

(19) **DANMARK**

(10) **DK/EP 3251059 T3**



(12) **Oversættelse af  
europæisk patentskrift**

Patent- og  
Varemærkestyrelsen

- 
- (51) Int.Cl.: **G 06 N 3/08 (2006.01)** **G 06 N 3/04 (2006.01)**
- (45) Oversættelsen bekendtgjort den: **2019-04-01**
- (80) Dato for Den Europæiske Patentmyndigheds bekendtgørelse om meddelelse af patentet: **2018-12-05**
- (86) Europæisk ansøgning nr.: **16704121.9**
- (86) Europæisk indleveringsdag: **2016-01-28**
- (87) Den europæiske ansøgnings publiceringsdag: **2017-12-06**
- (86) International ansøgning nr.: **US2016015476**
- (87) Internationalt publikationsnr.: **WO2016123409**
- (30) Prioritet: **2015-01-28 US 201562108984 P**
- (84) Designerede stater: **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**
- (73) Patenthaver: **Google LLC, 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA**
- (72) Opfinder: **IOFFE, Sergey, 1600 Amphitheatre Parkway, Mountain View, California 94043, USA**  
**CORTES, Corinna, 1600 Amphitheatre Parkway, Mountain View, California 94043, USA**
- (74) Fuldmægtig i Danmark: **Marks & Clerk (Luxembourg) LLP, 44 rue de la Vallée, B.P. 1775, L-1017 Luxembourg, Luxembourg**
- (54) Benævnelse: **BATCH-NORMALISERINGSLAG**
- (56) Fremdragne publikationer:  
**CN-A- 1 945 602**  
**US-A1- 2003 236 662**  
**Äzagliar GülÄşehre ET AL: "Knowledge Matters: Importance of Prior Information for Optimization", , 13 July 2013 (2013-07-13), pages 1-37, XP055266282, Retrieved from the Internet:  
URL:http://arxiv.org/pdf/1301.4083v6.pdf [retrieved on 2016-04-18]**  
**LECUN YANN A ET AL: "Efficient BackProp", 1 January 1901 (1901-01-01), CORRECT SYSTEM DESIGN; [LECTURE NOTES IN COMPUTER SCIENCE; LECT.NOTES COMPUTER], SPRINGER INTERNATIONAL PUBLISHING, CHAM, PAGE(S) 9 - 48, XP047292571, ISSN: 0302-9743 ISBN: 978-3-642-23953-3 page 13 - page 17 page 30 - page 35 page 40 - page 46**  
**Sergey Ioffe ET AL: "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift", , 2 March 2015 (2015-03-02), pages 1-11, XP055266268, Retrieved from the Internet:  
URL:http://arxiv.org/pdf/1502.03167v3.pdf [retrieved on 2016-04-18]**



# DESCRIPTION

## BACKGROUND

**[0001]** This specification relates to processing inputs through the layers of neural networks to generate outputs.

**[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. Gulcehre et al: "Knowledge Matters: Importance of Prior Information for Optimization", 13 July 2013, <http://arxiv.org/pdf/1301.4083v6.pdf> described the use of normalization at the output layer of a neural network to encourage sparse activations.

## SUMMARY

**[0003]** In general, one innovative aspect of the subject matter described in this specification can be embodied in a neural network system implemented by one or more computers that includes a batch normalization layer between a first neural network layer and a second neural network layer, wherein the first neural network layer generates first layer outputs having a plurality of components, where the batch normalization layer is configured to, during training of the neural network system on a batch of training examples: receive a respective first layer output for each training example in the batch; compute a plurality of normalization statistics for the batch from the first layer outputs; normalize each component of each first layer output using the normalization statistics to generate a respective normalized layer output for each training example in the batch; generate a respective batch normalization layer output for each of the training examples from the normalized layer outputs; and provide the batch normalization layer output as an input to the second neural network layer.

**[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]** Particular embodiments of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. A neural network system that includes one or more batch normalization layers can be trained more quickly than

an otherwise identical neural network that does not include any batch normalization layers. For example, by including one or more batch normalization layers in the neural network system, problems caused by the distribution of a given layer's inputs changing during training can be mitigated. This may allow higher learning rates to be effectively used during training and may reduce the impact of how parameters are initialized on the training process. Additionally, during training, the batch normalization layers can act as a regularizer and may reduce the need for other regularization techniques, e.g., dropout, to be employed during training. Once trained, the neural network system that includes one normalization layers can generate neural network outputs that are as accurate, if not more accurate, than the neural network outputs generated by the otherwise identical neural network system.

**[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. The invention is defined by the appended claims. Embodiments which do not fall within the scope of the claims do not describe part of the present invention.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

**[0007]**

FIG. 1 shows an example neural network system.

FIG. 2 is a flow diagram of an example process for processing an input using a batch normalization layer during training of the neural network system.

FIG. 3 is a flow diagram of an example process for processing an input using a batch normalization after the neural network system has been trained.

