
Parallelizing Legendre Memory Unit Training

Narsimha Chilkuri¹ Chris Eliasmith^{1,2}

Abstract

Recently, a new recurrent neural network (RNN) named the Legendre Memory Unit (LMU) was proposed and shown to achieve state-of-the-art performance on several benchmark datasets. Here we leverage the linear time-invariant (LTI) memory component of the LMU to construct a simplified variant that can be parallelized during training (and yet executed as an RNN during inference), resulting in up to 200 times faster training. We note that our efficient parallelizing scheme is general and is applicable to any deep network whose recurrent components are linear dynamical systems. We demonstrate the improved accuracy of our new architecture compared to the original LMU and a variety of published LSTM and transformer networks across seven benchmarks. For instance, our LMU sets a new state-of-the-art result on psMNIST, and uses half the parameters while outperforming DistilBERT and LSTM models on IMDB sentiment analysis.

1. Introduction

LSTMs (Hochreiter & Schmidhuber, 1997), the most popular class of recurrent neural networks (RNNs), process information in an inherently sequential manner. This prevents parallelization within training examples, and thus, they cannot fully leverage today’s commodity GPU hardware. This sequential nature of RNNs is one of the critical reasons for the shift towards purely self-attention based architectures such as the transformer and its derivatives (Vaswani et al., 2017; Devlin et al., 2018; Brown et al., 2020), especially in the domain of Natural Language Processing (NLP). Parallelization makes training such networks far more efficient on GPUs, and thus allows them to be applied to enormous datasets – for example, the Colossal Clean Crawled Corpus (C4) which comprises 750GB of English text (Raffel et al., 2019). Hence, such models, via unsupervised pre-training,

make it possible for us to exploit resources such as the internet,¹ which produces 20TB of text data each month. A feat such as this, from the training perspective, would be unimaginable using LSTM or other RNN based models.

In this paper, we exploit the idea that linear recurrence relations can be *solved*. More specifically, a linear time-invariant (LTI) system’s state update equation can be written in a non-sequential fashion (Åström & Murray, 2010). This allows for the computation of the hidden state of an LTI system to be done in parallel, thus overcoming the fundamental constraint that traditional RNNs suffer from. Despite this desirable property, to our knowledge, such systems have remained underexplored. This perhaps has to do with the general belief that nonlinear dynamics are critical for solving complicated problems, especially AI level tasks. Here, we challenge that notion by focussing our attention on a specific LTI system, the memory component of the recently proposed Legendre Memory Unit (Voelker et al., 2019), and demonstrating superior performance of our model on a variety of tasks, including predicting a nonlinear, chaotic dynamical system.

The Legendre Memory Unit (LMU) is an RNN cell that is constructed by coupling an LTI system to a nonlinear one. The LTI system, known as the Delay Network (Voelker & Eliasmith, 2018), is the core component of the LMU that projects a sliding window of length θ of the input signal onto the Legendre basis – hence the name Legendre Memory Unit. This construction, which employs the Delay Network (DN) to act as the memory and a nonlinear recurrent layer to compute arbitrary functions across time, has been shown to outperform LSTMs and other RNNs on various tasks. Of special interest to us is the ability of the LMU to handle temporal dependencies across 100,000 time-steps, which is orders of magnitude better than the LSTM.

In the following pages, we start by simplifying the LMU such that recurrence exists only in the linear system. After showing that the training of this simplified variant can be parallelized, we turn to the question of its effectiveness. We do so by first considering the two synthetic tasks that the original LMU was validated on: psMNIST and Mackey-Glass. We show that our variant exceeds the original LMU’s performance on both of these tasks, while also establishing

¹Center for Theoretical Neuroscience, University of Waterloo
²Applied Brain Research. Correspondence to: Narsimha Chilkuri <narasimha.chilkuri@gmail.com>.

¹<https://commoncrawl.org>

a new state-of-the-art result for RNNs on psMNIST. We also test our models against LSTM models of comparable sizes on several NLP tasks. We look at sentiment classification (IMDB), semantic similarity (QQP), and natural language inference (SNLI), and show that our models achieve better performance while using significantly fewer parameters (up to 650x). We then briefly investigate the transfer learning abilities of our model by training a language model on the (unlabelled) Amazon Reviews dataset and using that to improve the IMDB score. Here, we show that it outperforms DistilBert, a transformer based model, while using 50% fewer parameters. We conclude our NLP experiments by demonstrating superior performance on language modelling (text8) and machine translation (IWSLT’15 En-Vi). Finally, we show that these architectural modifications result in training times that are up to 200 times faster, relative to the original LMU.

2. Related Work

Several recent works have explored the applicability and architecture of the LMU. For example, Blouw et al. (2020) used an LMU based architecture to achieve SotA results on the problem of efficient key-word spotting, and Gu et al. (2020) recently proposed a generalization of the original architecture, which they use to beat the SotA result set by the LMU on psMNIST.

