

(19)



(11)

EP 3 497 630 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent:
04.11.2020 Bulletin 2020/45

(51) Int Cl.:
G06N 3/04 (2006.01) G10L 15/16 (2006.01)

(21) Application number: **17780544.7**

(86) International application number:
PCT/US2017/050335

(22) Date of filing: **06.09.2017**

(87) International publication number:
WO 2018/048945 (15.03.2018 Gazette 2018/11)

(54) PROCESSING SEQUENCES USING CONVOLUTIONAL NEURAL NETWORKS

VERARBEITUNGSSEQUENZEN MIT VERWENDUNG VON NEURONALEN KONVOLUTIONSNETZWERKEN

SÉQUENCES DE TRAITEMENT UTILISANT DES RÉSEAUX NEURONAUX CONVOLUTIFS

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR

- **VINYALS, Oriol**
London N1C 4AG (GB)
- **ESPEHOLT, Lasse**
London N1C 4AG (GB)

(30) Priority: **06.09.2016 US 201662384123 P**

(74) Representative: **Martin, Philip John**
Marks & Clerk LLP
62-68 Hills Road
Cambridge
CB2 1LA (GB)

(43) Date of publication of application:
19.06.2019 Bulletin 2019/25

(60) Divisional application:
20192441.2

(56) References cited:
CA-A1- 2 810 457

(73) Proprietor: **Deepmind Technologies Limited**
London EC4A 3TW (GB)

- **Aäron Van Den Oord ET AL: "Pixel Recurrent Neural Networks", , 19 August 2016 (2016-08-19), pages 1-11, XP055435831, Retrieved from the Internet:
URL:https://arxiv.org/pdf/1601.06759.pdf [retrieved on 2017-12-18]**
- **AARON VAN DEN OORD ET AL: "WaveNet: A Generative Model for Raw Audio", ARXIV.ORG, CORNELL UNIVERSITY LIBRARY, 201 OLIN LIBRARY CORNELL UNIVERSITY ITHACA, NY 14853, 12 September 2016 (2016-09-12), XP080725972,**

(72) Inventors:
• **VAN DEN OORD, Aaron Gerard Antonius**
London N1C 4AG (GB)
• **DIELEMAN, Sander Etienne Lea**
London N1C 4AG (GB)
• **KALCHBRENNER, Nal Emmerich**
London N1C 4AG (GB)
• **SIMONYAN, Karen**
London N1C 4AG (GB)

EP 3 497 630 B1

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

BACKGROUND

[0001] This specification relates to processing sequences using neural networks.

[0002] Neural networks are machine learning models that employ one or more layers of nonlinear units to predict an output for a received input. Some neural networks include one or more hidden layers in addition to an output layer. The output of each hidden layer is used as input to the next layer in the network, i.e., the next hidden layer or the output layer. Each layer of the network generates an output from a received input in accordance with current values of a respective set of parameters.

SUMMARY

[0003] This specification describes how a system implemented as computer programs on one or more computers in one or more locations can generate a neural network output by processing a sequence including multiple inputs.

[0004] In one innovative aspect a neural network system implemented by one or more computers is configured to receive an input sequence comprising a plurality of inputs and to generate a neural network output from the input sequence. The neural network system may comprise a convolutional subnetwork comprising a plurality of dilated convolutional neural network layers and an output subnetwork. The convolutional subnetwork may be configured to, for each of the plurality of inputs: receive a current input sequence comprising the input and the inputs that precede the input in the input sequence, and process the current input sequence to generate an alternative representation for the input. This alternative representation may thus comprise a numeric representation, i.e., an ordered collection of numeric values, in which the current input sequence has been encoded by the convolutional subnetwork, for example encoding features of the current sequence. The output subnetwork may be configured to receive the alternative representations of the inputs and to process the alternative representations to generate the neural network output.

[0005] Some of the many advantages of such a system are described later.

[0006] The neural network output can comprise a plurality of sets of scores with each set of scores including a respective score for each of a plurality of possible outputs. The system can select outputs in accordance with the sets of scores to generate an output sequence comprising a plurality of outputs.

[0007] In some implementations the convolutional neural network layers are dilated causal convolutional neural network layers. As described in more detail later, a dilated convolutional neural network layer applies a convolution to non-adjacent values in a sequence, i.e., as defined by the outputs from a previous layer. This can

increase the receptive field of the convolutional subnetwork by orders of magnitude whilst preserving the input (time) resolution and maintaining computational efficiency.

[0008] In some implementations the convolutional neural network layers include multiple stacked blocks of dilated convolutional neural network layers. Each block may comprise multiple dilated convolutional neural network layers with increasing dilation. For example the dilation may be increased by a factor n for each successive layer up to a limit within each block. This can further increase the receptive field size.

[0009] In some implementations one or more of the convolutional neural network layers may have gated activation units. For example a rectified linear or other unit following a convolution implemented by a layer may be replaced by a gated activation unit. In a gated activation unit the output may be a combination of two (causal) convolutions, a main convolution and a gate convolution. The convolutions may each be applied to some or all of the same outputs from the previous layer. The combination may involve a non-linear activation function applied to the gate convolution, for example an activation with a (0,1) range such as a sigmoid. This may then multiply a value from the main convolution; a non-linear activation function may, but need not be, applied to the main convolution. Such an approach may assist in capturing more complex structure within the data.

[0010] The convolutional subnetwork may comprise residual connections, for example a connection from an input of a convolutional layer to a summer to sum this with an intermediate output of the layer. This effectively allows the network to be trained to skip or partially skip a layer, thus speeding up convergence and facilitating training of deeper models. The convolutional subnetwork may additionally or alternatively comprise skip connections, for example directly from each of one or more intermediate layers of the convolutional subnetwork to one or more operations that directly generate the alternative representation that is provided to the output layer.

[0011] In some implementations, e.g., implementations where the system is used for speech recognition, the output subnetwork comprises: a mean pooling layer configured to: receive a subset of the alternative representations, and perform mean pooling on the subset of alternative representations to generate a mean pooled representation for the subset; and one or more additional neural network layers configured to, for each subset of the alternative representations, receive the mean pooled representation for the subset and to generate a set of scores for a position in the output sequence corresponding to the subset.

[0012] In some implementations processing the current sequence using the convolutional subnetwork, to generate an alternative representation for an input, reuses values computed for previous time steps. The reused values may comprise values derived from application of a convolutional filter to the input data or data de-

rived therefrom. The re-used values may be stored at one time step and retrieved at a later time step when the same filter is applied to the same (or some of the same) input data or data derived therefrom. This can make the system computationally more efficient and hence faster, because there is no need to re-compute the stored values.

