Published in final edited form as:

Neurocomputing. 2019 February 28; 331: 281–288. doi:10.1016/j.neucom.2018.11.066.

Machine Learning on Sequential Data Using a Recurrent Weighted Average

Jared Ostmeyer* and Lindsay Cowell**

Department of Clinical Sciences UT Southwestern Medical Center 5323 Harry Hines Blvd. Dallas, TX 75390-9066, USA

Abstract

Recurrent Neural Networks (RNN) are a type of statistical model designed to handle sequential data. The model reads a sequence one symbol at a time. Each symbol is processed based on information collected from the previous symbols. With existing RNN architectures, each symbol is processed using only information from the previous processing step. To overcome this limitation, we propose a new kind of RNN model that computes a recurrent weighted average (RWA) over every past processing step. Because the RWA can be computed as a running average, the computational overhead scales like that of any other RNN architecture. The approach essentially reformulates the *attention mechanism* into a stand-alone model. The performance of the RWA model is assessed on the variable copy problem, the adding problem, classification of artificial grammar, classification of sequences by length, and classification of the MNIST images (*where the pixels are read sequentially one at a time*). On almost every task, the RWA model is found to fit the data significantly faster than a standard LSTM model.

Keywords

Recurrent Neural Network; Attention Mechanism; Sequences

1. Introduction

Types of information as dissimilar as language, music, and genomes can be represented as sequential data. The essential property of sequential data is that the order of the information is important, which is why statistical algorithms designed to handle this kind of data must be able to process each symbol in the order that it appears. Recurrent neural network (RNN) models have been gaining interest as a statistical tool for dealing with the complexities of sequential data. The essential property of a RNN is the use of feedback connections. The

*Principal corresponding author. **Corresponding author: Jared.Ostmeyer@UTSouthwestern.edu (Jared Ostmeyer), Lindsay.Cowell@UTSouthwestern.edu (Lindsay Cowell).

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Source code and experiments at <https://github.com/jostmeyer/rwa>

¹The RWA model was initially implemented incorrectly. The mistake was discovered by Alex Nichol. After correcting the mistake, the old results were discarded and the RWA model was run again on each classification task

sequence is read by the RNN one symbol at a time through the model's inputs. The RNN starts by reading the first symbol and processing the information it contains. The processed information is then passed through a set of feedback connections. Every subsequent symbol read into the model is processed based on the information conveyed through the feedback connections. Each time another symbol is read, the processed information of that symbol is used to update the information conveyed in the feedback connections. The process continues until every symbol has been read into the model (Fig. 1a). The processed information is passed along each step like in the game *telephone* (a.k.a. *Chinese whispers*). With each step, the RNN produces an output that serves as the model's prediction. The challenge of designing a working RNN is to make sure that processed information does not decay over the many steps. Error correcting information must also be able to *backpropagate* through the same pathways without degradation [1, 2]. Hochreiter and Schmidhuber were the first to solve these issues by equipping a RNN with what they called long short-term memory (LSTM) [3].

Since the introduction of the LSTM model, several improvements have been proposed. The *attention mechanism* is perhaps one of the most significant [4]. The attention mechanism is nothing more than a weighted average. At each step, the output from the RNN is weighted by an *attention model*, creating a weighted output. The weighted outputs are then aggregated together by computing a weighted average (Fig. 1b). The outcome of the weighted average is used as the model's result. The attention model controls the relative contribution of each output, determining how much of each output is "seen" in the results. The attention mechanism has since been incorporated into several other neural network architectures leading to a variety of new models each specifically designed for a single task (partial reference list: [4, 5, 6, 7, 8, 9]). Unfortunately, the attention mechanism is not defined in a recurrent manner. The recurrent connections must come from a separate RNN model, restricting where the attention mechanism can be used.

Inspired by the attention mechanism used to modify existing neural network architectures, we propose a new kind of RNN that is created by reformulating the attention mechanism into a stand-alone model. The proposed model includes feedback connections from every past processing step, not just the preceding step (Fig. 1c). The feedback connections from each processing step are weighted by an attention model. The weighted feedback connections are then aggregated together by computing a weighted average. The model is said to use a recurrent weighted average (RWA) because the attention model also uses the feedback connections. At each step, the weighted average can be computed as a running average. By maintaining a running average, the model scales like any other type of RNN.

2. Theory

2.1 Mathematical Description

The proposed model is defined recursively over the length of the sequence. Starting at the first processing step $t = 1$ a set of values h_0 , representing the initial state of the model, is required. The distribution of values for h_0 must be carefully chosen to ensure that the initial state resembles the output from the subsequent processing steps. This is accomplished by defining h_0 in terms of s_0 , a parameter that must be fitted to the data.

$$h_0 = f(s_0) \quad (1)$$

The parameters s_0 are passed through the model's activation function f to mimic the processes that generate the outputs of the later processing steps.

