from the frame representations of the second-stage clean signal estimate 213. In one example, this may be achieved by applying an inverse transform

from the processing domain back to the time domain, together with a subsequent combination of the time domain frames, e.g. by Inverse Fast Fourier Transform (IFFT) and an overlap add (OLA) operation, in one example embodiment.

5 [043] In some examples, it is envisaged that a so-called overlap-save (OLS) technique may be used, where the IFFT is being used and a portion of the resulting time-domain frame is appended to a respective previous portion of the frame, thereby producing a time-domain output.

10 [044] In some examples, it is envisaged that if the next processing step is an automatic speech recognizer (ASR), the IFFT and OLA/OLS schemes may be dropped, and a recognition of speech may be achieved by directly operating on a frequency-domain output (instead of a time domain second-stage estimate 215) of the invention. Thus, it is envisaged that both the first acoustic interference suppression circuit (e.g. stage) 216, followed by the second speech restoration circuit (e.g. stage) 217 may be equally well performed directly in the time domain, as compared to the described frequency domain.

15 [045] In some examples, and before performing both the first acoustic interference suppression circuit (e.g. stage) 216, followed by the second speech restoration circuit (e.g. stage) 217, the first neural network circuit (or stage) 207 and second neural network circuit (or stage) 212 may be trained on a specified dataset. In some examples, such training may be performed by adapting, say, internal parameters , e.g. the weights and biases defining the  
20 mapping that is carried out by the neural network, in order to minimize a loss function between the neural network outputs and desired output values (e.g. targets). In a practical neural network training implementation, it is known to use some form of gradient descent algorithm, in order to iteratively adapt the internal parameters, i.e. the weights and biases in the neural network, in order to minimize the selected loss function as far as possible. In each  
25 iteration of such gradient descent algorithm, a relatively small internal parameter adjustment is performed in the direction of steepest descent of the loss function on a mini-batch of data. This iterative training scheme can be initialized with random settings for the internal parameters, or with some predetermined values. The iterative training scheme is performed until one or more stopping criteria is met, e.g. one such criterion could be to stop iterative  
30 training when no more decrease is seen in the loss function when performing gradient descent

iterations on new mini-batches of data.

**[046]** This training procedure may be performed either by separately training the two different networks, or by jointly training the two neural network circuits (i.e. the two neural networks are effectively considered as a combined two-stage network from the start). In some examples, the training procedure may be performed following separate first training routines and thereafter using the obtained internal neural network parameters e.g. the weights and biases defining the mapping that is carried out by the neural network, to initialize both neural networks and subsequently train the combined networks jointly.

**[047]** In some examples, it is envisaged that a use of discriminative training and a mean squared error (MSE) loss function may be used when training the second stage neural network, noting that the training relies on a use of a certain loss function, as would be appreciated by a skilled artisan. In some examples, it is envisaged that different loss functions, such as a perceptually motivated loss for maximizing speech naturalness and/or other training schemes such as generative adversarial network (GAN) training, may be used for different embodiments and implementations.

**[048]** FIG. 3 schematically illustrates a block diagram of separate training schemes including training an acoustic interference suppression neural network circuit 322 and subsequently training a speech restoration neural network circuit 323, according to an example of the present invention. In one example, the training of the acoustic interference suppression neural network circuit 322 is performed first. The example training approach in FIG. 3 may use the same frame timing and FFT, as used in processing the noisy speech signal training data 301 and the clean speech signal data 302, which offers advantages with respect to the overall delay and quality of the stored training data.

**[049]** In examples of the invention, the training of the acoustic interference suppression neural network circuit 322 and subsequently training of the speech restoration neural network circuit 323 operates on input/target pairs of speech signal training data 301 (i.e. speech plus acoustic background interference) and clean speech signal data 302, which are aligned in time and level, i.e. they stem from the same time-frame of the same speech sentence. The speech signal training data 301 is input to feature extraction circuit 307. After the feature extraction circuit 307 has extracted one or more feature(s) of the noisy speech signal training data 301,

such as signal magnitudes, real and imaginary parts of the audio spectrum, mel-frequency-based spectral magnitudes, etc., a normalization of the extracted features is performed in normalization circuit 310 using, say, a vector of means 308 and/or a vector of standard deviations 309, extracted from a data storage element 330. In some examples, the two vectors may be computed over all samples in a training data set. The resulting normalized features 306 output from normalization circuit 310 may be used as an input to the first neural network circuit (or stage) 311. In some examples, the first neural network circuit (or stage) 311 may also receive a 'direct' version of the speech signal training data 301, as shown.