[0013] Particular examples of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. The neural network system as described in this specification can achieve results that outperform the state of the art on a variety of sequence processing tasks, e.g., speech recognition. By using convolutional neural network layers, e.g., dilated convolutional layers, instead of recurrent neural network layers, the neural network system can achieve these results while not needing as many computational resources to train as other systems that do include recurrent neural network layers, resulting in a reduced training time. By using convolutional neural network layers, e.g., causal convolutional layers, instead of recurrent neural network layers, e.g., instead of long short-term memory (LSTM) layers, the neural network system can achieve these advantageous results while not needing as many computational resources to train as other systems that do include recurrent neural network layers, resulting in a reduced training time. By employing convolutional layers rather than recurrent layers, the computation of the neural network system can be more easily batched and more easily parallelized, e.g., because the layers of the network do not have to be unrolled for each time step, allowing the computation of the system to be performed more efficiently. Additionally, by employing dilated causal convolutional layers, the receptive field of the convolutional subnetwork and, therefore, the quality of the output sequences generated by the system, can be improved without greatly increasing the computational cost of generating the output sequence. Moreover, the paths traversed by forwards and backward signals in the system between inputs and outputs are independent of the sequence distance between the inputs and outputs, which allows better signal propagation and long-range dependency learning, thereby improving output quality.

[0014] The details of one or more embodiments of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims. The embodiments of the invention are only defined by the subject-matter of the claims.

[0015] CA 2 810 457 A1 discloses the mitigation of acoustic variations of a speech signal using a convolutional neural network (CNN). Several frequency bands are input to a convolution layer of a CNN, the convolution layer including several convolution units that receive the input from the frequency bands. The output of the con-

volution layer is provided to a pooling layer of the CNN, the pooling layer being provided with several pooling units that are provided to receive the input from the convolution unit. The output of the pooling layer is a representation of an acoustic signal that is provided to mitigate the acoustic variations.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016]

FIG. 1 shows an example neural network system.
 FIG. 2 shows a visualization of an example block of dilated causal convolutional layers.
 FIG. 3 shows an example architecture for the convolutional subnetwork.
 FIG. 4 is a flow diagram of an example process for processing a current input sequence.

[0017] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0018] FIG. 1 shows an example neural network system 100. The neural network system 100 is an example of a system implemented as computer programs on one or more computers in one or more locations, in which the systems, components, and techniques described below can be implemented.

[0019] The neural network system 100 is a system that performs a sequence processing task, i.e., generates an output sequence 152 from an input sequence 102. The system 100 can be configured to perform any of a variety of sequence processing tasks.

[0020] For example, the system 100 can be configured to perform a speech recognition task, where the input sequence is a sequence of audio data, e.g., raw audio data, and the system 100 generates a sequence of phonemes that is a phoneme representation of the input sequence or a sequence of graphemes that is a grapheme representation of the input sequence.

[0021] As another example, the system 100 can be configured to perform a language modelling task, where the input sequence is a sequence of phonemes or graphemes and the output sequence generated by the system 100 is a sequence of words in a first natural language that represents the input sequence.

[0022] As another example, the system 100 can be configured to perform a machine translation task, where the input sequence is a sequence of words in a first natural language and the output sequence generated by the system 100 is a sequence of words in a second natural language that is a translation of the input sequence into the second natural language.

[0023] As another example, the system 100 can be configured to perform a natural language processing task, e.g., sentence summarization, sentence parsing,

part of speech tagging, and so on, where the input sequence is a sequence of words in a first natural language and the output sequence generated by the system 100 is a sequence that classifies the words in the input sequence according to the kind of output required by the natural language processing task.

[0024] As another example, the system 100 can be configured to perform a compression task, where the input sequence is a sequence of data drawn from an input sample, e.g., audio data or image data, and the output sequence is a compressed or encoded sample.

[0025] As another example, the system 100 can be configured to perform a medical analysis task, where the input sequence is a sequence of health data for a particular patient, and the output sequence includes data that characterizes the health of the patient or predicts the future health of the patient.

[0026] More specifically, the input sequence 102 includes a respective input at each of multiple input time steps and the neural network system 100 generates the output sequence 152 by processing a current input sequence at some or all of the time steps. The current input sequence at a given time step includes the input at the given time step and the inputs that precede the input in the input sequence. For example, for time step 140 in input sequence 102, the current input sequence 142 includes the input at time step 140 and the inputs at time steps before time step 140 in the input sequence 102.

[0027] In some implementations, the neural network system 100 processes a respective current input sequence for each time step in the input sequence. In some other implementations, the neural network system 100 does not process current input sequences for one or more initial inputs in the input sequence, i.e., the system 100 waits until a specified number of inputs have been seen before beginning processing.

[0028] The neural network system 100 includes a convolutional subnetwork 110 and an output subnetwork 120.

[0029] At a given time step during the generation of an output sequence, the convolutional subnetwork 110 is configured to receive the current input sequence, i.e., the sequence that includes the inputs at and before the given time step, and to process the current input sequence to generate an alternative representation for the time step. For example, for the time step 140 in the input sequence 102, the convolutional subnetwork 110 can receive the current input sequence 142 and process the current input sequence 142 to generate an alternative representation 144.

[0030] The output subnetwork 120 is configured to receive the alternative representations for each of the time steps and to generate, from the alternative representations, a respective score distribution over possible outputs for each position in the output sequence. For example, the output subnetwork can generate a score distribution 146 for a particular position in the output sequence 152. Each score distribution includes a respective score

for each of multiple possible outputs and the score distributions generated by the output subnetwork 120 during generation of the output sequence 152 are collectively referred to as the neural network output for the input sequence 102.

[0031] In particular, when the neural network system 100 is configured to generate a sequence of phonemes, each score distribution includes a respective score for each phoneme in a vocabulary of phonemes. When the neural network system 100 is configured to generate a sequence of graphemes, the score distribution includes a respective score for each grapheme in a vocabulary of graphemes. When the neural network system 100 is configured to generate a sequence of words, the score distribution includes a respective score for each word in a vocabulary of words. In some implementations, the corresponding vocabulary, i.e., the vocabulary of graphemes, phonemes, or words, includes one or more special outputs, e.g., a blank output that indicates that the next output in the output sequence is not yet ready to be generated.

[0032] Once the output subnetwork 146 has generated a score distribution for a given position in the output sequence, the neural network system 100 can select an output to be included in the output sequence at the position from the multiple possible outputs in accordance with the score distribution. For example, the neural network system 100 can select an output by sampling from the score distribution, i.e., sampling from the possible outputs in accordance with the scores in the score distribution, or can select the possible outputs having the highest score according to the score distribution.

[0033] The convolutional subnetwork 110 generally includes multiple convolutional neural network layers. More specifically, the convolutional neural network layers include multiple causal convolutional layers.

[0034] A causal convolutional layer is a convolutional layer that operates on an input sequence that has a respective input at each of multiple time steps by, for each time step, generating an output that depends only on the inputs at the time step and at the time steps before the time step in the input sequence, i.e., and not on any inputs at any time steps after the time step in the input sequence. In some cases, the causal convolutional layers are implemented by applying a normal convolution and then shifting each output of the normal convolution by a few time steps, i.e., shifting each output forward by (filter length - 1) time steps, prior to applying the activation function for the convolutional layer, where "filter length" is the length of the filter of the convolution that is being applied.

[0035] To increase the receptive field of the causal convolutional layers without requiring an excessive number of layers or filters of excessive length, some or all of the convolutional layers can be dilated causal convolutional layers. A dilated convolution is a convolution where the filter is applied over an area larger than its length by skipping input values with a certain step that is defined by the dilation value for the dilated convolution. By incorpo-

rating dilated causal convolutions, the convolutional neural network layers effectively operate on their inputs with a coarser scale than with a normal convolution.

