

AN OPTIMIZED PARALLEL IMPLEMENTATION OF NON-ITERATIVELY TRAINED RECURRENT NEURAL NETWORKS

Julia El Zini, Yara Rizk and Mariette Awad

*Department of Electrical and Computer Engineering
American University of Beirut*

E-mail: {jwe04,yar01,mariette.awad}@aub.edu.lb

Submitted: 7th May 2020; Accepted: 14th September 2020

Abstract

Recurrent neural networks (RNN) have been successfully applied to various sequential decision-making tasks, natural language processing applications, and time-series predictions. Such networks are usually trained through back-propagation through time (BPTT) which is prohibitively expensive, especially when the length of the time dependencies and the number of hidden neurons increase. To reduce the training time, extreme learning machines (ELMs) have been recently applied to RNN training, reaching a 99% speedup on some applications. Due to its non-iterative nature, ELM training, when parallelized, has the potential to reach higher speedups than BPTT.

In this work, we present Opt-PR-ELM, an optimized parallel RNN training algorithm based on ELM that takes advantage of the GPU shared memory and of parallel QR factorization algorithms to efficiently reach optimal solutions. The theoretical analysis of the proposed algorithm is presented on six RNN architectures, including LSTM and GRU, and its performance is empirically tested on ten time-series prediction applications. Opt-PR-ELM is shown to reach up to 461 times speedup over its sequential counterpart and to require up to 20x less time to train than parallel BPTT. Such high speedups over new generation CPUs are extremely crucial in real-time applications and IoT environments.

Keywords: GPU implementation, parallelization, Recurrent Neural Network (RNN), Long-short Term Memory (LSTM), Gated Recurrent Unit (GRU), Extreme Learning Machines (ELM), non-iterative training

1 Introduction

Recurrent neural networks (RNN) are a type of neural networks that have been successfully applied to many problems in machine learning [22]. They have proven their ability to exceed human performance in time series prediction and sequential decision-making [31]. RNNs' training is usually based on gradient descent methods, specifically back-propagation through time (BPTT) [40], and real-time recurrent learning [41] which require a

substantial amount of iterations before converging. Moreover, when unfolded through time, RNNs become even deeper [1] and their training becomes even more expensive since the number of learned weights grows exponentially with the number of hidden neurons and the length of time dependency.

Non-iterative training algorithms have been investigated in the literature [32, 1, 35] to reduce the training cost of neural networks. Recently, Ertugrul et al. [10] proposed a non-iterative train-

ing algorithm for Jordan RNNs[19]. Then, Rizk et al. [30] extended it to different RNN architectures, including Elman, fully connected RNN, and Long Short-Term Memory (LSTM). Their algorithm was tested on time-series and sequential decision-making problems and achieved a speedup of up to 99% over iterative training.

Although they only need one iteration to obtain near-optimal solutions, non-iterative training algorithms minimize their cost function by computing a Moore-Penrose pseudo-inverse which requires ample computational resources, especially for large matrices. To the best of our knowledge, no attempts have been made in the literature to parallelize non-iterative training algorithms for RNNs. Fortunately, such algorithms hold great potential for parallelization due to their non-sequential nature.

In this work, we propose Basic-PR-ELM, a basic parallel version of ELM training applied on six RNN architectures: Elman, Jordan, NARMAX, fully connected, LSTM, and GRU. Basic-PR-ELM relies on parallel QR factorization to solve the pseudo-inverse required in ELM training algorithms. Then, the memory access patterns were studied and led to Opt-PR-ELM, an optimized version of parallel ELM training that utilizes the GPU shared memory to speedup the training process further.

The proposed algorithms, Basic-PR-ELM and Opt-PR-ELM, are tested on 10 publicly available time-series prediction applications and on different GPU architectures to empirically show their scalability, robustness, portability, speedup potentials, and energy efficiency. Compared to the sequential version proposed by Rizk et al. in [30], Basic-PR-ELM and Opt-PR-ELM achieve a speedup of up to 311 and 461, respectively on the LSTM architecture. Notably, Opt-PR-ELM is shown to train LSTM networks 20 times faster than the parallel iterative training algorithms (BPTT).

The rest of the paper is organized as follows: Section 2 presents the background on ELM-training and the RNN architectures. Section 3 summarizes the related work on RNN training and the parallel training algorithms. Section 4 presents the proposed algorithms Basic-PR-ELM and Opt-PR-ELM and Section 5 theoretically analyzes their memory and floating-point operations. Then, Sections 6 discusses the experimental setup and Sec-

tion 7 reports the empirical results. Finally, Section 8 concludes with final remarks.

2 Background

2.1 Extreme Learning Machine

Extreme Learning Machine (ELM) is a non-iterative training algorithm introduced by Huang et al. [16] for single hidden layer feedforward neural networks (SLFNs). Given n arbitrary distinct training samples (x_j, y_j) where $x_j \in \mathbb{R}^m, y_j \in \mathbb{R}$, M hidden nodes and g as activation function, the predicted output O_j can be written as

$\sum_{i=1}^M \beta_i g(w_i^T x_j + b_i)$ where $w_i \in \mathbb{R}^m$ is the weight vector connecting the i th hidden node and the input nodes, $\beta \in \mathbb{R}^M$ is the weight vector connecting all the hidden nodes and the output node and b_i is the bias of the i th hidden node. Throughout the training, the input weights w_{ij} are randomly generated and fixed and the output weights $\beta_1 \dots \beta_M$ are analytically computed. The goal is to minimize the error between the predicted and the true output as

$$\min_{\beta} \sum_{j=1}^n \|O_j - t_j\|^2 = \sum_{j=1}^n \left\| \sum_{i=1}^M \beta_i g(w_i^T x_j + b_i) - t_j \right\|. \quad (1)$$

Defining H and T as:

$$H_{(n \times M)} = \begin{bmatrix} g(w_1^T x_1 + b_1) & \dots & g(w_M^T x_1 + b_M) \\ \vdots & \ddots & \vdots \\ g(w_1^T x_n + b_1) & \dots & g(w_M^T x_n + b_M) \end{bmatrix} \quad (2)$$

$$T_{(n \times 1)} = [t_1, t_2, \dots, t_n]^T, \quad (3)$$

one can compactly write the problem in Equation 1 as minimizing $\|H\beta - T\|^2$. The solution of this problem is given as: $\beta = H^\dagger T$, where $H^\dagger = (H^T H)^{-1} H$ is the Moore-penrose generalized inverse of the matrix H .

2.2 RNN architectures

RNNs are one of the most powerful neural networks that are best suitable to model long-term dependencies in time-series applications [31]. RNN architectures differ in the way cycles are introduced in the network. In this work, we consider six RNN architectures, illustrated in Figure 1: Elman [9], Jordan [19], NARMAX [8],

fully connected RNN, LSTM [15] and GRU [5].

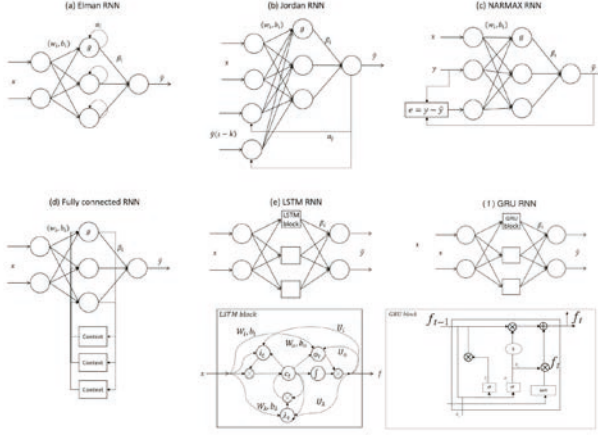


Figure 1. RNN architectures adapted from prior work in [30]

In Figure 1 and throughout this work, $x \in S \times Q$ is the input to the network, M is the number of hidden neurons, $w_i \in \mathbb{R}^S$ is the vector of weights connecting the input to the i^{th} neuron, $\alpha_{ik} \in \mathbb{R}$ is the weight from the neuron i to itself from the k^{th} previous time step and b_i is i^{th} bias.