10 [050] In a parallel signal path of the acoustic interference suppression neural network 322, the clean speech signal data 302 is input to a target preparation circuit 303, which is configured to restructure the clean speech signal data 302 into a suitable vector representation form 304 for network learning, which matches the dimension and domain of the desired neural network output. In one example embodiment, training is performed in the target preparation circuit 303 using standard backpropagation with the mean spectral approximation (MSA) loss function, as shown in equ.[1]:

$$J_{MSA}(\Theta^{(1)}) = \frac{1}{|\mathcal{B}_1|} \sum_{\ell' \in \mathcal{B}_1} \|\hat{\mathbf{m}}_{\ell'} \odot \mathbf{Y}_{\ell'} - \mathbf{S}_{\ell'}^{(1)}\|_2^2 \quad [1].$$

[051] With:  $\|\cdot\|_2$  being the L2-norm,

$\hat{\mathbf{m}}_{\ell'}$  being the network output vector,

20  $\hat{\mathbf{m}}_{\ell'}$  and  $\mathbf{S}_{\ell'}^{(1)}$  being the matching vector representations of noisy speech signal training data 301 and clean speech signal data 302 respectively, and

$\mathcal{B}_1$  and  $\Theta^{(1)}$  being the set of frame indices for the current mini-batch and the trainable parameters of the first stage network, respectively.

[052] In some examples, the vector representations  $\mathbf{Y}_{\ell'}$  and  $\mathbf{S}_{\ell'}^{(1)}$  may contain real and imaginary parts of spectral bins resulting in a complex MSA loss, or they may contain spectral magnitudes resulting in a magnitude MSA loss. Alternatively, in other examples, it is envisaged that other loss functions may be used, such as a magnitude or complex variant of

MSA with square-root values, as shown in equ.[2]:

$$J_{\text{MSA}_{\text{MSQ}}}(\Theta^{(1)}) = \frac{1}{|\mathcal{B}_1|} \sum_{\ell' \in \mathcal{B}_1} \|\hat{\mathbf{m}}_{\ell'} \odot \mathbf{Y}_{\ell'} - \mathbf{S}_{\ell'}^{(1)}\|_1 \quad [2],$$

with  $\|\cdot\|_1$  being the L1-norm.

[053] In other examples, it is envisaged that a generative adversarial network (GAN) loss function may be used, which, as opposed to the loss functions mentioned above, does not directly compare the network outputs to a target clean speech representation. Instead, an adversarial training scheme may be employed, where an additional binary classification network (referred to as a ‘discriminator’) with the trainable parameters  $\Theta_D^{(1)}$  is employed to distinguish between data of the target domain  $\mathbf{S}_{\ell'}^{(1)}$  and the masked output data  $\hat{\mathbf{S}}_{\ell'}^{(1)} = \hat{\mathbf{m}}_{\ell'} \odot \mathbf{Y}_{\ell'}$  of the first neural network circuit (or stage) 311. In this example, during training with a GAN loss function, separate update steps may be employed for the newly defined binary classifier and the first neural network circuit 311, which are based on the losses computed according to equ.[3] and equ.[4]:

$$J_{D,\ell'}(\Theta_D^{(1)}) = \frac{1}{2}(D(\mathbf{S}_{\ell'}^{(1)}) - 1)^2 + \frac{1}{2}(D(\hat{\mathbf{S}}_{\ell'}^{(1)}))^2 \quad [3]$$

15 and

$$J_{G,\ell'}(\Theta^{(1)}) = \frac{1}{2}(D(\hat{\mathbf{S}}_{\ell'}^{(1)}) - 1)^2 \quad [4].$$

[054] Where:  $D(\cdot)$  is an output of the classifier network, representing the probability of its input signal belonging to the target domain.

[055] During a first step, the classifier of the first neural network circuit (or stage) 311 is trained, according to  $J_{D,\ell'}$  [3] in order to distinguish between the clean signal  $\mathbf{S}_{\ell'}^{(1)}$  and an enhanced signal  $\hat{\mathbf{S}}_{\ell'}^{(1)}$ . The second step trains the first neural network circuit (or stage) 311 according to  $J_{G,\ell'}$  [4] to output an enhanced signal  $\hat{\mathbf{S}}_{\ell'}^{(1)}$ , that cannot be distinguished from the target clean speech  $\mathbf{S}_{\ell'}^{(1)}$  by the discriminator. It is envisaged that this example training scheme is one of many envisaged alternative ways of employing GAN loss functions, with

minor changes to the training scheme described above, that can also be used for training of the first neural network circuit (or stage) 311.

[056] In some examples, it is envisaged that the first neural network circuit (or stage) 311 may also consist of two or more equal or unequal networks (referred to as ‘subnetworks’), in order to achieve a better acoustic interference attenuation. This may mean that, in the first neural network circuit (or stage) 311, only one final loss function is defined. In other examples, a loss function may be defined after each of the ‘subnetworks’. In some examples, it is envisaged that such subnetworks may be trained one after the other, or they may be trained jointly.

10 [057] The inventors have recognised and appreciated that if the same training dataset is used for training both the first neural network circuit (or stage) 322 and the second stage speech restoration neural network 323, the first stage acoustic interference suppression may result in an unrealistic suppression level for the data that the second stage will encounter, i.e. the residual noise level following the first stage acoustic interference suppression may be too low. Thus, in some examples of the invention, it is envisaged that the training dataset used for training the first neural network circuit (or stage) 322 may be configured to be different to the training dataset used for training the second stage speech restoration neural network 323.