Our model, however, is more directly inspired by the success of self-attention. Self-attention based architectures have come to replace RNN based approaches for problems such as language modelling, machine translation, and a slew of other NLP tasks (Radford et al., 2018; Raffel et al., 2019). Three properties that make self-attention desirable over RNNs are: (1) it is better at handling the challenging problem of long-range dependencies; (2) it is purely feed-forward; and (3) when the sequence length is smaller than the dimension of representation, which is not uncommon in NLP applications, self-attention is computationally cheaper than an RNN.

Of these three desirable properties, our model inherits the first one from the original LMU, and satisfies properties (2) and (3) by construction. Additionally, the capability to run our model in a recurrent manner during inference can be advantageous in situations where there are memory constraints or where low-latency is critical.

3. Background and LMU Variants

In this section we start by introducing the delay problem and show how a delay is optimally realized by the DN, an LTI system. We then introduce the LMU which employs a single-input DN coupled to a nonlinear dynamical system to process sequential data. Finally, we introduce our model,

which is obtained by simplifying the original architecture. We show how this simplification allows for training to be parallelized (often reducing the computation complexity), while also retaining the ability to handle long range dependencies of the original formulation.

3.1. Delay Network

Ideal Delay A system is said to be an ideal delay system if it takes in an input, $u(t)$, and outputs a function, $y(t)$, which is the delayed version of the input. Mathematically, this can be described in the following manner:

$$y(t) = \mathcal{D}[u(t)] = \begin{cases} 0 & t < \theta \\ u(t - \theta) & t \geq \theta \end{cases}, \quad (1)$$

where \mathcal{D} is the ideal delay operator and $\theta \in \mathbb{R}$ is the length of the delay. There are two things of note here: (1) the ideal delay system is linear, i.e., for any two functions, $f(t)$ and $g(t)$, and any $a, b \in \mathbb{R}$, it respects the following equation:

$$\mathcal{D}[af(t) + bg(t)] = a\mathcal{D}[f(t)] + b\mathcal{D}[g(t)]; \quad (2)$$

and (2) although this problem looks deceptively simple, in fact it takes a system with infinite memory to take in a continuous stream of input (with unspecified frequency content), store it for θ seconds and then reproduce the entire input without error.

These two considerations tell us that the optimal system that implements delay must be linear and that even the most optimal physical implementation can only approximately realize the delay.

Approximating Delay If we are interested in constructing a dynamical system that implements delay, thanks to the linearity constraint, we can narrow our search space from the general system of ODEs of the form:

$$\dot{\mathbf{m}} = \mathbf{f}(\mathbf{m}, u), \quad (3)$$

$$y = \mathbf{g}(\mathbf{m}, u), \quad (4)$$

to just finding the four matrices ($\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$) that define an LTI system:

$$\dot{\mathbf{m}} = \mathbf{A}\mathbf{m} + \mathbf{B}u, \quad (5)$$

$$y = \mathbf{C}\mathbf{m} + \mathbf{D}u. \quad (6)$$

Following the derivation of the Delay Network in Voelker & Eliasmith (2018), we start by considering the transfer function of the delay system, which for a SISO system is defined as:

$$G(s) = \frac{y(s)}{u(s)} = e^{-\theta s}, \quad (7)$$

where $y(s)$ and $u(s)$ are found by taking the Laplace transform of the input and output functions in time. As expected,

this defines an infinite dimensional transfer function, capturing the intuitive difficulty of constructing a continuous delay.

The transfer function can be converted to a finite, causal state space realization *if and only if* it can be written as a proper² ratio of finite dimensional polynomials in s (Brogan, 1991). $G(s)$, however, is irrational, and it cannot be written as a proper, finite dimensional ratio. Therefore, making an approximation is necessary.

We can achieve an optimal convergence rate (in the least-square error sense) for rational approximants by means of Padé approximants (Partington, 2004). Choosing the order of the numerator to be one less than the order of the denominator and accounting for numerical issues in the state-space realization (see Voelker & Eliasmith (2018) for details), gives the following canonical realization:

$$A_{i,j} = \frac{(2i+1)}{\theta} \begin{cases} -1 & i < j \\ (-1)^{i-j+1} & i \geq j \end{cases}, \quad (8)$$

$$B_i = \frac{(2i+1)(-1)^i}{\theta}, \quad (9)$$

$$C_i = (-1)^i \sum_{l=0}^i \binom{i}{l} \binom{i+l}{j} (-1)^l, \quad (10)$$

$$D = 0, \quad i, j \in [0, d-1], \quad (11)$$

where we refer to d as the order of the system. The LTI system $(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D})$ is what is known as a Delay Network (DN). The order of the system, d , and the delay length, θ are the main hyperparameters to choose when using a DN. Higher order systems require more resources, but provide a more accurate emulation of the ideal delay. Because we have used Padé approximants, each order is optimal for that dimension of state vector \mathbf{m} .

