



- (51) International Patent Classification: **G06N 3/04** (2006.01)
- (21) International Application Number: **PCT/US2016/068913**
- (22) International Filing Date: **28 December 2016 (28.12.2016)**
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data: **62/290,624** 3 February 2016 (03.02.2016) **US**
- (71) Applicant: **GOOGLE INC.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).
- (72) Inventors: **ALSHARIF, Ouais**; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). **PRABHAVALKAR, Rohit Prakash**; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). **MCGRRAW, Ian C.**; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). **BRUGUIER, Antoine Jean**; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).
- (74) Agents: **PORTNOV, Michael** et al.; Fish & Richardson P.C., P.O. Box 1022, 3300 RBC Plaza, Minneapolis, Minnesota 55440-1022 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (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).

[Continued on next page]

(54) Title: COMPRESSED RECURRENT NEURAL NETWORK MODELS

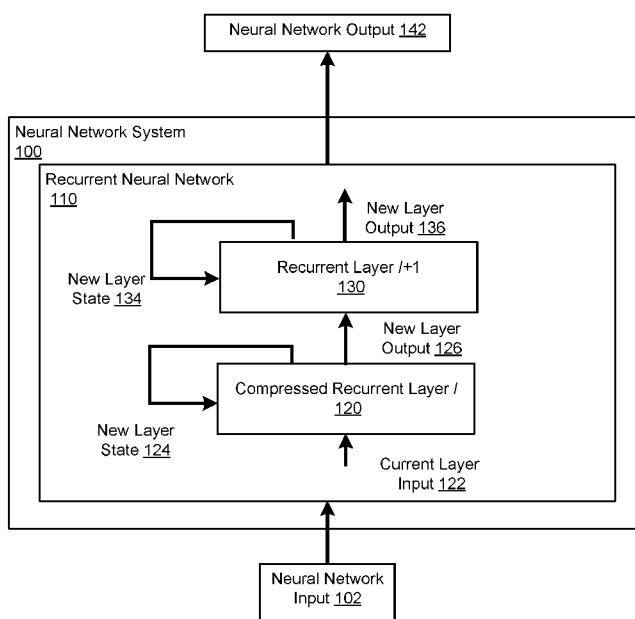


FIG. 1

(57) Abstract: Methods, systems, and apparatus, including computer programs encoded on computer storage media, for implementing a compressed recurrent neural network (RNN). One of the systems includes a compressed RNN, the compressed RNN comprising a plurality of recurrent layers, wherein each of the recurrent layers has a respective recurrent weight matrix and a respective inter-layer weight matrix, and wherein at least one of recurrent layers is compressed such that a respective recurrent weight matrix of the compressed layer is defined by a first compressed weight matrix and a projection matrix and a respective inter-layer weight matrix of the compressed layer is defined by a second compressed weight matrix and the projection matrix. Compression can provide a number of technical advantages, including the ability to process data more efficiently and use less data storage.

WO 2017/136070 A1

Declarations under Rule 4.17:

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

Published:

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

COMPRESSED RECURRENT NEURAL NETWORK MODELS

BACKGROUND

[0001] This specification relates to neural network architectures and compressing 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. Some neural networks, e.g., those that are designed for time series problems or sequence-to-sequence learning (recurrent neural networks (RNNs)), incorporate recurrent loops which permit memory, in the form of a hidden state variable, to persist within a layer between data inputs. A variation of RNNs, long short-term memory (LSTM) neural networks include multiple gates within each layer to control the persistence of data between data inputs. Some neural networks, e.g., those that are designed for time series problems or sequence-to-sequence learning, incorporate recurrent loops which permit memory, in the form of a hidden state variable, to persist within a layer between data inputs.

SUMMARY

[0003] This specification describes technologies that relate to recurrent neural network architectures. In general, a recurrent neural network includes at least one recurrent neural network layer that is compressed. In particular, the recurrent weight matrix and the inter-layer weight matrix for the compressed recurrent layer are jointly compressed using a shared projection matrix.

[0004] 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.

[0005] The subject matter described in this specification can be implemented in particular embodiments so as to realize one or more of the following advantages. By compressing the weight matrices of one or more of the recurrent layers in a recurrent neural network, the recurrent neural network is configured to be able to process data more efficiently and use less data storage. In particular, a recurrent neural network having one or compressed recurrent layers can be effectively trained to achieve performance that is comparable to full size, e.g., uncompressed, recurrent neural networks, while using less data storage and being able to process inputs faster by virtue of the compressed weight matrices of the compressed recurrent layers having fewer parameters than the weight matrices of the corresponding layers in the uncompressed recurrent neural network. In fact, because the compressed recurrent neural network has a smaller computational footprint, the compressed network may be able to be effectively implemented to process inputs in real-time on a mobile device having limited storage and processing power even when the uncompressed network could not be run on the mobile device.

[0006] The details of one or more embodiments of the subject matter of 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.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 shows an example neural network system.

[0008] FIG. 2 is a flow diagram of an example process for compressing a recurrent neural network.

