




(86) Date de dépôt PCT/PCT Filing Date: 2020/03/06
(87) Date publication PCT/PCT Publication Date: 2020/10/09
(45) Date de délivrance/Issue Date: 2021/07/06
(85) Entrée phase nationale/National Entry: 2020/10/22
(86) N° demande PCT/PCT Application No.: CA 2020/050303
(87) N° publication PCT/PCT Publication No.: 2020/176994
(30) Priorités/Priorities: 2019/03/06 (US62/814,767);
2019/05/06 (US62/844,090)

(51) Cl.Int./Int.Cl. *G06K 9/62* (2006.01),
G06N 3/063 (2006.01), *G06N 3/08* (2006.01)
(72) Inventeurs/Inventors:
VOELKER, AARON R., CA;
ELIASMITH, CHRISTOPHER D., CA
(73) Propriétaire/Owner:
APPLIED BRAIN RESEARCH INC., CA
(74) Agent: ELAN IP INC.

(54) Titre : UNITES DE MEMOIRE DE LEGENDRE DANS DES RESEAUX NEURONAUX RECURRENTS
(54) Title: LEGENDRE MEMORY UNITS IN RECURRENT NEURAL NETWORKS



(57) Abrégé/Abstract:

Neural network architectures, with connection weights determined using Legendre Memory Unit equations, are trained while optionally keeping the determined weights fixed. Networks may use spiking or non-spiking activation functions, may be stacked or recurrently coupled with other neural network architectures, and may be implemented in software and hardware. Embodiments of the invention provide systems for pattern classification, data representation, and signal processing, that compute using orthogonal polynomial basis functions that span sliding windows of time.

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

(19) World Intellectual Property

Organization

International Bureau

(43) International Publication Date

10 September 2020 (10.09.2020)



(10) International Publication Number

WO 2020/176994 A1

(51) International Patent Classification:

G06K 9/62 (2006.01)*G06N 3/08* (2006.01)*G06N 3/063* (2006.01)

(21) International Application Number:

PCT/CA2020/050303

(22) International Filing Date:

06 March 2020 (06.03.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/814,767

06 March 2019 (06.03.2019)

US

62/844,090

06 May 2019 (06.05.2019)

US

(71) Applicant: **APPLIED BRAIN RESEARCH INC.**

[CA/CA]; 118 Woodbend Cres, Waterloo, Ontario N2T 1G9 (CA).

(72) Inventors: **VOELKER, Aaron R.**; 16 Jonathan PackStreet, Stittsville, Ontario K2S 0R5 (CA). **ELIASMITH, Christopher David**; 118 Woodbend Crescent, Waterloo, Ontario N2T 1G9 (CA).(74) Agent: **ELAN IP INC.**; 3500-2 Bloor St E, Toronto, M4W

1A8 (CA).

(81) Designated States (unless otherwise indicated, for every

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

(84) Designated States (unless otherwise indicated, for every

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

(54) Title: LEGENDRE MEMORY UNITS IN RECURRENT NEURAL NETWORKS




FIG. 4

(57) **Abstract:** Neural network architectures, with connection weights determined using Legendre Memory Unit equations, are trained while optionally keeping the determined weights fixed. Networks may use spiking or non-spiking activation functions, may be stacked or recurrently coupled with other neural network architectures, and may be implemented in software and hardware. Embodiments of the invention provide systems for pattern classification, data representation, and signal processing, that compute using orthogonal polynomial basis functions that span sliding windows of time.

WO 2020/176994 A1

Published:

- *with international search report (Art. 21(3))*
- *in black and white; the international application as filed contained color or greyscale and is available for download from PATENTSCOPE*

Legendre Memory Units in Recurrent Neural Networks

TECHNICAL FIELD

[0001] The invention relates generally to artificial intelligence and deep learning, and more particularly to a recurrent neural network architecture that may be implemented in software and in hardware. This application claims priority to provisional application number 62/814,767, filed March 6, 2019 and provisional application number 62/844,090, filed May 6, 2019.

BACKGROUND

[0002] Deep learning has undoubtedly brought about many rapid and impressive advances to the field of artificial intelligence. Due to its black-box nature, neither domain expertise nor understanding of the neural network's internal function are required in order to achieve state-of-the-art performance on a large number of important problems, including: image recognition, speech recognition, natural language understanding, question answering, and language translation (see Y. LeCun, Y. Bengio, and G. Hinton, Deep learning. Nature, vol. 521, no. 7553, pp. 436–444, May 2015). The basic recipe is as follows: install a software library for deep learning, select a network architecture, set its hyperparameters, and then train using as much data as the hardware (e.g., graphics processing unit) can hold in memory.