Legendre Polynomials Suppose we construct a system using the $(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D})$ matrices defined above, and provide it with an input signal, $u(t)$. Given the state \mathbf{m}_t , we can use \mathbf{C} to decode $u(t-\theta)$ to a degree of accuracy determined by the order of our system, i.e.,

$$u(t-\theta) \approx \mathbf{C}^T \mathbf{m}_t. \quad (12)$$

Intuitively, given \mathbf{m}_t , it seems possible to decode not only $u(t-\theta)$ but also $u(t-\theta')$ $\forall 0 \leq \theta' \leq \theta$. This can be done using a slightly modified \mathbf{C} for a given θ' :

$$u(t-\theta') \approx \mathbf{C}(\theta')^T \mathbf{m}_t, \quad (13)$$

²A ratio $\frac{a(s)}{b(s)}$ is said be proper if the order of the numerator does not exceed the order of the denominator.

where

$$C_i(\theta') = (-1)^i \sum_{l=0}^i \binom{i}{l} \binom{i+l}{j} \left(-\frac{\theta'}{\theta}\right)^l, \quad 0 \leq \theta' \leq \theta, \quad (14)$$

and $\mathbf{C}(\theta' = \theta)$ corresponds to the \mathbf{C} defined in equation (10). Interestingly, the functions in (14) turn out to be the shifted Legendre polynomials.

3.2. Legendre Memory Unit

The LMU is obtained by coupling a single-input delay network to a nonlinear dynamical system. The DN orthogonalizes the input signal across a sliding window of length θ , whereas the nonlinear system relies on this memory to compute arbitrary functions across time. The state update equations that define the LMU are given below:

$$u_t = \mathbf{e}_x^T \mathbf{x}_t + \mathbf{e}_h^T \mathbf{h}_{t-1} + \mathbf{e}_m^T \mathbf{m}_{t-1}, \quad (15)$$

$$\mathbf{m}_t = \bar{\mathbf{A}} \mathbf{m}_{t-1} + \bar{\mathbf{B}} u_t, \quad (16)$$

$$\mathbf{h}_t = f(\mathbf{W}_x \mathbf{x}_t + \mathbf{W}_h \mathbf{h}_{t-1} + \mathbf{W}_m \mathbf{m}_t), \quad (17)$$

where $\bar{\mathbf{A}}$ and $\bar{\mathbf{B}}$ are the discretized versions³ of \mathbf{A} and \mathbf{B} . These matrices are usually frozen during training, although they need not be. The input to the LTI system, $u(t)$, is computed by projecting the input to the RNN cell, $\mathbf{x}_t \in \mathbb{R}^{d_x}$, the hidden state, $\mathbf{h}_t \in \mathbb{R}^{d_h}$, and the memory state, $\mathbf{m}_t \in \mathbb{R}^d$, onto their respective encoding weights $(\mathbf{e}_x, \mathbf{e}_h, \text{ and } \mathbf{e}_m)$. The memory state, hidden state and the input to the cell are then combined using the weight matrices $(\mathbf{W}_x, \mathbf{W}_h, \text{ and } \mathbf{W}_m)$ and passed through the nonlinearity, f . The encoding vectors and the kernel matrices are learned during training.

3.3. Our Model

In this paper, we propose a modified LMU architecture by making two changes to the equations defined above. First, we remove certain connections in (15) and (17), namely $(\mathbf{e}_h^T, \mathbf{e}_m^T)$ and (\mathbf{W}_h) ; this is done to aid parallelization (see Table (7) for a simple ablation study). Second, instead of projecting the input to the RNN cell, \mathbf{x}_t down to a scalar as in (15), we implement a general affine transformation followed by an element-wise nonlinearity; the original architecture is better suited for dealing with 1D or low-dimensional inputs, and this is a straightforward generalization of the encoding equation for higher dimensional inputs. Additionally, adding a bias term to equation (17), we end up with the following:

$$\mathbf{u}_t = f_1(\mathbf{U}_x \mathbf{x}_t + \mathbf{b}_u), \quad (18)$$

$$\mathbf{m}_t = \bar{\mathbf{A}} \mathbf{m}_{t-1} + \bar{\mathbf{B}} \mathbf{u}_t, \quad (19)$$

$$\mathbf{o}_t = f_2(\mathbf{W}_m \mathbf{m}_t + \mathbf{W}_x \mathbf{x}_t + \mathbf{b}_o). \quad (20)$$

³Using zero-order hold and $dt = 1$, exact discretization gives $\bar{\mathbf{A}} = e^{\mathbf{A}}$ and $\bar{\mathbf{B}} = \mathbf{A}^{-1}(e^{\mathbf{A}} - \mathbf{I})\mathbf{B}$.

Note that equations (18) and (20) are equivalent to having *time-distributed* dense layers before and after equation (19). In general, we expect our model to be modified and used in combination with other feed-forward layers, including self-attention. For example, we found a gated architecture (Srivastava et al., 2015) where equation (18) is modified to

$$\mathbf{u}_t = f_1(\mathbf{W}_u \mathbf{x}_t + \mathbf{b}_u) \cdot \mathbf{g} + \mathbf{x}_t \cdot (1 - \mathbf{g}),$$

to work well for the addition problem (results not shown). The gate is defined as $\mathbf{g} = \sigma(\mathbf{W}_g \mathbf{x}_t + \mathbf{b}_g)$, i.e., a sigmoid-activated affine transformation where the bias is initialized to -1.