[0009] FIG. 3 is a flow diagram of an example process for compressing an inter-layer weight matrix and a recurrent weight matrix for a particular recurrent layer.

[0010] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0011] 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 are implemented.

[0012] The neural network system 100 is a machine learning system that receives a respective neural network input at each of multiple time steps and generates a respective neural network output at each of the time steps. That is, at each of the multiple time steps, the neural network system 100 receives a neural network input and processes the neural network input to generate a neural network output. For example, at a given time step t , the neural network system 100 can receive a neural network input 102 and generate a neural network output 142.

[0013] The neural network system 100 can store the generated neural network outputs in an output data repository or provide the neural network outputs for use for some other immediate purpose.

[0014] The neural network system 100 can be configured to receive any kind of digital data input and to generate any kind of score or classification output based on the input.

[0015] For example, if the inputs to the neural network system 100 are images or features that have been extracted from images, the output generated by the neural network system 100 for a given image may be scores for each of a set of object categories, with each score representing an estimated likelihood that the image contains an image of an object belonging to the category.

[0016] As another example, if the inputs to the neural network system 100 are Internet resources (e.g., web pages), documents, or portions of documents or features extracted from Internet resources, documents, or portions of documents, the output generated by the neural network system 100 for a given Internet resource, document, or portion of a document may be a score for each of a set of topics, with each score representing an estimated likelihood that the Internet resource, document, or document portion is about the topic.

[0017] As another example, if the inputs to the neural network system 100 are features of a personalized recommendation for a user, e.g., features characterizing the context for the recommendation, e.g., features characterizing previous actions taken by the user, the output

generated by the neural network system 100 may be a score for each of a set of content items, with each score representing an estimated likelihood that the user will respond favorably to being recommended the content item. In some of these examples, the neural network system 100 is part of a reinforcement learning system that provides content recommendations to users.

[0018] As another example, if the input to the neural network system 100 is text in one language, the output generated by the neural network system 100 may be a score for each of a set of pieces of text in another language, with each score representing an estimated likelihood that the piece of text in the other language is a proper translation of the input text into the other language.

[0019] As another example, if the input to the neural network system 100 is features of a spoken utterance, the output generated by the neural network system 100 may be a score for each of a set of pieces of text, each score representing an estimated likelihood that the piece of text is the correct transcription for the utterance.

[0020] As another example, if the inputs to the neural network system 100 are images, the output generated by the neural network system 100 may be a score for each of a set of pieces of text, each score representing an estimated likelihood that the piece of text is text that is present in the input image.

[0021] In particular, the neural network system 100 includes a recurrent neural network 110 which, in turn, includes multiple recurrent layers, i.e., at least a compressed recurrent layer l 120 and a recurrent layer $l+1$ 130. The recurrent neural network 110 is configured to, at each of the time steps, receive the neural network input at the time step and to process the neural network input to generate the neural network output at the time step.

[0022] In addition to the compressed recurrent layer 120 and the recurrent layer 130, the recurrent neural network 110 may include one or more other components, e.g., other recurrent layers, other non-recurrent neural network layers, and so on.

[0023] For example, the recurrent neural network 100 may be a deep recurrent network that includes multiple recurrent layers including the compressed recurrent layer 120 and the recurrent layer 130 arranged in an ordered stack one on top of one another, and an output layer that, at each time step, receives the layer output from the highest recurrent layer in the

stack and, optionally, other recurrent layers in the stack, and processes the layer output to generate the neural network output 142 at the time step.

[0024] The compressed recurrent layer 120 is configured to, at each of the time steps, receive a current layer input 122 and to process the current layer input 122, a current layer state of the recurrent layer 120, and a current layer output of the recurrent layer 120 to generate a new layer output 126 and to update the current layer state to generate a new layer state 124.

[0025] Depending on the configuration of the recurrent neural network 110, the current layer input 122 may be the neural network input 102 or an output generated by a different component of the recurrent neural network 110.

[0026] Additionally, for each time step after the first step, the current layer state is the new layer state generated at the preceding time step. For the first time step, the current layer state may be a predetermined initial layer state.

[0027] The recurrent layer 130 is configured to, at each of the time steps, receive the new layer output 126 and to process the new layer output 126 and a current layer state of the recurrent layer 130 to generate a new layer output 136 and to update the current layer state to generate a new layer state 134.

[0028] Depending on the configuration of the recurrent neural network 110, the new layer output 126 may be provided as input to another recurrent layer in the recurrent neural network 110, as input to a different type of neural network component, e.g., to an output layer or a different type of neural network layer, or may be provided as the neural network output 142 of the recurrent neural network 110.

[0029] Each recurrent layer in the recurrent neural network 110 has two corresponding weight matrices: a recurrent weight matrix and an inter-layer weight matrix. Generally, during processing at a given time step, the recurrent weight matrix for a given recurrent layer is applied to the layer output generated by the recurrent layer at the preceding time step while the inter-layer weight matrix is applied to layer outputs generated by the recurrent layer at the given time step. Thus, the recurrent weight matrix for a given recurrent layer will generally be applied by the given recurrent layer while the inter-layer weight matrix will generally be applied by the next layer that receives the layer output generated by the given recurrent layer at the time step, e.g., the next layer above the given layer in the stack.