For every processing step that follows, a weighted average computed over every previous step is passed through the activation function f to generate a new output h_t for the model. The equation for the model is given below.

$$h_t = f \left(\frac{\sum_{i=1}^t z(x_i, h_{i-1}) \odot e^{a(x_i, h_{i-1})}}{\sum_{j=1}^t e^{a(x_j, h_{j-1})}} \right) \quad (2)$$

The weighted average consists of two models: z and a . The model z encodes the features x_i for each symbol in the sequence. Its recurrence relations, represented by h_{i-1} , provide the context necessary to encode the features in a sequence dependent manner. The model a serves as the attention model, determining the relative contribution of z at each processing step. The exponential terms of model a are normalized in the denominator to form a proper weighted average. The recurrent relations in a , represented by h_{i-1} , are required to compose the weighted average recursively. Because of the recurrent terms in a the model is said to use a RWA.

There are several models worth considering for z in future studies, but only one is considered here. Because z encodes the features, the output from z should ideally be dominated by the values in x_i and not h_{i-1} . This can be accomplished by separating the model for z into an unbounded component containing only x_i and a bounded component that includes the recurrent terms h_{i-1} .

$$z(x_i, h_{i-1}) = u(x_i) \odot \tanh g(x_i, h_{i-1}) \quad (3)$$

The model for u contains only x_i and encodes the features. With each processing step, information from u can accumulate in the RWA. The model for g contains the recurrent relations and is bounded between $[-1,1]$ by the tanh function. This model can control the sign of z but cannot cause the absolute magnitude of z to increase. Having a separate model for controlling the sign of z ensures that information encoded by u does not just accumulate but can negate information encoded from previous processing steps. Together the models u and g encode the features in a sequence dependent manner (see Appendix A for further discussion about the choice of z).

The terms u , g , and a can be modelled as feed-forward linear networks.

$$\begin{aligned}
 u(x_i) &= W_u \cdot x_i + b_u \\
 g(x_i, h_{i-1}) &= W_g \cdot [x_i, h_{i-1}] + b_g \\
 a(x_i, h_{i-1}) &= W_a \cdot [x_i, h_{i-1}]
 \end{aligned} \tag{4}$$

The matrices W_u , W_g , and W_a represent the weights of the feed-forward networks, and the vectors b_u and b_g represent the bias terms. The bias term for a would cancel when factored out of the numerator and denominator, which is why the term is omitted.

While running through a sequence, the output h_t from each processing step can be passed through a fully connected neural network layer to predict a label. Gradient descent based methods can then be used to fit the model parameters, minimizing the error between the true and predicted label.

2.2. Running Average

The RWA in equation (2) is recalculated from the beginning at each processing step. The first step to reformulate the model as a running average is to separate the RWA in equation (2) as a numerator term n_t and denominator term d_t .

$$\begin{aligned}
 n_t &= \sum_{i=1}^t z(x_i, h_{i-1}) \odot e^{a(x_i, h_{i-1})} \\
 d_t &= \sum_{j=1}^t e^{a(x_j, h_{j-1})}
 \end{aligned}$$

Because any summation can be rewritten as a recurrence relation, the summations for n_t and d_t can be defined recurrently (see Appendix B). Let $n_0 = 0$ and $d_0 = 0$.

$$\begin{aligned}
 n_t &= n_{t-1} + z(x_t, h_{t-1}) \odot e^{a(x_t, h_{t-1})} \\
 d_t &= d_{t-1} + e^{a(x_t, h_{t-1})}
 \end{aligned} \tag{5}$$

By saving the previous state of the numerator n_{t-1} and denominator d_{t-1} , the values for n_t and d_t can be efficiently computed using the work done during the previous processing step. The output h_t from equation (2) can now be obtained from the relationship listed below.

$$h_t = f\left(\frac{n_t}{d_t}\right) \tag{6}$$

Using this formulation of the model, the RWA can efficiently be computed dynamically.

2.3. Equations for Implementation

The RWA model can be implemented using equations (1) and (3)—(6), which are collected together and written below.

$$h_0 = f(s_0), n_0 = 0, d_0 = 0 \quad (7)$$

$$u(x_t) = W_u \cdot x_t + b_u$$

$$g(x_t, h_{t-1}) = W_g \cdot [x_t, h_{t-1}] + b_g$$

$$a(x_t, h_{t-1}) = W_a \cdot [x_t, h_{t-1}]$$

$$z(x_t, h_{t-1}) = u(x_t) \odot \tanh g(x_t, h_{t-1})$$

$$n_t = n_{t-1} + z(x_t, h_{t-1}) \odot e^{a(x_t, h_{t-1})}$$

$$d_t = d_{t-1} + e^{a(x_t, h_{t-1})}$$

$$h_t = f\left(\frac{n_t}{d_t}\right)$$

Starting from the initial conditions, the model is run recursively over an entire sequence. The features for every symbol in the sequence are contained in x_t , and the parameters s_0 , W_u , b_u , W_g , b_g , and W_a are determined by fitting the model to a set of training data. Because the model is differentiable, the parameters can be fitted using gradient optimization techniques.