Additionally, now that the input to the LTI system in this case is a vector, $\mathbf{u}_t \in \mathbb{R}^{1 \times d_u}$, we have that $\mathbf{m}_t \in \mathbb{R}^{d \cdot d_u}$, and equation (16) can be thought of as implementing the following:

$$\mathbf{m}_t = \text{reshape}(\bar{\mathbf{A}} \text{reshape}(\mathbf{m}_{t-1}, (d, d_u)) + \bar{\mathbf{B}} \mathbf{u}_t, d \cdot d_u). \quad (21)$$

Parallel Training One of the motivations for the above mentioned architectural changes is that the model now has only one recurrent connection: \mathbf{m}_t 's dependence on itself from the past in equation (19). But because this is an LTI system, standard control theory (Åström & Murray, 2010) gives us a non-iterative way of evaluating this equation as shown below⁴

$$\mathbf{m}_t = \sum_{j=1}^t \bar{\mathbf{A}}^{t-j} \bar{\mathbf{B}} \mathbf{u}_j. \quad (22)$$

Defining $\mathbf{H} = [\bar{\mathbf{A}}^0 \bar{\mathbf{B}} \quad \bar{\mathbf{A}} \bar{\mathbf{B}} \quad \dots] \in \mathbb{R}^{d \times n}$ and

$$\mathbf{U} = \begin{bmatrix} u_1 & u_2 & u_3 & \dots & u_n \\ & u_1 & u_2 & \dots & u_{n-1} \\ & & u_1 & \dots & u_{n-2} \\ & & & \ddots & \vdots \\ & & & & u_1 \end{bmatrix} \in \mathbb{R}^{n \times n}, \quad (23)$$

the above convolution equation can alternatively be written as a matrix multiplication:

$$\mathbf{m}_{1:n} = \mathbf{H} \mathbf{U}, \quad (24)$$

where n is the sequence length. Given that \mathbf{H} is the impulse response of the LTI system, in practice we compute \mathbf{H} by feeding in an impulse to the RNN version of the DN (equation (19)). Note that in our method the $\bar{\mathbf{A}}$ and $\bar{\mathbf{B}}$ matrices are frozen during training, so the impulse response need only be computed once. In case of multi-dimensional inputs, we would repeat the above computation several times,

⁴Restricting ourselves to 1D input for the purposes of illustration.

Table 1. Complexity per layer and minimum number of sequential operations of various architectures. n is the sequence length, d_x is the input dimension, d is the order, and k is the size of the kernel. First three rows are as reported in Vaswani et al. (2017).

Layer Type	Complexity	Sequential Ops
RNN	$O(n \cdot d_x^2)$	✓
Convolution	$O(k \cdot n \cdot d_x^2)$	×
Attention	$O(n^2 \cdot d_x)$	×
DN (19)	$O(n \cdot d^2 \cdot d_x)$	✓
DN (24)	$O(n^2 \cdot d \cdot d_x)$	×
DN (25)	$O(n \cdot d \cdot d_x)$	×
DN (26)	$O(n \cdot \log n \cdot d \cdot d_x)$	×

once for each dimension of the input. It is also evident from the structure of the \mathbf{U} matrix that although this reformulation turns the DN into a feedforward layer, it still respects causality. In other words, the state \mathbf{m}_t depends only on the inputs seen until that point of time ($u_i : i \leq t$).

Complexity With the new formulation, we immediately see that it is computationally (in terms of the number of operations) advantageous in situations where we only need the final state (`return_sequences=False` in Keras terminology). Instead of using (19) to explicitly simulate the first $n - 1$ steps in-order to compute the state at the time-step n , we can instead just compute

$$\mathbf{m}_n = \mathbf{H} \mathbf{U}_{:n}, \quad (25)$$

thus reducing the complexity of computing the final state, \mathbf{m}_n , from $O(n \cdot d^2 \cdot d_x)$ to $O(n \cdot d \cdot d_x)$, where n is the sequence length, d is the order, and d_x is the dimension of the input. We show in Section (4.6) that using this implementation results in up to 200x speedup.

The more general computation (24), although parallelizable, results in a complexity of $O(n^2 \cdot d \cdot d_x)$. This can be made more efficient by employing the convolution theorem which gives us an equivalent way of evaluating the convolution in the Fourier space as⁵

$$\mathbf{m}_{1:n} = \mathcal{F}^{-1}\{\mathcal{F}\{\mathbf{H}\} \cdot \mathcal{F}\{\mathbf{U}_{:n}\}\}. \quad (26)$$

Thanks to the fast Fourier transform, the above operation has a complexity of $O(n \cdot \log n \cdot d \cdot d_x)$. These, other algorithms, and their complexities are reported in Table (1).

It was argued in Vaswani et al. (2017) that a self-attention layer is cheaper than an RNN when the representation dimension of the input, d_x , is much greater than the length

⁵Assuming a padded Fourier transform across the appropriate axis and automatic broadcasting when computing element-wise multiplication.

of the sequence, n , which is seen in NLP applications. For example, standard word or sub-word based machine translation systems work with sequence lengths of about 100 and representation dimension ranging from 300 to 512. The same argument holds for the DN layer too, since we have found the inequality $\log_2 n \cdot d \ll d_x$ to hold in all our word-based NLP experiments – this excludes text8, which works at the level of characters.