[0030] In some implementations, the recurrent neural network 110 is a standard recurrent neural network and the state of each recurrent layer is therefore also used as the layer output of the recurrent layer. That is, the updated state of the layer for a given time step is also used as the layer output for the layer for the given time step. Thus, new layer output 136 is the same as new layer state 134 and new layer output 126 is the same as new layer state 124.

[0031] In these implementations, if neither of the compressed recurrent layer 120 or the recurrent layer 130 were compressed, the compressed recurrent layer 120 would be configured to generate a layer output h_t^l 126 (and new layer state 124) at time step t that satisfied:

$$h_t^l = \sigma(\mathbf{W}_x^{l-1}h_{t-1}^{l-1} + \mathbf{W}_h^l h_{t-1}^l + b^l),$$

where \mathbf{W}_x^{l-1} is an inter-layer weight matrix for the layer before the layer 120 in the recurrent neural network 110, h_{t-1}^{l-1} is the layer output of the layer before the layer 120, \mathbf{W}_h^l is the recurrent weight matrix for the uncompressed version of the recurrent layer 120, h_{t-1}^l is the current layer state (and the layer output from the preceding time step), b^l is the bias vector for the layer 120, and $\sigma(\cdot)$ denotes a non-linear activation function.

[0032] The recurrent layer 130 would be configured to generate a layer output h_t^{l+1} 136 (and new layer state 134) at time step t that satisfies:

$$h_t^{l+1} = \sigma(\mathbf{W}_x^l h_t^l + \mathbf{W}_h^{l+1} h_{t-1}^{l+1} + b^{l+1}),$$

where \mathbf{W}_x^l is an inter-layer weight matrix for the uncompressed version of the layer 120, h_t^l is the layer output 126 of the recurrent layer 120, \mathbf{W}_h^{l+1} is the recurrent weight matrix for the recurrent layer 130, h_{t-1}^{l+1} is the current layer state (and the layer output from the preceding time step) for the recurrent layer 130, b^{l+1} is the bias vector for the recurrent layer 130, and $\sigma(\cdot)$ denotes a non-linear activation function.

[0033] However, because at least the compressed recurrent layer 120 has been compressed, the inter-layer and recurrent weight matrices for the compressed recurrent layer 120 have been modified.

[0034] In particular, a compressed recurrent layer is a recurrent layer for which the recurrent and inter-layer matrices have each been replaced by a respective lower-rank approximation. That is, the recurrent weight matrix for the compressed recurrent layer has been replaced by a matrix that has a lower rank than the recurrent weight matrix and the

inter-layer weight matrix has been replaced by a matrix that has a lower rank than the inter-layer weight matrix. In so doing, the number of parameters in the recurrent and inter-layer weight matrices have been reduced.

[0035] In particular, using the compression scheme described in this specification, the compressed recurrent layer 120 has been compressed by replacing the recurrent weight matrix, \mathbf{W}_h^l , and inter-layer weight matrix, \mathbf{W}_x^l , of the compressed layer 120 with respective first and second compressed weight matrices, \mathbf{Z}_h^l , and \mathbf{Z}_x^l , and a corresponding projection matrix, \mathbf{P}^l . Specifically, the recurrent weight matrix, \mathbf{W}_h^l , and inter-layer weight matrix, \mathbf{W}_x^l , are jointly compressed by determining a projection matrix such that \mathbf{W}_h^l is defined by $\mathbf{Z}_h^l \mathbf{P}^l$ and \mathbf{W}_x^l is defined by $\mathbf{Z}_x^l \mathbf{P}^l$. The first and second compressed weight matrices and the projection matrix each have a rank that is lower than the rank of the inter-layer weight matrix and the recurrent weight matrix. Techniques for compressing a recurrent layer to reduce the number of parameters while maintaining high performance is described in more detail below with reference to FIG. 3.

[0036] Accordingly, after the compression, when the recurrent neural network 110 is a standard recurrent neural network, the layer output generated by the compressed recurrent layer 120 satisfies:

$$h_t^l = \sigma(\mathbf{W}_x^{l-1} h_{t-1}^{l-1} + \mathbf{Z}_h^l \mathbf{P}^l h_{t-1}^l + b^l) ,$$

while the layer output generated by the recurrent layer 130 satisfies:

$$h_t^{l+1} = \sigma(\mathbf{Z}_x^l \mathbf{P}^l h_t^l + \mathbf{W}_h^{l+1} h_{t-1}^{l+1} + b^{l+1}).$$

[0037] In some other implementations, the recurrent layers in the recurrent neural network 110 are long short-term memory (LSTM) layers and the state of a given LSTM layer and the layer output of the LSTM layer are different. To generate a layer output, a given LSTM layer applies multiple gates to the current layer input and the current layer state, to generate the new layer output and to update the current layer state to generate the new layer state. Thus, as part of generating a layer output, the LSTM will generally apply multiply different weight matrices to both the current layer input and the current layer state. The operation of LSTM layers is described in more detail in H. Sak, A. Senior, and F. Beaufays, “Long short-term memory recurrent neural network architectures for large scale acoustic modeling,” in Proc. of Interspeech, 2014, pp. 338–342.