[0003] Deep learning architectures, such as the multi-layer perceptron, excel at constructing static vector functions, that generalize to new examples, by automatically discovering the “latent representations” (i.e., hidden features) that are most relevant to the task at hand. However, the opacity of its optimization procedure comes as a double-edged sword: while it is easy to apply deep learning to many problems with minimal hand-engineering, it is unclear even to experts what effect most hyperparameter changes will have in advance on overall performance.

[0004] Despite its breakthroughs, the field is well-aware that a feed-forward architecture is incapable of learning relationships that span arbitrarily across the input data in time, which is necessary for tasks involving video, speech, and other sequential time-series data with long-

range temporal dependencies. Regardless of the depth of the network, a feed-forward network will always have some finite input response, which leaves a finite “memory” of previous inputs within the state of the network. In other words, the functions that are computable with such a network cannot access inputs that go beyond the depth of the network. The most general solution to overcome this problem is to introduce recurrent connections into the network, which transmit current state information back to itself, thus allowing the network to capture information about previous inputs and reuse it in the future. These networks are called Recurrent Neural Networks (RNNs).

[0005] The RNN is the most computationally powerful brand of neural network that we know how to physically implement. By using recurrent connections to persist state information through time, thus endowing the network with an internal memory, RNNs are able to compute functions outside the computational class afforded by deep feed-forward networks: dynamical systems – functions whose state evolves nonlinearly according to the history of its inputs. This enables the network to exploit patterns in the input that span time along arbitrary temporal scales.

[0006] Specifically, RNNs serve as a universal approximator to any finite-dimensional, causal, dynamical system in the discrete-time domain (see A. M. Schäfer and H. G. Zimmermann, Recurrent neural networks are universal approximators. In International Conference on Artificial Neural Networks, Springer, pp. 632–640, Sept. 2006) and in the continuous-time domain (see K. Funahashi and Y. Nakamura, Approximation of dynamical systems by continuous time recurrent neural networks. Neural Networks, vol. 6, no. 6, pp. 801–806, Nov. 1992). In practice, RNNs are often the best model for tasks that involve sequential inputs, such as recognizing speech, translating language, processing video, generating captions, and decoding human emotions.

[0007] A longstanding challenge with RNNs pertains to the difficulty in training initially random recurrent weights such that they are able to exploit long-range temporal dependencies (see Y. Bengio, P. Simard, and P. Frasconi, Learning long-term dependencies with gradient descent is difficult. IEEE Transactions on Neural Networks, vol. 5, no. 2, pp. 157–166, Mar. 1994). Many architectural solutions have been proposed, with the most

historically successful being the Long Short-Term Memory (LSTM; see S. Hochreiter and J. Schmidhuber, Long short-term memory. *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997). A variety of more recent, yet closely related, alternatives also exist, for instance the Gated Recurrent Unit (GRU; see J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv:1412.3555, Dec. 2014) and Non-Saturating Recurrent Unit (NRU; see S. Chandar, C. Sankar, E. Vorontsov, S.E. Kahou, and Y. Bengio, Towards non-saturating recurrent units for modelling long-term dependencies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 1, pp. 3280–3287, Jul. 2017).

[0008] The LSTM, GRU, NRU, and other related alternatives, are all specific RNN architectures that aim to mitigate the difficulty in training RNNs, by providing methods of configuring the connections between nodes in the network. These architectures typically train to better levels of accuracy than randomly initialized RNNs of the same size. Nevertheless, these architectures are presently incapable of learning temporal dependencies that span more than about 100–5,000 time-steps, which severely limits the scalability of these architectures to applications involving longer input sequences. There thus remains a need for improved RNN architectures that can be trained to accurately maintain longer (i.e., longer than 100–5,000 steps in a sequential time-series) representations of temporal information, which motivates the proposed Legendre Memory Unit (LMU).

SUMMARY OF THE INVENTION

[0009] In one embodiment of the invention, there is disclosed method for generating recurrent neural networks having Legendre Memory Unit (LMU) cells including defining a node response function for each node in the recurrent neural network, the node response function representing state over time, wherein the state is encoded into one of binary events or real values; each node having a node input and a node output; defining a set of connection weights with each node input; defining a set of connection weights with each node output; defining one or more LMU cells having a set of recurrent connections defined as a matrix that determine node connection weights based on the formula:

$$A = [a]_{ij} \in \mathbb{R}^{q \times q}$$

$$\text{where } a_{ij} = (2i + 1) \begin{cases} -1 & i < j \\ (-1)^{i-j+1} & i \geq j \end{cases}$$

[0010] where q is an integer determined by the user, i and j are greater than or equal to zero.

[0011] In one aspect of the invention, the set of input connection weights are defined as a matrix that determine node connection weights based on the formula: $B = [b]_i \in \mathbb{R}^{q \times 1}$ where $b_i = (2i + 1)(-1)^i$.