[058] In some examples, the training scheme of the second stage speech restoration neural network 323 may be configured as essentially the same as for the first acoustic interference suppression neural network circuit (or stage) 322, with the exception of using the output 312 of the first acoustic interference suppression neural network circuit (or stage) 322 as input to the second feature extraction circuit 317 and the usage of the MSE loss function, as shown in equ.[5]:

$$J_{\text{MSE}}(\Theta^{(2)}) = \frac{1}{|\mathcal{B}_2|} \sum_{\ell' \in \mathcal{B}_2} \|\hat{\mathbf{S}}_{\ell'}^{(2)} - \mathbf{S}_{\ell'}^{(2)}\|_2^2 \quad [5].$$

25 [059] With:

$\hat{\mathbf{S}}_{\ell'}^{(2)}$  and  $\mathbf{S}_{\ell'}^{(2)}$  being the network output representation of estimated speech and the matching representation of clean speech respectively; and

$\mathcal{B}_2$  and  $\Theta^{(2)}$  being the set of frame indices for the current mini-batch

and the trainable parameters for the second speech restoration neural network circuit (or stage) 323 respectively.

[060] In some examples, as mentioned above, the vector representations  $\hat{\mathbf{S}}_{\ell'}^{(2)}$  and  $\mathbf{S}_{\ell'}^{(2)}$  may be configured to contain real and imaginary parts of spectral bins resulting in a complex MSE loss. Alternatively, in some examples, the vector representations  $\hat{\mathbf{S}}_{\ell'}^{(2)}$  and  $\mathbf{S}_{\ell'}^{(2)}$  may be configured to contain spectral magnitudes resulting in a magnitude MSE loss. Alternatively, in other examples, other loss functions may be used, such as the magnitude or complex variant of the mean absolute error (MAE) loss function, as shown in equ.[6]:

$$J_{\text{MAE}}(\Theta^{(2)}) = \frac{1}{|\mathcal{B}_2|} \sum_{\ell' \in \mathcal{B}_2} \|\hat{\mathbf{S}}_{\ell'} - \mathbf{S}_{\ell'}^{(2)}\|_1 \quad [6].$$

[061] A further example option for the second speech restoration neural network circuit (or stage) 323 is to use a generative adversarial network (GAN) loss function. As described above, the second speech restoration neural network circuit (or stage) 323 may be configured so as to not directly compare the network outputs to a target clean speech representation. Instead, an adversarial training scheme may be employed in some examples, where an additional binary classification network (referred to as a ‘discriminator’) with the trainable parameters  $\Theta_D^{(2)}$  may be employed to distinguish between data of the target domain  $\mathbf{S}_{\ell'}^{(2)}$  and the output data  $\hat{\mathbf{S}}_{\ell'}^{(2)}$  of the second speech restoration neural network circuit (or stage) 321, where the frame index is for mini-batches used during training (i.e. a subset of the whole data grouped during neural network training). During training with a GAN loss, separate update steps may be employed for the newly defined binary classifier, and the second neural network circuit (or stage) 321, which are based on the losses computed according to, equ.[7] and equ.[8]:

$$J_{D,\ell'}(\Theta_D^{(2)}) = \frac{1}{2}(D(\mathbf{S}_{\ell'}^{(2)}) - 1)^2 + \frac{1}{2}(D(\hat{\mathbf{S}}_{\ell'}^{(2)}))^2 \quad [7],$$

and

$$J_{G,\ell'}(\Theta^{(2)}) = \frac{1}{2}(D(\hat{\mathbf{S}}_{\ell'}^{(2)}) - 1)^2 \quad [8].$$

[062] Where:

$D(\cdot)$  is the output of the classifier network, representing the probability of its input signal belonging to the target domain.

[063] In a first step, the classifier of the second neural network circuit (or stage) 321 is trained, according to  $J_{D,u}$  in equ.[7] in order to distinguish between the clean signal  $\mathbf{s}_u^{(2)}$  and enhanced signal  $\hat{\mathbf{s}}_u^{(2)}$ . The second step trains the second neural network circuit (or stage) 321 according to  $J_{G,u}$  in equ.[8] in order to output an enhanced signal  $\hat{\mathbf{s}}_u^{(2)}$ , that cannot be distinguished from the target clean speech  $\mathbf{s}_u^{(2)}$  by the discriminator. It is envisaged that this example training scheme is one of many envisaged alternative ways of employing GAN loss functions, with minor changes to the training scheme described above, that can also be used for training of the second neural network circuit (or stage) 321.

[064] Again, in some examples, it is envisaged that the second neural network circuit (or stage) 321 may also consist of two or more equal or unequal networks (referred to as 'subnetworks'), in order to achieve a better acoustic interference attenuation. This may mean that, in the second neural network circuit (or stage) 321, only one final loss function is defined. In other example, a loss function may be defined after each of the 'subnetworks'. In some examples, it is envisaged that such subnetworks may be trained one after the other, or they may be trained jointly.