2.2.1 Elman

Elman RNNs are single hidden layer networks where context neurons introduce recurrence by feeding back signals as internal state of the network. At time step t , the output is

$$\hat{y} = \sum_{i=1}^M \beta_i f_i(t), \quad (4)$$

where $f_i(t) = g\left(w_i^T x(t) + \sum_{k=1}^Q \alpha_{ik} f_i(t-k) + b_i\right)$ is the output of neuron i at time t .

2.2.2 Jordan

Jordan networks are similar to Elman's except for the way recurrence is introduced. In the Jordan architecture, signals are fed back from the predicted output of the previous time step. Consequently, such networks are more suitable for time series prediction where dependencies are on current input and previous outputs. Specifically, the output at time step t is described by Equation 4 with

$$f_i(t) = g\left(w_i^T x(t) + \sum_{k=1}^Q \alpha_{ik} \hat{y}(t-k) + b_i\right).$$

2.2.3 NARMAX

The Nonlinear Autoregressive Moving Average model with exogenous inputs (NARMAX) represents a wide class of nonlinear systems [2]. NARMAX networks, have been proposed for non-linear time series prediction using artificial neural networks and are described by $\hat{y}(t) = \sum_{i=1}^M \beta_i g\left(w_i^T x(t) + \sum_{l=1}^F w'_{il} y(t-l) + \sum_{l=1}^R w''_{il} e(t-l) + b_i\right)$, where F and R are the lengths of the time dependency of the output and the error feedbacks respectively, $e(t) = y(t) - \hat{y}(t)$, $w'_{il} \in \mathbb{R}$ ($w''_{il} \in \mathbb{R}$ resp.) is the weight from the output (error resp.) at the l^{th} time step to the i^{th} hidden neuron.

2.2.4 Fully Connected RNN

A fully connected RNN is the most general RNN architecture in which signals are fed back from all hidden neurons at previous time steps. Specifically, the output at time step t is described by Equation 4 with $f_i(t) = g\left(w_i^T x(t) + \sum_{l=1}^Q \sum_{m=1}^M \alpha_{ilm} f_m(t-k) + b_i\right)$. In this case, $\alpha_{ilm} \in \mathbb{R}$ is the weight connecting the neuron i to neuron l from the k^{th} previous time step.

2.2.5 LSTM

LSTMs were introduced by [15] to solve the vanishing gradient problem in BPTT. LSTMs have been successfully applied to a wide variety of applications including speech recognition [12, 13], machine translation [4, 42] and human action recognition [23, 24]. An LSTM unit is composed of the main cell, an input, output and forget gates which regulate the flow of information into and out of the cell through forgetting factors and weights. This formulation gives the network the ability to decide on which information to remember. The output of LSTM is described by Equation 4 with $f(t) = o(t) \circ g_f(c(t))$, \circ is the Hadamard product of two matrices and $o(t)$, $c(t)$, $\lambda(t)$ and $in(t)$ are given by

$$\begin{aligned} o(t) &= g_o(W_o x(t) + U_o f(t-1) + b_o) \\ c(t) &= \lambda(t) \circ c(t-1) + in(t) \circ g_c(W_c x(t) + U_c f(t-1) + b_c) \\ \lambda(t) &= g_\lambda(W_\lambda x(t) + U_\lambda f(t-1) + b_\lambda) \end{aligned}$$

$$in(t) = g_{in}(W_{in}x(t) + U_{in}f(t-1) + b_{in}).$$

2.2.6 GRU

GRUs are introduced in [5] as a gating mechanism for RNNs. They resemble LSTMs but have only two gates and fewer parameters. GRUs expose their state at each time step and do not have any mechanism to control the degree to which their state is exposed [7]. They exhibit good performances on small datasets [7] and are widely used in speech recognition [34, 6] and sequence modeling [7]. GRUs' output is described by Equation 4 while $f(t)$ is given by

$$f(t) = (1 - z(t)) \circ f(t-1) + z(t) \circ g_f(W_f x(t) + U_f(r_t \circ f(t-1) + b_f)), \quad (5)$$

where $z(t) = g_z(W_z x(t) + U_z f(t-1) + b_z)$ and $r(t) = g_r(W_r x(t) + U_r f(t-1) + b_r)$.

3 Related Work

This work focuses on the parallelization of a non-iterative training algorithm for RNNs. In what follows, we first discuss the basic training methods of RNNs while focusing on the non-iterative ones. Then, we report the parallelization attempts for training algorithms.

3.1 RNN Training

3.1.1 Iterative RNN Training

Training RNNs has been mainly done iteratively through BPTT [40] which unfolds the recurrence through time to transform the RNN into a feedforward network trained using gradient descent. BPTT is susceptible to local minima and suffers from the vanishing and exploding gradient problems with long time dependencies. BPTT can also be slow, given that it is applied iteratively in batch mode. Other iterative algorithms include, but are not limited to, Hessian free optimization [25], extended Kalman filters [39] and genetic algorithms (GA) [3]. Although successful, these algorithms are computationally expensive and require manually tuning of many hyper-parameters.

3.1.2 Non-Iterative RNN Training

Different non-iterative training algorithms have been proposed to reduce the computational cost of training neural networks in general. For instance, the authors in [32, 35, 28, 16] proposed ELM, a non-iterative method to train single hidden layer feedforward networks by randomly assigning input weights and computing output weights using the least-squares method. These methods were later extended to RNN architectures when Ertugrul implemented a non-iterative training for the Jordan RNN architecture in electricity load forecasting applications [10]. Later, Park et al. extended it to online RNNs [29] and Rizk et al. generalized the approach to more powerful RNN architectures [30].

Although these methods achieved high speedups (up to 99% in [30]), they heavily rely on stencil operations and on the computation of the generalized inverse of matrices which are CPU intensive operations and could be further optimized using parallel algorithms.

3.2 Parallelizing Training Algorithms

Several frameworks have been developed to solve challenges of high performance computing in the big data area [43], including parallelizing training algorithms. This is the first attempt to parallelize non-iterative training of RNNs; thus we describe previous work on the parallelization of RNN *iterative* training algorithms and on the *parallel non-iterative* training for neural networks - not exclusively RNN.

3.2.1 Parallelizing Iterative Training Algorithms For RNN

Parallelizing RNN training is mostly based on parallelizing the back-propagation algorithm (BP). For instance, Sierra et al. parallelized BP on CUBLAS and achieved a speedup of 63. In [44], data is distributed on multiple GPUs achieving a speedup of up to 51 [33]. In [38], parallel scan algorithm improves the step complexity of BP from $O(n)$ to $O(\log n)$. Khomenko et al. parallelized their data on multiple GPUs and relied on batch bucketing by input sequence length to accelerate RNN training achieving a speedup of up to 4 [20]. In [27], a semantic correlation-based data pre-fetch framework is implemented to break the dependency

in the input to parallelize the training of cognitive applications [27]. Their work is tested on LSTMs using image captioning, speech recognition, and language processing applications showing a speedup of 5.1, 44.9 and 1.53, respectively. Recently, GA is introduced into the Elman architecture to accelerate the training and prevent the local minima problem [18]. GA-Elman outperforms traditional training algorithms in terms of convergence speed and accuracy.