Recurrent Inference Machine learning algorithms are usually optimized for training rather than deployment (Crankshaw, 2019), and because of that models need to be modified, sometimes non-trivially, to be more suitable for inference. For example, edge applications such as keyword spotting (KWS) often have strong memory and power constraints. Transformers have shown faster training times and better performance than RNNs on speech related tasks, but since they employ (global) self-attention which needs to store the entire input in memory to output a higher level representation and has a computational dependency that is quadratic with respect to the sequence length, they are not natural fits for edge applications (Zhang et al., 2020; Miao et al., 2020). While our model can be trained in parallel, because of the equivalence of equations (19) and (26), it can also be run in an iterative manner during inference, and hence can process data in an online and memory-efficient fashion during inference.

4. Experiments

In the following experiments we compare our model against the LMU, LSTMs and transformers. With these experiments, we focus on benchmarking rather than establishing new state-of-the-art results. Hence, we stick to simple architectures, constrain ourselves to train all the models, with the exception of text8, using the Adam optimizer (Kingma & Ba, 2014) with all the default parameter settings. For text8, we found it helpful to reduce the learning rate by a factor of 10 halfway into training. Although unusual, we use the default optimization settings even when transfer learning.

With models we compare against, we use results found in published work, even when they make use of more sophisticated architectures, learning schedules or regularization schemes.

We do not consider the capacity task from the original LMU paper here as they employ just the delay layer without non-linear dynamics in order to deal with this problem, which makes their architecture essentially the same as ours.

4.1. psMNIST

The permuted sequential MNIST (psMNIST) dataset was introduced by Le et al. (2015) to test the ability of RNNs

Table 2. psMNIST results. The first four rows are from Voelker et al. (2019), and the fifth row is from Gu et al. (2020).

Model	Accuracy
LSTM	89.86
SRU	92.49
NRU	95.38
LMU	97.15
HiPPO-LegS	98.3
Our Model	98.49

Table 3. Mackey-Glass results.

Model	NRMSE
LSTM	0.059
LMU	0.049
Hybrid	0.045
Our Model	0.044

to model complex long term dependencies. psMNIST, as the name suggests, is constructed by permuting and then flattening the (28×28) MNIST images. The permutation is chosen randomly and is fixed for the duration of the task, and in order to stay consistent, we use the same permutation as Chandar et al. (2019) and Voelker et al. (2019). The resulting dataset is of the form (samples, 784, 1), which is fed into a recurrent network sequentially, one pixel at a time. We use the standard 50k/10k/10k split.

Architecture As was pointed out in Voelker et al. (2019), in order to facilitate fair comparison, RNN models being tested on this task should not have access to more than $28^2 = 784$ internal variables. Keeping that in mind, and in order to make direct comparisons to the original LMU model, we consider a model with $d = 468$ dimensions for memory. We set the dimension of the output state to 346 and use $\theta = 784$. Our model uses 165k parameters, the same as all the models reported in Table (2), except for the original LMU model, which uses 102k parameters, and the HiPPO-LegS model, which is reported to use 512 hidden dimensions (number of parameters is unknown).

Results & Discussion Test scores of various models on this dataset are reported in Table (2). Our model not only surpasses the LSTM model, but also beats the current state-of-the-art result of 98.3% set by HiPPO-LegS (Gu et al., 2020) recently. Thus, our model sets a new state-of-the-art result for RNNs of 98.49% on psMNIST. It is interesting that our model, despite being simpler than the original LMU, outperforms it on this dataset. This suggests that the main advantage of the LMU over past models is the quality of its temporal memory, implemented by the DN.

4.2. Mackey-Glass

Mackey-Glass equations are a set of two nonlinear differential equations that were originally developed to model the quantity of mature cells in the blood. The second of these equations is interesting because it can result in chaotic

attractors and is used to construct the dataset at hand. This is a time-series prediction task where we need to predict 15 time-steps into the future.

Architecture Voelker & Eliasmith (2018) compare three architectures on this dataset. The first one uses 4 layers of LSTMs ($h = 28$), the second one replaces LSTMs with LMUs ($d = 4, \theta = 4$), and the final one replaces the first and third layers in the first model with LMUs. We use a relatively simple architecture where we combine our model (single layer) with an additional dense layer. We use $d = 40$, $\theta = 50$, and 1 and 140 units in the input and output layers, with the additional dense layer containing 80 units. We did not try other variations. All the models contain about 18k parameters and are run for 500 epochs.

Results & Discussion NRMSE scores on the test set are presented in Table (3). We see that our model outperforms the other three models in accuracy and training time. In the original paper Voelker & Eliasmith (2018) hypothesize that the superior performance of the hybrid model is due to the presence of gates, but given that our model lacks gating mechanisms, we think that it might have to do with the LMU model being misparametrized.

4.3. Sentiment, Semantics and Inference

In this section, we explore the supervised and semi-supervised capabilities of our model. More specifically, we first look at the tasks of sentiment analysis (IMDB), semantic similarity (QQP) and natural language inference (SNLI), and then improve upon the IMDB score using a language model pre-trained on the (unlabelled) Amazon Reviews dataset (Ni et al., 2019).