In most cases, the numerator and denominator will need to be rescaled to prevent the exponential terms from becoming too large or small. The rescaling equations are provided in Appendix C.

3. Methods (Implementation)

A RWA model is implemented in *TensorFlow* using the equations in (7) [10]. The model is trained and tested on five different classification tasks each described separately in the following subsections.

The same configuration of the RWA model is used on each dataset. The activation function is $f(x) = \tanh x$ and the model contains 250 units. Following general guidelines for initializing the parameters of any neural network, the initial weights in W_u , W_g and W_a are drawn at random from the uniform distribution $\left[-\sqrt{\frac{3}{(N_{\text{in}} + N_{\text{out}})/2}}, \sqrt{\frac{3}{(N_{\text{in}} + N_{\text{out}})/2}}\right]$ and the bias terms b_u and b_g are initialized to 0's [11]. The initial state s_0 for the RWA model is drawn at random from a normal distribution according to $s_0 \sim \mathcal{N}(\mu = 0, \sigma^2 = 1)$. To avoid not-

a-number (NaN) and divide-by-zero errors, the numerator and denominator terms in equations (7) are rescaled using equations (B.1) in Appendix C, which do not alter the model's output.

The datasets are also scored on a LSTM model that contains 250 cells to match the number of units in the RWA model. Following the same guidelines used for the RWA model, the initial values for all the weights are drawn at random from the uniform distribution

$$\left[-\sqrt{\frac{3}{(N_{\text{in}} + N_{\text{out}})/2}}, \sqrt{\frac{3}{(N_{\text{in}} + N_{\text{out}})/2}} \right] \text{ and the bias terms are initialized to 0's except for the}$$

forget gates [11]. The bias terms of the forget gates are initialized to 1's, since this has been shown to enhance the performance of LSTM models [12, 13]. All initial cell states of the LSTM model are 0.

A fully connected neural network layer transforms the output from the 250 units into a predicted label. The error between the true label and predicted label is then minimized by fitting the model's parameters using Adam optimization [14]. All values for the ADAM optimizer follow published recommended settings. A step size of 0.001 is used throughout this study, and the other optimizer settings are $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. The same learning optimizer settings and learning rate are used both for the RWA and LSTM models and the models are run under identical conditions. Each parameter update consists of a batch of 100 training examples. Gradient clipping is not used.

Each model is immediately scored on the test set and no negative results are omitted¹. No hyperparameter search is done and no regularization is tried. At every 100 steps of training, a batch of 100 test samples are scored by each model to generate values that are plotted in the figures for each of the tasks described below. See Appendix D for a table summarizing the performance of models on each task. The code and results for each experiment may be found online (see: <https://github.com/jostmey/rwa>).

4. Results

4.1. Classifying Artificial Grammar

It is important that a model designed to process sequential data exhibit a sensitivity to the order of the symbols. For this reason, the RWA model is tasked with proofreading the syntax of sentences generated by an artificial grammatical system. Whenever the sentences are valid with respect to the artificial grammar, the model must return a value of 1, and whenever a typo exists the model must return a value of 0. This type of task is considered especially easy for standard RNN models, and is included here to show that the RWA model also performs well at this task [3].

The artificial grammar generator is shown in Figure 2a. The process starts with the arrow labeled *B*, which is always the first letter in the sentence. Whenever a node is encountered, the next arrow is chosen at random. Every time an arrow is used the associated letter is added to the sentence. The process continues until the last arrow *E* is used. All valid sentences end with this letter. Invalid sentences are constructed by randomly inserting a typo along the sequence. A typo is created by an invalid jump between unconnected arrows. No

more than one typo is inserted per sentence. Typos are inserted into approximately half the sentences. Each RNN model must perform with greater than 50% accuracy to demonstrate it has learned the task.

A training set of 100,000 samples are used to fit the model, and a test set of 10,000 samples are used to evaluate model performance. The RWA model does remarkably well, achieving 100% accuracy in 600 training steps. The LSTM model also learns to identify valid sentences, but requires 1000 training steps to achieve the same performance level (Fig 2b). This task demonstrates that the RWA model can classify patterns based on the order of information.

4.2. Classifying by Sequence Length

Classifying sequences by length requires that RNN models track the number of symbols contained in each sequence. For this task, the length of each sequence is randomly drawn from a uniform distribution over every possible length 0 to T , where T is the maximum possible length of the sequence. Each step in the sequence is populated with a random number drawn from a unit normal distribution (*i.e.* $\mu = 0$ and $\sigma^2 = 1$). Sequences greater than length $T/2$ are labeled with 1 while shorter sequences are labeled with 0. The goal is to predict these labels, which indicates if a RNN model has the capacity to classify sequences by length. Because approximately half the sequences will have a length above $T/2$, each RNN model must perform with greater than 50% accuracy to demonstrate it has learned the task.