3.2.2 Parallelizing Non-Iterative Training Algorithms

Non-iterative training algorithms for RNNs are shown to require less training time than iterative methods [30, 10, 29]. However, even with non-iterative training, large datasets require costly computations, especially when increasing the number of neurons or when model selection is performed to avoid over-fitting [36]. Parallelizing non-iterative training has been explored in single layer feedforward networks by [14]. Their approach is based on a Map-Reduce and achieves a speedup of up to 5.6 when tested on 32 cores. Following a similar approach, Wang et al. [37] developed a parallel implementation of online ELM and achieved a speedup of 3.5 when trained on 120K instances with 120 attributes. Huang et al. extended their approach to the ensemble online sequential ELM which was tested on real and synthetic data with 5120K training instances and 512 attributes and achieved a speedup of 40 on a cluster with 80 cores [17]. In [36], Van et al. attempted to parallelize ELM on Flink with multi hidden layer feedforward network and achieved a speedup of 17.

To the best of our knowledge, our work is the first attempt to parallelize non-iterative training for different RNN architectures.

4 Methodology

Before proposing our methods, we present the nomenclature that will be used throughout this paper in Table 1.

Table 1. Nomenclature

Symbol	Definition
n	Number of training samples
M	Number of hidden neurons
Q	Max number of time dependencies
S	Dimension of input
$x_j \in \mathbb{R}^{S \times Q}$	j^{th} Input instance
$y_j \in \mathbb{R}$	j^{th} Output instance
$X \in \mathbb{R}^{n \times S \times Q}$	Input matrix
$Y \in \mathbb{R}^n$	Output matrix
$W \in \mathbb{R}^{S \times L}$	Weight matrix connecting the input to the hidden neurons
$\alpha \in \mathbb{R}^{L \times Q}$	Weight matrix connection the hidden neuron to itself for previous time steps
$b \in \mathbb{R}^L$	Bias vector for the hidden neurons
$\beta \in \mathbb{R}^L$	Weight vector connecting hidden neurons to output layer
S-R-ELM	Sequential ELM for RNN training
Basic-PR-ELM	Basic parallel ELM RNN training
Opt-PR-ELM	Optimal parallel ELM RNN training
BPTT	Back-propagation through time
P-BPTT	Parallel Back-propagation through time
BS	Block size
TW	Tile width

In this work, a parallel version of ELM-trained RNNs will be formalized and implemented. The sequential version of our approach, denoted by S-RELM, is summarized in Algorithm 1 and is adopted from our previous work in [30].

Algorithm 1 S-R-ELM algorithm

- 1: Randomly assign W, α, b
- 2: Compute $H(t), t = 1 \dots Q$ according to the corresponding RNN architecture
- 3: Compute $\beta = H(Q)^\dagger Y$ using the generalized Moore–Penrose pseudoinverse

$H(t)$ at row i and column j is referred to as $h_{ij}[t]$ in this paper and is computed as in Equations 6, 7, 8, 9, 10 and 11 for the Elman, Jordan, NARMAX, fully connected, LSTM and GRU architectures respectively.

$$h_{ij}[t] = g(W[:, j].X[i, :, t] + b_i + \sum_{k=1}^Q \alpha[j, k]h_{ij}[t - k]) \quad (6)$$

$$h_{ij}[t] = g(W[:, j].X[i, :, t] + b_i + \sum_{k=1}^Q \alpha[j, k]\hat{y}(t - k)) \quad (7)$$

$$h_{ij}[t] = g(W[:, j].X[i, :, t] + b_i + \sum_{l=1}^F W'[i, l]y(t - l) + \sum_{l=1}^R W''[i, l]e(t - l)) \quad (8)$$

$$h_{ij}[t] = g\left(W[:, j].X[i, :, t] + b_i + \sum_{k=1}^Q \sum_{l=1}^M \alpha[j, l, k]h_{ij}[t - k]\right) \quad (9)$$

$$h_{ij}[t] = o[i, j, t] \circ g_f(c[i, j, t]) \quad (10)$$

$$h_{ij}[t] = \left(1 - z[i, j, t]\right) \circ h_{ij}[t - 1] + z[i, j, t] \circ g_f\left(W_f[:, j].X[i, :, t] + U_f(r[i, j, t] \circ h_{ij}[t - 1] + b_i)\right) \quad (11)$$

Considering Algorithm 1, one can see that the running time of the ELM training mainly consists of two CPU intensive operations: computing H and computing β by solving the linear system using the Moore–Penrose pseudo-inverse. Thus, those two operations are the main target when optimizing the performance of non-iterative training.

4.1 H Computation**4.1.1 Basic Parallel Implementation (Basic-PR-ELM)**

For all RNN architectures, the computation of $H(t)$ at row i and column j is independent of the computation of $H(t)$ at row i_2 and column j_2 , $\forall i_2 \neq i, j_2 \neq j$; it only depends on $H(t_2)$ at row i and column j for $t_2 < t$. Given only this dependency, a parallel H computation can be done as follows: each thread (i, j) can independently compute $H(t)$ at row i and column j for $t = 1, \dots, Q$. We describe the basic implementation of the computation of H for the Elman architecture in Algorithm 2.

Algorithm 2 Basic-PR-ELM by thread (i, j)

- 1: $tx \leftarrow \text{threadIdx.x}$
- 2: $ty \leftarrow \text{cuda.threadIdx.y}$
- 3: $Row \leftarrow tx + \text{blockIdx.x} \times \text{blockDim.x}$
- 4: $Col \leftarrow ty + \text{blockIdx.y} \times \text{blockDim.y}$
- 5: **for** $t = 1 \mapsto Q$ **do**
- 6: $h_{ij} \leftarrow W[:, Col].X[Row, :, t]$
- 7: $h_{ij} \leftarrow h_{ij} + b_{Col}$
- 8: **for** $t_{prev} = 1 \mapsto t$ **do**
- 9: $h_{ij} \leftarrow h_{ij} + \alpha[j, t_{prev}] \times H[Row, Col, t_{prev}]$
- 10: **end for**
- 11: $H[Row, Col, t] \leftarrow h_{ij}$
- 12: **end for**


4.1.2 Optimized Parallel Implementation (Opt-PR-ELM)

Figure 2. Basic-PR-ELM memory access patterns on Elman

Figure 2 illustrates the memory access patterns of Basic-PR-ELM on the Elman architecture. One can clearly see that threads in the same row access the same elements of X and threads in the

same column access the same elements of W and α . Thus, the tiling concept can be applied to utilize the shared memory to speed up the computation of H . Moreover, we notice that b_{Col} can be preloaded and used efficiently by other threads.

Algorithm 3 describes how these optimizations can be applied for the Elman architecture.

Algorithm 3 Opt-PR-ELM by thread (i, j)

```

1:  $tx \leftarrow threadIdx.x$ 
2:  $ty \leftarrow cuda.threadIdx.y$ 
3:  $Row \leftarrow tx + blockDim.x \times blockDim.x$ 
4:  $Col \leftarrow ty + blockDim.y \times blockDim.y$ 
5:  $H_{loc} \leftarrow t$ -dimensional array in the local memory of thread  $(i, j)$ 
6: for  $t \leftarrow 1 \mapsto Q$  do
7:    $h_{ij} \leftarrow 0$ 
8:   for  $tile = 1 \mapsto num\_tiles$  do
9:      $W_{shared} \leftarrow W[tx + tile \times TW :, Col]$ 
10:     $X_{shared} \leftarrow X[Row, ty + tile * TW, t]$ 
11:     $synch()$ 
12:     $h_{ij} \leftarrow h_{ij} + W_{shared} \cdot X_{shared}$ 
13:  end for
14:   $synch()$ 
15:  if  $tx = 0$  and  $ty = 0$  then
16:     $b_{shared} \leftarrow b[Col]$ 
17:  end if
18:   $synch()$ 
19:   $h_{ij} \leftarrow h_{ij} + b_{shared}$ 
20:  for  $tile \leftarrow 1 \mapsto \lceil \frac{t}{TW} \rceil$  do
21:     $\alpha_{shared} \leftarrow \alpha[Col, tx + tile \times TW]$ 
22:     $synch()$ 
23:     $h_{ij} \leftarrow h_{ij} + \alpha_{shared} \cdot H_{loc}[t_{prev}]$ 
24:  end for
25:   $synch()$ 
26:   $H_{loc}[t] \leftarrow h_{ij}$ 
27:   $H[Row, Col, t] \leftarrow h_{ij}$ 
28: end for

```

First, in the dot product $W[:, Col] \cdot X[Row, :, t]$, each thread can load only one element of W and one element of X into the shared memory. Once the threads synchronize, then all needed elements of W and X are loaded, and the dot product can be computed efficiently. Second, only one thread can load $b[j]$ that is needed by all the threads in the same column of the block. The same tiling concept used to compute $W[:, Col] \cdot X[Row, :, t]$ can be used to speed up the computation of $\alpha[j, t_{prev}] \times$

$H(t_{prev})[Row, Col]$. Lastly, each thread can save the values of $H(t)[Row, Col]$ in its register file to reduce the time taken to read from the global memory in line 8 of Algorithm 2. If these values do not fit in the registers, they are read from the global memory.