[0036] In some implementations, the convolutional neural network layers include a stack of multiple blocks of dilated causal convolutional layers. Each block in the stack can include multiple dilated convolutional neural network layers with increasing dilation. For example, within a block, the dilation can double for each layer starting from an initial dilation, and then return to the initial dilation for the first layer in the next block. As an illustrative example, the dilations of the dilated convolutional layers in a block can be, in order: 1, 2, 4, ..., 512. A simplified example of a block of dilated causal convolutional layers is described below with reference to FIG. 2.

[0037] In some implementations, the convolutional subnetwork includes residual connections, skip connections, or both. An example architecture of the convolutional subnetwork that includes both residual connections and skip connections is described below with reference to FIG. 3.

[0038] Generally, the output subnetwork 120 can be any neural network-based subnetwork that is appropriate for generating the kind of neural network output that the system 100 is configured to generate using the alternative representations generated by the convolutional subnetwork 110, e.g., a recurrent neural network - based subnetwork, a fully-connected subnetwork, a single output layer, and so on.

[0039] In some particular implementations, e.g., in implementations where the system 100 is configured to perform speech recognition, the output subnetwork 120 includes a mean pooling layer that is configured to receive a subset of the alternative representations, i.e., the alternative representations for a specified number of time steps in the input sequence, and perform mean pooling on the alternative representations to generate a mean pooled representation for the alternative representation that aggregates the alternative representations to a coarser frame. In these implementations, the output subnetwork 120 also includes one or more additional layers, e.g., one or more non-causal convolutional layers followed by a softmax output layer, configured to receive the mean pooled representation and to generate a score distribution for an output in the output sequence at a position corresponding to the coarser frame. That is, each output position in the output sequence corresponds to a different subset of the alternative representations generated by the convolutional subnetwork 110.

[0040] FIG. 2 shows a visualization 200 of an example block of dilated causal convolutional layers. In particular, the example block includes a dilated causal convolutional layer 204 with dilation one, a dilated causal convolutional layer 206 with dilation two, a dilated causal convolutional layer 208 with dilation four, and a dilated causal convolutional layer 210 with dilation eight.

[0041] In the visualization 200, the block of dilated causal convolutional layers are operating on a current

input sequence 202 to generate an output sequence. In particular, the visualization 200 visualizes using bold arrows how the block generates the output 212 that is the output at the time step that is currently the last time step in the current input sequence 202 and the output sequence.

[0042] As can be seen from the visualization 200, because each layer in the block is a causal convolutional layer, the output 212 depends only on outputs that are at the last current time step or time steps before the last current time step in the various sequences operated on by the layers in the block.

[0043] Additionally, as can be seen from the visualization 200, the layers in the block are arranged in order of increasing dilation, with the first layer in the block, i.e., dilated causal convolutional layer 204, having dilation one and the last layer in the block, i.e., dilated causal convolutional layer 210, having dilation eight. In particular, as is shown by the bold arrows in the visualization 200, because the dilated causal convolutional layer 204 has dilation one, the filter of the layer 204 is applied to adjacent inputs in the current input sequence 202. Because the dilated causal convolutional layer 206 has dilation two, the filter of the layer 206 is applied to outputs that are separated by one output in the output sequence generated by the layer 204. Because the dilated causal convolutional layer 208 has dilation four, the filter of the layer 208 is applied to outputs that are separated by three outputs in the output sequence generated by the layer 206. Because the dilated causal convolutional layer 210 has dilation eight, the filter of the layer 210 is applied to outputs that are separated by seven outputs in the output sequence generated by the layer 208.

[0044] FIG. 3 shows an example architecture 300 for the convolutional subnetwork 110 of FIG. 1. As described above, in the example architecture 300, the dilated causal convolutional layers that are in the convolutional subnetwork have residual connections and skip connections.

[0045] In particular, in the architecture 300, the convolutional subnetwork 110 includes a causal convolutional layer 302 that processes the current input sequence 142.

[0046] The convolutional subnetwork 110 then processes the output of the causal convolutional layer 302 through a stack of dilated causal convolutional layers.

[0047] Each dilated causal convolutional layer 304 in the stack applies a dilated causal convolution 308 to the input 306 to the dilated causal convolutional layer 304. As described above, in some implementations, the dilated causal convolutional layers in the stack are arranged in blocks, with the dilation of the dilated causal convolutions applied by each layer increasing within a given block and then restarting at the initial value for the first layer in the next block.

[0048] In some implementations, the dilated causal convolutional layers in the stack have a gated activation unit in which the output of an element-wise non-linearity, i.e., of a conventional activation function, is element-wise multiplied by a gate vector. In some of these implemen-

tations, the dilated causal convolution 308 includes two dilated causal convolutions on the layer input 302 - a first dilated causal convolution between a main filter for the layer 304 and the layer input 306 and another dilated causal convolution between a gate filter for the layer 304 and the layer input 306. In others of these implementations, dilated causal convolution 308 is a single dilated causal convolution and half of the output of the single convolution is provided as the output of the dilated causal convolution between the main filter for the layer 304 and the layer input 306 and the other half of the output of the single convolution is provided as the output of the dilated causal convolution between the gate filter for the layer 304 and the layer input 306.

[0049] The dilated causal convolutional layer 304 then determines the output of the activation unit of the layer 304 using the outputs of the dilated causal convolution.

[0050] In particular, when the activation function is a gated activation unit, the layer 304 applies an element-wise non-linear function 310 which, in the example of FIG. 3 is the *tanh* function, to the output of the dilated convolution with the main filter and applies an element-wise gating function which, in the example of FIG. 3, is the *sigmoid* function, to the output of the dilated convolution with the gate filter. The layer 304 then performs an element-wise multiplication 314 between the output of the non-linear function 310 and the output of the gating function 312 to generate the activation unit output.

[0051] More specifically, when the element-wise non-linearity is *tanh* and the element-wise gating function is the *sigmoid* function, the output of the activation unit \mathbf{z} for a layer k satisfies:

$$\mathbf{z} = \tanh(W_{f,k} * \mathbf{x}) \odot \sigma(W_{g,k} * \mathbf{x}),$$

where $W_{f,k}$ is the main filter for the layer k , \mathbf{x} is the layer input, $*$ denotes a causal dilated convolution, \odot denotes element-wise multiplication, and $W_{g,k}$ is the gate filter for the layer k .

[0052] Because the architecture 300 includes skip connections and residual connections for the dilated causal convolutional layers, the layer 304 then performs a 1×1 convolution 316 on the activation unit output.

[0053] The layer 304 provides the output of the 1×1 convolution as the skip output 318 of the layer and adds the residual, i.e., the layer input 306, and the output of the 1×1 convolution to generate the final output 320 of the layer 304.

[0054] In some implementations, the layer 304 performs two 1×1 convolutions on the activation unit output, one with a residual filter and the other with a skip filter. In these implementations, the layer 304 provides the output of the convolution with the skip filter as the skip output 318 of the layer and adds the residual and the output of the 1×1 convolution with the residual filter to generate the final output 320 of the layer 304.

