
Efficient Orthogonal Parametrisation of Recurrent Neural Networks Using Householder Reflections

Zakaria Mhammedi^{1,2} Andrew Hellicar² Ashfaqur Rahman² James Bailey¹

Abstract

The problem of learning long-term dependencies in sequences using Recurrent Neural Networks (RNNs) is still a major challenge. Recent methods have been suggested to solve this problem by constraining the transition matrix to be unitary during training which ensures that its norm is equal to one and prevents exploding gradients. These methods either have limited expressiveness or scale poorly with the size of the network when compared with the simple RNN case, especially when using stochastic gradient descent with a small mini-batch size. Our contributions are as follows; we first show that constraining the transition matrix to be unitary is a special case of an orthogonal constraint. Then we present a new parametrisation of the transition matrix which allows efficient training of an RNN while ensuring that the matrix is always orthogonal. Our results show that the orthogonal constraint on the transition matrix applied through our parametrisation gives similar benefits to the unitary constraint, without the time complexity limitations.

1. Introduction

Recurrent Neural Networks (RNNs) have been successfully used in many applications involving time series. This is because RNNs are well suited for sequential data as they process inputs one element at a time and store relevant information in their hidden state. In practice, however, training simple RNNs (sRNN) can be challenging due to the problem of exploding and vanishing gradients (Hochreiter et al., 2001). It has been shown that exploding gradients can occur when the transition matrix of an RNN has a spectral norm larger than one (Glorot & Bengio, 2010). This results

in an error surface, associated with some objective function, having very steep walls (Pascanu et al., 2013). On the other hand, when the spectral norm of the transition matrix is less than one, the information at one time step tend to vanish quickly after a few time steps. This makes it challenging to learn long-term dependencies in sequential data.

Different methods have been suggested to solve either the vanishing or exploding gradient problem. The LSTM has been specifically designed to help with the vanishing gradient (Hochreiter & Schmidhuber, 1997). This is achieved by using gate vectors which allow a linear flow of information through the hidden state. However, the LSTM does not directly address the exploding gradient problem. One approach to solving this issue is to clip the gradients (Mikolov, 2012) when their norm exceeds some threshold value. However, this adds an extra hyperparameter to the model. Furthermore, if exploding gradients can occur within some parameter search space, the associated error surface will still have steep walls. This can make training challenging even with gradient clipping.

Another way to approach this problem is to improve the shape of the error surface directly by making it smoother, which can be achieved by constraining the spectral norm of the transition matrix to be less than or equal to one. However, a value of exactly one is best for the vanishing gradient problem. A good choice of the activation function between hidden states is also crucial in this case. These ideas have been investigated in recent works. In particular, the unitary RNN (Arjovsky et al., 2016) uses a special parametrisation to constrain the transition matrix to be unitary, and hence, of norm one. This parametrisation and other similar ones (Hyland & Rättsch, 2017; Wisdom et al., 2016) have some advantages and drawbacks which we will discuss in more details in the next section.

The main contributions of this work are as follows:

- We first show that constraining the search space of the transition matrix of an RNN to the set of unitary matrices $U(n)$ is equivalent to limiting the search space to a subset of $O(2n)$ ($O(2n)$ is the set of $2n \times 2n$ orthogonal matrices) of a new RNN with twice the hidden size. This suggests that it may not be necessary to

¹The University of Melbourne, Parkville, Australia ²Data61, CSIRO, Australia. Correspondence to: Zakaria Mhammedi <zak.mhammedi@data61.csiro.au>.

work with complex matrices.

- We present a simple way to parametrise orthogonal transition matrices of RNNs using Householder matrices, and we derive the expressions of the back-propagated gradients with respect to the new parametrisation. This new parametrisation can also be used in other deep architectures.
- We develop an algorithm to compute the back-propagated gradients efficiently. Using this algorithm, we show that the worst case time complexity of one gradient step is of the same order as that of the sRNN.

2. Related Work

Throughout this work we will refer to elements of the following sRNN architecture.

$$h^{(t)} = \phi(W_h h^{(t-1)} + V x^{(t)}), \quad (1)$$

$$o^{(t)} = Y h^{(t)}, \quad (2)$$

where W , V and Y are the hidden-to-hidden, input-to-hidden, and hidden-to-output weight matrices. $h^{(t-1)}$ and $h^{(t)}$ are the hidden vectors at time steps $t-1$ and t respectively. Finally, ϕ is a non-linear activation function. We have omitted the bias terms for simplicity.

Recent research explored how the initialisation of the transition matrix W influences training and the ability to learn long-term dependencies. In particular, initialisation with the identity or an orthogonal matrix can greatly improve performance (Le et al., 2015). In addition to these initialisation methods, one study also considered removing the non-linearity between the hidden-to-hidden connections (Henaff et al., 2016), i.e. the term $W_h h^{(t-1)}$ in Equation (1) is outside the activation function ϕ . This method showed good results when compared to the LSTM on pathological problems exhibiting long-term dependencies.