[0012] In another aspect of the invention, the LMU node connection weights are determined based on the equation:

$$f(A; \theta, t)$$

where f is a function of A , θ is a predetermined parameter and t is time.

[0013] In another aspect of the invention, the predetermined parameter is one of selected by a user or determined using the output of a node in the neural network.

[0014] In another aspect of the invention, the LMU node connection weights are determined based on the equation:

$$\overline{f}(A; \theta, t, \Delta t)$$

where f is a function of A , θ is a predetermined parameter, t is time, and Δt is a predetermined parameter.

[0015] In another aspect of the invention, each of the predetermined parameters is one of selected by a user or determined using the output of the neural network.

[0016] In another aspect of the invention, one or more connection weights from node outputs are determined by evaluating Legendre polynomials.

[0017] In another aspect of the invention, the LMU cells are stacked, wherein each LMU cell is connected to the next using either a connection weight matrix or another neural network.

[0018] In another aspect of the invention, one or more LMU cells include connections to the inputs and from the outputs of other network architectures selected from LSTM cells, GRU cells, NRU cells, other LMU cells, multi-layer perceptrons, sigmoidal layers, and other linear or nonlinear layers.

[0019] In another aspect of the invention, the network is trained as a neural network by updating a plurality of its parameters.

[0020] In another aspect of the invention, the network is trained as a neural network by fixing one or more parameters while updating the remaining parameters.

[0021] According to another embodiment of the invention, there is provided a system for pattern classification, data representation, or signal processing in neural networks, the system including one or more input layers presenting a vector of one or more dimensions, wherein each dimension is provided to the network either by external input or by using previous outputs from the network; one or more intermediate layers coupled via weight matrices to at least one of the input, other intermediate, or output layers; one or more output layers generating a vector representation of the data presented at the input layer or computing a function of that data at one or more discrete points in time or continuously over time; wherein the system generates a recurrent neural network as herein described.

[0022] According to another embodiment of the invention, there is provided a circuit implemented in hardware with one or more recurrent connections that determine node connection weights as herein described.

BRIEF DESCRIPTION OF THE DRAWINGS

[0023] The invention is illustrated in the figure of the accompanying drawings which are meant to be exemplary and not limiting, in which like references are intended to refer to like or corresponding parts, and in which:

FIG. 1 illustrates an embodiment in software for a feed-forward network that determines the connection weights such that each layer encodes a more progressively lowpass filtered version of the input signal.

FIG. 2 illustrates an embodiment in software for a recurrent network that determines the connection weights in order to undo the effects of a lowpass filter at each layer according to embodiments of the invention.

FIG. 3 illustrates a circuit embodiment that implements the continuous-time LMU equations for six-dimensional recurrent and input weights.

FIG. 4 illustrates a method according to one embodiment of the invention.

FIG. 5 is a schematic diagram of an exemplary neural network on which embodiments of the invention may be implemented.

DETAILED DESCRIPTION OF THE INVENTION

[0024] Having summarized the invention above, certain exemplary and detailed embodiments will now be described below, with contrasts and benefits over the prior art being more explicitly described.

[0025] It will be apparent to one of skill in the art that other configurations, hardware etc. may be used in any of the foregoing embodiments of the products, methods, and systems of this invention. It will be understood that the specification is illustrative of the present invention and that other embodiments suggest themselves to those skilled in the art.

[0026] The embodiments of the systems and methods described herein may be implemented in hardware or software, or a combination of both. These embodiments may be implemented in computer programs executing on programmable computers, each computer including at least one processor, a data storage system (including volatile memory or non-volatile memory or other data storage elements or a combination thereof), and at least one communication interface.

[0027] In this invention, neural network architectures, with connection weights determined using Legendre Memory Unit (LMU) equations, are trained while optionally keeping the determined weights fixed. Networks may use spiking or non-spiking activation functions, may be stacked or recurrently coupled with other neural network architectures, and may be implemented in software and hardware. Embodiments of the invention provide systems for pattern classification, data representation, and signal processing, that compute using orthogonal polynomial basis functions that span sliding windows of time. Recurrent neural networks are well known in the art and their description and operation are assumed to be known in this application. The invention provides for an improved method and system by which recurrent network node weights are determined using Legendre Memory Unit (LMU) approaches and algorithms. Each node having an LMU approach applied is also referred to herein as an LMU cell.

[0028] Neural network architectures, with connection weights determined using Legendre Memory Unit equations, are trained while optionally keeping the determined weights fixed. Networks may use spiking or non-spiking activation functions, may be stacked or recurrently coupled with other neural network architectures, and may be implemented in software and hardware. Embodiments of the invention provide systems for pattern classification, data representation, and signal processing, that compute using orthogonal polynomial basis functions that span sliding windows of time.