[065] Referring now to FIG. 4, a block diagram 400 illustrates a joint neural network training scheme, to jointly train a first acoustic interference suppression neural network circuit (or stage) and second speech restoration neural network circuit (or stage), according to some examples of the present invention. One example of a mechanism to perform joint training of the neural networks for both circuits or stages, includes a feature extraction circuit 402 configured to extract one or more feature(s) of a noisy speech signal training data 401. A normalization of the extracted features is performed in first normalization circuit 403 using, say, a vector of means 404 and/or a vector of standard deviations 405, extracted from a storage element 430. In some examples, the two vectors may be computed over all samples in the training data set. The resulting normalized features 406 output from normalization circuit

403 may be used as an input to the joint, two-stage, neural network circuit (or stage) 416. The joint, two-stage, neural network circuit (or stage) 416 includes a first neural network circuit 408, configured to receive the first pre-trained parameters 407.

[066] Inside of the joint, two-stage, neural network circuit (or stage) 416, the processing operations then include an estimation 410 of the mask vector  $\hat{\mathbf{m}}_\ell$  being performed in masking circuit 409, following the mask vector  $\hat{\mathbf{m}}_\ell$  being output from the first neural network circuit 408, resulting in  $\hat{\mathbf{S}}_\ell^{(1)} = \hat{\mathbf{m}}_\ell \odot \mathbf{Y}_\ell$ . The estimation 410 of the mask vector  $\hat{\mathbf{m}}_\ell$  is then input to a second normalization circuit 417 using, say, mean and variance value statistics 411, 412 computed during separate training, and processing of the second-stage neural network circuit 415. The second-stage neural network circuit 415 is configured to receive second pre-trained parameters 414 and is configured to reconstruct the speech signal. The network parameters  $\Theta^{(1)}$  and  $\Theta^{(2)}$  as defined above, for the first neural network circuit 408 and the second-stage neural network circuit 415, are initialized with the first pre-trained parameters 407 and the configured to receive the second pre-trained parameters 414, notably obtained from separate training of the respective networks in this example.

[067] In a parallel signal processing path to the joint, two-stage, neural network circuit (or stage) 416, the targets for backpropagation training of the joint, two-stage, neural network circuit (or stage) may be computed from a framed frequency domain representation of the corresponding clean speech signal 418. This computation may be performed during target preparation 419, which is configured to restructure the clean speech signal data 418 into a suitable vector representation form 420 for network learning, which matches the dimension and domain of the desired neural network output. In one example embodiment, training is performed in the target preparation circuit 419 using standard backpropagation with the mean spectral approximation (MSA) loss function. Depending on the used loss function, either only the target representation 421 from second-stage training or only the target representation 420 from first-stage training is used. In another example, target representation from joint training may be used) Accordingly, the loss function can be either  $J_{\text{MSE}}(\Theta^c)$  where  $\Theta^c$  represents all trainable parameters 416 of, or a combined loss, as shown in equ.[9]:

[068]  $J_{\text{combined}}(\Theta^c) = \beta J_{\text{MSA}}(\Theta^{(1)}) + (1 - \beta) J_{\text{MSE}}(\Theta^c)$  [9],

including the intermediate loss  $J_{\text{MSA}}(\Theta^{(1)})$  and the weighting parameter  $\beta$ .

[069] In other example embodiments, it is envisaged that it is also possible to use the described joint training procedure either with  $J_{\text{MSE}}(\Theta^c)$  or  $J_{\text{combined}}(\Theta^c)$  using a random initialization for  $\Theta^c$ , and excluding the normalization circuit 417 to perform joint training from the outset.

[070] The aforementioned circuit topologies and neural network processing domains, e.g. relating to FIG. 1 for the first neural network circuit (or stage) 102, configured to suppress acoustic interference in a degraded input speech signal 101, and the second neural network circuit (or stage) 103, configured to restore natural sounding speech, have to be chosen according to their specific tasks for the final two-stage enhancement to perform well. That is, the first neural network circuit (or stage) 102 is configured to suppress acoustic interferences in a degraded speech signal, and the second neural network circuit (or stage) 103 is configured to restore natural sounding speech and potentially suppress residual acoustic interferences based on the outputs of the first neural network circuit (or stage) 102.

[071] For acoustic interference suppression of speech signals, and in reference to FIG. 2, the preferred circuit model for the first neural network circuit (or stage) 207 is a long short-term memory (LSTM) network, which in some examples comprises two LSTM layers that enable a modelling of temporal dependencies in the noisy speech and a fully-connected feedforward output layer, as would be understood by a skilled artisan.

[072] Referring now to FIG. 5, a flowchart 500 illustrates an example of a topology for an acoustic interference suppression stage neural network, in accordance with some examples of the invention. The flowchart comprises a series of feedforward layers and a series of long short-term memory (LSTM) network operations. The example of FIG. 5 illustrates a feedforward layer 510, followed by the two LSTM layers 511 and 512, two feedforward layers 513 and 514, employing skip connections 520 and 521, and a final output layer 515. It is envisaged that many alternative topologies may be employed for the acoustic interference suppression stage neural network, with the illustrated example of three feedforward layers

510, 513, 514 and two LSTM network operations 511, 512 just one possible example. Thus, in other examples, any number of feedforward layers and/or LSTM network operations may be employed. In some examples, the LSTM layers 511, 512 may comprise of standard LSTM cells without ‘peephole’ connections, and the feedforward layers 510, 513, 514 may use  
5 rectified linear unit (ReLU) activation functions, and the output layer 515 may use a hyperbolic tangent activation function, for example.