[0055] The convolutional subnetwork 110 then provides the final output 320 as the layer input to the next dilated convolutional layer in the stack. For the last layer in the stack, because there is no next layer, the convolutional subnetwork 110 can either discard the final output 320 generated by the last layer or can refrain from computing a final output, i.e., can refrain from performing the 1×1 convolution and the residual sum for the last layer in the stack.

[0056] Once the processing of all of the layers 304 in the stack of dilated convolutional layers has been completed, the convolutional subnetwork 110 sums 322 the skip outputs generated by the layers 304 and, optionally, applies one or more additional operations on the sum to generate the alternative representation 144. In particular, in the example of FIG. 3, the convolutional subnetwork 110 applies an element-wise non-linearity 324, e.g., a ReLU, to generate the alternative representation 144.

[0057] As described above, the output subnetwork 120 then processes the alternative representation 144 to generate the score distribution 146. While illustrated in FIG. 3 as generating a score distribution 146 from a single alternative representation, in implementations where the output sequence has fewer positions than there are inputs in the input sequence, the output subnetwork 120 may instead aggregate multiple alternative representations to generate the score distribution 146, e.g., by employing a mean pooling layer as described above.

[0058] FIG. 4 is a flow diagram of an example process 400 for processing a current input sequence. For convenience, the process 400 will be described as being performed by a system of one or more computers located in one or more locations. For example, a neural network system, e.g., the neural network system 100 of FIG. 1, appropriately programmed, can perform the process 400.

[0059] The system provides a current input sequence as of a given time step as input to the convolutional subnetwork (step 402). The current input sequence is a sequence that includes the input at the given time step and the inputs at time steps before the given time step in the input sequence. As described above, the convolutional subnetwork includes convolutional neural network layers, e.g., dilated causal convolutional layers, and is configured to process the current sequence to generate an alternative representation for the given time step.

[0060] The system provides the alternative representation as input to an output subnetwork (step 404). The output subnetwork is configured to process the alternative representation and, optionally, alternative representations from other time steps to generate a score distribution over possible outputs for a position in the output sequence.

[0061] The system selects an output for inclusion in the output sequence in accordance with the score distribution (step 406). For example, the system can sample a possible output in accordance with the score distribution or can select the possible output having the highest

score.

[0062] The system can be trained on training data that includes source sequence - target sequence pairs. In some implementations, at training time, i.e., during the training of the convolutional subnetwork and the output subnetwork to determine trained values of the filters of the convolutional layers and any other parameters of the system, the system trains the networks with one loss term that measures how well the networks classify the frame. In some other implementations, the system trains the networks with two loss terms, one to predict the next sample and one to classify the frame.

[0063] Additionally, because the system processes values from the input sequence in multiple passes, in some implementations, the convolutional subnetwork re-uses values computed for previous time steps when computing the alternative representation for the given time step. In particular, because the same inputs are provided as input to the subnetwork more than once, i.e., as part of multiple current input sequences, some of the computation performed by the convolutional subnetwork will be the same at multiple different time steps. In these implementations, rather than re-compute these computations each time step, the convolutional subnetwork can store the output values of the computation the first time that the computation is performed and then re-use the stored output values at subsequent time steps. As a simple example, the first convolutional layer in the convolutional subnetwork will apply the same filter or filters multiple times to the same inputs during the generation of an output sequence. Rather than re-compute the output of these filter applications at each time step, the system can re-use outputs computed at previous time steps.

[0064] In some implementations, as another way to increase the receptive field, one stack of dilated causal convolutional layers with a very large (long) receptive field, but preferably fewer units per layer, may be employed to condition another (larger) stack with a smaller receptive field. The larger stack may process a shorter part of the input sequence, for example cropped at the end.

[0065] This specification uses the term "configured" in connection with systems and computer program components. For a system of one or more computers to be configured to perform particular operations or actions means that the system has installed on it software, firmware, hardware, or a combination of them that in operation cause the system to perform the operations or actions. For one or more computer programs to be configured to perform particular operations or actions means that the one or more programs include instructions that, when executed by data processing apparatus, cause the apparatus to perform the operations or actions.

[0066] Examples of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specifica-

tion and their structural equivalents, or in combinations of one or more of them. Examples of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions encoded on a tangible non transitory storage medium for execution by, or to control the operation of, data processing apparatus. The computer storage medium can be a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them. Alternatively or in addition, the program instructions can be encoded on an artificially generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus.

[0067] The term "data processing apparatus" refers to data processing hardware and encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can also be, or further include, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit). The apparatus can optionally include, in addition to hardware, code that creates an execution environment for computer programs, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[0068] A computer program, which may also be referred to or described as a program, software, a software application, an app, a module, a software module, a script, or code, can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages; and it can be deployed in any form, including as a stand alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub programs, or portions of code. A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a data communication network.

[0069] In this specification, the term "database" is used broadly to refer to any collection of data: the data does not need to be structured in any particular way, or structured at all, and it can be stored on storage devices in one or more locations. Thus, for example, the index database can include multiple collections of data, each of which may be organized and accessed differently.

[0070] Similarly, in this specification the term "engine" is used broadly to refer to a software-based system, sub-system, or process that is programmed to perform one or more specific functions. Generally, an engine will be implemented as one or more software modules or components, installed on one or more computers in one or more locations. In some cases, one or more computers will be dedicated to a particular engine; in other cases, multiple engines can be installed and running on the same computer or computers.

[0071] The processes and logic flows described in this specification can be performed by one or more programmable computers executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by special purpose logic circuitry, e.g., an FPGA or an ASIC, or by a combination of special purpose logic circuitry and one or more programmed computers.

[0072] Computers suitable for the execution of a computer program can be based on general or special purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a central processing unit for performing or executing instructions and one or more memory devices for storing instructions and data. The central processing unit and the memory can be supplemented by, or incorporated in, special purpose logic circuitry. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

[0073] Computer readable media suitable for storing computer program instructions and data include all forms of non volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks; and CD ROM and DVD-ROM disks.

[0074] To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback

provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's device in response to requests received from the web browser. Also, a computer can interact with a user by sending text messages or other forms of message to a personal device, e.g., a smartphone that is running a messaging application, and receiving responsive messages from the user in return.

[0075] Data processing apparatus for implementing machine learning models can also include, for example, special-purpose hardware accelerator units for processing common and compute-intensive parts of machine learning training or production, i.e., inference, workloads.

[0076] Machine learning models can be implemented and deployed using a machine learning framework, .e.g., a TensorFlow framework, a Microsoft Cognitive Toolkit framework, an Apache Singa framework, or an Apache MXNet framework.

[0077] Examples of the subject matter described in this specification can be implemented in a computing system that includes a back end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface, a web browser, or an app through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network (LAN) and a wide area network (WAN), e.g., the Internet.

[0078] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other. In some embodiments, a server transmits data, e.g., an HTML page, to a user device, e.g., for purposes of displaying data to and receiving user input from a user interacting with the device, which acts as a client. Data generated at the user device, e.g., a result of the user interaction, can be received at the server from the device.