[0029] We define the LMU cell as follows. Let $q \geq 1$ be an integer, provided by the user. Let $A = [a]_{ij} \in \mathbb{R}^{q \times q}$ be a square $q \times q$ matrix ($0 \leq i, j \leq q - 1$), with the following coefficients:

$$[0030] \ a_{ij} = (2i + 1) \begin{cases} -1 & i < j \\ (-1)^{i-j+1} & i \geq j \end{cases}$$

[0031] The output of each node may be defined as follows. Let $B = [b]_i \in \mathbb{R}^{q \times 1}$ be a $q \times 1$ matrix, with the following coefficients:

$$[0032] \ b_i = (2i + 1)(-1)^i.$$

[0033] Let θ be a parameter that is provided by the user or determined using the output of a node in the neural network.

[0034] Let t be either a continuous point in time, or a discrete point in time. For the discrete-time case, let Δt be a parameter that is provided by the user or determined using the output of a node in the neural network.

[0035] The LMU recurrent connections determine node connection weights by evaluating the following equation:

[0036] $f(A; \theta, t)$ in the continuous-time case; or $\bar{f}(A; \theta, t, \Delta t)$ in the discrete-time case, (1)

[0037] where f is a function of A that is parameterized by θ, t , in the continuous-time case, and additionally Δt for the function $\bar{f}$ in the discrete-time case.

[0038] The LMU connection weights to node inputs are optionally determined by evaluating the following equation:

[0039] $g(B; \theta, t)$ in the continuous-time case; or $\bar{g}(B; \theta, t, \Delta t)$ in the discrete-time case, (2)

[0040] where g is a function of B that is parameterized by θ, t , in the continuous-time case, and additionally Δt for the function $\bar{g}$ in the discrete-time case.

[0041] The LMU connections from node outputs are optionally determined by evaluating the Legendre polynomials (see A.M. Legendre, Recherches sur l'attraction des sphéroïdes homogènes. Mémoires de Mathématiques et de Physique, présentés à l'Académie Royale des Sciences, pp. 411–435, 1782).

[0042] This approach to determining recurrent connection weights is novel in the art, and as discussed below provides an improved recurrent neural network.

Derivation of A and B Matrices

[0043] To derive equations 1 and 2, let $\mathbf{x} \in \mathbb{R}^{q \times k}$ correspond to some subset of the state vector represented by some LMU cell, and $\mathbf{u} \in \mathbb{R}^{1 \times k}$ correspond to some subset of the vector

provided as input to the aforementioned LMU cell. Given our choice of (A, B) matrices, we define the following continuous-time dynamical system:

$$[0044] \quad \theta \frac{dx}{dt} = A\mathbf{x}(t) + B\mathbf{u}(t). \quad (3)$$

[0045] This dynamical system represents a memory of $\mathbf{u}$ across a sliding time-window of length θ using the orthogonal Legendre basis with coefficients given by the state $\mathbf{x}$. This provides computational benefits that are not available in any other RNN architecture.

[0046] An example of (A, B) for $q = 6$ is the following:

$$[0047] \quad A = \begin{pmatrix} -1 & -1 & -1 & -1 & -1 & -1 \\ 3 & -3 & -3 & -3 & -3 & -3 \\ -5 & 5 & -5 & -5 & -5 & -5 \\ 7 & -7 & 7 & -7 & -7 & -7 \\ -9 & 9 & -9 & 9 & -9 & -9 \\ 11 & -11 & 11 & -11 & 11 & -11 \end{pmatrix}$$

$$[0048] \quad B = \begin{pmatrix} 1 \\ -3 \\ 5 \\ -7 \\ 9 \\ -11 \end{pmatrix}$$

Example of Determining Continuous-Time Recurrent and Input Weights

[0049] Equation 3 then corresponds to the following continuous-time system of q ordinary differential equations (ODEs):

$$[0050] \quad \frac{dx}{dt} = f(A; \theta, t)\mathbf{x}(t) + g(B; \theta, t)\mathbf{u}(t),$$

[0051] where, for example, we define the following function to determine recurrent weights:

$$[0052] \quad f(A; \theta, t) = \frac{A}{\theta},$$

[0053] and the following function to determine input weights:

[0054] $g(B; \theta, t) = \frac{B}{\theta}.$

Example of Determining Discrete-Time Recurrent and Input Weights

[0055] Equation 3 corresponds to the following discrete-time dynamical system of q ODEs, discretized to a time-step of Δt :

[0056] $\mathbf{x}_{t+\Delta t} = \bar{f}(A; \theta, t, \Delta t)\mathbf{x}_t + \bar{g}(B; \theta, t, \Delta t)\mathbf{u}_t,$