Algorithms 2 and 3 could be easily extended to other architectures when Equation 6 is replaced by Equations 7, 8, 9, 10 or 11.

4.2 Computing β

β is the solution of the following system: $H\beta = Y$. Instead of computing the pseudo-inverse H^\dagger and then multiplying it by Y , one can perform a QR factorization of H as $H = QR$, then compute $z = Q^T Y$. Having that, β would be the solution of $R\beta = z$ by back substitution since R will be an upper triangular matrix. In this work, we make use of *Numba* [21] and *Numpy* [26] libraries which provide an efficient implementation of this method in *Python*.

5 Theoretical Analysis

We analyze the memory read and write operations and the floating point operations (FLOPS) for the proposed algorithms: Basic-PR-ELM and Opt-PR-ELM. For the Elman architecture, Basic-PR-ELM performs $Q(2S + Q + 2)$ read operations divided as follows:

- $2 \times SQ$ to read the values needed in line 6
- Q reads for b_{Col} in line 7
- $2 \times (Q \frac{Q+1}{2})$ reads in the loop at line 8


Moreover, only Q write operations are needed (in line 11) and $Q(2S + Q + 2)$ FLOPS are performed as follows:

- $2 \times SQ$ to perform the dot product at line 6
- Q FLOPS for the addition in line 7
- $2 \times (Q \frac{Q+1}{2})$ to perform the loop at line 8

The memory operations to FLOPS ratio is $\frac{2S+Q+3}{2S+Q+2} > 1$ which might limit the performance of Basic-PR-ELM. This ratio improves with Opt-PR-ELM as it minimizes the memory operations while keeping the same number of FLOPS. Specifically, Opt-PR-ELM decreases the number of reads to $\frac{1}{TW^2} (2 \times SQ + \frac{Q(Q+1)}{2}) + 1$ divided as follows:

Table 2. Number of memory operations and FLOPS for each RNN architecture for Basic-PR-ELM

Architecture	# Read Operations	# Write Operations	FLOPS
Elman	$Q(2S + Q + 2)$	Q	$Q(2S + Q + 2)$
Jordan	$Q(2S + 1 + (Q + 1)(1/2 + M))$	Q	$Q(2S + 1 + \frac{Q+1}{2}(2SM + M))$
NARMAX	$Q(2S + 1) + 2(2F + M + R)$	Q	$Q(2S + 1 + 2F + R(2 + 2SM + M))$
Fully Connected	$Q(2S + 1 + 2MQ)$	Q	$Q(2S + Q + 2QM)$
LSTM	$Q(5S + 13)$	$5Q$	$Q(8S + 18)$
GRU	$Q(4S + 8)$	$3Q$	$Q(3S + 17)$

**Figure 3.** Speedup of Basic-PR-ELM and Opt-PR-ELM for the different architectures when $M = 50$

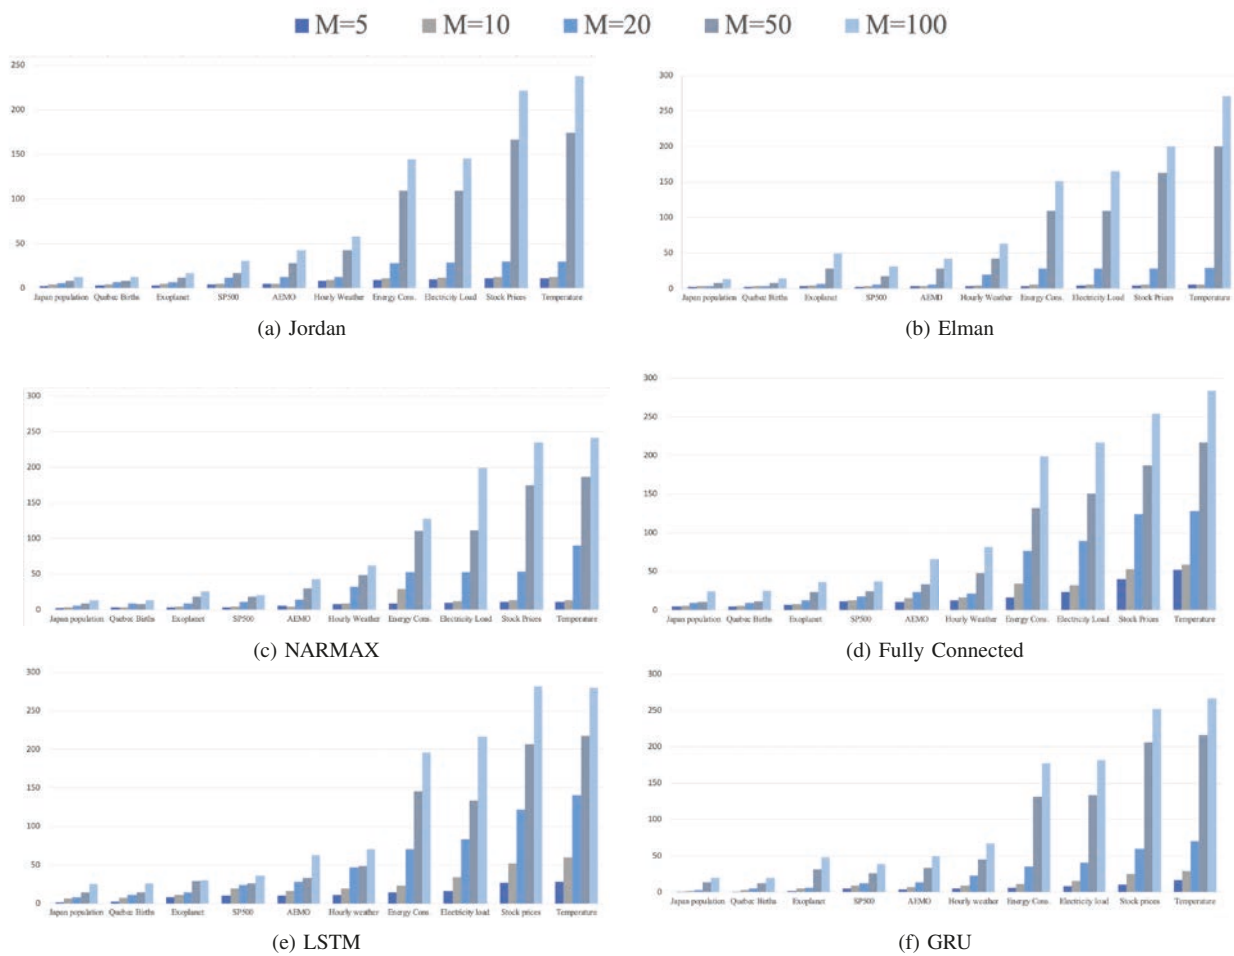


Figure 4. Speedup of Opt-PR-ELM for the different architectures when the number of hidden neurons increases from 5 to 100