[0079] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or on the scope of what may be claimed, but rather as descriptions of features that may be specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination

in a single embodiment.

[0080] Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0081] Particular embodiments of the subject matter have been described. Other examples are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In some cases, multitasking and parallel processing may be advantageous.

Claims

1. A neural network system implemented by one or more computers, wherein the neural network system is configured to receive an input sequence comprising a plurality of inputs and to generate a neural network output from the input sequence, and wherein the neural network system comprises:
 - a convolutional subnetwork comprising a plurality of dilated convolutional neural network layers, wherein the convolutional subnetwork is configured to, for each of the plurality of inputs:
 - receive a current input sequence comprising the input and the inputs that precede the input in the input sequence, and
 - process the current input sequence to generate an alternative representation for the input; and
 - an output subnetwork, wherein the output subnetwork is configured to receive the alternative representations of the inputs and to process the alternative representations to generate the neural network output;
- wherein the neural network output comprises a plurality of sets of scores, wherein each set of scores includes a respective score for each of a plurality of

possible outputs, and

wherein the neural network system further comprises:

a subsystem configured to select outputs in accordance with the sets of scores to generate an output sequence comprising a plurality of outputs; and wherein the dilated convolutional neural network layers are causal convolutional neural network layers; and wherein

i) the input sequence is a sequence of audio data and the output sequence is a sequence of graphemes that is a grapheme representation of the input sequence, or

ii) the input sequence is a sequence of audio data and the output sequence is a sequence of phonemes that is a phoneme representation of the input sequence, or

iii) the input sequence is a sequence of phonemes or graphemes and the output sequence is a sequence of words in a first natural language that represents the input sequence, or

iv) wherein the input sequence is a sequence of words in a first natural language and the output sequence is a sequence of words in a second natural language that is a translation of the input sequence into the second natural language.

2. The neural network system of claim 1, wherein the output subnetwork comprises:

a mean pooling layer configured to:

receive a subset of the alternative representations, and perform mean pooling on the subset of alternative representations to generate a mean pooled representation for the subset; and

one or more additional neural network layers configured to, for each subset of the alternative representations, receive the mean pooled representation for the subset and to generate a set of scores for a position in the output sequence corresponding to the subset.

3. The neural network system of any one of claims 1-2, wherein the dilated convolutional neural network layers include multiple blocks of dilated convolutional neural network layers, wherein each block comprises multiple dilated convolutional neural network layers with increasing dilation.
4. The neural network system of any one of claims 1-3, wherein one or more of the dilated convolutional neural network layers have gated activation units.

5. The neural network system of any one of claims 1-4, wherein the convolutional subnetwork comprises residual connections.
6. The neural network system of any one of claims 1-5, wherein the convolutional subnetwork comprises skip connections.
7. A computer-implemented method of generating a neural network output from an input sequence that comprises a plurality of inputs, the method comprising:
- for each of the inputs, providing a current input sequence that comprises the input and the inputs preceding the input in the input sequence to a convolutional subnetwork comprising a plurality of dilated convolutional neural network layers, wherein the convolutional subnetwork is configured to, for each of the plurality of inputs:
- receive the current input sequence for the input, and
process the current input sequence to generate an alternative representation for the input; and
- providing the alternative representations to an output subnetwork, wherein the output subnetwork is configured to receive the alternative representations and to process the alternative representations to generate the neural network output;
- wherein the neural network output comprises a plurality of sets of scores, wherein each set of scores includes a respective score for each of a plurality of possible outputs, and
wherein the method further comprises:
selecting outputs in accordance with the sets of scores to generate an output sequence comprising a plurality of outputs; and
wherein the dilated convolutional neural network layers are causal convolutional neural network layers, and wherein
- i) the input sequence is a sequence of audio data and the output sequence is a sequence of graphemes that is a grapheme representation of the input sequence, or
ii) the input sequence is a sequence of audio data and the output sequence is a sequence of phonemes that is a phoneme representation of the input sequence, or
iii) the input sequence is a sequence of phonemes or graphemes and the output sequence is a sequence of words in a first natural language that represents the input sequence, or
iv) wherein the input sequence is a sequence of words in a first natural language and the output sequence is a sequence of words in a second natural language that is a translation of the input sequence into the second natural language.
8. The method of claim 7, wherein the output subnetwork comprises:
- a mean pooling layer configured to:
- receive a subset of the alternative representations, and
perform mean pooling on the subset of alternative representations to generate a mean pooled representation for the subset; and
- one or more additional neural network layers configured to, for each subset of the alternative representations, receive the mean pooled representation for the subset and to generate a subset output that defines an output in the output sequence at a position corresponding to the subset.
9. The method of any one of claims 7-8, wherein the dilated convolutional neural network layers include multiple blocks of dilated convolutional neural network layers, wherein each block comprises multiple dilated convolutional neural network layers with increasing dilation.
10. The method of any one of claims 7-9, wherein one or more of the dilated convolutional neural network layers have gated activation units.
11. The method of any one of claims 7-10, wherein the convolutional subnetwork comprises residual connections.
12. The method of any one of claims 7-11, wherein the convolutional subnetwork comprises skip connections.
13. One or more computer storage media encoded with instructions that when executed by one or more computers cause the one or more computers to perform the operations of the method of any one of claims 7-12.
14. The neural network system, computer-implemented method, or computer storage media of any preceding claim wherein each of the dilated convolutional neural network layers is configured to apply a convolution to non-adjacent values in a sequence input to

the layer.

Patentansprüche

1. Neuronales Netzwerksystem, implementiert durch einen oder mehrere Computer, wobei das neuronale Netzwerksystem konfiguriert ist zum Empfangen einer Eingangssequenz, umfassend eine Vielzahl von Eingängen, und zum Erzeugen eines Ausgangs des neuronalen Netzwerks aus der Eingangssequenz und wobei das neuronale Netzwerksystem umfasst: ein Konvolutionsteilnetzwerk, umfassend eine Vielzahl von erweiterten neuronalen Konvolutionsnetzwerkschichten, wobei das Konvolutionsteilnetzwerk für jede der Vielzahl von Eingängen konfiguriert ist zum:

Empfangen einer gegenwärtigen Eingangssequenz, umfassend den Eingang und die Eingänge, die dem Eingang in der Eingangssequenz vorhergehen, und

Verarbeiten der gegenwärtigen Eingangssequenz, um eine alternative Repräsentation für den Eingang zu erzeugen; und

ein Ausgangsteilnetzwerk, wobei das Ausgangsteilnetzwerk konfiguriert ist zum Empfangen der alternativen Repräsentationen der Eingänge und zum Verarbeiten der alternativen Repräsentationen, um den Ausgang des neuronalen Netzwerks zu erzeugen;

wobei der Ausgang des neuronalen Netzwerks eine Vielzahl von Mengen von Bewertungsziffern umfasst, wobei jede Menge von Bewertungsziffern eine jeweilige Bewertungsziffer für jeden einer Vielzahl von möglichen Ausgängen enthält, und

wobei das neuronale Netzwerksystem ferner umfasst:

ein Teilsystem, konfiguriert zum Auswählen von Ausgängen gemäß den Mengen von Bewertungsziffern, um eine Ausgangssequenz, umfassend eine Vielzahl von Ausgängen, zu erzeugen; und

wobei die erweiterten neuronalen Konvolutionsnetzwerkschichten neuronale Kausal-konvolutionsnetzwerkschichten sind; und wobei