After training a model for a few iterations using gradient descent, nothing guarantees that the initial structures of the transition matrix will be held. In fact, its spectral norm can deviate from one, and exploding and vanishing gradients can be a problem again. It is possible to constrain the transition matrix to be orthogonal during training using special parametrisations (Arjovsky et al., 2016; Hyland & Rätsch, 2017), which ensure that its spectral norm is always equal to one. The unitary RNN (uRNN) (Arjovsky et al., 2016) is one example where the hidden matrix $W \in \mathbb{C}^{n \times n}$ is the product of elementary matrices, consisting of reflection, diagonal, and Fourier transform matrices. When the size of hidden layer is equal to n , the transition matrix has a total of only $7n$ parameters. Another advantage of this parametrisation is computational efficiency - the matrix-vector product Wv , for some vector v , can be calculated in time complexity $\mathcal{O}(n \log n)$. However, it has been shown that this

parametrisation does not allow the transition matrix to span the full unitary group (Wisdom et al., 2016) when the size of the hidden layer is greater than 7. This may limit the expressiveness of the model.

Another interesting parametrisation (Hyland & Rätsch, 2017) has been suggested which takes advantage of the algebraic properties of the unitary group $\mathbf{U}(n)$. The idea is to use the corresponding matrix Lie algebra $\mathfrak{u}(n)$ of skew hermitian matrices. In particular, the transition matrix can be written as $W = \exp \left[\sum_{i=1}^{n^2} \lambda_i T_i \right]$, where \exp is the exponential matrix map and $\{T_i\}_{i=1}^{n^2}$ are predefined $n \times n$ matrices forming a bases of the Lie algebra $\mathfrak{u}(n)$. The learning parameters are the weights $\{\lambda_i\}$. The fact that the matrix Lie algebra $\mathfrak{u}(n)$ is closed and connected ensures that the exponential mapping from $\mathfrak{u}(n)$ to $\mathbf{U}(n)$ is surjective. Therefore, with this parametrisation the search space of the transition matrix spans the whole unitary group. This is one advantage over the original unitary parametrisation (Arjovsky et al., 2016). However, the cost of computing the matrix exponential to get W is $\mathcal{O}(n^3)$, where n is the dimension of the hidden state.

Another method (Wisdom et al., 2016) performs optimisation directly of the Stiefel manifold using the Cayley transformation. The corresponding model was called full-capacity unitary RNN. Using this approach, the transition matrix can span the full set of unitary matrices. However, this method involves a matrix inverse as well as matrix-matrix products which have time complexity $\mathcal{O}(n^3)$. This can be problematic for large neural networks when using stochastic gradient descent with a small mini-batch size.

A more recent study (Vorontsov et al., 2017) investigated the effect of soft versus hard orthogonal constraints on the performance of RNNs. The soft constraint was applied by specifying an allowable range for the maximum singular value of the transition matrix. To this end, the transition matrix was factorised as $W = USV'$, where U and V are orthogonal matrices and S is a diagonal matrix containing the singular values of W . A soft orthogonal constraint consists of specifying small allowable intervals around 1 for the diagonal elements of S . Similarly to (Wisdom et al., 2016), the matrices U and V were updated at each training iteration using the Cayley transformation, which involves a matrix inverse, to ensure that they remain orthogonal.

All the methods discussed above, except for the original unitary RNN, involve a step that requires at least a $\mathcal{O}(n^3)$ time complexity. All of them, except for one, require the use of complex matrices. Table 1 summarises the time complexities of various methods, including our approach, for one stochastic gradient step. In the next section, we show that imposing a unitary constraint on a transition matrix $W \in \mathbb{C}^{n \times n}$ is equivalent to imposing a special orthog-

onal constraint on a new RNN with twice the hidden size. Furthermore, since the norm of orthogonal matrices is always one, using the latter has the same theoretical benefits as using unitary matrices when it comes to the exploding gradient problem.

3. Complex unitary versus orthogonal

We can show that when the transition matrix $W \in \mathbb{C}^{n \times n}$ of an RNN is unitary, there exists an equivalent representation of this RNN involving an orthogonal matrix $\hat{W} \in \mathbb{R}^{2n \times 2n}$.

In fact, consider a complex unitary transition matrix $W = A + iB \in \mathbb{C}^{n \times n}$, where A and B are now real-valued matrices in $\mathbb{R}^{n \times n}$. We also define the following new variables

$$\forall t, \hat{h}^{(t)} = \begin{bmatrix} \Re(h^{(t)}) \\ \Im(h^{(t)}) \end{bmatrix}, \hat{V} = \begin{bmatrix} \Re(V) \\ \Im(V) \end{bmatrix}, \hat{W} = \begin{bmatrix} A & -B \\ B & A \end{bmatrix},$$

where \Re and \Im denote the real and imaginary parts of a complex number. Note that $\hat{h}^{(t)} \in \mathbb{R}^{2n}$, $\hat{W} \in \mathbb{R}^{2n \times 2n}$, and $\hat{V} \in \mathbb{R}^{2n \times n_x}$, where n_x is the dimension of the input vector $x^{(t)}$ in Equation (1).