**[0008]** Like reference numbers and designations in the various drawings indicate like elements.

#### **DETAILED DESCRIPTION**

**[0009]** This specification describes a neural network system implemented as computer programs on one or more computers in one or more locations that includes a batch normalization layer.

**[0010]** 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.

**[0011]** The neural network system 100 includes multiple neural network layers that are arranged in a sequence from a lowest layer in the sequence to a highest layer in the sequence. The neural network system generates neural network outputs from neural network inputs by processing the neural network inputs through each of the layers in the sequence.

**[0012]** 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.

**[0013]** 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.

**[0014]** 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.

**[0015]** As another example, if the inputs to the neural network system 100 are features of an impression context for a particular advertisement, the output generated by the neural network system 100 may be a score that represents an estimated likelihood that the particular advertisement will be clicked on.

**[0016]** 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.

**[0017]** 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.

**[0018]** As another example, if the input to the neural network system 100 is a spoken utterance, a sequence of spoken utterances, or features derived from one of the two, 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 transcript for the utterance or sequence of utterances.

**[0019]** As another example, the neural network system 100 can be part of an autocompletion system or part of a text processing system.

**[0020]** As another example, the neural network system 100 can be part of a reinforcement learning system and can generate outputs used for selecting actions to be performed by an agent interacting with an environment.

**[0021]** In particular, each of the layers of the neural network is configured to receive an input and generate an output from the input and the neural network layers collectively process neural network inputs received by the neural network system 100 to generate a respective neural network output for each received neural network input. Some or all of the neural network layers in the sequence generate outputs from inputs in accordance with current values of a set of parameters for the neural network layer. For example, some layers may multiply the received input by a matrix of current parameter values as part of generating an output from the received input.

**[0022]** The neural network system 100 also includes a batch normalization layer 108 between a neural network layer A 104 and a neural network layer B 112 in the sequence of neural network layers. The batch normalization layer 108 is configured to perform one set of operations on inputs received from the neural network layer A 104 during training of the neural network system 100 and another set of operations on inputs received from the neural network layer A 104 after the neural network system 100 has been trained.

**[0023]** In particular, the neural network system 100 can be trained on multiple batches of training examples in order to determine trained values of the parameters of the neural network layers. A batch of training examples is a set of multiple training examples. For example, during training, the neural network system 100 can process a batch of training examples 102 and generate a respective neural network output for each training example in the batch 102. The neural network outputs can then be used to adjust the values of the parameters of the neural network layers in the sequence, e.g., through conventional gradient descent and backpropagation neural network training techniques.

**[0024]** During training of the neural network system 100 on a given batch of training examples, the batch normalization layer 108 is configured to receive layer A outputs 106 generated by the neural network layer A 104 for the training examples in the batch, process the layer A outputs 106 to generate a respective batch normalization layer output 110 for each training example in the batch, and then provide the batch normalization layer outputs 110 as an input to the neural network layer B 112. The layer A outputs 106 include a respective output generated by the neural network layer A 104 for each training example in the batch. Similarly, the batch normalization layer outputs 110 include a respective output generated by the batch normalization layer 108 for each training example in the batch.

**[0025]** Generally, the batch normalization layer 108 computes a set of normalization statistics for the batch from the layer A outputs 106, normalizes the layer A outputs 106 to generate a respective normalized output for each training example in the batch, and, optionally, transforms each of the normalized outputs before providing the outputs as input to the neural network layer B 112.

**[0026]** The normalization statistics computed by the batch normalization layer 108 and the manner in which the batch normalization layer 108 normalizes the layer A outputs 106 during training depend on the nature of the neural network layer A 104 that generates the layer A outputs 106.

**[0027]** In some cases, the neural network layer A 104 is a layer that generates an output that includes multiple components indexed by dimension. For example, the neural network layer A 104 may be a fully-connected neural network layer. In some other cases, however, the neural network layer A 104 is a convolutional layer or other kind of neural network layer that generates an output that includes multiple components that are each indexed by both a feature index and a spatial location index. Generating the batch normalization layer output during training of the neural network system 100 in each of these two cases is described in more detail below with reference to FIG. 2.

**[0028]** Once the neural network system 100 has been trained, the neural network system 100 may receive a new neural network input for processing and process the neural network input through the neural network layers to generate a new neural network output for the input in accordance with the trained values of the parameters of the components of the neural network system 100. The operations performed by the batch normalization layer 108 during the processing of the new neural network input also depend on the nature of the neural network layer A 104. Processing a new neural network input after the neural network system 100 has been trained is described in more detail below with reference to FIG. 3.

**[0029]** The batch normalization layer 108 may be included at various locations in the sequence of neural network layers and, in some implementations, multiple batch normalization layers may be included in the sequence.