- i) die Eingangssequenz eine Sequenz von Audiodaten ist und die Ausgangssequenz eine Sequenz von Graphemen ist, die eine Graphemrepräsentation der Eingangssequenz ist, oder
 ii) die Eingangssequenz eine Sequenz von Audiodaten ist und die Ausgangs-

sequenz eine Sequenz von Phonemen ist, die eine Phonemrepräsentation der Eingangssequenz ist, oder

iii) die Eingangssequenz eine Sequenz von Phonemen oder Graphemen ist und die Ausgangssequenz eine Sequenz von Wörtern in einer ersten natürlichen Sprache ist, die die Eingangssequenz repräsentiert, oder

iv) wobei die Eingangssequenz eine Sequenz von Wörtern in einer ersten natürlichen Sprache ist und die Ausgangssequenz eine Sequenz von Wörtern in einer zweiten natürlichen Sprache ist, die eine Übersetzung der Eingangssequenz in die zweite natürliche Sprache ist.

2. Neuronales Netzwerksystem nach Anspruch 1, wobei das Ausgangsteilnetzwerk umfasst: eine Schicht für gemittelte Zusammenfassung, konfiguriert zum:

Empfangen einer Teilmenge der alternativen Repräsentationen und

Durchführen von gemittelter Zusammenfassung an der Teilmenge von alternativen Repräsentationen, um eine gemittelte zusammengefasste Repräsentation für die Teilmenge zu erzeugen; und

eine oder mehrere zusätzliche neuronale Netzwerkschichten, konfiguriert zum, für jede Teilmenge der alternativen Repräsentationen, Empfangen der gemittelten zusammengefassten Repräsentation für die Teilmenge und zum Erzeugen einer Menge von Bewertungsziffern für eine Position in der Ausgangssequenz korrespondierend mit der Teilmenge.

3. Neuronales Netzwerksystem nach einem der Ansprüche 1-2, wobei die erweiterten neuronalen Konvolutionsnetzwerkschichten mehrere Blöcke von erweiterten neuronalen Konvolutionsnetzwerkschichten enthalten, wobei jeder Block mehrere erweiterte neuronale Konvolutionsnetzwerkschichten mit zunehmender Erweiterung umfasst.

4. Neuronales Netzwerksystem nach einem der Ansprüche 1-3, wobei eine oder mehrere der erweiterten neuronalen Konvolutionsnetzwerkschichten gattergesteuerte Aktivierungseinheiten aufweisen.

5. Neuronales Netzwerksystem nach einem der Ansprüche 1-4, wobei das Konvolutionsteilnetzwerk restliche Verbindungen umfasst.

6. Neuronales Netzwerksystem nach einem der Ansprüche 1-5, wobei das Konvolutionsteilnetzwerk

ausgelassene Verbindungen umfasst.

7. Computerimplementiertes Verfahren zum Erzeugen eines neuronalen Netzwerkausgangs aus einer Eingangssequenz, die eine Vielzahl von Eingängen umfasst, das Verfahren umfassend:
- 5 für jeden der Eingänge Bereitstellen einer gegenwärtigen Eingangssequenz, die den Eingang und die Eingänge, die dem Eingang in der Eingangssequenz vorhergehen, umfasst, an ein Konvolutionsteilnetzwerk, umfassend eine Vielzahl von erweiterten neuronalen Netzwerkschichten, wobei das Konvolutionsteilnetzwerk für jeden der Vielzahl von Eingängen konfiguriert ist zum:

Empfangen der gegenwärtigen Eingangssequenz für den Eingang und

Verarbeiten der gegenwärtigen Eingangssequenz, um eine alternative Repräsentation für den Eingang zu erzeugen; und

20 Bereitstellen der alternativen Repräsentationen an ein Ausgangsteilnetzwerk, wobei das Ausgangsteilnetzwerk konfiguriert ist zum Empfangen der alternativen Repräsentationen und zum Verarbeiten der alternativen Repräsentationen, um den Ausgang des neuronalen Netzwerks zu erzeugen;

25 wobei der Ausgang des neuronalen Netzwerks eine Vielzahl von Mengen von Bewertungsziffern umfasst,

30 wobei jede Menge von Bewertungsziffern eine jeweilige Bewertungsziffer für jeden einer Vielzahl von möglichen Ausgängen enthält und wobei das Verfahren ferner umfasst:

Auswählen von Ausgängen gemäß den Mengen von Bewertungsziffern, um eine Ausgangssequenz, umfassend eine Vielzahl von Ausgängen, zu erzeugen; und

40 wobei die erweiterten neuronalen Konvolutionsteilnetzwerkschichten neuronale Kausal-konvolutionsteilnetzwerkschichten sind, und wobei

- i) die Eingangssequenz eine Sequenz von Audiodaten ist und die Ausgangssequenz eine Sequenz von Graphemen ist, die eine Graphemrepräsentation der Eingangssequenz ist, oder
- 50 ii) die Eingangssequenz eine Sequenz von Audiodaten ist und die Ausgangssequenz eine Sequenz von Phonemen ist, die eine Phonemrepräsentation der Eingangssequenz ist, oder
- 55 iii) die Eingangssequenz eine Sequenz von Phonemen oder Graphemen ist und die Ausgangssequenz eine Sequenz von Wörtern in einer ersten na-

türlichen Sprache ist, die die Eingangssequenz repräsentiert, oder

iv) wobei die Eingangssequenz eine Sequenz von Wörtern in einer ersten natürlichen Sprache ist und die Ausgangssequenz eine Sequenz von Wörtern in einer zweiten natürlichen Sprache ist, die eine Übersetzung der Eingangssequenz in die zweite natürliche Sprache ist.

8. Verfahren nach Anspruch 7, wobei das Ausgangsteilnetzwerk umfasst:
- 15 eine Schicht für gemittelte Zusammenfassung, konfiguriert zum:

Empfangen einer Teilmenge der alternativen Repräsentationen, und

Durchführen von gemittelter Zusammenfassung an der Teilmenge von alternativen Repräsentationen, um eine gemittelte zusammengefasste Repräsentation für die Teilmenge zu erzeugen; und

eine oder mehrere zusätzliche neuronale Netzwerkschichten, konfiguriert zum, für jede Teilmenge der alternativen Repräsentationen, Empfangen der gemittelten zusammengefassten Repräsentation für die Teilmenge und zum Erzeugen eines Teilmengenausgangs, der einen Ausgang in der Ausgangssequenz bei einer Position korrespondierend mit der Teilmenge definiert.

9. Verfahren nach einem der Ansprüche 7-8, wobei die erweiterten neuronalen Konvolutionsnetzwerkschichten mehrere Blöcke von erweiterten neuronalen Konvolutionsnetzwerkschichten enthalten, wobei jeder Block mehrere erweiterte neuronale Konvolutionsnetzwerkschichten mit zunehmender Erweiterung umfasst.