Table 3. Benchmarks Description

Database		Size			Output Statistics			
Category	Name	# of instances	Q	% Train	Mean	Std Dev	Min	Max
Small	Japan population	2,540	10	80	1.40E+06	1.40E+06	1.00E+05	1.03E+08
	Quebec Births	5,113	10	80	2.51E+02	4.19E+01	-2.31E+01	3.66E+02
	Exoplanet	5,657	3197	80	-3.01E+02	1.45E+04	-6.43E+05	2.11E+05
Medium	SP500	17,218	10	80	8.99E+08	1.53E+09	1.00E+06	1.15E+10
	AEMO	17,520	10	80	7.98E+03	1.19E+03	5.11E+03	1.38E+04
	Hourly weather	45,300	50	80	2.79E+02	3.78E+01	0.00E+00	3.07E+02
Large	Energy Consumption	119,000	10	70	1.66E+03	3.02E+02	0.00E+00	3.05E+03
	Electricity load	280,514	10	70	2.70E+14	2.60E+14	0.00E+00	9.90E+14
	Stock prices	619,000	50	70	4.48E+06	1.08E+07	0.00E+00	2.06E+09
	Temperature	998,000	50	70	5.07E+01	2.21E+01	4.00E+00	8.10E+01

– $\frac{2}{TW^2} \times SQ$ to read the values needed in line 12

– at most 1 read for b_{Col} in line 16

– $\frac{1}{TW^2} (Q \frac{Q+1}{2})$ reads in the loop at line 20

where TW is the tile width which is set to block size in this work. The new memory operations to FLOPS ratio is $\frac{\frac{1}{TW^2} (2 \times SQ + \frac{Q(Q+1)}{2}) + 1 + Q}{Q(2S+Q+2)}$ which is less than the ratio of Basic-PR-ELM by a factor of $\approx TW^2$. Specifically, Opt-PR-ELM minimizes the number of read operations by a factor of 256 (1024 resp.) when the tile width is set to 16 (32 resp.).

Table 2 reports the number of memory operations and FLOPS needed by Basic-PR-ELM for each RNN architecture. The values of Opt-PR-ELM are omitted as it requires the same number of write operations and FLOPS and less read operations by a factor of $\approx TW^2$.

¹kaggle.com/jd1325/japan-population-data

²datamarket.com/data/list/?q=provider%3Atdsl

³kaggle.com/keplersmachines/kepler-labelled-time-series-data

⁴kaggle.com/benjibb/sp500-since-1950

⁵aemo.com.au/

⁶kaggle.com/selfishgene/historical-hourly-weather-data

6 Experimental Setup

6.1 Setup

Serial algorithms were run on an Intel 64-bit core-i7 machine with a memory of 16 GB. Parallel algorithms were run on NVidia Tesla K20m GPU with 2688 CUDA cores and 723MHz GPU core clock speed. The GPU main memory is 6GB and bandwidth of 250 GB/s between the host and the device. All experiments are repeated 5 times, and the average value is reported.

6.2 Time Series Prediction Benchmarks

Basic-PR-ELM and Opt-PR-ELM were validated on time series prediction problems. Table 3 presents the characteristics of the datasets ordered according to the number of instances. According to their size, we split the databases into three categories: small datasets containing less than 10K instances, medium datasets with multiples of 10K instances and large dataset consisting of multiples of 100K instances. Japan population¹ tracks the population of various Japanese regions, while the Quebec Births² tracks the number of births in Quebec

Table 4. Average RMSE (\pm standard deviation) of S-R-ELM and Opt-PR-ELM (BS=32) showing that both algorithm achieve accuracies within the same range for different RNN architectures on all the datasets.

Dataset	Algorithm	Architecture						
		Elman	Jordan	NARMAX	Fully Connected	LSTM	GRU	
Japan pop.	S-R-ELM	3.97E-2 \pm 4.67E-2	1.12E-1 \pm 3.75E-1	6.54E-1 \pm 3.32E-2	5.43E-3 \pm 3.89E-5	2.45E-1 \pm 2.36E-1	4.46E-1 \pm 3.35E-4	
	Opt-PR-ELM	3.74E-2 \pm 7.17E-8	1.23E-1 \pm 2.89E-2	6.23E-1 \pm 2.31E-2	6.23E-3 \pm 2.65E-4	2.46E-1 \pm 4.56E-2	4.75E-2 \pm 5.81E-5	
Quebec Births	S-R-ELM	4.06E-3 \pm 7.68E-5	1.01E-1 \pm 5.00E-3	3.42E-1 \pm 5.05E-3	2.02E-2 \pm 4.99E-7	1.01E-1 \pm 5.76E-1	1.01E+0 \pm 5.16E-4	
	Opt-PR-ELM	2.02E-3 \pm 4.89E-5	4.35E-1 \pm 5.32E-4	3.46E-1 \pm 3.79E-3	2.42E-2 \pm 7.07E-1	1.49E-2 \pm 1.46E-4	1.16E+0 \pm 3.56E-3	
Exoplanets	S-R-ELM	5.40E+0 \pm 3.03E-1	2.87E+0 \pm 7.91E-3	2.01E-1 \pm 2.98E-3	3.46E-1 \pm 1.01E-2	5.45E-1 \pm 2.31E-1	4.32E+0 \pm 4.56E-1	
	Opt-PR-ELM	5.42E+0 \pm 3.05E-1	2.34E+0 \pm 7.34E-2	2.53E-1 \pm 1.98E-3	3.42E-1 \pm 1.51E-2	3.65E-1 \pm 2.31E-5	5.21E+0 \pm 3.76E-2	
SP500	S-R-ELM	1.69E-1 \pm 7.78E-3	1.32E-1 \pm 3.75E-4	9.01E-1 \pm 8.70E-4	1.96E+0 \pm 4.32E-1	1.01E-1 \pm 5.16E-2	7.84E+0 \pm 5.55E-2	
	Opt-PR-ELM	2.34E-1 \pm 7.98E-4	4.01E-1 \pm 6.36E-5	9.11E-1 \pm 8.32E-5	1.36E+0 \pm 1.90E-2	1.24E-1 \pm 3.14E-2	7.83E+0 \pm 5.53E-1	
AEMO	S-R-ELM	1.26E-1 \pm 1.45E-3	3.30E-2 \pm 7.16E-3	9.61E-2 \pm 8.79E-3	5.00E-2 \pm 1.32E-5	1.36E-2 \pm 5.33E-4	2.33E-1 \pm 2.23E-5	
	Opt-PR-ELM	1.34E-1 \pm 1.25E-4	1.12E-2 \pm 5.16E-2	3.23E-3 \pm 1.01E-2	5.36E-2 \pm 1.12E-4	1.22E-2 \pm 5.67E-3	2.01E-1 \pm 2.13E-6	
Hourly Weather	S-R-ELM	1.98E-1 \pm 5.17E+0	3.14E-1 \pm 2.07E-3	8.06E-1 \pm 7.63E-5	7.39E-2 \pm 6.03E-2	2.10E-2 \pm 2.24E-5	3.21E-1 \pm 9.61E-3	
	Opt-PR-ELM	1.52E-1 \pm 3.34E+0	3.98E-1 \pm 5.67E-4	2.00E-1 \pm 7.03E-4	3.79E-2 \pm 5.03E-3	1.02E-2 \pm 2.14E-5	4.32E-1 \pm 9.16E-3	
Energy Cons.	S-R-ELM	1.83E-4 \pm 1.98E-3	2.21E-3 \pm 3.43E-1	2.22E-4 \pm 5.26E-3	3.56E-3 \pm 5.56E-4	1.56E-3 \pm 9.96E-4	2.34E-2 \pm 2.22E-5	
	Opt-PR-ELM	1.38E-4 \pm 2.45E-3	3.48E-3 \pm 3.03E-2	6.44E-5 \pm 5.16E-4	2.65E-3 \pm 5.16E-5	2.56E-3 \pm 5.326E-5	3.24E-3 \pm 2.12E-5	
Elec. Load	S-R-ELM	2.56E+0 \pm 7.93E+0	2.40E+0 \pm 3.90E-1	8.64E+0 \pm 9.81E+0	4.16E-1 \pm 3.45E-1	8.32E+0 \pm 8.05E+0	1.12E+0 \pm 5.16E-1	
	Opt-PR-ELM	2.34E+0 \pm 7.03E-1	4.76E+0 \pm 2.20E-2	4.86E+0 \pm 8.91E-1	4.64E-1 \pm 3.97E-2	2.84E+0 \pm 8.13E-1	2.98E+0 \pm 5.06E+0	
Stock Prices	S-R-ELM	6.41E-1 \pm 7.93E-1	1.10E-1 \pm 9.09E-5	4.80E+0 \pm 3.87E-1	2.13E-2 \pm 3.89E-1	4.00E-1 \pm 1.09E-3	2.62E-1 \pm 3.82E-4	
	Opt-PR-ELM	3.41E-1 \pm 3.35E-2	1.56E-1 \pm 9.23E-5	4.81E+0 \pm 3.32E-2	2.03E-3 \pm 1.92E-4	4.94E-1 \pm 5.69E-4	6.28E-1 \pm 3.28E-3	
Temp.	S-R-ELM	4.32E-4 \pm 9.85E-5	5.65E-3 \pm 6.79E-9	3.56E-4 \pm 7.10E-6	2.91E-5 \pm 3.72E-9	4.92E-4 \pm 6.02E-5	3.54E-4 \pm 2.95E-6	
	Opt-PR-ELM	4.12E-4 \pm 9.67E-4	5.03E-3 \pm 6.19E-2	3.15E-4 \pm 9.25E-6	9.21E-5 \pm 3.02E-5	8.17E-4 \pm 6.92E-4	3.19E-3 \pm 5.29E-5	