**[0030]** In the example of FIG. 1, in some implementations, the neural network layer A 104 generates outputs by modifying inputs to the layer in accordance with current values of a set of parameters for the first neural network layer, e.g., by multiplying the input to the layer by a matrix of the current parameter values. In these implementations, the neural network layer B 112 may receive an output from the batch normalization layer 108 and generate an output by applying a non-linear operation, i.e., a non-linear activation function, to the batch normalization layer output. Thus, in these implementations, the batch normalization layer 108 is inserted within a conventional neural network layer, and the operations of the conventional neural network layer are divided between the neural network layer A 104 and the neural network layer B 112.

**[0031]** In some other implementations, the neural network layer A 104 generates the outputs by modifying layer inputs in accordance with current values of a set of parameters to generate a modified first layer inputs and then applying a non-linear operation to the modified first layer inputs before providing the output to the batch normalization layer 108. Thus, in these implementations, the batch normalization layer 108 is inserted after a conventional neural network layer in the sequence.

**[0032]** FIG. 2 is a flow diagram of an example process 200 for generating a batch normalization layer output during training of a neural network on a batch of training examples. 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 batch normalization layer included in a neural network system, e.g., the batch normalization layer 108 included in the neural network system 100 of FIG.1, appropriately programmed, can perform the process 200.

**[0033]** The batch normalization layer receives lower layer outputs for the batch of training examples (step 202). The lower layer outputs include a respective output generated for each training example in the batch by the layer below the batch normalization layer in the sequence of neural network layers.

**[0034]** The batch normalization layer generates a respective normalized output for each training example in the batch (step 204). That is, the batch normalization layer generates a respective normalized output from each received lower layer output.

**[0035]** In some cases, the layer below the batch normalization layer is a layer that generates an output that includes multiple components indexed by dimension.

**[0036]** In these cases, the batch normalization layer computes, for each dimension, the mean and the standard deviation of the components of the lower layer outputs that correspond to the dimension. The batch normalization layer then normalizes each component of each of the lower level outputs using the means and standard deviations to generate a respective normalized output for each of the training examples in the batch. In particular, for a given component of a given output, the batch normalization layer normalizes the component using the mean and the standard deviation computed for the dimension corresponding to the component. For example, in some implementations, for a component  $x_{k,i}$  corresponding to the k-th dimension of the i-th lower layer output from a batch  $\beta$ , the normalized output  $\hat{x}_{k,i}$  satisfies:

$$\hat{x}_{k,i} = \frac{x_{k,i} - \mu_B}{\sigma_B},$$

where  $\mu_B$  is the mean of the components corresponding to the k-th dimension of the lower layer outputs in the batch  $\beta$  and  $\sigma_B$  is the standard deviation of the components corresponding to the k-th dimension of the lower layer outputs in the batch  $\beta$ . In some implementations, the standard deviation is a numerically stable standard deviation that is equal to  $(\sigma_B^2 + \epsilon)^{1/2}$ , where  $\epsilon$  is a constant value and  $\sigma_B^2$  is the variance of the components corresponding to the k-th dimension of the lower layer outputs in the batch  $\beta$ .



**[0037]** In some other cases, however, the neural network layer below the batch normalization layer is a convolutional layer or other kind of neural network layer that generates an output that includes multiple components that are each indexed by both a feature index and a spatial location index.

**[0038]** In some of these cases, the batch normalization layer computes, for each possible feature index and spatial location index combination, the mean and the variance of the components of the lower layer outputs that have that feature index and spatial location index. The batch normalization layer then computes, for each feature index, the average of the means for the feature index and spatial location index combinations that include the feature index. The batch normalization layer also computes, for each feature index, the average of the variances for the feature index and spatial location index combinations that include the feature index. Thus, after computing the averages, the batch normalization layer has computed a mean statistic for each feature across all of the spatial locations and a variance statistic for each feature across all of the spatial locations.

**[0039]** The batch normalization layer then normalizes each component of each of the lower level outputs using the average means and the average variances to generate a respective normalized output for each of the training examples in the batch. In particular, for a given component of a given output, the batch normalization layer normalizes the component using the average mean and the average variance for the feature index corresponding to the component, e.g., in the same manner as described above when the layer below the batch normalization layer generates outputs indexed by dimension.

**[0040]** In others of these cases, the batch normalization layer computes, for each feature index the mean and the variance of the components of the lower layer outputs that correspond to the feature index, i.e., that have the feature index.

**[0041]** The batch normalization layer then normalizes each component of each of the lower level outputs using the means and the variances for the feature indices to generate a respective normalized output for each of the training examples in the batch. In particular, for a given component of a given output, the batch normalization layer normalizes the component using the mean and the variance for the feature index corresponding to the component, e.g., in the same manner as described above when the layer below the batch normalization layer generates outputs indexed by dimension.

**[0042]** Optionally, the batch normalization layer transforms each component of each normalized output (step 206).

**[0043]** In cases where the layer below the batch normalization layer is a layer that generates an output that includes multiple components indexed by dimension, the batch normalization layer transforms, for each dimension, the component of each normalized output in the dimension in accordance with current values of a set of parameters for the dimension. That is,

the batch normalization layer maintains a respective set of parameters for each dimension and uses those parameters to apply a transformation to the components of the normalized outputs in the dimension. The values of the sets of parameters are adjusted as part of the training of the neural network system. For example, in some implementations, the transformed normalized output  $y_{k,i}$  generated from the normalized output  $\hat{x}_{k,i}$  satisfies:

$$y_{k,i} = \gamma_k \hat{x}_{k,i} + A_k,$$

where  $\gamma_k$  and  $A_k$  are the parameters for the  $k$ -th dimension.