[0057] where, for example considering zero-order hold (ZOH; see W.L. Brogan, Modern Control Theory. 3rd Edition, Pearson, Oct. 1990) discretization, we define the following function to determine recurrent weights:

[0058] $\bar{f}(A; \theta, t, \Delta t) = e^{A\frac{\Delta t}{\theta}},$ (4)

[0059] and the following function to determine input weights:

[0060] $\bar{g}(B; \theta, t, \Delta t) = A^{-1} \left(e^{A\frac{\Delta t}{\theta}} - I \right) B.$ (5)

Additional Considerations for Equations 1 and 2

[0061] We permit other possible definitions of $f(A; \theta, t)$ or $\bar{f}(A; \theta, t, \Delta t)$ and $g(B; \theta, t)$ or $\bar{g}(B; \theta, t, \Delta t)$ to determine the connection weights when evaluating equations 1 and 2, respectively. Examples include, but are not limited to, the use of alternative methods of numerically integrating differential equations, and transforming θ as a function of t and Δt .

[0062] Input sequences with irregular intervals (a.k.a., “unevenly spaced time series”) are supported by providing Δt using an input node in the neural network.

[0063] If the output of equation 1 or 2 are constant (i.e., if none of its parameters are variable, nor depend on the outputs of any nodes in the neural network), then they only need to be evaluated once (e.g., to initialize the weights). Otherwise, they may be reevaluated whenever their parameters change.

Example of Determining Output Weights

[0064] To determine the output connection weights from one or more nodes using the Legendre polynomials, we can for example evaluate the first q polynomials using the Rodrigues' formula (see O. Rodrigues, De l'attraction des sphéroïdes, Correspondence sur l'École Impériale Polytechnique. PhD Thesis, University of Paris, 1816) for the shifted Legendre polynomials:

$$[0065] \tilde{\mathcal{P}}_i(r) = (-1)^i \sum_{j=0}^i \binom{i}{j} \binom{i+j}{j} (-r)^j = \mathcal{P}_i(2r-1) \quad (6)$$

[0066] where $r \in [0,1]$, $0 \leq i \leq q-1$, and $\mathcal{P}_i$ is the Legendre polynomial of order i . To provide a specific example, we state the following property:

$$[0067] u(t - \theta') \approx \sum_{i=0}^{q-1} \tilde{\mathcal{P}}_i\left(\frac{\theta'}{\theta}\right) x_i(t).$$

[0068] For each connection projecting from the node representing x_i one can choose θ' ($0 \leq \theta' \leq \theta$) and then set $r = \frac{\theta'}{\theta}$ to evaluate equation 6 to determine its weight. More generally, one may compute any function of these polynomials (e.g., integral transforms such as the Fourier transform) in order to have the output nodes approximate functions of the sliding window of $\mathbf{u}$.

[0069] Referring to Figs. 1 and 2, we apply these methods to determine the output connection weights between layers by choosing $\theta' = 0$. In this example, equation 4 is used to determine recurrent connection weights, and equation 5 is used to determine input connection weights. In this exemplary embodiment, a lowpass filter is harnessed to implement the integration required by the dynamical system at each layer, and the choice of $\theta' = 0$ effectively undoes the temporal convolution performed by each lowpass filter. Consequently, with these LMU weights, the system propagates its input signal instantaneously through to the deepest layers as shown in Fig. 2. Without this choice of weights, the signal becomes progressively more lowpass filtered at each layer, as shown in Fig. 1.

Training the Neural Network

[0070] The parameters of the neural network can be trained using any available method, for example backpropagation through time (BPTT; see P.J. Werbos, Backpropagation through time: What it does and how to do it. Proceedings of the IEEE, vol. 78, no. 10, pp. 1550–1560, Oct. 1990).

[0071] During training, one or more of the weight parameters produced by evaluating equation 1 or 2 or the Legendre polynomials may be held fixed. Alternatively, one or more of the weights produced by evaluating equation 1 or 2 or the Legendre polynomials may be trained. In either case, when using BPTT, the error may be backpropagated through the multiply-accumulate operations implementing the connection weights.

[0072] Likewise, the parameters of equation 1 or 2 (e.g., θ or Δt , or the parameters of the neural network determining θ or Δt) may also be trained, for example by backpropagating the error through the gradients of equation 1 or 2 (also see T.Q. Chen, Y. Rubanova, J. Bettencourt, and D.K. Duvenaud, Neural Ordinary Differential Equations. In Advances in Neural Information Processing Systems, pp. 6571–6583, Dec. 2018).