10. Verfahren nach einem der Ansprüche 7-9, wobei eine oder mehrere der erweiterten neuronalen Konvolutionsnetzwerkschichten gattergesteuerte Aktivierungseinheiten aufweisen.

11. Verfahren nach einem der Ansprüche 7-10, wobei das Konvolutionsteilnetzwerk restliche Verbindungen umfasst.

12. Verfahren nach einem der Ansprüche 7-11, wobei das Konvolutionsteilnetzwerk ausgelassene Verbindungen umfasst.

13. Ein oder mehrere Computerspeichermedien, codiert mit Anweisungen, die, wenn sie durch einen oder mehrere Computer ausgeführt werden, bewirken, dass der eine oder die mehreren Computer die Ope-

rationen des Verfahrens nach einem der Ansprüche 7-12 durchführen.

14. Neuronales Netzwerksystem, computerimplementiertes Verfahren oder Computerspeichermedien nach einem der vorstehenden Ansprüche, wobei jede der erweiterten neuronalen Konvolutionsnetzwerkschichten konfiguriert ist zum Anwenden einer Konvolution auf nicht angrenzende Werte in einer in die Schicht eingegebene Sequenz.

Revendications

1. Système de réseau neuronal mis en œuvre par un ou plusieurs ordinateurs, dans lequel le système de réseau neuronal est configuré pour recevoir une séquence d'entrée qui comprend une pluralité d'entrées et pour générer une sortie du réseau neuronal à partir de la séquence d'entrée ; et dans lequel le système de réseau neuronal comprend : un sous-réseau de convolution qui comprend une pluralité de couches de réseau neuronal de convolution dilatées, dans lequel le sous-réseau de convolution est configuré pour, pour chacune de la pluralité d'entrées :

recevoir une séquence d'entrée courante qui comprend l'entrée et les entrées qui précèdent l'entrée dans la séquence d'entrée ; et traiter la séquence d'entrée courante pour générer une représentation alternative pour l'entrée ; et un sous-réseau de sortie, dans lequel le sous-réseau de sortie est configuré pour recevoir les représentations alternatives des entrées et pour traiter les représentations alternatives pour générer la sortie du réseau neuronal ; dans lequel la sortie du réseau neuronal comprend une pluralité de jeux de scores, dans lequel chaque jeu de scores inclut un score respectif pour chacune d'une pluralité de sorties possibles ; et dans lequel le système de réseau neuronal comprend en outre :

un sous-système qui est configuré pour sélectionner des sorties conformément aux jeux de scores pour générer une séquence de sortie qui comprend une pluralité de sorties ; et dans lequel les couches de réseau neuronal de convolution dilatées sont des couches de réseau neuronal de convolution causales ; et dans lequel :

i) la séquence d'entrée est une séquence de données audio et la séquence de sortie est une séquence de graphèmes qui est une représentation par graphèmes de la séquence d'entrée ; ou
ii) la séquence d'entrée est une séquence de données audio et la séquence de sortie est une séquence de phonèmes qui est une représentation par phonèmes de la séquence d'entrée ; ou
iii) la séquence d'entrée est une séquence de phonèmes ou de graphèmes et la séquence de sortie est une séquence de mots dans une première langue naturelle qui représente la séquence d'entrée ; ou
iv) dans lequel la séquence d'entrée est une séquence de mots dans une première langue naturelle et la séquence de sortie est une séquence de mots dans une seconde langue naturelle qui est une traduction de la séquence d'entrée dans la seconde langue naturelle.

2. Système de réseau neuronal selon la revendication 1, dans lequel le sous-réseau de sortie comprend : une couche de pooling moyen qui est configurée pour :

recevoir un sous-jeu des représentations alternatives ; et réaliser un pooling moyen sur le sous-jeu de représentations alternatives pour générer une représentation à pooling moyen pour le sous-jeu ; et une ou plusieurs couches de réseau neuronal additionnelles qui sont configurées pour, pour chaque sous-jeu des représentations alternatives, recevoir la représentation à pooling moyen pour le sous-jeu et pour générer un jeu de scores pour une position dans la séquence de sortie qui correspond au sous-jeu.

3. Système de réseau neuronal selon l'une quelconque des revendications 1-2, dans lequel les couches de réseau neuronal de convolution dilatées incluent de multiples blocs de couches de réseau neuronal de convolution dilatées, dans lequel chaque bloc comprend de multiples couches de réseau neuronal de convolution dilatées présentant une dilatation croissante.

4. Système de réseau neuronal selon l'une quelconque des revendications 1-3, dans lequel une ou plusieurs des couches de réseau neuronal de convolution dilatées comportent des unités d'activation à déclenchement.

5. Système de réseau neuronal selon l'une quelconque des revendications 1-4, dans lequel le sous-réseau de convolution comprend des connexions résiduelles. 5
6. Système de réseau neuronal selon l'une quelconque des revendications 1-5, dans lequel le sous-réseau de convolution comprend des connexions de saut. 5
7. Procédé mis en œuvre par ordinateur consistant à générer une sortie du réseau neuronal à partir d'une séquence d'entrée qui comprend une pluralité d'entrées, le procédé comprenant : 10
pour chacune des entrées, la fourniture d'une séquence d'entrée courante qui comprend l'entrée et les entrées qui précèdent l'entrée dans la séquence d'entrée à un sous-réseau de convolution qui comprend une pluralité de couches de réseau neuronal de convolution dilatées, dans lequel le sous-réseau de convolution est configuré pour, pour chacune de la pluralité d'entrées : 15
recevoir la séquence d'entrée courante pour l'entrée ; et
traiter la séquence d'entrée courante pour générer une représentation alternative pour l'entrée ; et 25
la fourniture des représentations alternatives à un sous-réseau de sortie, dans lequel le sous-réseau de sortie est configuré pour recevoir les représentations alternatives et pour traiter les représentations alternatives pour générer la sortie du réseau neuronal ; 30
dans lequel la sortie du réseau neuronal comprend une pluralité de jeux de scores, dans lequel chaque jeu de scores inclut un score respectif pour chacune d'une pluralité de sorties possibles ; et 35
dans lequel le procédé comprend en outre : 40
la sélection de sorties conformément aux jeux de scores pour générer une séquence de sortie qui comprend une pluralité de sorties ; et 45
dans lequel les couches de réseau neuronal de convolution dilatées sont des couches de réseau neuronal de convolution causales ; et dans lequel : 50
i) la séquence d'entrée est une séquence de données audio et la séquence de sortie est une séquence de graphèmes qui est une représentation par graphèmes de la séquence d'entrée ; ou 55
ii) la séquence d'entrée est une séquence de données audio et la séquence de sortie est une séquence de phonèmes qui est une représentation par phonèmes de la séquence d'entrée ; ou
iii) la séquence d'entrée est une séquence de phonèmes ou de graphèmes et la séquence de sortie est une séquence de mots dans une première langue naturelle qui représente la séquence d'entrée ; ou
iv) dans lequel la séquence d'entrée est une séquence de mots dans une première langue naturelle et la séquence de sortie est une séquence de mots dans une seconde langue naturelle qui est une traduction de la séquence d'entrée dans la seconde langue naturelle.
8. Procédé selon la revendication 7, dans lequel le sous-réseau de sortie comprend : 5
une couche de pooling moyen qui est configurée pour : 20
recevoir un sous-jeu des représentations alternatives ; et
réaliser un pooling moyen sur le sous-jeu de représentations alternatives pour générer une représentation à pooling moyen pour le sous-jeu ; et
une ou plusieurs couches de réseau neuronal additionnelles qui sont configurées pour, pour chaque sous-jeu des représentations alternatives, recevoir la représentation à pooling moyen pour le sous-jeu et pour générer une sortie de sous-jeu qui définit une sortie dans la séquence de sortie au niveau d'une position qui correspond au sous-jeu.
9. Procédé selon l'une quelconque des revendications 7-8, dans lequel les couches de réseau neuronal de convolution dilatées incluent de multiples blocs de couches de réseau neuronal de convolution dilatées, dans lequel chaque bloc comprend de multiples couches de réseau neuronal de convolution dilatées présentant une dilatation croissante.
10. Procédé selon l'une quelconque des revendications 7-9, dans lequel une ou plusieurs des couches de réseau neuronal de convolution dilatées comportent des unités d'activation à déclenchement.
11. Procédé selon l'une quelconque des revendications 7-10, dans lequel le sous-réseau de convolution comprend des connexions résiduelles.
12. Procédé selon l'une quelconque des revendications 7-11, dans lequel le sous-réseau de convolution comprend des connexions de saut.
13. Un ou plusieurs supports de stockage d'ordinateur qui sont codés avec des instructions qui, lorsqu'elles