**[0044]** In cases where the layer below the batch normalization layer is a convolutional layer, the batch normalization layer transforms, for each component of each of the normalized outputs, the component in accordance with current values of a set of parameters for the feature index corresponding to the component. That is, the batch normalization layer maintains a respective set of parameters for each feature index and uses those parameters to apply a transformation to the components of the normalized outputs that have the feature index, e.g., in the same manner as described above when the layer below the batch normalization layer generates outputs indexed by dimension. The values of the sets of parameters are adjusted as part of the training of the neural network system.

**[0045]** The batch normalization layer provides the normalized outputs or the transformed normalized outputs as input to a layer above the batch normalization layer in the sequence (step 208).

**[0046]** After the neural network has generated the neural network outputs for the training examples in the batch, the normalization statistics are backpropagated through as part of adjusting the values of the parameters of the neural network, i.e., as part of performing the backpropagation training technique.

**[0047]** FIG. 3 is a flow diagram of an example process 300 for generating a batch normalization layer output for a new neural network input after the neural network has been trained. 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 batch normalization layer included in a neural network system, e.g., the batch normalization layer 108 included in the neural network system 100 of FIG. 1, appropriately programmed, can perform the process 300.

**[0048]** The batch normalization layer receives a lower layer output for the new neural network input (step 302). The lower layer output is an output generated for the new neural network input by the layer below the batch normalization layer in the sequence of neural network layers.

**[0049]** The batch normalization layer generates a normalized output for the new neural network input (step 304).

**[0050]** If the outputs generated by the layer below the batch normalization layer are indexed

by dimension, the batch normalization layer normalizes each component of the lower layer output using pre-computed means and standard deviations for each of the dimensions to generate a normalized output. In some cases, the means and standard deviations for a given dimension are computed from the components in the dimension of all of outputs generated by the layer below the batch normalization layer during the training of the neural network system.

**[0051]** In some other cases, however, the means and standard deviations for a given dimension are computed from the components in the dimension of the lower layer outputs generated by the layer below the batch normalization layer after training, e.g., from lower layer outputs generated during in a most recent time window of specified duration or from a specified number of lower layer outputs most recently generated by the layer below the batch normalization layer.

**[0052]** In particular, in some cases the distribution of network inputs and, accordingly, the distribution of lower layer outputs may change between the training examples used during training and the new neural network inputs used after the neural network system is trained, e.g., if the new neural network inputs are different kinds of inputs from the training examples. For example, the neural network system may have been trained on user images and may now be used to process video frames. The user images and the video frames likely have different distributions in terms of the classes pictured, image properties, composition, and so on. Therefore, normalizing the lower layer inputs using statistics from the training may not accurately capture the statistics of the lower layer outputs being generated for the new inputs. Thus, in these cases, the batch normalization layer can use normalization statistics computed from lower layer outputs generated by the layer below the batch normalization layer after training.

**[0053]** If the outputs generated by the layer below the batch normalization layer are indexed by feature index and spatial location index, the batch normalization layer normalizes each component of the lower layer output using pre-computed average means and average variances for each of the feature indices, to generate a normalized output. In some cases, as described above, the average means and average variances for a given feature index, are computed from the outputs generated by the layer below the batch normalization layer for all of the training examples used during training. In some other cases, as described above, the means and standard deviations for a given feature index are computed from the lower layer outputs generated by the layer below the batch normalization layer after training.

**[0054]** Optionally, the batch normalization layer transforms each component of the normalized output (step 306).

**[0055]** If the outputs generated by the layer below the batch normalization layer are indexed by dimension, the batch normalization layer transforms, for each dimension, the component of the normalized output in the dimension in accordance with trained values of the set of parameters for the dimension. If the outputs generated by the layer below the batch normalization layer are indexed by feature index and spatial location index, the batch

normalization layer transforms each component of the normalized output in accordance with trained values of the set of parameters for the feature index corresponding to the component. The batch normalization layer provides the normalized output or the transformed normalized output as input to the layer above the batch normalization layer in the sequence (step 308).

**[0056]** 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.

**[0057]** The term "data processing apparatus" 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 include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit). The apparatus can also include, in addition to hardware, code that creates an execution environment for the computer program in question, 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.

**[0058]** 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.

**[0059]** 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).

**[0060]** 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.

**[0061]** 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.

**[0062]** 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.

**[0063]** 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.

**[0064]** 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.

**[0065]** 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.

**[0066]** 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.

**[0067]** Particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. For example, some 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.