SUPERVISED

IMDB The IMDB dataset (Maas et al., 2011) is a standard sentiment classification task containing a collection of 50k highly polar reviews, with the training and testing sets containing 25k reviews each. We use the standard pre-processed dataset available from the Keras website,⁶ consider a vocabulary of 20k words, and set the maximum sequence length to 500 words.

QQP For Quora Question Pairs (QQP), given a pair of sentences, the task is to identify whether the two sentences are semantically similar. In this case, we experiment on two train/dev/test splits: 390k/8k/8k like in Shen et al. (2018), and 280k/80k/40k like in Sharma et al. (2019). We use a vocabulary of 20k words and truncate the sentences to be less than 25 words.

⁶<https://keras.io/api/datasets/imdb/>.

Table 4. IMDB, QQP and SNLI results. IMDB result is from (Gu et al., 2020), QQP results are from (Shen et al., 2018) and (Sharma et al., 2019) respectively, and SNLI is from (Bowman et al., 2015).

	LSTM		Our Model	
	Param.	Acc.	Param.	Acc.
IMDB	50k	87.29	301	89.10
QQP	-/800k	82.58/81.4	1201	86.95/85.36
SNLI	220k	77.6	3.6k	78.85

SNLI The Stanford Natural Language Inference⁷ was released to serve as a benchmark for evaluating machine learning systems on the task of natural language inference. Given a premise, the task is to determine whether a hypothesis is true (entailment), false (contradiction), or independent (neutral). We use the standard 550k/10k/10k split, consider a vocabulary of 20k words, and set the maximum sequence length to 25 words.

Architecture In our experiments, confusingly, we found the use of the DN, without any nonlinearities, to work well. Therefore, we construct parameter efficient models that employ just the DN layer, with $d = 1$ and $\theta = \text{maxlen}$. We use 300D Glove embeddings (840B Common Crawl; Pennington et al. (2014)) for all our models. For the IMDB task, which is a single sentence task, we encode the sentence and pass the encoded vector to the final classification layer. For two-sentence tasks, QQP and SNLI, we encode the two sentences to produce two vectors, and then pass the vector obtained by concatenating the two vectors, their absolute difference, and their element-wise product to the final classification layer. We compare our models against the LSTM models described in (Gu et al., 2020) for IMDB, both (Shen et al., 2018) and (Sharma et al., 2019) for QQP, and (Bowman et al., 2015) for SNLI. They all use at least an order of magnitude more parameters than our models.

Results & Discussion We present the results from the three experiments in Table (4). As we can see, our simple models based on the DN alone do indeed outperform the LSTM models. It is also noteworthy that our models use significantly fewer parameters than the LSTM models: 160x on IMDB, 650x on QQP and 60x on SNLI.

SEMI-SUPERVISED

Amazon Reviews NLP over the past couple of years has been dominated by transfer learning approaches, where language models pre-trained on large corpora are used to initialize general purpose fine-tuning architectures to help with

⁷Dataset and published results are available at <https://nlp.stanford.edu/projects/snli/>.

various downstream tasks. We therefore consider using a pre-trained language model on the Amazon Reviews dataset to improve the previously achieved score on IMDB. Here, we make direct comparisons to the LSTM model described in Radford et al. (2017). While they use the entire dataset (82 million reviews) for pre-training and train their model for a month on 4 GPUs, due to resource constraints, we use less than 5% (3.7 million reviews) and train our model for about 12 hours on a single GPU. We use a vocabulary of 30k words.

Architecture We have found the repeating block architecture, where each block is composed of our model, a few highway layers, and a dense layer, when used with skip connections to work well (see Figure in supplementary materials). Since the inputs, d_x , are high dimensional, using a large θ , which would have to be accompanied by a large order d to maintain resolution, would result in vectors that are $d_x \cdot d$ dimensional, which is not desirable. Therefore, we instead work with smaller θ (and hence smaller d) and use several repeating blocks to take long-term dependencies into account. This is similar to how convolutional networks are used: small kernel sizes but many convolutional layers. If θ_i is the setting we use with the DN in the i^{th} block, then the effective delay, $\theta_e = \sum_i \theta_i$. In this specific case, we have $\theta_i = 6 \forall i$, and $\theta_e = 30$. Thus, the model has access to the past 30 tokens when making a prediction. In terms of the fine-tuning performance, like Peters et al. (2018), we found that using deep representations, i.e., a linear combination of representations from all the blocks, to be better than using just the output of the top block. For fine-tuning, we compute a weighted sum of the outputs from the language model, and use that to classify a given review as expressing a positive or negative sentiment. Even during fine-tuning, we use the Adam optimizer with all the default settings. We did not experiment with smaller learning rates or other learning schedules.

Results & Discussion Results for this experiment are presented in Table (5). Despite using a much smaller pre-training dataset, training for just 12 hours, and using less than 50% of the parameters, our model outperforms the LSTM model described in Radford et al. (2017). We also include a self-attention based model, DistilBert (Sanh et al., 2019), for comparison; it must be noted however that DistilBert was trained on a more general yet much larger dataset (concatenation of English Wikipedia and Toronto Book Corpus).

4.4. Language Modelling

text 8 We evaluate our model on character level language modelling using the text8 dataset⁸. It contains 100MB of

⁸<http://mattmahoney.net/dc/textdata>.