[0073] To train the architecture with nodes consisting of spiking nonlinearities, one can use any available method of training spiking neural networks (see E. Hunsberger, C. Eliasmith, Spiking deep networks with LIF neurons, arXiv:1510.08829, Oct. 2015).

Software Architecture

[0074] Neural networks, with the aforementioned connection weights, may be implemented in software. Layers with one or more connection weights determined by evaluating equation 1 or equation 2 or the Legendre polynomials may be implemented using program code to create an LMU cell. These layers may be recurrently coupled with other neural network architectures. These layers may also be stacked by using connection weights or other neural networks to connect each layer to the next.

[0075] Program code is applied to input data to perform the functions described herein and to generate output information. The output information is applied to one or more output devices, in known fashion.

[0076] Each program may be implemented in a high-level procedural or object-oriented programming or scripting language, or both, to communicate with a computer system. Alternatively the programs may be implemented in assembly or machine language, if desired. The language may be a compiled or interpreted language. Each such computer program may be stored on a storage media or a device (e.g., read-only memory (ROM), magnetic disk, optical disc), readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer to perform the procedures described herein. Embodiments of the system may also be considered to be implemented as a non-transitory computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner to perform the functions described herein.

[0077] Furthermore, the systems and methods of the described embodiments are capable of being distributed in a computer program product including a physical, non-transitory computer readable medium that bears computer useable instructions for one or more processors. The medium may be provided in various forms, including one or more diskettes, compact disks, tapes, chips, magnetic and electronic storage media, and the like. Non-transitory computer-readable media comprise all computer-readable media, with the exception being a transitory, propagating signal. The term non-transitory is not intended to exclude computer readable media such as a volatile memory or random access memory (RAM), where the data stored thereon is only temporarily stored. The computer useable instructions may also be in various forms, including compiled and non-compiled code.

[0078] Fig. 5 shows a schematic of a neural network 500 that may be implemented in hardware or in software, having an input layer 508, one or more intermediate layers 512 and an output layer 516. The input layer has a plurality of nodes 508, 530, 536. The intermediate layers have recurrent nodes 532 that loop in the intermediate layer, with input weights 518 and output weights 520 coupling the nodes of each of the layers. Recurrent weights provide the feedback loop within the nodes of the intermediate layers. The output layers have nodes 534. The input to the input layer is shown as either an external input 502 or an input from a previous output 504 (derived from 528), for example.

Hardware Architecture

[0079] Neural networks, with the aforementioned connection weights, may be implemented in hardware including neuromorphic, digital, or analog hardware and/or hybrids thereof. More specifically, this architecture may be implemented in an application-specific integrated circuit (ASIC), field-programmable gate array (FPGA), graphics processing unit (GPU), or using configurations of analog components and other physical primitives including but not limited to transistors, and/or other parallel computing systems.

[0080] Referring to Fig. 3, we illustrate an exemplary circuit 300 implementing a neural network according to the invention in which connection weights are determined by evaluating equations 1 and 2 by module 300 with $q = 6$ in the continuous-time case. Large circles correspond to each dimension of $\mathbf{x}$. Small circles indicate elements that add (arrow head) or subtract (circular head) their inputs. The i 'th dimension temporally integrates and scales its input (triangular head) by $(2i + 1)/\theta$.

[0081] This design exploits the alternation of signs, and reuses the intermediate computations within the upper and lower triangles of A , by decomposing them into two separate cascading chains of summations that are then combined by a feedback loop. These same computations are also reused to implement the connection weights of B by supplying $\mathbf{u}$ to the appropriate intermediate nodes.

[0082] Increasing the dimensionality of the system by one requires appending $O(1)$ wires, adders, and state-variables, to the existing circuitry. In total, this circuit requires $O(q)$ wires, adders, and state-variables, thus making the circuit linearly scalable in both space and time.

Simulation Results

[0083] We consider a set of experiments that are designed to evaluate the memory capacity of stacked LSTMs relative to stacked LMUs with equivalent resource usage. For this, we use an off-the-shelf Keras implementation of a stacked LSTM, and construct 3 layers with 50 cells each. Each layer is fully connected to the next, and uses all of the default settings (e.g., tanh activations). The final layer likewise consists of a tanh activation unit for each output. To evaluate the continuous-time memory capacity, the input data is white noise, bandlimited

to 30 Hz, starting at 0, and normalized to an absolute range of $[-1, 1]$. The output data is a 50-dimensional vector representing a uniform arrangement of delayed inputs between 0–0.2 s. The data set consists of 256 samples, each 1 s long. This data is randomly partitioned into 50% training and 50% testing. The training data is further partitioned into a separate random 25% sample used to report validation accuracy during training. Backpropagation through time is carried out using the Adam optimizer with respect to the mean-squared error (MSE) loss function. Training is parallelized using Keras and TensorFlow across four Nvidia Titan Xp GPUs (12 GB each).