For this task, $T = 1,000$ (*The task was found to be too easy for both models for $T = 100$*). A training set of 100,000 samples are used to fit the model, and a test set of 10,000 samples are used to evaluate model performance. The RWA model does remarkably well, learning to correctly classify sequences by their length in fewer than 100 training steps. The LSTM model also learns to correctly classify sequences by their length, but requires over 2,000 training steps to achieve the same level of performance (Fig. 3).

4.3. Variable Copy Problem

The variable copy problem, proposed by Henaff et al. [15], requires that a RNN memorize a random sequence of symbols and recall it only when prompted. The input sequence starts with the recall sequence. The recall sequence for the RNN to memorize consists of S many symbols drawn at random from $\{a_i\}_{i=1}^K$. After that, the input sequence contains a stretch of T blank spaces, denoted by the symbol a_{K+1} . Somewhere along this stretch of blank spaces, one of the symbols is replaced by a delimiter a_{K+2} . The delimiter indicates when the recall sequence should be repeated back. The input sequence is then padded with an additional stretch of S blank spaces, providing sufficient time for the model to repeat the recall sequence. The goal is to train the model so that its output is always blank except for the spaces immediately following the delimiter, in which case the output must match the recall sequence.

An example of the task is shown in Figure 4a. The recall pattern drawn at random from symbols A through H is $DEAE EBHGBH$. Blank spaces represented by $*$ fill the rest of the

sequence. One blank space is chosen at random and replaced by X , which denotes the delimiter. After X appears the model must repeat the recall pattern in the output. The naïve strategy is to always guess that the output is $*$ because this is the most common symbol. Each RNN must perform better than the naïve strategy to demonstrate that it has learned the task. Using this naïve strategy, the expected cross-entropy error between the true output and predicted output is $\frac{S \cdot \ln k}{2 \cdot S + T}$, represented by the dashed line in Figures 4b, c. This is the baseline to beat.

For this challenge, $K = 8$, $S = 10$, and models are trained and evaluated on two separate cases where $T = 100$ and $T = 1,000$. For both $T = 100$ and $T = 1,000$, the training set contains 100,000 examples and the test set contains 10,000 examples. For the case of $T = 100$, the RWA requires roughly 1,000 training steps to beat the baseline score, whereas the LSTM model requires over 10,000 training steps to achieve the same level of performance. The RWA model scales well to $T = 1,000$, requiring only 3,000 training steps to beat the baseline score. The LSTM model is only barely able to beat the baseline error after 50,000 training steps (Fig. 4c). The RWA model appears to scale much better as the sequence length increases.

4.4. Adding Problem

The adding problem, first proposed by Hochreiter and Schmidhuber [3], tests the ability of a RNN model to form long-range connections across a sequence. The task requires that the model learn to add two numbers randomly spaced apart on a sequence. The input sequence consists of two values at each step. The first value serves as an indicator marking the value to add while the second value is the actual number to be added and is drawn at random from a uniform distribution over $[0,1]$. Whenever the indicator has a value of 1, the randomly drawn number must be added to the sum, and whenever the indicator has a value of 0, the randomly drawn number must be ignored. Only two steps in the entire sequence will have an indicator of 1, leaving the indicator 0 everywhere else.

An example of the adding problem is shown in Figure 5a. The top row contains the indicator values and the bottom row contains randomly drawn numbers. The two numbers that have an indicator of 1 must be added together. The numbers in this example are 0.5 and 0.8, making the target output 1.3. Because the two numbers being added together are uniformly sampled over $[0,1]$, the naïve strategy is to always guess that the target output is 1. Each RNN must perform better than the naïve strategy to demonstrate that it has learned the task. Using this naïve strategy, the expected mean square error (MSE) between the true answer and the prediction is approximately 0.167, represented by the dashed line in Figures 5b,c. This is the baseline to beat.

The adding problem is repeated twice, first with sequences of length 100 and again with sequences of length 1,000. In both cases, a training set of 100,000 samples are used to fit the model, and a test set of 10,000 samples are used to evaluate the model's performance. When the sequences are of length 100, the RWA model requires fewer than 1,000 training steps to beat the baseline score while the LSTM model requires around 3,000 steps (Fig. 5b). When the sequences are of length 1,000, the RWA model requires approximately 1,000

training steps to beat the baseline score, while the LSTM model requires over 15, 000 training steps on the same task (Fig. 5c). The RWA model appears to scale much better as the sequence length increases.

4.5. Classifying MNIST Images (Pixel by Pixel)

The MNIST dataset contains 28×28 pixel images of handwritten digits 0 through 9. The goal is to predict the digit represented in the image [16]. Using the same setup suggested by Le, Jaitly, and Hinton [17], the images are arranged into a sequence of pixels. The length of each sequence is $28 \times 28 = 784$ pixels. Each RNN model reads the sequence one pixel at a time and must predict the digit being represented in the image from this sequence.