[0038] When the recurrent layers are LSTM layers, the recurrent weight matrix for a given LSTM layer can be considered to be the vertical concatenation of the weight matrices that the given LSTM layer applies to the current layer state. The inter-layer weight matrix for a given LSTM layer can be considered to be the vertical concatenation of the weight matrices that a next LSTM layer applies to the layer outputs generated by the given LSTM layer.

[0039] While in the example of FIG. 1 only the compressed recurrent layer 120 is compressed, in some cases more than one of the recurrent layers or even all of the recurrent layers in the recurrent neural network 110 can each be compressed as described below with reference to FIG. 3.

[0040] In some implementations, jointly compressing the recurrent weight matrix and the inter-layer weight matrix such that the projection matrix is shared across the recurrent and inter-layer weight matrices as described in this specification may allow for more efficient parameterization of the weight matrices. In some implementations, where the recurrent neural network 110 is a speech recognition model, the above described techniques may be used to compress the recurrent neural network 110 by at least 68% while achieving a word error rate that is within 5% of the uncompressed model.

[0041] FIG. 2 is a flow diagram of an example process 200 for compressing a recurrent neural network. For convenience, the process 200 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 in accordance with this specification, can perform the process 200.

[0042] The system trains an uncompressed recurrent neural network on training data (step 202) to determine trained values of the parameters in the weight matrices of the layers of the uncompressed recurrent neural network. For example, the system can train the recurrent neural network 110 of FIG. 1 before any of the recurrent layers in the network have been compressed. The system can train the uncompressed recurrent neural network using conventional recurrent neural network training techniques, e.g., stochastic gradient descent with backpropagation through time.

[0043] The system compresses one or more of the recurrent layers in the recurrent neural network (step 204). In particular, for each recurrent layer that is to be compressed, the system generates a first compressed weight matrix, Z_h^1 , and a projection matrix, P^1 so that the

product of the first compressed weight matrix and the projection matrix approximates the recurrent weight matrix \mathbf{W}_h^l of the recurrent layer and generates a second compressed weight matrix, \mathbf{Z}_x^l , based on the first compressed weight matrix, \mathbf{Z}_h^l , and the projection matrix, \mathbf{P}^l so that the product of the second compressed weight matrix and the projection matrix approximates the inter-layer weight matrix of the recurrent layer. Compressing a particular recurrent layer is described in more detail below with reference to FIG. 3. Because the compression of a given layer depends only on the recurrent weight matrix and the inter-layer weight matrix of the layer, i.e., and not on the matrices of any other layers, the system can compress a single recurrent layer, multiple recurrent layers, or all of the recurrent layers in the recurrent neural network.

[0044] The system re-configures the recurrent neural network with the compressed weight matrices (step 206). That is, for each recurrent layer that was compressed, the system replaces the recurrent weight matrix for the layer with the product of the first compressed weight matrix and the projection matrix and the inter-layer weight matrix for the layer with the product of the second compressed weight matrix and the projection matrix. Because the product of the first compressed weight matrix and the projection matrix is of a lower rank than the recurrent weight matrix and the product of the second compressed weight matrix and the projection matrix the inter-layer weight matrix of the layer, the matrices include fewer parameters than their corresponding matrices in the uncompressed neural network.

[0045] Optionally, after re-configuring the recurrent neural network, the system can fine-tune the performance of the compressed neural network by training the re-configured neural network on additional training data to further adjust the trained values of the parameters while maintaining the ranks of the projection matrices and the compressed weight matrices, i.e., constraining the ranks of the projection matrices and the compressed weight matrices to not increase.

[0046] The system stores the weight matrices of the re-configured neural network for use in instantiating a trained neural network, i.e., a trained recurrent neural network that can effectively be used to process neural network inputs (step 208). In some cases, in addition to or instead of storing the weight matrices, the system can transmit the weight matrices and other data defining the configuration of the neural network to another system for use in implementing a trained recurrent neural network. For example, the system can transmit the

configuration data to a mobile device to allow the compressed recurrent neural network to be implemented on the mobile device.

[0047] FIG. 3 is a flow diagram of an example process 300 for compressing an inter-layer weight matrix and a recurrent weight matrix for a particular recurrent layer. For convenience, the process 300 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 in accordance with this specification, can perform the process 300.