## **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.

**Non-patent literature cited in the description**

- **GULCEHRE** Knowledge Matters: Importance of Prior Information for Optimization, 2013, [\[0002\]](#)

## BATCH-NORMALISERINGSLAG

## PATENTKRAV

1. Neuralt netværkssystem (100), der er implementeret af én eller flere computere, hvilket neuralt netværkssystem omfatter:
  - 5 et batch-normaliseringslag (108) mellem et første neuralt netværkslag (104) og et andet neuralt netværkslag (112), hvor det første neurale netværkslag genererer første lag-output med en flerhed af komponenter, og hvor batch-normaliseringslaget er konfigureret til, under træning af det neurale netværkssystem på en batch af træningseksempler, at:
    - modtage et tilsvarende første lag-output for hvert træningseksempel i batchen;
    - 10 beregne en flerhed af normaliseringsstatistikker for batchet ud fra de første lag-outputs;
    - normalisere hver komponent af hvert første lag-output ved anvendelse af normaliseringsstatistikkerne til at generere et tilsvarende normaliseret lag-output for hvert træningseksempel i batchen;
    - generere et tilsvarende batch-normaliseringslag-output for hvert af træningseksemplerne ud fra de
    - 15 normaliserede lag-outputs; og
    - tilvejebringe batch-normaliseringslag-outputtet som et input til det andet neurale netværkslag;
    - hvor flerheden af komponenter af det første lag-output er indekseret efter dimension, og hvor beregning af en flerhed af normaliseringsstatistikker for de første lag-outputs omfatter:
      - beregning, for hver af dimensionerne, af en middelværdi for komponenterne af de første lag-
      - 20 outputs i dimensionen; og
      - beregning, for hver af dimensionerne, af en standardafvigelse for komponenterne af de første lag-outputs i dimensionen.
  2. Neuralt netværkssystem ifølge krav 1, hvor normalisering af hver komponent af hvert lag-output omfatter:
    - 25 normalisering af komponenten ved anvendelse af den beregnede middelværdi og beregnede standardafvigelse for den dimension, der svarer til komponenten.
  3. Neuralt netværkssystem ifølge krav 1 eller 2, hvor batch-normaliseringslaget bevarer et tilsvarende sæt af parametre for hver dimension, og hvor generering af det tilsvarende batch-normaliseringslag-output for hvert af træningseksemplerne ud fra de normaliserede lag-outputs omfatter:
    - 30 transformering, for hver dimension, af komponenten af det normaliserede lag-output for træningseksemplet i dimensionen i overensstemmelse med aktuelle værdier af et sæt af parametre for dimensionen.
  4. Neuralt netværkssystem ifølge krav 3, hvor batch-normaliseringslaget er konfigureret til, efter at det neurale netværkssystem er blevet trænet til at bestemme trænede værdier af parametrene for hver af
  - 35 dimensionerne, at:
    - modtage et nyt første lag-output genereret af det første neurale netværkslag for et nyt neuralt netværks-input;



- 2 -

- normalisere hver komponent af det nye første lag-output ved anvendelse af forudberegnete  
middel- og standardafvigelsesstatistikker for dimensionerne for at generere et nyt normaliseret lag-output;
- generere et nyt batch-normaliseringslag-output ved transformering, for hver dimension, af  
komponenten af det nye normaliserede lag-output for træningseksemplet i dimensionen i overensstemmelse  
5 med trænedede værdier af sættet af parametre for dimensionen; og
- tilvejebringe batch-normaliseringslag-outputtet som et nyt lag-input til det andet neurale  
netværkslag.
5. Neuralt netværkssystem ifølge krav 4, hvor de forudberegnete middel- og  
standardafvigelsesstatistikker for dimensionerne er beregnet ud fra de første lag-outputs genereret af det  
10 første neurale netværkslag under træning af det neurale netværkssystem eller er beregnet ud fra nye første  
lag-outputs genereret af det første neurale netværkslag, efter at det neurale netværk er blevet trænet.
6. Neuralt netværkssystem implementeret af én eller flere computere, hvilket neuralt netværkssystem  
omfatter:
- et batch-normaliseringslag mellem et første neuralt netværkslag og et andet neuralt netværkslag,  
15 hvor det første neurale netværkslag genererer første lag-outputs med en flerhed af komponenter, og hvor  
batch-normaliseringslaget er konfigureret til, under træning af det neurale netværkssystem på en batch af  
træningseksempler, at:
- modtage et tilsvarende første lag-output for hvert træningseksempel i batchen;  
beregne en flerhed af normaliseringsstatistikker for batchet fra de første lag-outputs;
- 20 normalisere hver komponent af hvert første lag-output ved anvendelse af  
normaliseringsstatistikkerne for at generere et tilsvarende normaliseret lag-output for hvert  
træningseksempel i batchen;
- generere et tilsvarende batch-normaliseringslag-output for hvert af træningseksemplerne ud fra de  
normaliserede lag-outputs; og
- 25 tilvejebringe batch-normaliseringslag-outputtet som et input til det andet neurale netværkslag;  
hvor det første neurale netværkslag er et konvolutionslag, hvor flerheden af komponenter af det  
første lag-output er indekseret efter karakteristikindeks og rumligt lokalitetsindeks, og hvor beregning af en  
flerhed af normaliseringsstatistikker for de første lag-outputs omfatter, for hvert af karakteristikindekserne:
- beregning af en middelværdi for komponenterne af de første lag-outputs, der svarer til  
30 karakteristikindekset; og
- beregning af en varians for komponenterne af de første lag-outputs, der svarer til  
karakteristikindekset.
7. Neuralt netværkssystem ifølge krav 6, hvor normalisering af hver komponent af hvert lag-output  
omfatter:
- 35 normalisering af komponenten ved anvendelse af middelværdien og variansen for det  
karakteristikindeks, der svarer til komponenten.
8. Neuralt netværkssystem ifølge krav 6 eller 7, hvor generering af det tilsvarende batch-  
normaliseringslag-output for hvert af træningseksemplerne fra de normaliserede lag-outputs omfatter:

- transformering af hver komponent af det normaliserede lag-output i overensstemmelse med aktuelle værdier af et sæt af parametre for det karakteristikeindeks, der svarer til komponenten.
9. Neuralt netværkssystem ifølge krav 8, hvor batch-normaliseringslaget er konfigureret til, efter at det neurale netværk er blevet trænet til at bestemme trænedede værdier af parametrene for hver af dimensionerne, at:
- 5 modtage et nyt første lag-input genereret fra et nyt neuralt netværks-input;
- normalisere hver komponent af det nye første lag-output ved anvendelse af forbergnede middel- og standardafvigelsesstatistikker for karakteristikeindekserne til at generere et nyt normaliseret lag-output;
- 10 generere et nyt batch-normaliseringslag-output ved transformering af hver komponent af det normaliserede lag-output i overensstemmelse med trænedede værdier af sættet af parametre for det karakteristikeindeks, der svarer til komponenten; og
- tilvejebringe det nye batch-normaliseringslag-output som et nyt lag-input til det andet neurale netværkslag.
10. Neuralt netværkssystem ifølge krav 6, hvor flerheden af komponenter af det første lag-output er indekseret efter karakteristikeindeks og rumligt lokalitetsindeks, og hvor beregning af en flerhed af normaliseringsstatistikker for de første lag-outputs endvidere omfatter:
- 15 beregning, for hver kombination af karakteristikeindeks og rumligt lokalitetsindeks, af en middelværdi for komponenterne af de første lag-outputs med karakteristikeindekset og rumligt lokalitetsindeks;
- 20 beregning, for hvert karakteristikeindeks, af et gennemsnit af middelværdierne for kombinationerne, der indbefatter karakteristikeindekset;
- beregning, for hver kombination af karakteristikeindeks og rumligt lokalitetsindeks, af en varians for komponenterne af de første lag-outputs med karakteristikeindekset og rumligt lokalitetsindeks; og
- 25 beregning, for hvert karakteristikeindeks, af et gennemsnit af varianserne for kombinationerne, der indbefatter karakteristikeindekset.
11. Neuralt netværkssystem ifølge krav 6, hvor normalisering af hver komponent af hvert lag-output omfatter:
- normalisering af komponenten ved anvendelse af gennemsnittet af middelværdierne og gennemsnittet af varianserne for det karakteristikeindeks, der svarer til komponenten.
- 30 12. Neuralt netværkssystem ifølge et hvilket som helst af kravene 1-11, hvor det første neurale netværkslag genererer de første lag-outputs ved modificering af de første lag-inputs i overensstemmelse med aktuelle værdier af et sæt af parametre for det første neurale netværkslag og hvor det andet neurale netværkslag genererer andet lag-outputs ved anvendelse af en ikke-lineær operation til batch-normaliseringslag-outputs; eller hvor det første neurale netværkslag genererer de første lag-outputs ved modificering af første lag-inputs i overensstemmelse med aktuelle værdier af et sæt af parametre for det første neurale netværkslag for at generere modificerede første lag-inputs og derefter anvendelse af en ikke-lineær operation til de modificerede første lag-inputs.
- 35