Examples of MNIST digits with the correct label are shown in Figure 6a. The pixels at the top and bottom of each image are empty. When the images are arranged into a sequence of pixels, all the important pixels will be in the middle of the sequence. To utilize these pixels, each RNN model will need to form long-range dependencies that reach the middle of each sequence. The model will have formed the necessary long-range dependencies when it outperforms a naive strategy of randomly guessing each digit. A naive strategy will achieve an expected accuracy of 10%, represented by the dashed line in Figures 6b. This is the baseline to beat.

For this challenge, the standard training set of 60,000 images is used to fit the model, and the standard test set of 10,000 images is used to evaluate the model's performance on unseen data. The RWA model fits the dataset in under 20, 000 steps, considerably faster than the LSTM model (Fig. 6b). After a quarter million training steps, the RWA model achieves an accuracy of 98.1% with an error of 0.175 bits, while LSTM model achieves an accuracy of 99.0% with an error of 0.077 bits. In this example, the LSTM model generalizes better to the unseen data.

A separate and more challenging task is to randomly permute the pixels, leaving the pixels out of order, as described in Le et al. [17]. The same permutation mapping must be used on each image to keep the data consistent between images. As before, a naive strategy of randomly guessing the answer will achieve an expected accuracy of 10%, represented by the dashed line in Figures 6c. This is the baseline to beat.

The classification task is repeated with the pixels randomly permuted. This time the LSTM model fit the dataset faster than the RWA model (Fig. 6c). After a quarter million training steps, the RWA model achieves an accuracy of 93.5% with an error of 0.561 bits, while LSTM model achieves an accuracy of 93.6% with an error of 0.577 bits. Neither model generalizes noticeably better to the unseen data.

5. Comparison to Related Work

The purpose of the RWA is to aggregate information across every previous processing step. Other RNN architectures have been considered for aggregating information from multiple past processing steps. Jan Koutnik et al proposed a clockwork RNN architecture with separate modules of recurrent connections for different past processing steps [18]. The

architecture represents a brute force approach by maintaining recurrent connections to multiple past processing steps, which requires separate weight matrices for the recurrent connections of each module. To reduce computational costs, Jan Koutník et al proposed updating the i th module of recurrent connections every 2^{i-1} steps. When a set of recurrent connections are not in use, its weights are treated as being 0. By evaluating modules only periodically, the architecture efficiently scales up across every symbol in long sequences. That said, the number of past processing steps that can be used is limited by the number of recurrent terms that can fit into computer memory. The connections to the past processing steps are also indirect, forming only when a module undergoes a periodic evaluation.

To include information from every past processing step without requiring separate modules of recurrent connections to each past processing step, an attention layer can be added to aggregate the output of a RNN model at each processing step into a single output. The attention layer essentially calculates a weighted average over all the past processing steps from the RNN model. The weight terms are determined by an attention model a that controls the relative contribution of each processing step in the output. Unfortunately, an attention layer does not take into account the order of the information in the sequence. This is because the terms in the attention layer can be permuted without changing the output of the aggregation. We have addressed this deficiency by introducing the concept of a recurrent weighted average.

Because the RWA model computes a weighted average across every symbol in a sequence, it is not limited to information from the previous processing step. This is both the strength and weakness of the model. The RWA model does not give preferential treatment to the most recent information and instead aggregates information across the entire sequence as a weighted average. Unfortunately, the RWA model is unable to discount information from past steps. This becomes a serious issue on tasks where only the most recent information is important. On the toy-problems presented in this manuscript it has not been an issue. However, on complex tasks in natural language processing, the RWA model fails to fit the training data [18].

To overcome the limitations of the RWA model, a modification has been proposed by Maginnis and Richemond that enables past information to be discounted [18]. The proposed model is called a Recurrent Discounted Attention (RDA) model, and it essentially adds a forget gate into the RWA model enabling it to erase information from past states. Maginnis and Richemond compared the performance of the RWA, RDA, GRU, and LSTM models on a variety of challenging tasks (including a character prediction task based on wikipedia text). They report that the RWA model failed to fit the training data on several complex tasks but that their RDA model achieved state of the art or near state of the art performance on each task. The RDA model represents an exciting new approach to modelling sequential data, and it combines the strengths of the RWA model with the strengths of popular RNN architectures like the LSTM.

6. Discussion

The RWA model reformulates the attention mechanism into a stand-alone model that can be optimized using gradient descent based methods. Given that the attention mechanism has been shown to work well on a wide range of problems, the robust performance of the RWA model on the five classification tasks in this study is not surprising [4, 5, 6, 7, 8, 9]. Moreover, the RWA model did not require a hyperparameter search to tailor the model to each task. The same configuration successfully generalized to unseen data on every task. Clearly, the RWA model can form long-range dependencies across the sequences in each task and does not suffer from the vanishing or exploding gradient problem that affects other RNN models [1, 2].

On each task, the RWA model requires less clock time and fewer parameters than a LSTM model with the same number of units. On almost every task, the RWA model beat the baseline score using fewer training steps. The number of training steps could be further reduced using a larger step size for the parameter updates. It is unclear how large the step size can become before the convergence of the RWA model becomes unstable (*a larger step size may or may not require gradient clipping, which was not used in this study*). The RWA model also uses over 25% fewer parameters per unit than a LSTM model. Depending on the computer hardware, the RWA model can either run faster or contain more units than a LSTM model on the same computer.