[0048] The system determines a singular value decomposition (SVD) of the recurrent weight matrix \mathbf{W}_h for the particular recurrent layer (step 302). The singular value decomposition of the recurrent weight matrix is a decomposition of the matrix \mathbf{W}_h into a first unitary matrix \mathbf{U} , a rectangular diagonal matrix $\mathbf{\Sigma}$, and a second unitary matrix \mathbf{V} . In particular, the SVD satisfies:

$$\mathbf{W}_h = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T.$$

[0049] The system can determine the SVD of the recurrent weight matrix using a known SVD decomposition technique. For example, the system can first reduce the recurrent weight matrix to a bidiagonal matrix and then compute the SVD of the bidiagonal matrix using an iterative method, e.g., a variant of the QR algorithm.

[0050] The system truncates the SVD to generate the first compressed weight matrix \mathbf{Z}_h^l and the projection matrix \mathbf{P}^l (step 304). In particular, the system truncates the SVD by (i) retaining the top, i.e., highest, l values in the rectangular diagonal matrix $\mathbf{\Sigma}$ and setting the remaining values to zero, (ii) retaining the top l values in each singular vector, i.e., column, of the first unitary matrix \mathbf{U} and setting the remaining values to zero, and (iii) retaining the top l values in each singular vector, i.e., column, of the second unitary matrix \mathbf{V} and setting the remaining values to zero.

[0051] Generally, l is a value that is less than the dimensionality of the weight matrices and that has been configured to control the degree of compression applied to the recurrent layer. That is, the smaller the value of l , the higher the degree of compression that is applied.

[0052] In some implementations, l is a predetermined value.

[0053] In some other implementations, the system determines l so that the truncated SVD retains at most a predetermined threshold fraction τ of the explained variance in the SVD

operation. In particular, the system can set l to be the value for which the ratio of (i) the sum of the squares of the top, i.e., highest, l values in the rectangular diagonal matrix Σ to (ii) the sum of the squares of all of the values in the rectangular diagonal matrix Σ is greatest while still being less than τ .

[0054] In these implementations, if the system is compressing multiple recurrent layers within the recurrent neural network, the value of l will likely differ between the multiple layers, i.e., because a different l will satisfy the above criteria for SVDs of different matrices.

[0055] Once the system has truncated the SVD to generate a truncated first unitary matrix \tilde{U} , a truncated rectangular diagonal matrix $\tilde{\Sigma}$, and a truncated second unitary matrix \tilde{V} , the system can set the first compressed weight matrix \mathbf{Z}_h^1 equal to $\tilde{U}\tilde{\Sigma}$ and the projection matrix \mathbf{P}^1 equal to \tilde{V}^T .

[0056] The system determines the second compressed weight matrix \mathbf{Z}_x^1 from the first compressed weight matrix \mathbf{Z}_h^1 and the projection matrix \mathbf{P}^1 (step 306). In particular, the system determines the second compressed weight matrix by solving the following least-squares problem:

$$\mathbf{Z}_x^1 = \arg \min_Y (\|Y\mathbf{P}^1 - \mathbf{W}_x^1\|_F^2),$$

where $\|X\|_F$ denotes the Frobenius norm of the matrix X .

[0057] Embodiments 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 specification and their structural equivalents, or in combinations of one or more of them. Embodiments 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 program carrier for execution by, or to control the operation of, data processing apparatus. 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. 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.

[0058] 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.

[0059] A computer program (which may also be referred to or described as a program, software, a software application, 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 computer 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 communication network.

[0060] 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, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit).

[0061] Computers suitable for the execution of a computer program include, by way of example, 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. 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.

[0062] 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. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0063] 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 client device in response to requests received from the web browser.

[0064] Embodiments 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 or a Web

browser 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.

[0065] 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.

[0066] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or 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. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[0067] 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.

[0068] Particular embodiments of the subject matter have been described. Other embodiments 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 certain implementations, multitasking and parallel processing may be advantageous.

[0069] What is claimed is:

CLAIMS

1. A system comprising:
 - a compressed recurrent neural network (RNN) implemented by one or more computers, the compressed RNN comprising a plurality of recurrent layers,
 - wherein each of the plurality of recurrent layers of the RNN is configured to, for each of a plurality of time steps, receive a respective layer input for the time step and process the layer input for the time step to generate a respective layer output for the time step,
 - wherein each of the plurality of recurrent layers has a respective recurrent weight matrix, \mathbf{W}_h , and a respective inter-layer weight matrix, \mathbf{W}_x , and
 - wherein at least one of the plurality of recurrent layers is compressed such that a respective recurrent weight matrix, \mathbf{W}_h^l , of the compressed layer is defined by a first compressed weight matrix, \mathbf{Z}_h^l , and a projection matrix, \mathbf{P}^l , and a respective inter-layer weight matrix, \mathbf{W}_x^l , of the compressed layer is defined by a second compressed weight matrix, \mathbf{Z}_x^l , and the projection matrix, \mathbf{P}^l .
2. The system of claim 1, wherein each of the plurality of recurrent layers is configured to, for each time step:
 - generate the respective layer output for the time step by applying an inter-layer weight matrix for a previous layer to a current input to the layer and applying a recurrent weight matrix for the layer to a recurrent input to the layer.
3. The system of any preceding claim, wherein the at least one of the plurality of recurrent layers that is compressed is configured to, for each time step:
 - generate the respective layer output for the time step in part by applying the first compressed weight matrix and the projection matrix to a respective recurrent input to the layer.
4. The system of any preceding claim, wherein the respective recurrent weight matrix is defined by the product of the first compressed weight matrix and the projection matrix and wherein the respective inter-layer weight matrix is defined by the product of the second compressed matrix and the projection matrix.