Table 5. IMDB results with pre-training. First row is from Radford et al. (2017), and the second row is from Sanh et al. (2019).

Model	# parameters (Millions)	Accuracy
LSTM	75	92.88
DistilBERT	66	92.82
Our Model	34	93.20

clean text from Wikipedia and has an alphabet size of 27. As is standard, we use the first 90MB as the training set, the next 5MB as the validation set and the final 5MB as the test set. Following Mikolov et al. (2012) and Zhang et al. (2016), we set the sequence length to 180.

Architecture This architecture is similar to the language model used in the above section, except for the use of deep representations; we simply work with the output from the top block.⁹ For text8, minor changes were made to adapt the model to deal with longer sequences, which is a consequence of modelling at the level of characters, and to parameter match it to the LSTM model in Zhang et al. (2016), which uses around 3.2 million weights. We employ three blocks in this case and use all three DNs with the setting $\theta = 15$.

Results & Discussion We report the scores in bits per character in Table (6). We see that our model performs better than the LSTM model of similar size. A couple of observations regarding model optimization: 1) we noticed that the training loss plateaus around the twelfth epoch, so we found it helpful it to decrease the learning rate by a factor of 10 around then (this is the only dataset for which we alter the optimization parameters); 2) despite that, the training slows down after a few more epochs, and we stop the optimization after 20 epochs; we believe that a more carefully designed learning schedule might be able to help with faster convergence. We also observed that using a self-attention layer after the final block helps with generalization; this perhaps points to the fact that the model needs a context that is longer than 45 tokens in order to make predictions.

4.5. Translation

IWSLT’15 En-Vi IWSLT’15 English to Vietnamese is a medium-resource language translation task containing 133k sentence pairs. Following Luong & Manning (2015), we do not perform any pre-processing other than replacing words that occur less frequently than 5 by the `<unk>` token, which leaves us with vocabulary sizes of 17k and 7.7k for English and Vietnamese, respectively. We use the TED tst2012 as the validation set and TED tst2013 as the test set. We use

⁹We did not test the use of deep representations in this context.

Table 6. Language modelling and translation results. text8 score is from Zhang et al. (2016), and IWSLT score is from Luong & Manning (2015). ^a(case sensitive), ^b(case insensitive).

Model	text8	IWSLT'15
LSTM	1.65	23.3
Our Model	1.61	25.5^a/26.2^b

300D representation for source and target embeddings.

Architecture For translation, we use a standard encoder-decoder architecture inspired by the Amazon reviews language model, and we also employ an attention layer to help with translation. Our model’s about the same size as the the 24 million LSTM model described in Luong & Manning (2015). Due to time constraints, the architecture and hyper-parameters for this problem were relatively underexplored.

Results & Discussion Case sensitive BLEU scores are reported in Table (6). Our model brings in an improvement of 2.2 BLEU over the LSTM model. When tested on lower-cased text (without altering the training procedure), we obtained a higher score of 26.2 BLEU. One major limiting factor of this analysis is that we use a small, fixed vocabulary (17k and 7.7k words), with no way of dealing with out-of-vocabulary words. For future work, we intend to experiment with an open-vocabulary encoding scheme such as byte pair encoding (Sennrich et al., 2015).

4.6. Training Time

Here we explore the effect of parallelization on training time. We refer to architectures that implement the DN using equation (19) as the *LTI version* and the ones that use either (25) or (26), depending on whether `return_sequences` is false or true, as the *parallel version*.

Results are presented in Figure (1). On the left, we compare the speedup we obtain by going from the original LMU to our model, in both LTI and parallel forms, on psMNIST and Mackey-Glass. We first notice that switching from the LMU to the LTI version results in non-trivial gains. This is solely due to the reduction in the number of recurrent connections. On top of this, switching to the parallel version, owing to long sequences (784 and 5000) and, in the case of psMNIST, a reduction in the computational burden, results in substantial improvements in training times of 220x and 64x respectively.¹⁰

On the right of Figure (1), we look at how varying the

¹⁰We note that for Mackey-Glass we use a (parameter matched) one layer LMU model instead of the 4 layer model used in Voelker et al. (2019); the difference with respect to the original architecture is more drastic, with our model (parallel) being about 200x faster.

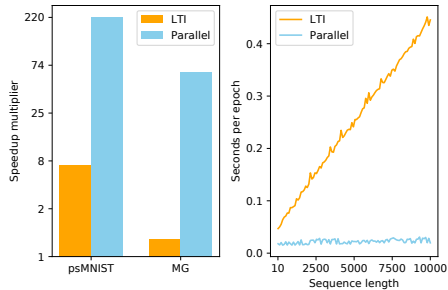


Figure 1. (left) The plot shows the speedup we obtain from switching from the LMU to the LTI (in orange) and parallel implementation (in blue) of our model. (right) This show the effect of increasing the sequence length on the LTI and parallel versions of our model. All results were measured using a single GTX 1080.

sequence length of the input affects the time it takes to complete one epoch. We see that the LTI version exhibits a linear growth whereas the parallel one is nearly constant over this scale.