Assuming that the activation function ϕ applies to the real and imaginary parts separately, it is easy to show that the update equation of the complex hidden state $h^{(t)}$ of the unitary RNN has the following real space representation

$$\hat{h}^{(t)} = \phi(\hat{W}\hat{h}^{(t-1)} + \hat{V}x^{(t)}). \quad (3)$$

Even when the activation function ϕ does not apply to the real and imaginary parts separately, it is still possible to find an equivalent representation in the real space. Consider the activation function proposed by (Arjovsky et al., 2016)

$$\sigma_{\text{modReLU}}(z) = \begin{cases} (|z| + b) \frac{z}{|z|}, & \text{if } |z| + b > 0 \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where b is a bias vector. For a hidden state $\hat{h} \in \mathbb{R}^{2n}$, the equivalent activation function in the real space representation is given by

$$\left[\hat{\phi}(a) \right]_i = \begin{cases} \frac{\sqrt{a_i^2 + a_{k_i}^2} + b_{\tilde{k}_i}}{\sqrt{a_i^2 + a_{k_i}^2}} a_i, & \text{if } \sqrt{a_i^2 + a_{k_i}^2} + b_{\tilde{k}_i} > 0 \\ 0, & \text{otherwise} \end{cases}$$

where $k_i = ((i + n) \bmod 2n)$ and $\tilde{k}_i = (i \bmod n)$ for all $i \in \{1, \dots, 2n\}$. The activation function $\hat{\phi}$ is no longer applied to hidden units independently.

Now we will show that the matrix \hat{W} is orthogonal. By definition of a unitary matrix, we have $WW^* = I$ where the $*$ represents the conjugate transpose. This implies that $AA' + BB' = I$ and $BA' - AB' = 0$. And since we have

$$\hat{W}\hat{W}' = \begin{bmatrix} AA' + BB' & AB' - BA' \\ BA' - AB' & AA' + BB' \end{bmatrix}, \quad (5)$$

it follows that $\hat{W}\hat{W}' = I$. Also note that \hat{W} has a special structure - it is a block-matrix.

The discussion above shows that using a complex, unitary transition matrix in $\mathbb{C}^{n \times n}$ is equivalent to using an orthogonal matrix, belonging to a subset of $\mathbf{O}(2n)$, in a new RNN with twice the hidden size. This is why in this work we focus mainly on parametrising orthogonal matrices.

4. Parametrisation of the transition matrix

Before discussing the details of our parametrisation, we first introduce a few notations. For $n, k \in \mathbb{N}$ and $2 \leq k \leq n$, let $\mathcal{H}_k : \mathbb{R}^k \rightarrow \mathbb{R}^{n \times n}$ be defined as

$$\mathcal{H}_k(\mathbf{u}) = \begin{bmatrix} I_{n-k} & 0 \\ 0 & I_k - 2 \frac{\mathbf{u}\mathbf{u}'}{\|\mathbf{u}\|^2} \end{bmatrix}, \quad (6)$$

where I_k denotes the k -dimensional identity matrix. For $\mathbf{u} \in \mathbb{R}^k$, $\mathcal{H}_k(\mathbf{u})$ is the *Householder Matrix* in $\mathbf{O}(n)$ representing the reflection about the hyperplane orthogonal to the vector $(\mathbf{0}'_{n-k}, \mathbf{u}') \in \mathbb{R}^n$ and passing through the origin, where $\mathbf{0}_{n-k}$ denotes the zero vector in \mathbb{R}^{n-k} .

We also define the mapping $\mathcal{H}_1 : \mathbb{R} \rightarrow \mathbb{R}^{n \times n}$ as

$$\mathcal{H}_1(u) = \begin{bmatrix} I_{n-1} & 0 \\ 0 & u \end{bmatrix}. \quad (7)$$

Note that $\mathcal{H}_1(u)$ is not necessarily a Householder reflection. However, when $u \in \{1, -1\}$, $\mathcal{H}_1(u)$ is orthogonal.

Finally, for $n, k \in \mathbb{N}$ and $1 \leq k \leq n$, we define

$$\begin{aligned} \mathcal{M}_k : \mathbb{R}^k \times \dots \times \mathbb{R}^n &\rightarrow \mathbb{R}^{n \times n} \\ (\mathbf{u}_k, \dots, \mathbf{u}_n) &\mapsto \mathcal{H}_n(\mathbf{u}_n) \dots \mathcal{H}_k(\mathbf{u}_k). \end{aligned}$$

We propose to parametrise the transition matrix W of an RNN using the mappings $\{\mathcal{M}_k\}$. When using m reflection vectors $\{\mathbf{u}_i\}$, the parametrisation can be expressed as

$$\begin{aligned} W &= \mathcal{M}_{n-m+1}(\mathbf{u}_{n-m+1}, \dots, \mathbf{u}_n) \\ &= \mathcal{H}_n(\mathbf{u}_n) \dots \mathcal{H}_{n-m+1}(\mathbf{u}_{n-m+1}), \end{aligned} \quad (8)$$

where $\mathbf{u}_i \in \mathbb{R}^i$ for $i \in \{n-m+1, \dots, n\}$.

For the particular case where $m = n$ in the above parametrisation, we have the following result.

Theorem 1. *The image of \mathcal{M}_1 includes the set of all $n \times n$ orthogonal matrices, i.e. $\mathbf{O}(n) \subset \mathcal{M}_1[\mathbb{R} \times \dots \times \mathbb{R}^n]$.*

Note that Theorem 1 would not be valid if $\mathcal{H}_1(\cdot)$ was a standard Householder reflection. In fact, in the two-dimensional case, for instance, the matrix $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ cannot be expressed as the product of exactly two standard Householder matrices.