Unlike previous implementations of the attention mechanism that read an entire sequence before generating a result, the RWA model computes an output in real time with each new input. With this flexibility, the RWA model can be deployed anywhere existing RNN models can be used. Several architectures are worth exploring. Bidirectional RWA models for interpreting genomic data could be created to simultaneously account for information that is both upstream and downstream in a sequence. Multi-layered versions could also be created to handle XOR classification problems at each processing step. In addition, RWA models could be used to autoencode sequences or to learn mappings from a fixed set of features to a sequence of labels. The RWA model offers a compelling framework for performing machine learning on sequential data. Already, enhancements to the RWA model are being proposed to improve its performance on tasks such as natural language processing [18].

7. Conclusion

In contrast to many popular RNN architectures, the RWA model aggregates all information in a sequence without giving preference to the last few symbols in the sequence. This aspect of the model compliments existing RNN architectures making it worth considering the RWA model as part of a tool-kit of RNN models for performing statistical classification of sequential data. Many problems may benefit from the RWA model being used either as a stand alone model or as part of an ensemble of RNN models. The RWA model represents an exciting new direction in research on RNN architectures.

Acknowledgements

Special thanks are owed to Elizabeth Harris for proofreading and editing the manuscript. She brought an element of clarity to the manuscript that it would otherwise lack. Alex Nichol also needs to be recognized. Alex identified a flaw in equations (B.1), which are used to compute a numerically stable update of the numerator and denominator terms. Without Alex's careful examination of the manuscript, all the results for the RWA model would be incorrect. The department of Clinical Sciences at the University of Texas South-western Medical Center also needs to be acknowledged. Ongoing research projects at the medical center highlighted the need to find better ways to process genomic data and provided the inspiration behind the development of the RWA model.

Funding

This work was supported by a training grant from the Cancer Prevention and Research Institute of Texas [RP1601570] and an R01 from the National Institute of Allergy and Infectious Diseases [R01AI097403].

Appendix A.

There are several models worth considering for z that may be useful depending on the task. For the purpose of this study, we use equation 3 for z because we believe it provides functionality that is lacking and non-redundant with the rest of the RWA model. To understand why, we will first discuss alternative models for z .

We could use the original RNN definition for z .

$$z(x_i, h_{i-1}) = W_z \cdot [x_i, h_{i-1}] + b_z$$

This creates a second set of recurrent terms h in z that appear along with the recurrent terms in the attention model a . The recurrent terms in both z and a can control the magnitude of each term in the RWA model leading to redundant functionality. Alternatively, we could use just the features x in z without any recurrent terms.

$$z(x_i) = W_z \cdot x_i + b_z$$

However, the exponential terms of the attention model a are always positive. If the same features reappeared at different steps, the information would only be able to add together and could not cancel out. For example, consider what happens if the model must learn to map the sequence $\{x_1 = 1\}$ to 1 and the sequence $\{x_1 = 1, x_2 = 1\}$ to 0. The values for z will always be $z(1)$ because both $x_1 = 1$ and $x_2 = 1$. The RWA for both $\{x_1 = 1\}$ and $\{x_1 = 1, x_2 = 1\}$ would be the same, and it would be impossible for the model to map the first sequence to 1 and the second sequence to 0.

The problem is solved by adding a term like in equation (3) that controls the sign of z at each step in the sequence. When $g > 0$ the sign of z will remain the same and when $g < 0$ the sign of z will flip. By flipping the sign between the first and second term, the $g < 0$ the sign of z will flip. By flipping the sign between the first and second term, the first term in the sequence $\{x_1 = 1\}$ can map to 1 and the second term can cancel out the first term so that $\{x_1 = 1, x_2 = 1\}$ can map to 0.

One way to think about equation (3) for z is that the exponential terms of the attention model control the magnitude of each term while g controls the sign *or phase* of each term. The recurrent terms h in equation (3) are restricted to only controlling the sign of z , minimizing redundant functionality that the recurrent terms in g share with the recurrent terms in a .

Appendix B.

Any summation of the form $y_n = \sum_{i=1}^n r_i$ can be written as a recurrent relation. Let the initial values be $y_0 = 0$.

$$y_n = \sum_{i=1}^n r_i = \sum_{i=1}^{n-1} r_i + r_n = y_{n-1} + r_n$$

The summation is now defined recursively.

Appendix C.

The exponential terms in equations (7) are prone to underflow and overflow. The underflow condition can cause all the terms in the denominator to become zero, leading to a divide-by-zero error. The overflow condition leads to not-a-number (NaN) errors. To avoid both kinds of errors, the numerator and denominator can be multiplied by an *exponential scaling factor*. Because the numerator and denominator are scaled by the same factor, the quotient remains unchanged and the output of the model is identical. The idea is similar to how a softmax function must sometimes be rescaled to avoid the same errors.