**[073]** In other examples, it is envisioned that types of recurrent neural networks (RNNs) differing from LSTMs can be used for the acoustic interference suppression stage. RNNs are regarded as the class of neural networks that employ internal network states from previous  
10 time steps in a temporal sequence of input data (in addition to the feature vector) for the computation of the network output. This enables RNNs to model temporal dynamics of the underlying sequential data.

**[074]** For performing speech restoration, and in reference to FIG. 2, the preferred circuit model for the second neural network circuit (or stage) 212 is a convolutional neural network  
15 (CNN) that is configured to use an encoder-decoder structure and skip connections, thereby enabling the model to focus on local context in the input features, whilst still being able to use high resolution data in the last model layers.

**[075]** Examples of the invention also described a method of speech processing that includes a first neural network circuit configured to suppress acoustic interference of an input  
20 speech signal and a second neural network circuit configured to restore at least one clean speech component of the input speech signal. In examples of the invention, the acoustic interference suppression aspect of the method includes receiving an input speech signal that comprises clean speech and acoustic interference; suppressing the acoustic interference in the input speech signal by a first neural network circuit; outputting a representation of an  
25 estimate of a clean speech signal to a second neural network circuit. In examples of the invention, the clean speech restoration aspect of the method includes restoring at least one speech component of the clean speech in the input speech signal by the second neural network circuit; and outputting a reconstructed speech signal.

**[076]** FIG. 6 illustrates a flowchart 600 of an example of a topology for a speech  
30 restoration stage neural network, in accordance with some examples of the invention. The

flowchart describes a method for speech processing that includes receiving an acoustic background interference suppressed input speech signal from a first neural network circuit. The second speech restoration stage is illustrated with an encoder 610/decoder 620 structure. The encoder 610 does not necessarily perform only acoustic interference suppression. Likewise the decoder 620 does not necessarily perform speech restoration only. In accordance with examples of the invention, the speech restoration stage neural network restores at least one speech component of the estimated clean speech signal and outputs a reconstructed speech signal.

[077] In particular, in the example of FIG. 6, the dimension of one example of the feature representation for each layer is shown next to the arrows. Convolutional layers are denoted by  $\text{Conv}(N_F, S_1 \times S_2)$  with  $N_F$  being the number of kernels used in the layer and  $S_1$  and  $S_2$  being the kernel sizes on the feature and frame axis respectively. This example embodiment exclusively uses one-dimensional convolutions with  $S_2 = 1$ , as shown, and therefore only needs features extracted from the current frame. In this example, all convolutional layers 652, 654, 658, 660, 664, 668, 670, 676, 678, 682 use leaky rectified linear units (ReLU) activation functions. It is noted in FIG. 6 that convolutional layer 664 is a bottleneck layer between the encoder 610 and decoder 620. In other examples, it is envisaged that alternative options may be used, as would be appreciated by a skilled artisan, such as:  $\tanh()$  = tangens, hyperbolicus, sigmoid functions, exponential linear units (ELUs), linear function, etc ..., with the last layer being an exception using a linear activation function. Further layers that may be utilised include a maximum pooling layers 656, 662 with kernel size  $2 \times 1$  and up-sampling layers 666, 674 that perform up-sampling of a factor 2 on the feature axis. Although this example illustrates one or more maximum pooling layers 656, 662 it is envisaged that the one or more pooling layers may be non-maximum, for example average pooling layer(s), in other implementations. Although this example illustrates one or more up-sampling layers 666, 674 it is envisaged that inverse convolution layers with a stride larger than one may be employed in other implementations.

[078] Referring now to FIG. 7, a flowchart 700 illustrates an alternative example of a topology for a second stage neural network, which can be employed to restore at least one speech component, in accordance with some examples of the invention. In this alternative

example, the encoder-decoder structure of FIG. 6 is replicated, with the exception of including an LSTM-based layer 764 in a bottleneck layer between the encoder 710 and decoder 720. Thus, the identical operations to FIG. 6 will not be further explained, to avoid obfuscating the invention. In some examples of the invention, the LSTM-based layer 764 may be a convolutional LSTM layer 764, denoted by ConvLSTM (or for other examples a fully-connected LSTM layer), and may account for temporal dependencies in the estimation task of the second-stage network 700. It is envisaged that in some examples, other convolutional layers in 700, e.g., 778 or 782, may be exchanged with LSTM-based layers in a similar manner.