13. Neuralt netværkssystem ifølge et hvilket som helst af kravene 1-12, hvor, under træningen af det neurale netværk, det neurale netværkssystem er konfigureret til at tilbagepropagere normaliseringsstatistikkerne som en del af justering af værdier af det neurale netværks parametre.
14. Fremgangsmåde til behandling af data ved anvendelse af det neurale netværkssystem ifølge et hvilket som helst af kravene 1-13, hvor fremgangsmåden omfatter udførelse af handlinger med et batch-normaliseringslag mellem et første neuralt netværkslag og et andet neuralt netværkslag, hvor det første neurale netværkslag genererer første lag-output med en flerhed af komponenter, hvor fremgangsmåden omfatter anvendelse af batch-normaliseringslaget til, under træning af det neurale netværkssystem på en batch af træningseksempler, at:
- 10 modtage et tilsvarende første lag-output for hvert træningseksempel i batchen;  
beregne en flerhed af normaliseringsstatistikker for batchet fra de første lag-outputs;  
normalisere hver komponent af hvert første lag-output ved anvendelse af normaliseringsstatistikkerne til at generere et tilsvarende normaliseret lag-output for hvert træningseksempel i batchen;
- 15 generere et tilsvarende batch-normaliseringslag-output for hvert af træningseksemplerne ud fra de normaliserede lag-outputs; og  
tilvejebringe batch-normaliseringslag-outputtet som et input til det andet neurale netværkslag;  
hvor flerheden af komponenter af det første lag-output er indekseret efter i) dimension eller ii) hvor det første neurale netværkslag er et konvolutionslag og hvor flerheden af komponenter af det første lag-output er indekseret efter karakteristikiindex og rumligt lokalitetsindex, og hvor beregning af en flerhed af normaliseringsstatistikker for de første lag-outputs endvidere omfatter, for (i):  
beregning, for hver af dimensionerne, af en middelværdi for komponenterne af de første lag-outputs i dimensionen; og  
beregning, for hver af dimensionerne, af en standardafvigelse for komponenterne af de første lag-outputs i dimensionen; og for (ii):  
beregning, for hvert af karakteristikiindexerne, af en middelværdi for komponenterne af de første lag-outputs, der svarer til karakteristikiindexet; og  
beregning, for hvert af karakteristikiindexerne, af en varians for komponenterne af de første lag-outputs der svarer til karakteristikiindexet.
- 30 15. Et eller flere ikke-transitoriske computerlagringsmedier, der er kodet med et computerprogram, hvilket computerprogram omfatter instruktioner, der, når de udføres af én eller flere computere, bevirker, at den ene eller flere computere implementerer fremgangsmåden ifølge krav 14.