The exponential scaling factor is determined by the largest value in a among every processing step. Let $a_t^{\max}$ represent the largest observed value. The initial value for $a_0^{\max}$ needs to be less than any value that will be observed, which can be accomplished by using an extremely large negative number.

$$a_0^{\max} = -10^{38} \quad (\text{B.1})$$

$$\begin{aligned} a_t^{\max} &= \text{Max}\{a_{t-1}^{\max}, a(x_t, h_{t-1})\} \\ n_t &= n_{t-1} \odot e^{\frac{a_{t-1}^{\max} - a_t^{\max}}{a_{t-1}^{\max} - a_t^{\max}}} + z(x_t, h_{t-1}) \odot e^{\frac{a(x_t, h_{t-1}) - a_t^{\max}}{a_{t-1}^{\max} - a_t^{\max}}} \\ d_t &= d_{t-1} \odot e^{\frac{a_{t-1}^{\max} - a_t^{\max}}{a_{t-1}^{\max} - a_t^{\max}}} + e^{\frac{a(x_t, h_{t-1}) - a_t^{\max}}{a_{t-1}^{\max} - a_t^{\max}}} \end{aligned}$$

The first equation sets the initial value for the exponential scaling factor to one of the largest numbers that can be represented using single-precision floating-point numbers. Starting with this value avoids the underflow condition. The second equation saves the largest value observed in a . The final two equations compute an updated numerator and denominator. The equations scale the numerator and denominator accounting for the exponential scaling factor

used during previous processing steps. The results from equations (B.1) can replace the results for n_t and d_t in equations (7) without affecting the model's output $h_t = f\left(\frac{n_t}{d_t}\right)$.

Appendix D.

Task	Training Step	Model	Accuracy	Error
Reber Grammar	2000	RWA	99%	0.0238
		LSTM	100%	0.0102
Classifying Sequence Length	5000	RWA	100%	0.0131
		LSTM	100%	0.0168
Variable Copy Problem 100	50000	RWA	N/A	4.736e-06
		LSTM	N/A	5.709e-04
Variable Copy Problem 1000	50000	RWA	N/A	0.001417
		LSTM	N/A	0.009594
Adding Problem 100	50000	RWA	N/A	2.008e-06
		LSTM	N/A	7.338e-06
Adding Problem 1000	50000	RWA	N/A	6.266e-06
		LSTM	N/A	1.093e-05
MNIST pixel by pixel	125000	RWA	99%	0.01936
		LSTM	98%	0.05430
MNIST pixel by pixel (permuted)	125000	RWA	87%	0.7946
		LSTM	96%	0.2589

The above table summarizes the performance of each model on every task in this study. The column “training step” indicates the last training step as it appears in the figure, which is the point when the models are evaluated. The columns “Accuracy” and “Error” show the model performance calculated over a batch of 100 random test samples. There is considerable variability because the results are calculated from only 100 random test samples.

Biography

Jared Ostmeyer received his B.S. degree from the University of Arkansas in 2008 and his Ph.D. from the University of Chicago in 2016. He is currently a postdoctoral researcher at the University of Texas Southwestern Medical Center in the department of Clinical Sciences. His research focus is on applying machine learning and bioinformatics techniques to help better understand the role of the adaptive immune system in fighting disease.

Lindsay Cowell received a M.S. in Biomathematics with a minor in Mathematics in 1995 from North Carolina State University. In 2000, she received a Ph.D. in Biomathematics with a minor in Immunology, also from North Carolina State University. She spent three years as a postdoctoral fellow in the Department of Immunology at Duke University Medical Center and then became an Assistant Professor in the Department of Biostatistics and Bioinformatics. In September 2010, she joined the Biomedical Informatics Division in the Department of Clinical Sciences at UT Southwestern. Dr. Cowell is broadly interested in understanding the mechanisms of adaptive immunity and their role in infectious diseases, autoimmune diseases, cancer immunology, and vaccine responses. Her methodologic focus has centered on the development of probabilistic models and the use of formal logics for representing and computing with descriptive information.