The parametrisation in (8) has the following advantages:

Methods	Constraint on the transition matrix	Time complexity of one online gradient step	Search space of the transition matrix
uRNN (Arjovsky et al., 2016)	$\ W\ = 1$	$\mathcal{O}(Tn \log(n))$	A subset of $\mathbf{U}(n)$ when $n > 7$
Full-capacity uRNN (Wisdom et al., 2016)	$\ W\ = 1$	$\mathcal{O}(Tn^2 + n^3)$	The full $\mathbf{U}(n)$ set
Unitary RNN (Hyland & Rätsch, 2017)	$\ W\ = 1$	$\mathcal{O}(Tn^2 + n^3)$	The full $\mathbf{U}(n)$ set
oRNN (Our approach)	$\ W\ = 1$	$\mathcal{O}(Tnm)$ where $m \leq n$	The full $\mathbf{O}(n)$ set when $m = n$

Table 1. Table showing the time complexities associated with one stochastic gradient step (mini-batch size = 1) for different methods, when the size of the hidden layer of an RNN is n and the input sequence has length T .

1. The parametrisation is smooth*, which is convenient for training with gradient descent. It is also flexible - a good trade-off between expressiveness and speed can be found by tuning the number of reflection vectors.
2. The time and space complexities involved in one gradient calculation are, in the worst case, the same as that of the sRNN with the same number of hidden units. This is discussed in the following subsections.
3. When $m < n$, the matrix W is always orthogonal, as long as the reflection vectors are nonzero. For $m = n$, the only additional requirement for W to be orthogonal is that $u_1 \in \{-1, 1\}$.
4. When $m = n$, the transition matrix can span the whole set of $n \times n$ orthogonal matrices. In this case, the total number of parameters needed for W is $n(n+1)/2$. This results in only n redundant parameters since the orthogonal set $\mathbf{O}(n)$ is a $n(n-1)/2$ manifold.

4.1. Back-propagation algorithm

Let $\mathbf{u}_i \in \mathbb{R}^i$. Let $U := (\mathbf{u}_n | \dots | \mathbf{u}_{n-m+1}) \in \mathbb{R}^{n \times m}$ be the parameter matrix constructed from the reflection vectors $\{\mathbf{u}_i\}$. In particular, the j -th column of U can be expressed using the zero vector $\mathbf{0}_{j-1} \in \mathbb{R}^{j-1}$ as

$$U_{*,j} = \begin{bmatrix} \mathbf{0}_{j-1} \\ \mathbf{u}_{n-j+1} \end{bmatrix} \in \mathbb{R}^n, \quad 1 \leq j \leq m. \quad (9)$$

Let \mathcal{L} be a scalar loss function and $C^{(t)} := Wh^{(t-1)}$, where W is constructed using the $\{\mathbf{u}_i\}$ vectors following Equation (8). In order to back-propagate the gradients through time, we need to compute the following partial derivatives

$$\frac{\partial \mathcal{L}}{\partial U^{(t)}} := \left[\frac{\partial C^{(t)}}{\partial U} \right]' \frac{\partial \mathcal{L}}{\partial C^{(t)}}, \quad (10)$$

$$\frac{\partial \mathcal{L}}{\partial h^{(t-1)}} = \left[\frac{\partial C^{(t)}}{\partial h^{(t-1)}} \right]' \frac{\partial \mathcal{L}}{\partial C^{(t)}}, \quad (11)$$

*except on a subset of zero Lebesgue measure.

at each time step t . Note that in Equation (10) $h^{(t-1)}$ is taken as a constant with respect to U . Furthermore, we have $\frac{\partial \mathcal{L}}{\partial U} = \sum_{t=1}^T \frac{\partial \mathcal{L}}{\partial U^{(t)}}$, where T is the length of the input sequence. The gradient flow through the RNN at time step t is shown in Figure 1.




Figure 1. Gradient flow through the RNN at time step t . Note that we have $\mathcal{L} = \sum_{t=1}^T \mathcal{L}_t(o_t, y_t)$, where $\{y_t\}$ are target outputs.

Before describing the algorithm to compute the back-propagated gradients $\frac{\partial \mathcal{L}}{\partial U^{(t)}}$ and $\frac{\partial \mathcal{L}}{\partial h^{(t-1)}}$, we first derive their expressions as a function of U , $h^{(t-1)}$ and $\frac{\partial \mathcal{L}}{\partial C^{(t)}}$ using the compact WY representation (Joffrain et al., 2006) of the product of Householder reflections.

Proposition 1. Let $n, m \in \mathbb{N}$ s.t. $m \leq n - 1$. Let $\mathbf{u}_i \in \mathbb{R}^i$ and $U = (\mathbf{u}_n | \dots | \mathbf{u}_{n-m+1})$ be the matrix defined in Equation (9). We have

$$T := \text{striu}(U'U) + \frac{1}{2} \text{diag}(U'U), \quad (12)$$

$$\mathcal{H}_n(\mathbf{u}_n) \dots \mathcal{H}_{n-m+1}(\mathbf{u}_{n-m+1}) = I - UT^{-1}U', \quad (13)$$

where $\text{striu}(U'U)$, and $\text{diag}(U'U)$ represent the strictly upper part and the diagonal of the matrix $U'U$, respectively.