# DRAWINGS

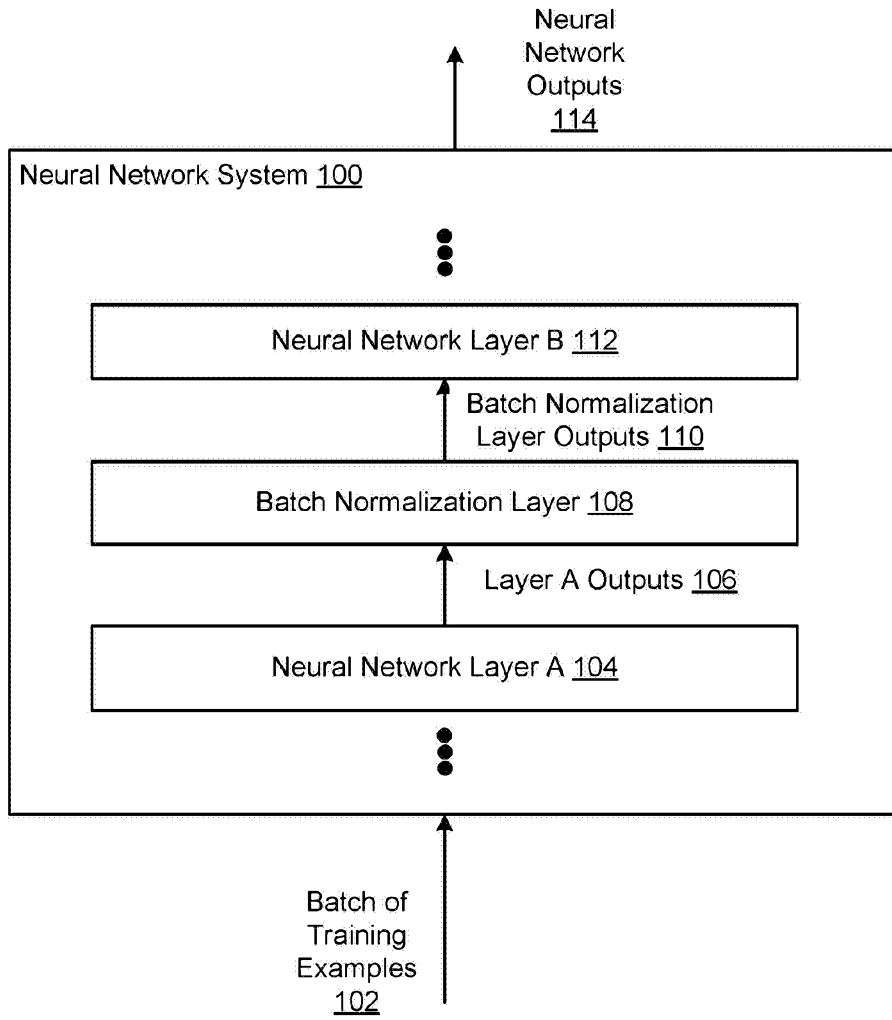


FIG. 1

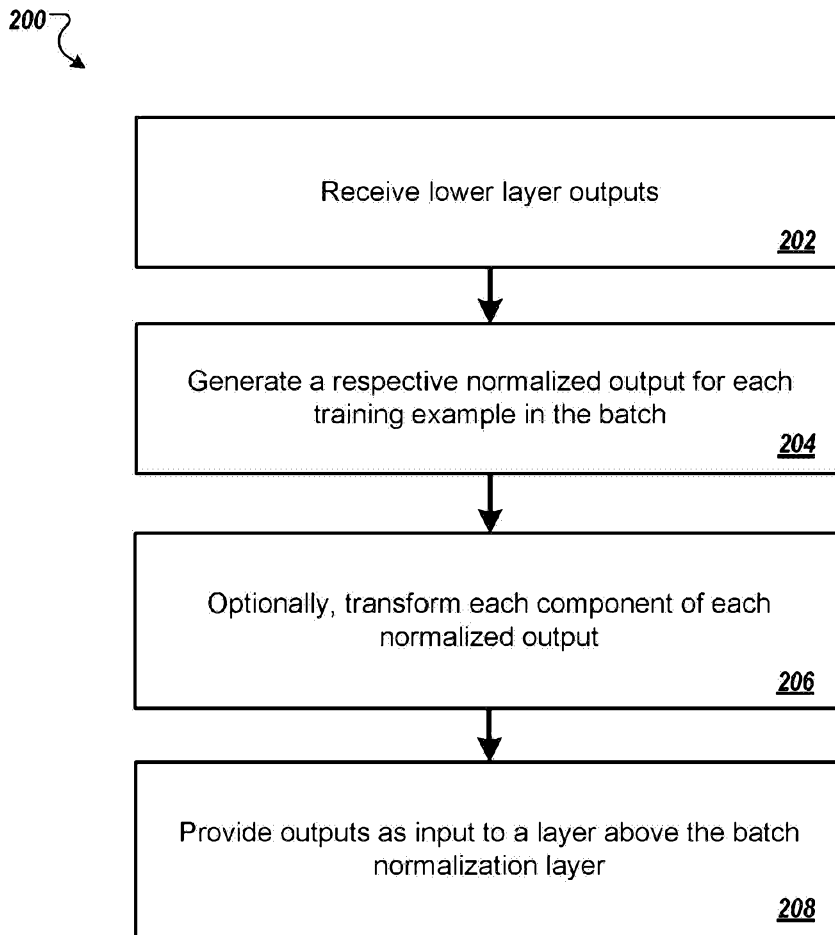


FIG. 2

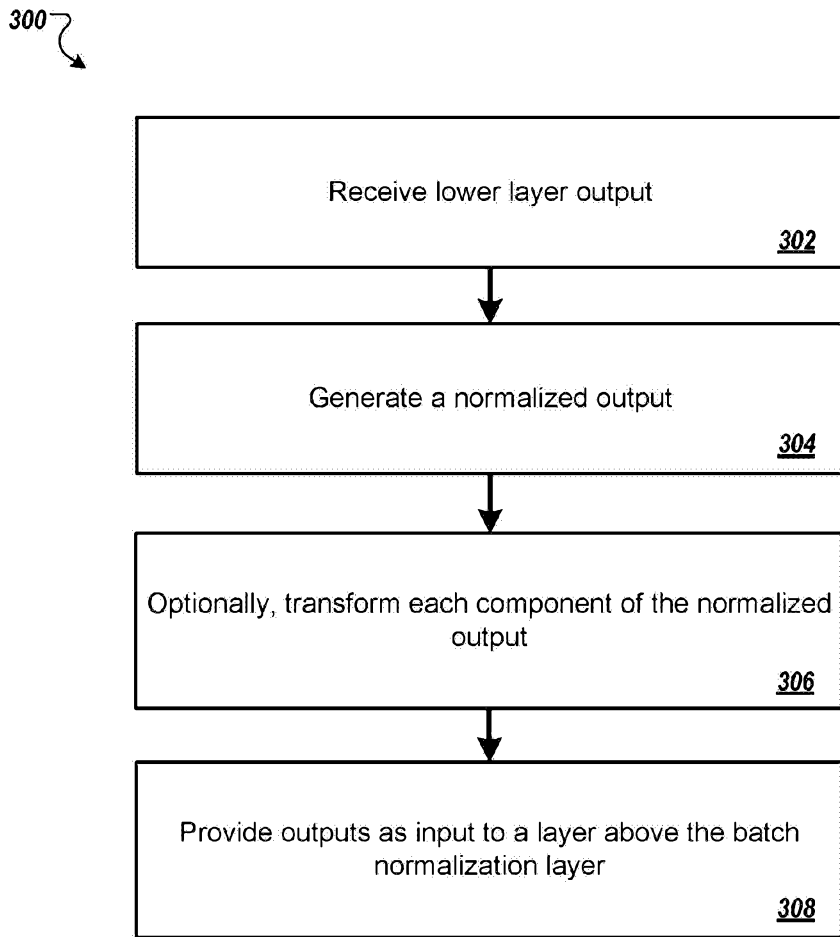


FIG. 3