10 **[079]** In some example, the preferred processing domain for both stages of the invention is the frequency domain, using spectral magnitudes, as shown in equ.[10]:

$$\hat{\mathbf{x}}_{\ell}^{(1)} = [|Y_{\ell}(0)|, |Y_{\ell}(1)|, \dots, |Y_{\ell}(K/2)|]^T \quad [10],$$

as features for the first neural network circuit (or stage) 207, 311, 408, where  $K$  is the FFT size. In other examples, it is envisaged that alternative exemplary options for feature

15 representations may be used, such as real and imaginary part of the spectrum and Mel-frequency-based spectral magnitudes. The output domain of the first neural network circuit (or stage) 207, 311, 408 is chosen to consist of real and imaginary parts of the spectral bins to enable the network to estimate the clean speech phase and provide even higher acoustic interference attenuation. This is achieved by setting the target representation to, as shown in equ.[11]:

$$\mathbf{S}_{\ell}^{(1)} = [\Re\{S_{\ell}(0)\}, \Re\{S_{\ell}(1)\}, \dots, \Re\{S_{\ell}(K/2)\}, \Im\{S_{\ell}(2)\}, \Im\{S_{\ell}(2)\}, \dots, \Im\{S_{\ell}(K/2-1)\}]^T \in \mathbb{R}^K$$

...[11],

constructing  $\mathbf{Y}_{\ell}$  accordingly from  $Y_{\ell}(k)$  and setting the output size of the first neural network circuit (or stage) 207, 311, 408 to  $K$ .

25 **[080]** In some examples, the second neural network circuit (or stage) 212, 321, 415 may also operate on real and imaginary parts of spectral bins, and accordingly the respective feature extraction circuit 210, 317 can be omitted, setting:

$$\hat{\mathbf{x}}_{\ell}^{(2)} = \hat{\mathbf{S}}_{\ell}^{(1)} = [\Re\{\hat{S}_{\ell}^{(1)}(0)\}, \Re\{\hat{S}_{\ell}^{(1)}(1)\}, \dots, \Re\{\hat{S}_{\ell}^{(1)}(K/2)\}, \Im\{\hat{S}_{\ell}^{(1)}(1)\}, \Im\{\hat{S}_{\ell}^{(1)}(2)\}, \dots, \Im\{\hat{S}_{\ell}^{(1)}(K/2-1)\}]^T \in \mathbb{R}^K$$

...[12] and

$$\mathbf{S}_f^{(2)} = \mathbf{S}_f^{(1)} \quad [13].$$

[081] Alternatively, the real and imaginary parts of the spectral bins can be presented to the second neural network circuit (or stage) 212, 321, 415 as separate feature maps for real and imaginary parts if the preferred CNN topology is chosen for this neural network.

5 [082] It is envisaged that the first and second stages of the two-stage neural network-based speech enhancement circuit 100, configured to provide acoustic interference suppression and speech reconstruction, may be implemented in hardware, software or firmware. Furthermore, it is envisaged that the neural network circuits employed by the acoustic interference suppression and speech reconstruction circuits may be performed in  
10 firmware, such as using field programmable logic/gate arrays (FPGAs), or programmed in software using a digital signal processor or may be implemented in hardware, although a firmware or software-based approach is typically the most economical. Also for example, in one embodiment, the illustrated examples may be implemented as circuitry located on a single integrated circuit or on an integrated circuit within a single device, such as a speech  
15 processing system. Alternatively, the examples may be implemented as any number of separate integrated circuits or separate devices interconnected with each other in a suitable manner.

[083] In this document, the terms ‘computer program product’ ‘computer-readable medium’ and the like may be used generally to refer to media such as, for example, memory,  
20 storage device, or storage unit. These and other forms of computer-readable media may store one or more instructions for use by a processor, to cause the speech processor to perform specified operations. Such instructions, generally referred to as ‘computer program code’ (which may be grouped in the form of computer programs or other groupings), when executed, enable the speech processing system 100 to perform functions of embodiments of  
25 the present invention. Note that the code may directly cause the speech processor to perform specified operations, be compiled to do so, and/or be combined with other software, hardware, and/or firmware elements (e.g., software libraries for performing standard functions) to do so. In an embodiment where the elements are implemented using software, the software may be stored in a computer-readable medium and loaded into speech

processing system 100 using, for example, removable storage unit, a drive or a communications interface. The control logic (in this example, software instructions or computer program code), when executed by the speech processor, causes the speech processor to perform the functions of the invention as described herein. In examples of the invention, the computer program product may be a non-transitory tangible computer program product comprising executable code stored therein for speech processing, wherein the code is operable for, when executed at a speech processor: receiving an input speech signal that comprises clean speech and acoustic interference; suppressing the acoustic interference in the input speech signal by a first neural network circuit; outputting a representation of an estimate of a clean speech signal to a second neural network circuit; restoring at least one speech component of the clean speech in the input speech signal by the second neural network circuit; and outputting a reconstructed speech signal.

**[084]** Although examples of the invention have been described with reference to two neural networks, a first configured for acoustic interference suppression followed by a second configured for speech restoration, it is envisaged that examples of the invention may be applied to all kinds of acoustic interferences, e.g. environmental noise, acoustic echo/feedback and reverberation or residuals of those acoustic interferences after preliminary enhancement. Therefore, in some examples, the neural networks may be applied in products using acoustic interference reduction techniques, for the mobile, automotive, wearable and Internet of Things (IoT) technology space. It is also envisaged that examples of the invention may also be applied in the preprocessing for automatic speech recognition systems, operating in noisy environments, such as in the automotive context, mobile devices or smart speaker systems.

**[085]** A skilled artisan will appreciate that the level of integration of circuits or components may be, in some instances, implementation-dependent.