Table 5. Speedup of Opt-PR-ELM (BS=32) when tested on the Tesla K20m and Quadro K2000 GPUS for different RNN architectures on various datasets when the number of hidden neurons M is 20.

Archi- tec- ture	GPU	Dataset									
		Japan pop.	Quebec Births	Exop.	SP500	AEMO	Hourly weather	Energy cons.	Elec. Load	Stock Prices	Temp.
Elman	Tesla K20m	12	12	18	26	42	64	163	164	251	261
	Quadro K2000	12	10	16	23	40	61	60	163	239	251
Jordan	Tesla K20m	12	13	42	26	42	64	163	165	244	300
	Quadro K2000	11	11	39	23	40	60	163	163	189	295
NAR-	Tesla K20m	13	12	29	29	45	72	167	168	263	281
MAX	Quadro K2000	11	11	28	26	42	71	162	162	257	273
Fully	Tesla K20m	17	18	35	36	50	73	198	226	281	326
Connec.	Quadro K2000	14	16	33	34	48	71	196	225	279	324
LSTM	Tesla K20m	21	21	43	39	50	74	219	201	310	327
	Quadro K2000	19	20	41	36	45	70	215	196	307	323
GRU	Tesla K20m	20	18	46	40	50	67	197	200	309	326
	Quadro K2000	15	14	42	35	47	58	192	187	300	320

Table 6. Runtime (seconds) of Opt-PR-ELM (BS=32) and the iterative training algorithm and the ratio= $\frac{BP}{Opt-PR-ELM}$

	Fully Connected			LSTM			GRU		
	Opt-PR-ELM	P-BPTT	Ratio	Opt-PR-ELM	P-BPTT	Ratio	Opt-PR-ELM	P-BPTT	Ratio
Japan pop.	0.23	3.52	15	0.38	7.41	20	0.38	6.59	17
Quebec Births	0.56	6.75	12	0.85	13.56	16	0.81	12.94	16
Exoplanet	10.03	24.98	2	15.23	54.32	4	13.14	43.12	3
SP500	3.56	20.66	6	7.77	37.55	5	5.61	35.65	6
AEMO	3.01	21.34	7	7.29	38.32	5	5.62	35.71	6
Hourly Weather	30.46	156.76	5	50.49	243.99	5	30.04	201.12	7
Energy Cons.	32.14	203.45	6	51.90	525.87	10	45.67	435.89	10
Elec. Load	36.70	256.89	7	53.60	572.74	11	51.7	532.31	10
Stock Prices	41.30	301.23	7	56.78	639.04	11	52.34	621.18	12
Temperature	45.45	354.99	8	62.00	678.11	11	59.32	641.09	11

and Exoplanet³ describes the change in the light intensity of several thousand stars. Additionally, SP 500⁴ records the stock prices since 1950 while AEMO⁵ reports the electricity load demand in Australia and hourly weather⁶ contains ≈ 5 years of temperature measures. The energy consumption dataset⁷ reports the hourly power consumption data in megawatts, the electricity load dataset⁸ reports the electricity demand at the MT166 and MT257 substations and the stock prices dataset⁹ consists of historical stock prices for all companies currently on the S&P 500 index. Finally, the temperature dataset¹⁰ reports sensor data collected from a permanent magnet synchronous motor (PMSM) deployed on a testbench where PMSM represents a german OEM's prototype model.

7 Experimental Results

7.1 Speedup

Figure 3 illustrates the speedups of Basic-PR-ELM and Opt-PR-ELM for the six architectures tested against the serial version when the number of hidden neurons M is 50. Opt-PR-ELM was tested with two different configurations: when the number of threads per block, block size BS, is 16 and 32, respectively.

Clearly, Basic-PR-ELM and Opt-PR-ELM achieve high speedups, especially when the size of the dataset increases. For instance, for the Elman architecture, Basic-PR-ELM achieves a speedup of 19 on the small Exoplanet dataset, 72 on the hourly energy consumption medium dataset, and up to 207 on the largest dataset (Temperature). Opt-PR-ELM achieves higher speedups that reach up to 311 with LSTM on the temperature dataset when $BS = 16$. The speedup increases to 461 when BS increases to 32.

⁷kaggle.com/selfishgene/historical-hourly-weather-data

⁸archive.ics.uci.edu/ml/index.php

⁹kaggle.com/camnugent/sandp500

¹⁰kaggle.com/wkirgsn/electric-motor-temperature

However, Opt-PR-ELM does not always achieve higher speedups. Specifically, Basic-PR-ELM and Opt-PR-ELM achieve similar speedups for the Japan population, Quebec births, SP500, AEMO, energy consumption, and the electricity load datasets. To investigate these results, we take a closer look at the characteristics of the datasets. When $Q = 10$, a thread is computing the dot product between a row of X and a column of W , and it is doing 2×10 memory read operations. Consequently, num_tiles will be only 1, and the loop at line 8 of Alg. 3 will only be executed once. In this case, the performance does not improve and might slightly decrease due to the thread synchronization in Opt-PR-ELM. However, Opt-PR-ELM achieves higher speedups when $Q > BS$ and when BS increases to 32. We notice that the speedup increases with more complex architectures, LSTM for example, since these architectures require more computations that can be better accelerated on a GPU.

7.2 Scalability

To test the scalability of our approach, we change the number of hidden neurons M , and we report the speedup of Opt-PR-ELM ($BS=32$) for the different architectures on the various datasets. Figure 4 illustrates that the speedup increases when M increases from 5 to 10, 20, 50, 100. Specifically, the speedup increases by a factor of 20 when M increases from 5 to 100 with a GRU on the energy consumption dataset. Thus, Opt-PR-ELM scales up well with more computationally expensive operations.