5. Conclusion

With the intention of alleviating the long standing problem of speeding up RNN training, we introduce a variant of the LMU architecture that can process data in a feedforward or sequential fashion. When implemented as a feedforward layer, it can utilize parallel processing hardware such as GPUs and thus is suitable for large scale applications, and when run as an RNN, it is useful in applications where the amount of available memory is a limitation. We demonstrate the effectiveness of this architecture by experimenting on a range of tasks, of varying scale and difficulty. We also briefly consider the question of computational complexity of our model, and argue for its suitability to large scale applications in the domain of NLP, a direction we will pursue in the future.

We note that our method of parallelization applies to all deep architectures with linear recurrent dependencies, and although we focus on a specific LTI system throughout, we hope that our analysis highlights the utility of linear systems for the purposes of machine learning. In sum, we believe that linear systems offer a scalable solution to many time series problems without sacrificing accuracy.

Acknowledgements

We thank Aaron Voelker, Vincenzo Heska, Alison Shi, Andreas Stockel, and Terry Stewart for discussions that helped improve this paper. This work was supported by CFI and OIT infrastructure funding as well as the Canada Research Chairs program, NSERC Discovery grant 261453, AFOSR grant FA9550-17-1-0026 and OGS graduate funding.

References

- Åström, K. J. and Murray, R. M. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.
- Blouw, P., Malik, G., Morcos, B., Voelker, A. R., and Elias-Smith, C. Hardware aware training for efficient keyword spotting on general purpose and specialized hardware. *arXiv preprint arXiv:2009.04465*, 2020.
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*, 2015.
- Brogan, W. L. *Modern control theory*. Pearson education india, 1991.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Chandar, S., Sankar, C., Vorontsov, E., Kahou, S. E., and Bengio, Y. Towards non-saturating recurrent units for modelling long-term dependencies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3280–3287, 2019.
- Crankshaw, D. *The Design and Implementation of Low-Latency Prediction Serving Systems*. PhD thesis, UC Berkeley, 2019.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Gu, A., Dao, T., Ermon, S., Rudra, A., and Re, C. Hippo: Recurrent memory with optimal polynomial projections. *arXiv preprint arXiv:2008.07669*, 2020.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Le, Q. V., Jaitly, N., and Hinton, G. E. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015.
- Luong, M.-T. and Manning, C. D. Stanford neural machine translation systems for spoken language domains. In *Proceedings of the International Workshop on Spoken Language Translation*, pp. 76–79, 2015.
- Maas, A., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pp. 142–150, 2011.
- Miao, H., Cheng, G., Gao, C., Zhang, P., and Yan, Y. Transformer-based online ctc/attention end-to-end speech recognition architecture. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6084–6088. IEEE, 2020.
- Mikolov, T., Sutskever, I., Deoras, A., Le, H.-S., Kombrink, S., and Cernocky, J. Subword language modeling with neural networks. *preprint (http://www.fit.vutbr.cz/imikolov/rnnlm/char.pdf)*, 8:67, 2012.
- Ni, J., Li, J., and McAuley, J. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 188–197, 2019.
- Partington, J. R. Some frequency-domain approaches to the model reduction of delay systems. *Annual Reviews in Control*, 28(1):65–73, 2004.
- Pennington, J., Socher, R., and Manning, C. D. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. Deep contextualized word representations. In *Proc. of NAACL*, 2018.
- Radford, A., Jozefowicz, R., and Sutskever, I. Learning to generate reviews and discovering sentiment. *arXiv preprint arXiv:1704.01444*, 2017.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. 2018.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.
- Sanh, V., Debut, L., Chaumond, J., and Wolf, T. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015.

- Sharma, L., Graesser, L., Nangia, N., and Evcı, U. Natural language understanding with the quora question pairs dataset. *arXiv preprint arXiv:1907.01041*, 2019.
- Shen, D., Wang, G., Wang, W., Min, M. R., Su, Q., Zhang, Y., Li, C., Henao, R., and Carin, L. Baseline needs more love: On simple word-embedding-based models and associated pooling mechanisms. *arXiv preprint arXiv:1805.09843*, 2018.
- Srivastava, R. K., Greff, K., and Schmidhuber, J. Training very deep networks. In *Advances in neural information processing systems*, pp. 2377–2385, 2015.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Voelker, A., Kajić, I., and Eliasmith, C. Legendre memory units: Continuous-time representation in recurrent neural networks. In *Advances in Neural Information Processing Systems*, pp. 15544–15553, 2019.
- Voelker, A. R. and Eliasmith, C. Improving spiking dynamical networks: Accurate delays, higher-order synapses, and time cells. *Neural computation*, 30(3):569–609, 2018.
- Zhang, Q., Lu, H., Sak, H., Tripathi, A., McDermott, E., Koo, S., and Kumar, S. Transformer transducer: A streamable speech recognition model with transformer encoders and rnn-t loss. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7829–7833. IEEE, 2020.
- Zhang, S., Wu, Y., Che, T., Lin, Z., Memisevic, R., Salakhutdinov, R. R., and Bengio, Y. Architectural complexity measures of recurrent neural networks. In *Advances in neural information processing systems*, pp. 1822–1830, 2016.