**[086]** In the foregoing specification, the invention has been described with reference to specific examples of embodiments of the invention. It will, however, be evident that various modifications and changes may be made therein without departing from the scope of the invention as set forth in the appended claims and that the claims are not limited to the specific examples described above. In particular, examples of the invention describe an integrated

circuit with an input that receives an input speech signal; and two respective neural networks. A first neural network circuit is configured to suppress input speech corrupted by acoustic interference and output an estimated clean speech signal; and a second neural network circuit is configured to receive the estimated clean speech signal and restore at least one speech component and output a reconstructed speech signal.

**[087]** Furthermore, because the illustrated embodiments of the present invention may, for the most part, be implemented using electronic components and circuits known to those skilled in the art, details will not be explained in any greater extent than that considered necessary as illustrated above, for the understanding and appreciation of the underlying concepts of the present invention and in order not to obfuscate or distract from the teachings of the present invention.

**[088]** Those skilled in the art will recognize that the boundaries between logic blocks are merely illustrative and that alternative embodiments may merge logic blocks or circuit elements or impose an alternate decomposition of functionality upon various logic blocks or circuit elements. Thus, it is to be understood that the architectures depicted herein are merely exemplary, and that in fact many other architectures can be implemented that achieve the same functionality.

**[089]** Any arrangement of components to achieve the same functionality is effectively 'associated', such that the desired functionality is achieved. Hence, any two components herein combined to achieve a particular functionality can be seen as being 'associated with' each other, such that the desired functionality is achieved, irrespective of architectures or intermediary components. Likewise, any two components so associated can also be viewed as being 'operably connected,' or 'operably coupled,' to each other to achieve the desired functionality.

**[090]** Furthermore, those skilled in the art will recognize that boundaries between the above described operations are merely illustrative. The multiple operations may be executed at least partially overlapping in time. Moreover, alternative example embodiments may include multiple instances of a particular operation, and the order of operations may be altered in various other embodiments.

**[091]** Also for example, in one embodiment, the illustrated examples may be

implemented as circuitry located on a single integrated circuit or within a same device. Alternatively, the examples may be implemented as any number of separate integrated circuits or separate devices interconnected with each other in a suitable manner.

[092] Also, examples of the invention are not limited to units implemented in non-programmable hardware but can also be applied in wireless programmable devices or units able to perform the desired device functions by operating in accordance with suitable program code. However, other modifications, variations and alternatives are also possible. The specifications and drawings are, accordingly, to be regarded in an illustrative rather than in a restrictive sense.

10 [093] In the claims, any reference signs placed between parentheses shall not be construed as limiting the claim. The word 'comprising' does not exclude the presence of other elements or steps than those listed in a claim. Furthermore, the terms 'a' or 'an,' as used herein, are defined as one, or more than one. Also, the use of introductory phrases such as 'at least one' and 'one or more' in the claims should not be construed to imply that the  
15 introduction of another claim element by the indefinite articles 'a' or 'an' limits any particular claim containing such introduced claim element to inventions containing only one such element, even when the same claim includes the introductory phrases 'one or more' or 'at least one' and indefinite articles such as 'a' or 'an.' The same holds true for the use of definite articles. Unless stated otherwise, terms such as 'first' and 'second' are used to arbitrarily  
20 distinguish between the elements such terms describe. Thus, these terms are not necessarily intended to indicate temporal or other prioritization of such elements. The mere fact that certain measures are recited in mutually different claims does not indicate that a combination of these measures cannot be used to advantage.

## CLAIMS

What is claimed is:

1. A speech processing system (100, 200) comprising:  
an input that receives an input speech signal (101, 201) that comprises clean  
5 speech and acoustic interference;  
a first neural network circuit (102, 216) configured to suppress the acoustic  
interference in the input speech signal (101, 201) and output (120) a representation of  
an estimate of a clean speech signal (209); and  
a second neural network circuit (103, 217) configured to receive the  
10 representation of the estimated clean speech signal (209) and restore at least one  
speech component of the clean speech in the input speech signal (101, 201) and output  
a reconstructed speech signal (104, 215).
2. The speech processing system (100, 200) Claim 1, wherein the second  
neural network circuit (103, 217) is configured to additionally suppress residual  
15 acoustic interference contained in the representation of the estimated clean speech  
signal (209).
3. The speech processing system (100, 200) of any preceding Claim,  
wherein the first neural network circuit (102, 216) comprises a recurrent neural  
network, RNN, topology.
- 20 4. The speech processing system (100, 200) of Claim 3 wherein the first  
neural network circuit (102, 216) comprises at least one long short-term memory,  
LSTM, layer.
5. The speech processing system (100, 200) of Claim 4 wherein the first  
neural network circuit (102, 216) comprises a combination of at least one LSTM,  
25 layer (511, 512) and at least one feedforward layer (510, 513, 514).
6. The speech processing system (100, 200) of Claim 4 or Claim 5 wherein  
at least one layer (510, 513, 514) of the first neural network circuit (102, 216)  
comprises at least one skip connection (520, 521).

7. The speech processing system (100, 200) of any preceding Claim, wherein the second neural network circuit (103, 217) comprises an encoder (610, 710) - decoder (620, 720) topology that includes at least one bottleneck layer (664, 764).