7.3 Robustness

Robustness, i.e. repeatability, is a key property for Opt-PR-ELM where random initialization might affect the solution. Moreover, floating-point computations might differ between the GPU and the CPU, which might affect the output. To ensure that such perturbations do not affect the performance of our parallel algorithm, we run S-R-ELM and Opt-PR-ELM ($BS=32$) five times, and we measure their root mean squared error (RMSE). Table 4 reports the average RMSE and its standard deviation when S-R-ELM and Opt-PR-ELM are tested on different

datasets with different RNN architectures. We select M according to the size of the problem; i.e. we used $M = 100$ for exoplanet where $Q = 5657$, $M = 20$ for hourly weather, stock prices and temperature where $Q = 50$ and $M = 10$ for the rest of the datasets that have $Q = 10$. Tables 3 and 4 show that the cases where the RMSE is high correspond to datasets with large outputs. For instance, having outputs ranging from 0 to 2.06×10^9 , the electricity load dataset has higher RMSE than other datasets. However, S-R-ELM and Opt-PR-ELM achieve accuracies in the same range for different RNN architectures on all the datasets, which means that GPU floating-point operations do not have a clear effect on the performance of our algorithm.

7.4 Portability

To verify that our algorithm is portable, we ran Opt-PR-ELM ($BS=32$) on an NVIDIA Quadro K2000 GPU while fixing the number of hidden nodes M at 20. It is important to check for portability to understand how much the proposed algorithm is architecture-dependent. Table 5 shows that Opt-PR-ELM also achieves high speedups on the Quadro K2000 GPUs for different RNN architectures on different datasets, but the speedups on the Tesla K20m GPU are constantly higher because of the computational capability of the latter. The speedups in Table 5 are reported with respect to the core-i7 CPU with 16 GB RAMs. Speedups with respect to sequential code on older generation CPUs (core-i5 with 8GB RAMs) show a speedup of up to 5 times higher. One can draw the following conclusion: increasing the number of cores of a CPU yields a speedup of up to only 5 times. However, parallelizing the code can yield a speedup of up to 326 with respect to sequential code on core-i7 CPUs and 651 on core-i5 CPUs. A rough estimation of the current pricing based on google search shows that GPU architectures cost between 500\$ and 7,000\$ for NVIDIA GTX 1080 and Tesla GPUs¹¹ respectively, while CPU architectures such as Intel Core i7-9700K with 8 cores cost 400\$¹². Considering the aforementioned speedups, one can conclude that investing in parallel architectures can be more profitable than upgrading the existing CPU architecture, especially in applications where real-time perfor-

¹¹<https://www.amazon.com/PNY-TCSV100MPCIE-PB-Nvidia-Tesla-v100/dp/B076P84525>

¹²<https://www.amazon.com/CPU-Processors-Memory-Computer-Add-Ons/b?ie=UTF8&node=229189>

mance and cost efficiency are essential such as general IoT applications.

7.5 Comparison with Parallel Iterative RNN Training

Although Opt-PR-ELM achieves high speedups compared its S-R-ELM, we need to show that its absolute training time is lower than the parallel version of the BPTT (P-BPTT) as implemented in [11]. We choose the architectures that [11] implements, i.e. fully connected, LSTM and GRU, and we report the training time of Opt-PR-ELM (BS=32) and P-BPTT when $M = 10$. P-BPTT is trained for 10 epochs with 64 as batch size, *mean squared error* (MSE) as loss function and *ADAM* as optimizer. We are interested in the absolute training times of the two parallel algorithms rather than their speedup over their sequential versions. Thus, we report the runtimes of Opt-PR-ELM and P-BPTT algorithms when tested on the same Tesla K20m GPU and the ratio between both training times. As Table 6 shows, Opt-PR-ELM runs up to 10x faster than P-BPTT when tested with LSTM on the energy consumption dataset. Figure 5 illustrates the MSE versus time for P-BPTT algorithms when tested with LSTM on the energy consumption with $M = 50$. For the same dataset and RNN architecture, Opt-PR-ELM reaches 2.56×10^{-3} as MSE, whereas P-BPTT reaches a lower MSE of 1.4×10^{-3} . However, Opt-PR-ELM took only 57 sec to reach its optimal MSE, whereas P-BPTT took 525 sec to reach its optimal MSE and 340 sec to reach the same MSE (1.1×10^{-3}).




Figure 5. MSE versus time (sec) for P-BPTT algorithms when tested on the energy consumption dataset with $M = 50$ and LSTM as architecture

Thus, Opt-PR-ELM could reach the same performance as P-BPTT 6 times faster. The sequential nature of iterative training explains the results: although one can attempt to parallelize each epoch,

the training needs to be done in a sequence of consecutive dependent epochs.

7.6 Opt-PR-ELM Runtime

One can argue that using memory streams or initializing the random weights on the GPU can lead to higher speedups. To investigate this, we study how the runtime of Opt-PR-ELM is decomposed between the parameters initialization, data transfer to and from the GPU and the actual computations for the six architectures. Figure 6 shows what portion each step takes from the runtime of Opt-PR-ELM when tested on the energy consumption dataset with $M = 50$. The initialization does not appear on the bar because it is less than 0.01% of the total runtime. Moreover, transfer data to the GPU consistently takes more time than the transfer back because the former deals with the following matrices: $X \in \mathbb{R}^{n \times S \times Q}$, $Y \in \mathbb{R}^n$, $W \in \mathbb{R}^{S \times L}$, $\alpha \in \mathbb{R}^{L \times Q}$ and $b \in \mathbb{R}^L$, while the latter only transfers $\beta \in \mathbb{R}^L$. The steps that take the major time portion are the computations of H and β . One can conclude that data streams or the GPU random initializations will not affect the speedup since initialization and data transfer are not a bottleneck in Opt-PR-ELM.




Figure 6. Time decomposition of Opt-PR-ELM on the energy consumption dataset with $M = 50$

8 Conclusion

In this work, we proposed Opt-PR-ELM, a parallel version of non-iteratively trained RNNs for time series prediction. Focusing on six RNN architectures: Elman, Jordan, NARMAX, fully connected RNN, LSTM and GRU, we first developed a basic version of the parallel algorithm. Then, we studied its memory access patterns to propose an optimized version that takes advantage of the shared memory of the GPU. In addition to performing a theoretical, computational analysis of Opt-PR-ELM on the various architectures, empirical validation was performed on 10 publicly available time series prediction datasets.

Opt-PR-ELM was shown to achieve a speedup of up to 461 over its sequential version and requires less time to train than the parallel BPTT by a factor of 20. Higher speedups are achieved on older generation CPUs which highlights the importance of investing in high-end parallel architectures, especially in IoT and machine learning applications that require accurate, cost-sensitive yet efficient solutions.

We further studied the portability and scalability of our proposed algorithm by changing the GPU architecture and the number of hidden neurons and reporting the speedup. Opt-PR-ELM showed higher speedups when the number of computations increases or the number of launched threads per block increases. Finally, Opt-PR-ELM was shown to reach similar accuracies as its sequential version.

Future work includes extending Opt-PR-ELM to RNNs with multiple layers and investigating its performance on applications that have multi-dimensional outputs such as machine translation and speech recognition.

Acknowledgment

This work was supported by the University Research Board at the American University of Beirut.