[0084] We found that, for a time-step of 2 ms, backpropagation could find adequate parameters to solve this task – that is, the LSTM could in fact accurately represent the entire delay interval consisting of $\theta = 100$ time-steps with a normalized root mean-squared error (NRMSE) of about 10%. However, after decreasing the time-step, by an order of magnitude, to 200 μ s—while increasing the length of data by the same factor so that the data still represents the exact same 1 s signals—the performance collapses; accuracy exponentially decays as a function of delay length across the $\theta = 1,000$ time-step window. In the worst case, the LSTM does no better than random chance, with an NRMSE of about 100%. Thus, even the most historically successful RNN architecture is clearly unable to represent increasingly long windows of time, which motivates the need for more capable RNN architectures.

[0085] We then took the exact same training code and network specification – but replaced each LSTM cell with a layer of LMU cells, where the (A, B) matrices for the continuous-time case are used (equivalent to using Euler’s method to discretize the system). These matrices are shared across each cell within the same layer (akin to weight-sharing in a convolutional neural network). Finally a plurality of tanh nonlinearities (one for each cell) are included that receive input from all state-variables across the same layer, thus supporting nonlinear computations across a mixture of scaled Legendre bases. For small values of q (e.g., 9), this network has comparable resource requirements to the aforementioned LSTM.

[0086] Each LMU cell receives a one-dimensional input. The trainable parameters are the weights between layers, and the delay lengths θ within each cell. In this experiment, we

disable training on the shared (A , B) weights. The overall architecture is consistent with the LSTM, as the LMU contains 50 cells stacked 3 times. The final output layer consists of linear activation units, since tanh has already been applied at this point. Finally, we set $q = 9$, initialize the encoding weights of each cell to 1 for the first layer and $1/50$ for all subsequent layers (i.e., the reciprocal of the fan-in), distribute θ values uniformly across $U[100, 1000]$, and set the weights projecting to each tanh by evaluating the Legendre polynomials at $r = 1$, with zero weights for all other state-variables from outside the cell. In other words, each cell is initialized to approximate $\tanh(u[t - \theta])$, where $u[\cdot]$ is the cell's mean input. Backpropagation then trains the values of θ and learns to mix weighted nonlinear combinations of inputs and outputs between layers.

[0087] Running the exact same code and analysis, on the exact same training, validation, and testing data, reveals a dramatic difference in training time between the two approaches. We found that the stacked LMU takes 52.5 s per epoch to train, compared to 102.6 s per epoch for the stacked LSTM. Furthermore, the LMU outperforms the LSTM in every measure of accuracy. Specifically, three orders of magnitude reduction in MSE across both training and validation, while converging much more rapidly to the ideal solution. The LMU architecture achieves consistent 3-4% error across the delay interval, while the equivalently-sized LSTM cell architecture approaches 100% error rates towards the end of the window. This illustrates that the stacked LSTM struggles to memorize low-frequency signals (relative to the time-step) across long intervals of time. In contrast, this task is natural for the stacked LMU, as its state represents a q -degree Legendre expansion of input history.

[0088] Backpropagation enables stacked LMUs to outperform stacked LSTMs even on tasks that are not readily supported by the initial configuration of the network. To assess the performance of each network on a continuous-time prediction task, we consider a synthetic dataset called Mackey-Glass (MG): a chaotic time-series described by a nonlinear delay-differential equation. The MG data is generated using a discrete time-delay of $\tau = 17$ (each time-step is 1 unit of time). The desired output is a lookahead (prediction) of 15 time-steps in advance (see Figure 6.15). We simulate this for 5,000 time-steps after removing the first 100 step transient. We repeat this 128 times, each time starting from initial random conditions.

The entire dataset is then centered to have a global mean of zero. Next, the dataset is randomly split into 32 training examples, 32 validation examples, and 64 testing examples.

[0089] We use the same networks from the previous experiment, but with 4 layers of 100 cells each. For the LMU cells, we make all parameters trainable (including the A , B matrices shared across cells within the same layer). We set $q = 6$ and initialize $\theta \in U[25, 50]$ to account for the shorter time-scale of this dataset. We initialize the remaining weights using standard Keras weight initializers. All three methods are trained across 500 epochs using the Adam optimizer. In this case, to minimize overfitting, we keep only the model from the epoch that has the highest validation score.

[0090] Test performance and training times are summarized as follows. The LSTM achieves 7.084% error using 282,101 parameters while taking 50.0 seconds per training epoch. The LMU achieves 6.783% error using 270,769 parameters while taking 30.5 seconds per training epoch. Thus, the LMU outperforms the LSTM in accuracy and training time. We posit that this is because the LMU more readily supports a delay-embedding within its 6-dimensional state. Moreover, the LMU provides improved scaling through time with respect to lower frequencies across longer continuous time-intervals.