Equation (13) is the compact WY representation of the product of Householder reflections. For the particular case where $m = n$, the RHS of Equation (13) should be replaced by $(I - UT^{-1}U') \mathcal{H}_1(u_1)$, where \mathcal{H}_1 is defined in (7) and $U = (\mathbf{u}_n | \dots | \mathbf{u}_2)$.

The following theorem gives the expressions of the gradients $\frac{\partial \mathcal{L}}{\partial U^{(t)}}$ and $\frac{\partial \mathcal{L}}{\partial h^{(t-1)}}$ when $m \leq n - 1$ and $h = h^{(t-1)}$.

Algorithm 1 Local forward and backward propagations at time step t . For a matrix A , $A_{*,k}$ denotes the k -th column.

- 1: **Inputs:** $h^{(t-1)}$, $\frac{\partial \mathcal{L}}{\partial C}$, $U = (\mathbf{u}_n | \dots | \mathbf{u}_{n-m+1})$.
- 2: **Outputs:** $\frac{\partial \mathcal{L}}{\partial U^{(t)}}$, $\frac{\partial \mathcal{L}}{\partial h^{(t-1)}}$, $C^{(t)} = Wh^{(t-1)}$
- 3: **Require:** $G \in \mathbb{R}^{n \times m}$, $g \in \mathbb{R}^n$, $H \in \mathbb{R}^{n \times (m+1)}$
- 4: $N \leftarrow (\|\mathbf{u}_n\|^2, \dots, \|\mathbf{u}_{n-m+1}\|^2)$
- 5: $H_{*,m+1} \leftarrow h^{(t-1)}$
- 6: $g \leftarrow \frac{\partial \mathcal{L}}{\partial C}$
- 7: **for** $k = m$ **to** 1 **do** {Local Forward Propagation}
- 8: $\tilde{h}_k \leftarrow \frac{2}{N_k} U'_{*,k} H_{*,k+1}$
- 9: $H_{*,k} \leftarrow H_{*,k+1} - \tilde{h}_k U_{*,k}$
- 10: **end for**
- 11: **for** $k = 1$ **to** m **do** {Local Backward Propagation}
- 12: $\tilde{C}_k \leftarrow \frac{2}{N_k} U'_{*,k} g$
- 13: $g \leftarrow g - \tilde{C}_k U_{*,k}$
- 14: $G_{*,k} \leftarrow -\tilde{h}_k g - \tilde{C}_k H_{*,k+1}$
- 15: **end for**
- 16: $C^{(t)} \leftarrow H_{*,1}$
- 17: $\frac{\partial \mathcal{L}}{\partial h^{(t-1)}} \leftarrow g$
- 18: $\frac{\partial \mathcal{L}}{\partial U^{(t)}} \leftarrow G$

Theorem 2. Let $n, m \in \mathbb{N}$ s.t. $m \leq n-1$. Let $U \in \mathbb{R}^{n \times m}$, $h \in \mathbb{R}^n$, and $C = (I - UT^{-1}U')h$, where T is defined in Equation (12). If \mathcal{L} is a scalar loss function which depends on C , then we have

$$\frac{\partial \mathcal{L}}{\partial U} = U[(\tilde{h}\tilde{C}') \circ B' + (\tilde{C}\tilde{h}') \circ B] - \frac{\partial \mathcal{L}}{\partial C} \tilde{h}' - h\tilde{C}', \quad (14)$$

$$\frac{\partial \mathcal{L}}{\partial h} = \frac{\partial \mathcal{L}}{\partial C} - U\tilde{C}, \quad (15)$$

where $\tilde{h} = T^{-1}U'h$, $\tilde{C} = (T')^{-1}U'\frac{\partial \mathcal{L}}{\partial C}$, and $B = \text{striv}(J) + \frac{1}{2}I$, with J being the $m \times m$ matrix of all ones and \circ the Hadamard product.

The proof of Equations (14) and (15) is provided in Appendix A. Based on Theorem 2, Algorithm 1 performs the one-step forward-propagation (FP) and back-propagation (BP) required to compute $C^{(t)}$, $\frac{\partial \mathcal{L}}{\partial U^{(t)}}$, and $\frac{\partial \mathcal{L}}{\partial h^{(t-1)}}$. See Appendix B for more detail about how this algorithm is derived using Theorem 2.

In the next section we analyse the time and space complexities of this algorithm.

4.2. Time and Space complexity

At each time step t , the flop count required by Algorithm 1 is $(13n+2)m$; $6nm$ for the one-step FP and $(7n+2)m$ for the one-step BP. Note that the vector N only needs to be calculated at one time step. This reduces the flop count at the remaining time steps to $(11n+3)m$. The fact that the matrix U has all zeros in its upper triangular part can be used to further reduce the total flop count to $(11n-3m+$

Model	Flop counts	
	FP	BP
oRNN (n, m)	$(4n-m+2)m$	$(7n-2m+3)m$
sRNN (n)	$2n^2-n$	$3n^2-n$

Table 2. Summary of the flop counts due to the computations of one-step FP and BP through the hidden-to-hidden connections. The BP flop counts for the oRNN case assumes that the H matrices (see Algorithm 1) are not locally generated during the BP steps. Otherwise, the flop count would be $(11n-3m+5)m$.