References

- [1] Yoshua Bengio, Patrice Simard, Paolo Frasconi, et al. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.
- [2] Stephen A Billings. *Nonlinear system identification: NARMAX methods in the time, frequency, and spatio-temporal domains*. John Wiley & Sons, 2013.
- [3] Armando Blanco, Miguel Delgado, and Maria C Pegalajar. A real-coded genetic algorithm for training recurrent neural networks. *Neural networks*, 14(1):93–105, 2001.
- [4] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.
- [5] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [6] Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio. Attention-based models for speech recognition. In *Advances in neural information processing systems*, pages 577–585, 2015.
- [7] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [8] Jerome T Connor, R Douglas Martin, and Les E Atlas. Recurrent neural networks and robust time series prediction. *IEEE transactions on neural networks*, 5(2):240–254, 1994.
- [9] Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990.
- [10] Ömer Faruk Ertugrul. Forecasting electricity load by a novel recurrent extreme learning machines approach. *International Journal of Electrical Power & Energy Systems*, 78:429–435, 2016.
- [11] Martín Abadi et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [12] Alex Graves, Navdeep Jaitly, and Abdel-rahman Mohamed. Hybrid speech recognition with deep bidirectional lstm. In *2013 IEEE workshop on automatic speech recognition and understanding*, pages 273–278. IEEE, 2013.
- [13] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6645–6649. IEEE, 2013.
- [14] Qing He, Tianfeng Shang, Fuzhen Zhuang, and Zhongzhi Shi. Parallel extreme learning machine for regression based on mapreduce. *Neurocomputing*, 102:52–58, 2013.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [16] Guang-Bin Huang, Qin-Yu Zhu, Chee-Kheong Siew, et al. Extreme learning machine: a new learning scheme of feedforward neural networks. *Neural networks*, 2:985–990, 2004.
- [17] Shan Huang, Botao Wang, Junhao Qiu, Jitao Yao, Guoren Wang, and Ge Yu. Parallel ensemble of online sequential extreme learning machine based on mapreduce. *Neurocomputing*, 174:352–367, 2016.

- [18] Weikuan Jia, Dean Zhao, Yuanjie Zheng, and Sujan Hou. A novel optimized ga-elman neural network algorithm. *Neural Computing and Applications*, 31(2):449–459, 2019.
- [19] Michael I Jordan. Serial order: A parallel distributed processing approach. In *Advances in psychology*, volume 121, pages 471–495. Elsevier, 1997.
- [20] Viacheslav Khomenko, Oleg Shyshkov, Olga Radyvonenko, and Kostiantyn Bokhan. Accelerating recurrent neural network training using sequence bucketing and multi-gpu data parallelization. In *IEEE First International Conference on Data Stream Mining & Processing*, pages 100–103. IEEE, 2016.
- [21] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings of the second Workshop on the LLVM Compiler Infrastructure in HPC*, pages 1–6. ACM, 2015.
- [22] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436, 2015.
- [23] Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *European Conference on Computer Vision*, pages 816–833. Springer, 2016.
- [24] Jun Liu, Gang Wang, Ling-Yu Duan, Kamila Abdiyeva, and Alex C Kot. Skeleton-based human action recognition with global context-aware attention lstm networks. *IEEE Transactions on Image Processing*, 27(4):1586–1599, 2017.
- [25] James Martens and Ilya Sutskever. Learning recurrent neural networks with hessian-free optimization. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 1033–1040. Citeseer, 2011.
- [26] Travis Oliphant. *Guide to NumPy*. 01 2006.
- [27] Peng Ouyang, Shouyi Yin, and Shaojun Wei. A fast and power efficient architecture to parallelize lstm based rnn for cognitive intelligence applications. In *Proceedings of the 54th Annual Design Automation Conference 2017*, pages 1–6. ACM, 2017.
- [28] Yoh-Han Pao, Gwang-Hoon Park, and Dejan J Sobajic. Learning and generalization characteristics of the random vector functional-link net. *Neurocomputing*, 6(2):163–180, 1994.
- [29] Jin-Man Park and Jong-Hwan Kim. Online recurrent extreme learning machine and its application to time-series prediction. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 1983–1990. IEEE, 2017.
- [30] Yara Rizk and Mariette Awad. On extreme learning machines in sequential and time series prediction: A non-iterative and approximate training algorithm for recurrent neural networks. *Neurocomputing*, 325:1–19, 2019.
- [31] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- [32] Wouter F Schmidt, Martin A Kraaijveld, and Robert PW Duin. Feedforward neural networks with random weights. In *11th IAPR International Conference on Pattern Recognition. Vol. II. Conference B: Pattern Recognition Methodology and Systems*, pages 1–4. IEEE, 1992.
- [33] Xavier Sierra-Canto, Francisco Madera-Ramirez, and Victor Uc-Cetina. Parallel training of a back-propagation neural network using cuda. In *2010 Ninth International Conference on Machine Learning and Applications*, pages 307–312. IEEE, 2010.
- [34] Zhiyuan Tang, Ying Shi, Dong Wang, Yang Feng, and Shiyue Zhang. Memory visualization for gated recurrent neural networks in speech recognition. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2736–2740. IEEE, 2017.
- [35] Hubert AB Te Braake and Gerrit Van Straten. Random activation weight neural net (rawn) for fast non-iterative training. *Engineering Applications of Artificial Intelligence*, 8(1):71–80, 1995.
- [36] Mark Van Heeswijk, Yoan Miche, Erkki Oja, and Amaury Lendasse. Gpu-accelerated and parallelized elm ensembles for large-scale regression. *Neurocomputing*, 74(16):2430–2437, 2011.
- [37] Botao Wang, Shan Huang, Junhao Qiu, Yu Liu, and Guoren Wang. Parallel online sequential extreme learning machine based on mapreduce. *Neurocomputing*, 149:224–232, 2015.
- [38] Shang Wang, Yifan Bai, and Gennady Pekhimenko. Scaling back-propagation by parallel scan algorithm. *arXiv preprint arXiv:1907.10134*, 2019.
- [39] Xiaoyu Wang and Yong Huang. Convergence study in extended kalman filter-based training of recurrent neural networks. *IEEE Transactions on Neural Networks*, 22(4):588–600, 2011.
- [40] Paul J Werbos et al. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.
- [41] Ronald J Williams and David Zipser. Gradient-based learning algorithms for recurrent. *Backpropagation: Theory, architectures, and applications*, 433, 1995.

- [42] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google's neural machine translation system: Bridging the gap between human and machine translation. arXiv preprint arXiv:1609.08144, 2016.
- [43] Feng Zhang, Jidong Zhai, Marc Snir, Hai Jin, Hironori Kasahara, and Mateo Valero. Guest edito-

rial: Special issue on network and parallel computing for emerging architectures and applications, 2019.

- [44] Shunlu Zhang, Pavan Gunupudi, and Qi-Jun Zhang. Parallel back-propagation neural network training technique using cuda on multiple gpus. In IEEE MTT-S International Conference on Numerical Electromagnetic and Multiphysics Modeling and Optimization, pages 1–3. IEEE, 2015.



Julia El Zini is a Ph.D. student enrolled in the electrical and computer engineering department at the American University of Beirut (AUB). She has received her B.S. and M.S. in computer science from AUB, Lebanon, in 2015 and 2017, respectively. Her research interests include distributed optimization, parallel computing, rein-

forcement learning, multi-task and transfer learning, and scalable machine learning applications.



Yara Rizk obtained her Ph.D. in Electrical and Computer Engineering from the American University of Beirut (AUB) in 2018. Prior, she received her BE in Computer and Communication Engineering from AUB, Lebanon, in 2012. Her research interests span robotics, multi-agent systems, machine learning, classification, clustering, and

artificial intelligence. Rizk has attended a technical internship (2013-2014) at Intel in Hillsboro, Oregon, USA and is an active researcher multiple peer-reviewed publications.



Mariette Awad obtained her Ph.D. in Electrical Engineering from the University of Vermont (2007). Her current research focuses on HMI, efficient artificial intelligence, applied machine learning and Internet of Things. Dr. Awad has received more than 25 grants to support her research including 2 multidisciplinary multi-million

dollar grants from the Qatar National Research Fund (QNRF) and Intel. Her work culminated in a book, *Efficient Machine Learning*, in 2015 as well as more than 100 conference, book chapter, and journal publications. Prior to her academic position, she was with IBM System and Technology group in Vermont for six years where her technical leadership and innovative spirit has earned her management recognition twice, two business awards, and 10 patents.