Exemplary Applications

[0091] These methods can be used to produce a system that uses neural networks for pattern classification, data representation, or signal processing in hardware and in software.

[0092] For example, automatic speech recognition (ASR), is a system for computer speech recognition that processes speech (as an audio input waveform) and produces text (as model output). The input can be preprocessed into audio features (e.g., Mel-frequency cepstral coefficients, FilterBANK coefficients, and feature space Maximum Likelihood Linear Regression coefficients; see M. Ravanelli, T. Parcollet, and Y. Bengio, The pytorch-kaldi speech recognition toolkit. In International Conference on Acoustics, Speech and Signal Processing, IEEE, pp. 6465–6469, May, 2019) and provided to a neural network consisting of layers with connection weights determined using the LMU cell equations, with the output node of the neural network being post-processed using available methods of generating text

(e.g., contextual beam search). This system can thus be trained as a neural network to build an ASR system.

[0093] To provide another example, we consider the application of anomaly detection, which is the identification of outliers, or “anomalies,” in a dataset. This data may be provided sequentially, one input vector at a time, to a neural network consisting of layers with connection weights determined using the LMU cell equations, with the output node of the neural network classifying the input as being either typical or anomalous. This system can thus be trained using available methods (e.g., using unsupervised, semi-supervised, or fully supervised learning rules) to build an anomaly detector.

We claim:

1. A method for generating recurrent neural networks having Legendre Memory Unit (LMU) cells comprising:

defining a node response function for each node in the recurrent neural network, the node response function representing state over time, wherein the state is encoded into one of binary events or real values; each node having a node input and a node output;

defining a set of connection weights with each node input;

defining a set of connection weights with each node output;

defining one or more LMU cells having a set of recurrent connections defined as a matrix that determine node connection weights based on the formula:

$$A = [a]_{ij} \in \mathbb{R}^{q \times q}$$

$$\text{where } a_{ij} = (2i + 1) \begin{cases} -1 & i < j \\ (-1)^{i-j+1} & i \geq j \end{cases}$$

where q is an integer determined by the user, i and j are greater than or equal to zero.

2. The method of claim 1, wherein the set of input connection weights are defined as a matrix that determine node connection weights based on the formula: $B = [b]_i \in \mathbb{R}^{q \times 1}$ where $b_i = (2i + 1)(-1)^i$.
3. The method of claim 1, wherein the LMU node connection weights are determined based on the equation:

$$f(A; \theta, t)$$

where f is a function of A , θ is a predetermined parameter and t is time.

4. The method of claim 3, wherein the predetermined parameter is one of selected by a user or determined using the output of a node in the neural network.
5. The method of claim 1, wherein the LMU node connection weights are determined based on the equation:

$$\bar{f}(A; \theta, t, \Delta t)$$

where $\bar{f}$ is a function of A , θ is a predetermined parameter, t is time, and Δt is a predetermined parameter.

6. The method of claim 5, wherein each of the predetermined parameters is one of selected by a user or determined using the output of a node in the neural network.
7. The method of claim 1, wherein one or more connection weights from node outputs are determined by evaluating Legendre polynomials.
8. The method of claim 1, wherein the LMU cells are stacked, wherein each LMU cell is connected to the next using either a connection weight matrix or another neural network.
9. The method of claims 1, wherein one or more LMU cells include connections to the inputs and from the outputs of other network architectures selected from LSTM cells, GRU cells, NRU cells, other LMU cells, multi-layer perceptrons, sigmoidal layers, and other linear or nonlinear layers.
10. The method of claims 1, wherein the network is trained as a neural network by updating a plurality of its parameters.
11. The method of claims 1, wherein the network is trained as a neural network by fixing one or more parameters while updating the remaining parameters.

12. A system for pattern classification, data representation, or signal processing in neural networks, the system comprising:

- one or more input layers presenting a vector of one or more dimensions, wherein each dimension is provided to the network either by external input or by using previous outputs from the network;
- one or more intermediate layers coupled via weight matrices to at least one of the input, other intermediate, or output layers;
- one or more output layers generating a vector representation of the data presented at the input layer or computing a function of that data at one or more discrete points in time or continuously over time;

wherein

the system generates a recurrent neural network using the method of claim 1.

13. A circuit implemented in hardware with one or more recurrent connections that determine node connection weights using the method of claim 1.

14. The circuit of claim 13, wherein one or more connection weights from node outputs are determined by evaluating the Legendre polynomials.




FIG. 1




FIG. 2

3/5





FIG. 3

4/5

400 ↘

**FIG. 4**

5/5




FIG. 5

400