$5)m$; $(4n-m+2)m$ for the one-step FP and $(7n-2m+3)m$ for the one-step BP. See Appendix C for more details.

Note that if the values of the matrices H , defined in Algorithm 1, are first stored during a ‘‘global’’ FP (i.e. through all time steps), then used in the BP steps, the time complexity[†] for a global FP and BP using one input sequence of length T are, respectively, $\approx 3n^2T$ and $\approx 5n^2T$, when $m \approx n$ and $n \gg 1$. In contrast with the sRNN case with n hidden units, the global FP and BP have time complexities $\approx 2n^2T$ and $\approx 3n^2T$. Hence, when $m \approx n$, the FP and BP steps using our parametrisation require only about twice more flops than the sRNN case with the same number of hidden units.

Note, however, that storing the values of the matrices H at all time steps requires the storage of mnT values for one sequence of length T , compared with nT when only the hidden states $\{h^{(t)}\}_{t=1}$ are stored. When $m \gg 1$ this may not be practical. One solution to this problem is to generate the matrices H locally at each BP step using U and $h^{(t-1)}$. This results in a global BP complexity of $(11n-3m+5)mT$. Table 2 summarises the flop counts for the FP and BP steps. Note that these flop counts are for the case when $m \leq n-1$. When $m = n$, the complexity added due to the multiplication by $\mathcal{H}_1(u_1)$ is negligible.

4.3. Extension to the Unitary case

Although we decided to focus on the set of real-valued orthogonal matrices, for the reasons given in Section 3, our parametrisation can readily be modified to apply to the general unitary case.

Let $\hat{\mathcal{H}}_k : \mathbb{C}^k \rightarrow \mathbb{C}^{n \times n}$, $2 \leq k \leq n$, be defined by Equation (6) where the transpose sign $'$ is replaced by the conjugate transpose $*$. Furthermore, let $\hat{\mathcal{H}}_1 : \mathbb{R}^n \rightarrow \mathbb{C}^{n \times n}$ be defined as $\hat{\mathcal{H}}_1(\theta) = \text{diag}(e^{i\theta_1}, \dots, e^{i\theta_n})$. With the new mappings $\{\hat{\mathcal{H}}_k\}_{k=1}^{k=n}$, we have the following corollary.

[†]We considered only the time complexity due to computations through the hidden-to-hidden connections of the network.

Corollary 1. Let $\hat{\mathcal{M}}_1$ be the mapping defined as

$$\hat{\mathcal{M}}_1 : \mathbb{R}^n \times \mathbb{C}^2 \times \dots \times \mathbb{C}^n \rightarrow \mathbb{C}^{n \times n}$$

$$(\theta, \mathbf{u}_2, \dots, \mathbf{u}_n) \mapsto \hat{\mathcal{H}}_n(\mathbf{u}_n) \dots \hat{\mathcal{H}}_2(\mathbf{u}_2) \hat{\mathcal{H}}_1(\theta).$$

The image of $\hat{\mathcal{M}}_1$ spans the full set of unitary matrices $U(n)$ and any point on its image is a unitary matrix.

5. Experiments

All RNN models were implemented using the python library *theano* (Theano Development Team, 2016). For efficiency, we implemented the one-step FP and BP algorithms described in Algorithm 1 using C code[‡]. We tested the new parametrisation on five different datasets all having long-term dependencies. We call our parametrised network oRNN (for orthogonal RNN). We set its activation function to the `leaky_ReLU` defined as $\phi(x) = \max(\frac{x}{10}, x)$. To ensure that the transition matrix of the oRNN is always orthogonal, we set the scalar u_1 to -1 if $u_1 \leq 0$ and 1 otherwise after each gradient update. Note that the parameter matrix U in Equation (9) has all zeros in its upper triangular part. Therefore, after calculating the gradient of a loss with respect to U (i.e. $\frac{\partial \mathcal{L}}{\partial U}$), the values in the upper triangular part are set to zero.

For all experiments, we used the *adam* method for stochastic gradient descent (Kingma & Ba, 2014). We initialised all the parameters using uniform distributions similar to (Arjovsky et al., 2016). The biases of all models were set to zero, except for the forget bias of the LSTM, which we set to 5 to facilitate the learning of long-term dependencies (Koutník et al., 2014).

5.1. Sequence generation

In this experiment, we followed a similar setting to (Koutník et al., 2014) where we trained RNNs to encode song excerpts. We used the track *Manyrista* from album *Musica Deposita* by *Cuprum*. We extracted five consecutive excerpts around the beginning of the song, each having 800 data points and corresponding to 18ms with a 44.1Hz sampling frequency. We trained an sRNN, LSTM, and oRNN for 5000 epochs on each of the pieces with five random seeds. For each run, the lowest Normalised Mean Squared Error (NMSE) during the 5000 epochs was recorded. For each model, we tested three different hidden sizes. The total number of parameters N_p corresponding to these hidden sizes was approximately equal to 250, 500, and 1000. For the oRNN, we set the number of reflection vectors to the hidden size for each case, so that the transition matrix is allowed to span the full set of orthogonal

[‡]Our implementation can be found at https://github.com/zmhammedi/Orthogonal_RNN.