5. The system of any preceding claim, wherein the RNN is an acoustic model.
6. The system of any one of claims 1-4, wherein the RNN is a speech recognition model.
7. The system of any preceding claim, wherein the RNN is at least a 68% compression of an uncompressed version of the RNN.
8. The system of claim 7, wherein a word error rate of the RNN is within 5% of a word error rate of an uncompressed version of the RNN.
9. A method for compressing a recurrent neural network (RNN), wherein the RNN is implemented by one or more computers and comprises a plurality of recurrent layers,
wherein each of the plurality of recurrent layers of the RNN is configured to, for each of a plurality of time steps, receive a respective layer input for the time step and process the layer input to generate a respective layer output for the time step, and
wherein each recurrent layer has a respective recurrent weight matrix, \mathbf{W}_h , and a respective inter-layer weight matrix, \mathbf{W}_x , the method comprising, for one of the plurality of recurrent layers:
generating a first compressed weight matrix, \mathbf{Z}_h^1 , and a projection matrix, \mathbf{P}^1 , wherein the product of the first compressed weight matrix and the projection matrix approximates the recurrent weight matrix of the one of the plurality of recurrent layers; and
generating a second compressed weight matrix, \mathbf{Z}_x^1 , based on the first compressed weight matrix, \mathbf{Z}_h^1 , and the projection matrix, \mathbf{P}^1 , wherein the product of the second compressed weight matrix and the projection matrix approximates the inter-layer weight matrix of the one of the plurality of recurrent layers.
10. The method of claim 9, further comprising, for the one of the plurality of recurrent layers:

replacing the respective recurrent weight matrix with the product of the first compressed weight matrix and the projection matrix; and

replacing the respective inter-layer weight matrix with the product of the second compressed weight matrix and the projection matrix.

11. The method of any one of claims 9 or 10, wherein the respective ranks \mathbf{Z}_h and \mathbf{P} of are less than the rank of \mathbf{W}_h .

12. The method of any one of claims 9-11, wherein the respective ranks \mathbf{Z}_x and \mathbf{P} of are less than the rank of \mathbf{W}_x .

13. The method of any one of claims 9-12, wherein generating the second compressed weight matrix, \mathbf{Z}_x , comprises inverting the projection matrix and multiplying the inverted projection matrix by the inter-layer weight matrix.

14. The method of any one of claims 9-13, wherein singular value decomposition (SVD) compression techniques are used to generate the first compressed weight matrix, \mathbf{Z}_h , and the projection matrix, \mathbf{P} .

15. The method of any one of claims 9-14, wherein the one of the plurality of layers is an l -th layer, and wherein the output for the l -th layer can be expressed by $h_t^l = \sigma(\mathbf{W}_x^{l-1}h_t^{l-1} + \mathbf{Z}_h^l\mathbf{P}^lh_{t-1}^l + b^l)$, wherein h_t^l represents a hidden layer activation output of the l -th layer at time t , \mathbf{W}_x^{l-1} represents an inter-layer weight matrix from a previous, $(l-1)$ -th, layer b^l represents an l -th layer bias vector, and $\sigma(\cdot)$ denotes a non-linear activation function.

16. The method of any one of claims 9-15, wherein the RNN is a long short-term memory (LSTM) RNN.

17. The method of any one of claims 9-16, further comprising training the RNN before generating the first compressed weight matrix, \mathbf{Z}_h , the projection matrix, \mathbf{P} , and the second

compressed weight matrix, \mathbf{Z}_x ; and re-training the RNN after generating the first compressed weight matrix, \mathbf{Z}_h , the projection matrix, \mathbf{P} , and the second compressed weight matrix, \mathbf{Z}_x .

18. The method of any one of claims 9-17, wherein the RNN is an acoustic model.

19. The method of any one of claims 9-17, wherein the RNN is a speech recognition model.

20. A system comprising one or more computers and one or more storage devices storing instructions that are operable, when executed by the one or more computers, to cause the one or more computers to perform the operations of the respective method of any one of claims 10-19.

21. A non-transitory computer storage medium encoded with instructions that are operable, when executed by one or more computers, cause the one or more computers to perform the operations of the respective method of any one of claims 10-19.