References

- [1]. Hochreiter S, Untersuchungen zu dynamischen neuronalen Netzen, Ph.D. thesis, diploma thesis, institut für informatik, lehrstuhl prof. brauer, technische universitat miinchen, 1991.
- [2]. Bengio Y, Simard P, Frasconi P, Learning long-term dependencies with gradient descent is difficult, IEEE transactions on neural networks 5 (2) (1994) 157–166. [PubMed: 18267787]
- [3]. Hochreiter S, Schmidhuber J, Long short-term memory, Neural computation 9 (8) (1997) 1735–1780. [PubMed: 9377276]
- [4]. D. Bahdanau, K. Cho, Y. Bengio, Neural machine translation by jointly learning to align and translate, arXiv preprint arXiv:1409.0473.
- [5]. Vinyals O, Toshev A, Bengio S, Erhan D, Show and tell: A neural image caption generator, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 3156–3164, 2015.
- [6]. Xu K, Ba J, Kiros R, Cho K, Courville AC, Salakhutdinov R, Zemel RS, Bengio Y, Show, Attend and Tell: Neural Image Caption Generation with Visual Attention., in: ICML, vol. 14, 77–81, 2015.
- [7]. Sonderby SK, Sonderby CK, Nielsen H, Winther O, Convolutional LSTM networks for subcellular localization of proteins, in: International Conference on Algorithms for Computational Biology, Springer, 68–80, 2015.
- [8]. W. Chan, N. Jaitly, Q. V. Le, O. Vinyals, Listen, attend and spell, arXiv preprint arXiv:1508.01211
- [9]. Vinyals O, Kaiser L, Koo T, Petrov S, Sutskever I, Hinton G, Grammar as a foreign language, in: Advances in Neural Information Processing Systems, 2773–2781, 2015.
- [10]. Abadi, et al., Tensorflow: Large-scale machine learning on heterogeneous distributed systems, arXiv preprint arXiv:1603.04467.
- [11]. Glorot X, Bengio Y, Understanding the difficulty of training deep feedforward neural networks., in: Aistats, vol. 9, 249–256, 2010.
- [12]. Gers FA, Schmidhuber J, Cummins F, Learning to forget: Continual prediction with LSTM, Neural computation 12 (10) (2000) 2451–2471. [PubMed: 11032042]
- [13]. Jozefowicz R, Zaremba W, Sutskever I, An Empirical Exploration of Recurrent Network Architectures, in: Proceedings of The 32nd International Conference on Machine Learning, 2342–2350, 2015.
- [14]. D. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980.

- [15]. Henaff M, Szlam A, LeCun Y, Recurrent Orthogonal Networks and Long-Memory Tasks, in: Proceedings of The 33rd International Conference on Machine Learning, 2034–2042, 2016.
- [16]. LeCun Y, Bottou L, Bengio Y, Haffner P, Gradient-based learning applied to document recognition, Proceedings of the IEEE 86 (11) (1998) 2278–2324.
- [17]. Q. V. Le, N. Jaitly, G. E. Hinton, A simple way to initialize recurrent networks of rectified linear units, arXiv preprint arXiv:1504.00941.
- [18]. B. Maginnis, P. H. Richmond, Efficiently applying attention to sequential data with the Recurrent Discounted Attention unit, arXiv preprint arXiv:1705.08480.

Figure 1:

Comparison of models for classifying sequential data. (a) Standard RNN architecture with LSTM requires that information contained in the first symbol x_1 pass through the feedback connections repeatedly to reach the output h_t , like in a game of *telephone* (a.k.a. *Chinese whispers*). (b) The attention mechanism aggregates the outputs into a single state by computing a weighted average. It is not recursively defined. (c) The proposed model incorporates pathways to every previous processing step using a recurrent weighted average (RWA). By maintaining a running average, the computational cost scales like that of other RNN models.

Figure 2:

(a) The generator function used to create each sentence. Examples of valid and invalid sentences are shown below. (b) A plot comparing the performance of the RWA and LSTM models. The traces show the accuracy of each model on the test data while the models are being fitted to the training data. The RWA model reaches 100% accuracy slightly before the LSTM model.

Figure 3:

Plot comparing the performance of the RWA and LSTM models when classifying sequences by length. The models must determine if a sequence is longer than 500 symbols, which is half the maximum possible sequence length. Each sequence is populated with numbers drawn at random from a unit normal distribution. The traces show the accuracy of each model on the test data while the models are being fitted to the training data. The RWA model achieves a classification accuracy near 100% before the LSTM model.

Figure 4:

(a) An example of the variable copy problem. The input and target sequences represent a single sample of data. (b) A plot comparing the performance of the RWA and LSTM models when the sequences are 100 symbols in length. The traces show the error of each model on the test data while the models are being fitted to the training data. The RWA model beats the baseline score (dashed line) before the LSTM model. (c) The same problem as before except the sequences are 1,000 symbols in length. The RWA model beats the baseline score (dashed line) before the LSTM model.

Figure 5:

(a) An example of the adding problem. The input sequence and target output represent a single sample of data. (b) A plot comparing the performance of the RWA and LSTM models when the sequences are of length 100. The traces show the error of each model on the test data while the models are being fitted to the training data. The RWA model beats the baseline score (dashed line) before the LSTM model. (c) The same problem as before except the sequences are of length 1,000. The RWA model beats the baseline score (dashed line) before the LSTM model.

Figure 6:

(a) Examples of the MNIST classification task. Each image of a handwritten digit must be classified by the value it represents. The images are feed into the RNNs as a sequence of pixels one at a time. (b) A plot comparing the performance of the RWA and LSTM models. The traces show the accuracy of each model on the test data while the models are being fitted to the training data. The RWA model beats the baseline score (dashed line) before the LSTM model. (c) Same task as before expect that the pixels have been randomly permuted. The LSTM model trained much faster on the permutation task while the RWA model took slightly longer.