Figure 2. Sequence generation task: The plots show the NMSE distributions for the different models with respect to the total number of parameters for the sequence generation task. The horizontal red lines represent the medians of the NMSE over 25 data points (i.e. five seeds for each of the five song excerpts). The solid rectangles and the dashed bars represent the [25% – 75%] and [9% – 91%] confidence intervals respectively.

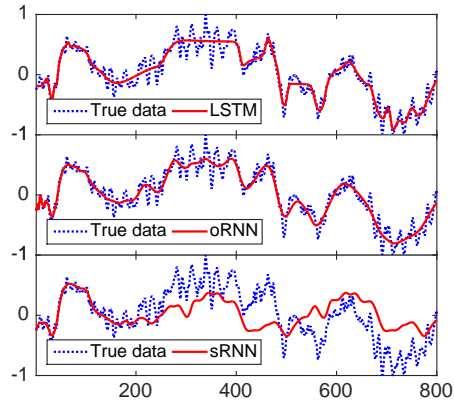


Figure 3. Sequence generation task: The RNN-generated sequences against the true data for one of the five excerpts used. We only displayed the best performing models for $N_p \simeq 1000$.

matrices. The results are shown in Figures 2 and 3. All the learning rates were set to 10^{-3} . The orthogonal parametrisation outperformed the sRNN and performed on average better than the LSTM.

5.2. Addition Task

In this experiment, we followed a similar setting to (Arjovsky et al., 2016), where the goal of the RNN is to output the sum of two elements in the first dimension of a two-dimensional sequence. The location of the two elements to be summed are specified by the entries in the second dimension of the input sequence. In particular, the first dimension of every input sequence consists of random numbers between 0 and 1. The second dimension has all zeros except for two elements equal to 1. The first unit entry

is located in the first half of the sequence, and the second one in the second half. We tested two different sequence lengths $T = 400, 800$. All models were trained to minimise the Mean Squared Error (MSE). The baseline MSE for this task is 0.167; for a model that always outputs one.




Figure 4. Addition task: For each lag T , the red curves represent two runs of the oRNN model with two different random initialisation seeds. LSTM and sRNN did not beat the baseline MSE.

We trained an oRNN with $n = 128$ hidden units and $m = 16$ reflections. We trained an LSTM and sRNN with hidden sizes 28 and 54, respectively, corresponding to a total number of parameters ≈ 3600 (i.e. same as the oRNN model). We chose a batch size of 50, and after each iteration, a new set of sequences was generated randomly. The learning rate for the oRNN was set to 0.01. Figure 4 displays the results for both lags.

The oRNN was able to beat the baseline MSE in less than 5000 iterations for both lags and for two different random initialisation seeds. This is in line with the results of the unitary RNN (Arjovsky et al., 2016).

5.3. Pixel MNIST

In this experiment, we used the MNIST image dataset. We split the dataset into training (55000 instances), validation (5000 instances), and test sets (10000 instances). We trained oRNNs with $n \in \{128, 256\}$ and $m \in \{16, 32, 64\}$, where n and m are the number of hidden units and reflection vectors respectively, to minimise the cross-entropy error function. We experimented with (mini-batch size, learning rate) $\in \{(1, 10^{-4}), (50, 10^{-3})\}$.

Table 3 compares the test performance of our best model against results available in the literature for unitary/orthogonal RNNs. Despite having fewer total number of parameters, our model performed better than three out the four models selected for comparison (all having $\geq 16K$ parameters). Figure 5 shows the validation accuracy as a function of the number of epochs of our oRNN model in Table 3.

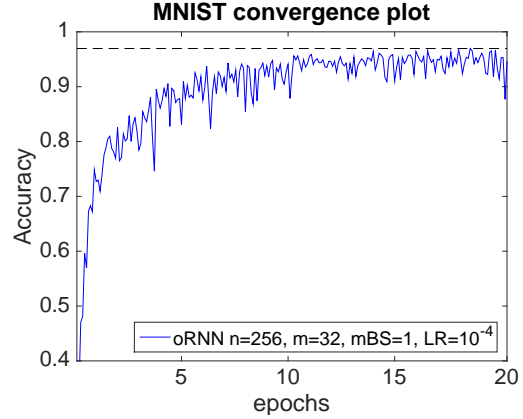


Figure 5. MNIST experiment: Validation accuracy of the oRNN in Table 3 as a function of the number of epochs. mBS and LR in the legend stand for mini-batch size and learning rate respectively.

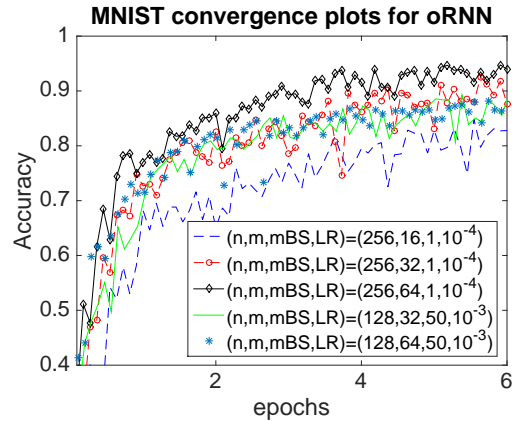


Figure 6. MNIST experiment: Effect of varying the number of reflection vectors m on the validation accuracy and speed of convergence. mBS (resp, LR) stands for mini-batch size (resp, learning rate).