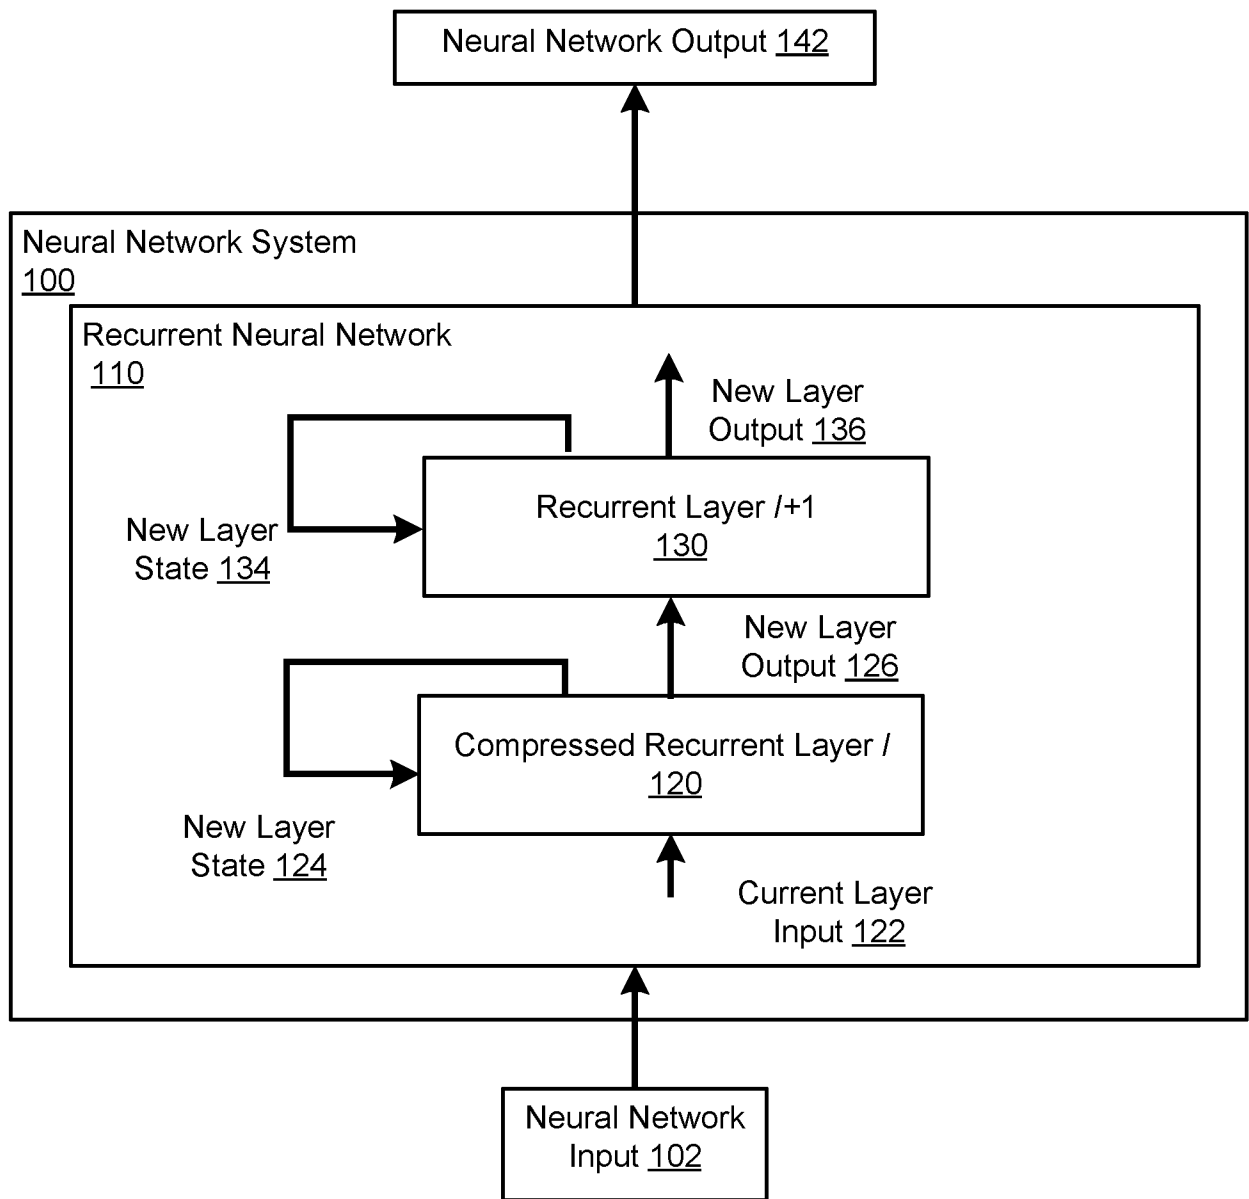


FIG. 1

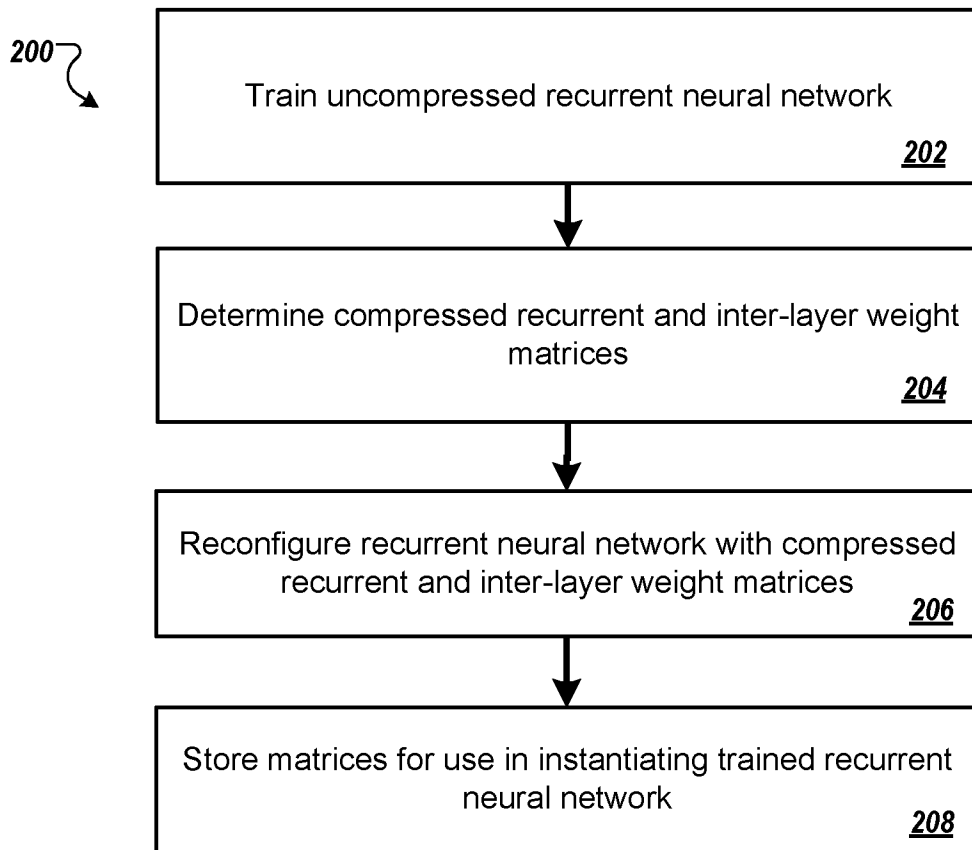


FIG. 2

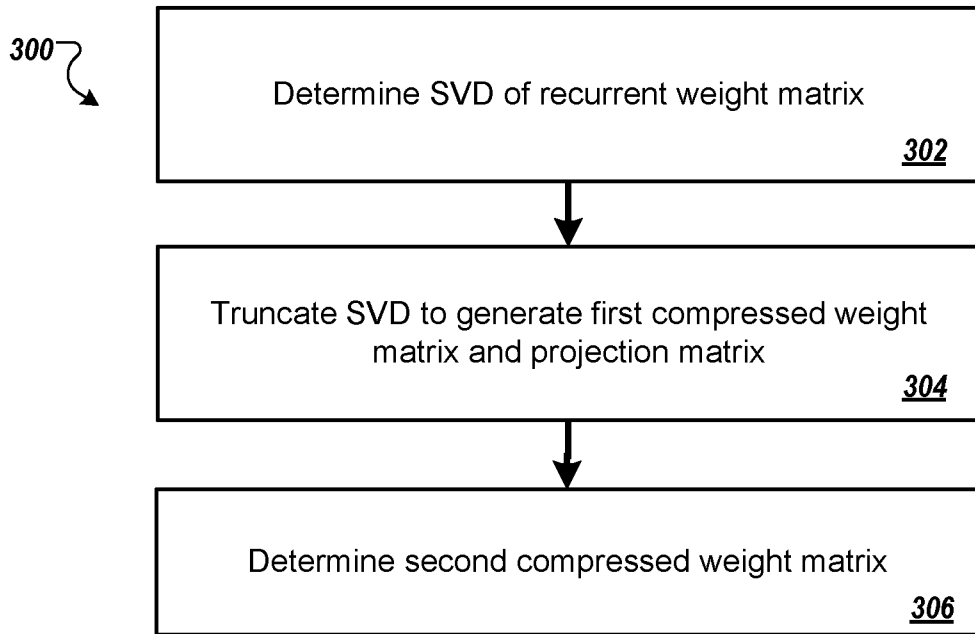


FIG. 3

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/068913

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06N3/04
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
G06N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>Hasim Sak ET AL: "Long Short-Term Memory Based Recurrent Neural Network Architectures for Large Vocabulary Speech Recognition", 5 February 2014 (2014-02-05), XP055274936, Retrieved from the Internet: URL:http://static.googleusercontent.com/media/research.google.com/en//pubs/archive/43895.pdf [retrieved on 2016-05-24] abstract page 1 - page 2 figure 1</p> <p style="text-align: center;">-----</p>	1-21

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

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

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

Date of the actual completion of the international search

16 March 2017

Date of mailing of the international search report

27/03/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Appeltant, Lennert