sont exécutées par un ou plusieurs ordinateurs, forcent les un ou plusieurs ordinateurs à réaliser les opérations du procédé selon l'une quelconque des revendications 7-12.

5

14. Système de réseau neuronal, procédé mis en œuvre par ordinateur ou supports de stockage d'ordinateur selon l'une quelconque des revendications précédentes, dans lequel chacune des couches de réseau neuronal de convolution dilatées est configurée pour appliquer une convolution à des valeurs non adjacentes dans une séquence qui est entrée sur la couche.

10

15

20

25

30

35

40

45

50

55

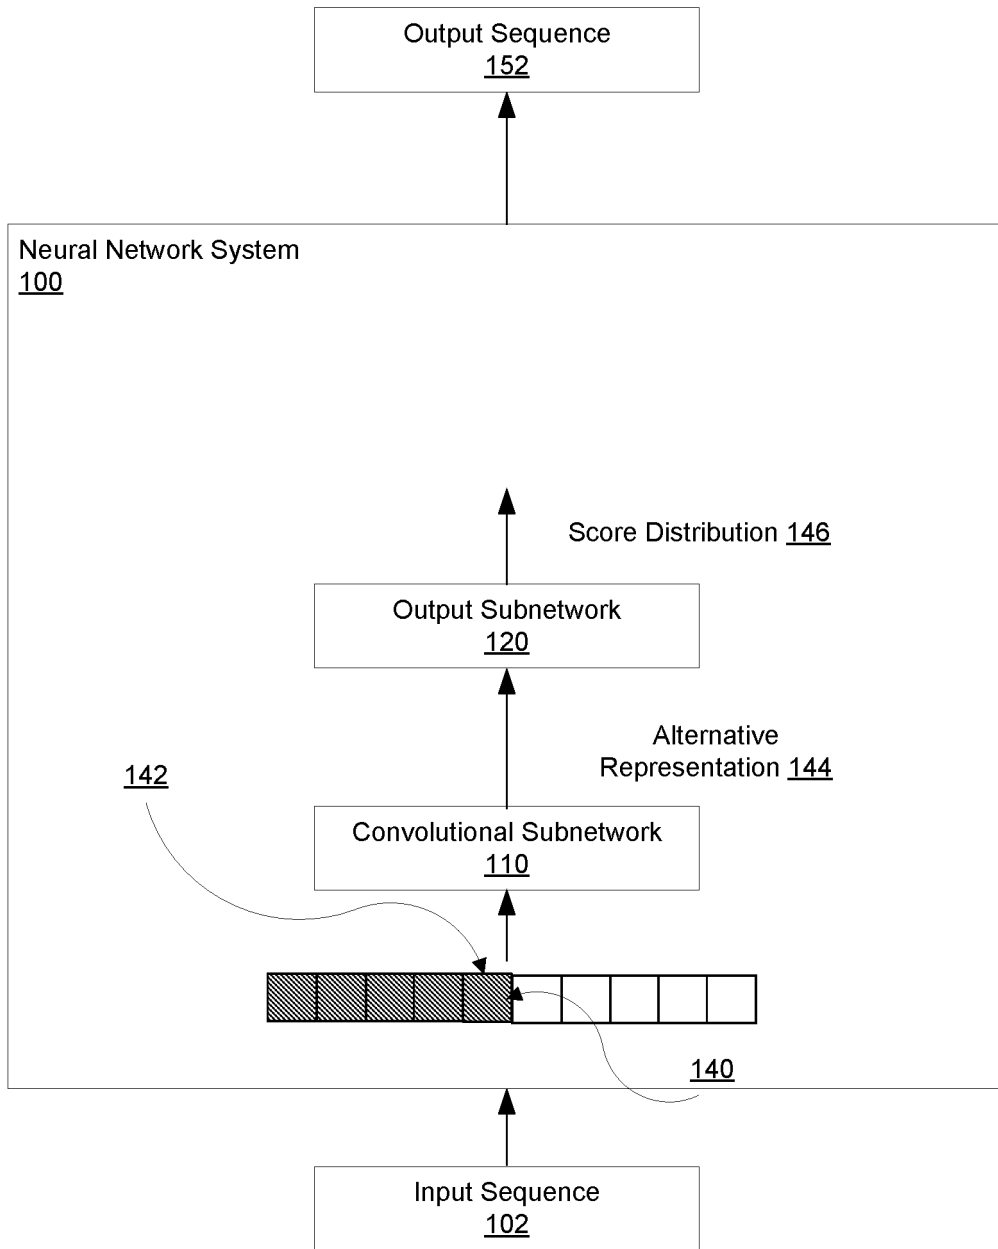


FIG. 1




FIG. 2


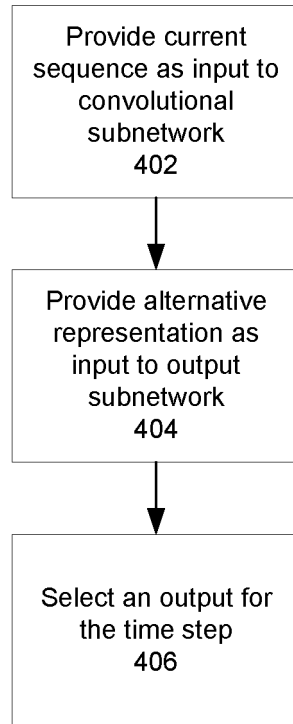


FIG. 3



400 ↗

FIG. 4

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- CA 2810457 A1 [0015]