8. The speech processing system (100, 200) of Claim 7 wherein the second  
5 neural network circuit (103, 217) comprises a convolutional neural network, CNN, (600) topology that comprises at least one convolutional layer (652) and at least one of: at least one pooling layer, at least one re-sampling layer, at least one skip connection.

9. The speech processing system (100, 200) of Claim 7 or Claim 8 wherein  
10 the bottleneck layer (764) comprises at least one of: at least one long short-term memory, LSTM, bottleneck layer, at least one convolutional LSTM bottleneck layer.

10. The speech processing system (100, 200) of any preceding Claim, wherein the input is configured to receive at least one of: at least one auxiliary input signal (105) that comprises acoustic interference or an acoustic interference estimate;  
15 at least one additional microphone signal (260).

11. The speech processing system (100, 200) of any preceding Claim, further comprising at least one pre-processing circuit coupled to the input and configured to provide the input speech signal (101, 201) with reduced interference to the first neural network circuit (102, 216).

20 12. The speech processing system (100, 200) of any preceding Claim, wherein the first neural network circuit (102, 216) and second neural network circuit (103, 217) are subjected to joint training of both circuits combined.

13. The speech processing system (100, 200) of any preceding Claim, wherein the second neural network circuit (103) is configured to additionally receive  
25 the input speech signal (101) as an additional input in generating an output representation of an enhanced speech signal (104).

14. A method for speech processing comprising:

receiving an input speech signal (101, 201) that comprises clean speech and acoustic interference;

30 suppressing the acoustic interference in the input speech signal (101, 201) by a

first neural network circuit (102, 216);

outputting (120) a representation of an estimate of a clean speech signal (209)  
to a second neural network circuit (103, 217);

restoring at least one speech component of the clean speech in the input speech  
5 signal (101, 201) by the second neural network circuit (103, 217); and

outputting a reconstructed speech signal.

15. A non-transitory tangible computer program product comprising  
executable code stored therein for speech processing, wherein the code is operable for,  
when executed at a speech processor:

10 receiving an input speech signal (101, 201) that comprises clean speech and  
acoustic interference;

suppressing the acoustic interference in the input speech signal (101, 201) by a  
first neural network circuit (102, 216);

outputting (120) a representation of an estimate of a clean speech signal (209)  
15 to a second neural network circuit (103, 217);

restoring at least one speech component of the clean speech in the input speech  
signal (101, 201) by the second neural network circuit (103, 217); and

outputting a reconstructed speech signal.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/CN2020/080959

<b>A. CLASSIFICATION OF SUBJECT MATTER</b>		
G10L 21/02(2013.01)i; G10L 25/30(2013.01)i; G10L 15/16(2006.01)i; G10L 17/18(2013.01)i		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b>		
Minimum documentation searched (classification system followed by classification symbols)		
G10L		
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)		
CNABS, CNTXT, VEN, IEEE, CNKI: speech, voice, enhance, process, neural, network, CNN, ANN, DNN, clean, restore, reconstruct, interference, noise		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	ZHAO, Yan et. al. "Two-Stage Deep Learning for Noisy-Reverberant Speech Enhancement" <i>IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING</i> , Vol. 27, No. 1, 31 January 2019 (2019-01-31), the introduction and algorithm description, pages 53-56, Fig. 1	1-15
X	ZHAO, Yan et. al. "A two-stage algorithm for noisy and reverberant speech enhancement" <i>2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , 31 December 2017 (2017-12-31), pages 5580 to 5582	1-15
A	ZHAO, Ziyue et. al. "A CNN postprocessing to enhance coded speech" <i>International Workshop on Acoustic Signal Enhancement (IWAENC2018)</i> , 30 September 2018 (2018-09-30), pages 406 to 409	1-15
A	CN 105261359 A (UNIV. NANJING NORMAL) 20 January 2016 (2016-01-20) the whole document	1-15
A	CN 107845389 A (UNIV. BEIJING TECHNOLOGY) 27 March 2018 (2018-03-27) the whole document	1-15
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> 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		Date of mailing of the international search report
15 June 2020		24 June 2020
Name and mailing address of the ISA/CN		Authorized officer
National Intellectual Property Administration, PRC 6, Xitucheng Rd., Jimen Bridge, Haidian District, Beijing 100088 China		ZHANG, Yaling
Facsimile No. (86-10)62019451		Telephone No. 86-(010)-62085722

Form PCT/ISA/210 (second sheet) (January 2015)

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/CN2020/080959

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	CN 104952448 A (ZHANG, Aiyong et, al.) 30 September 2015 (2015-09-30) the whole document	1-15
A	CN 108447495 A (TIANJIN UNIV.) 24 August 2018 (2018-08-24) the whole document	1-15

**INTERNATIONAL SEARCH REPORT**  
**Information on patent family members**

International application No.

**PCT/CN2020/080959**

Patent document cited in search report			Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
CN	105261359	A	20 January 2016	CN	105261359	B	09 November 2018
CN	107845389	A	27 March 2018	None			
CN	104952448	A	30 September 2015	None			
CN	108447495	A	24 August 2018	None			

Form PCT/ISA/210 (patent family annex) (January 2015)