Figure 6 shows the effect of varying the number of reflection vectors m on the performance.

5.4. Penn Tree Bank

In this experiment, we tested the oRNN on the task of character level prediction using the Penn Tree Bank Corpus. The data was split into training (5017K characters), validation (393K characters), and test sets (442K characters). The total number of unique characters in the corpus was 49. The vocabulary size was 10K and any other words were replaced by the special token $\langle \text{unk} \rangle$. The number of characters per instance (i.e. char/line) in the training data ranged between 2 and 518 with an average of 118 char/line. We trained an oRNN and LSTM with hidden units 512 and 183 respectively, corresponding to a total of $\approx 180K$ parameters, for 20 epochs. We set the number of reflections to 510

Model	hidden size (# reflections)	Number of parameters	validation accuracy	test accuracy
oRNN	256 (m=32)	≈11K	97.0 %	97.2 %
RNN (Vorontsov et al., 2017)	128	≈35K	-	94.1 %
uRNN (Arjovsky et al., 2016)	512	≈16K	-	95.1 %
RC uRNN (Wisdom et al., 2016)	512	≈16K	-	97.5 %
FC uRNN (Wisdom et al., 2016)	116	≈16K	-	92.8 %

Table 3. Results summary for the MNIST digit classification experiment and comparison with the uRNN results available in the literature. ‘FC’ and ‘RC’ stand for Full-Capacity and Restricted Capacity respectively. The oRNN with $n = 256$ was trained using a mini-batch size of 1 and a learning rate of 10^{-4} .

	Model (n_h)	# params	valid cost	test cost
1-step	oRNN (512)	183K	1.73	1.68
	LSTM (184)	181K	1.69	1.68
5-step	oRNN (512)	183K	3.85	3.85
	LSTM (183)	180K	3.81	3.8

Table 4. Results summary for the Penn Tree bank experiment.

for the oRNN. The learning rate was set to 0.0001 for both models with a mini-batch size of 1.

Similarly to (Pascanu et al., 2013) we considered two tasks: one where the model predicts one character ahead and the other where it predicts a character five steps ahead. It was suggested that solving the later task would require the learning of longer term correlations in the data rather than the shorter ones. Table 4 summarises the test results. The oRNN and LSTM performed similarly to each other on the one-step head prediction task. Whereas on the five-step ahead prediction task, the LSTM was better. The performance of both models on this task was close to the state of the art result for RNNs 3.74 bpc (Pascanu et al., 2013).

Nevertheless, our oRNN still outperformed the results of (Vorontsov et al., 2017) which used both soft and hard orthogonality constraints on the transition matrix. Their RNN was trained on 99% of the data (sentences with ≤ 300 characters) and had the same number of hidden units as the oRNN in our experiment. The lowest test cost achieved was 2.20(bpc) for the one-step-ahead prediction task.

5.5. Copying task

We tested our model on the copy task described in details in (Gers et al., 2001; Arjovsky et al., 2016). Using an oRNN with the `leaky_ReLU` we were not able to reproduce the same performance as the uRNN (Arjovsky et al., 2016; Wisdom et al., 2016). However, we were able to achieve a comparable performance when using the `OPLU` activation function (Chernodub & Nowicki, 2016), which is a norm-preserving activation function. In order to explore whether the poor performance of the oRNN was only due to the activation function, we tested the same activation as the uRNN

(i.e. the real representation of `modReLU` defined in Equation (4)) on the oRNN. This did not improve the performance compared to the `leaky_ReLU` case suggesting that the block structure of the uRNN transition matrix, when expressed in the real space (see Section 3), may confer special benefits in some cases.

6. Discussion

In this work, we presented a new parametrisation of the transition matrix of a recurrent neural network using Householder reflections. This method allows an easy and computationally efficient way to enforce an orthogonal constraint on the transition matrix which then ensures that exploding gradients do not occur during training. Our method could also be applied to other deep neural architectures to enforce orthogonality between hidden layers. Note that a “soft” orthogonal constraint could also be applied using our parametrisation by, for example, allowing u_1 to vary continuously between -1 and 1.

It is important to note that our method is particularly advantageous for stochastic gradient descent when the mini-batch size is close to 1. In fact, if B is the mini-batch size and T is the average length of the input sequences, then a network with n hidden units trained using other methods (Vorontsov et al., 2017; Wisdom et al., 2016; Hyland & Rättsch, 2017) that enforce orthogonality (see Section 2), would have time complexity $\mathcal{O}(BTn^2 + n^3)$. Clearly when $BT \gg n$ this becomes $\mathcal{O}(BTn^2)$, which is the same time complexity as that of the sRNN and oRNN (with $m = n$). In contrast with the case of fully connected deep forward networks with no weight sharing between layers (# layer = L), the time complexity using our method is $\mathcal{O}(BLnm)$ whereas other methods discussed in this work (see Section 2) would have time complexity $\mathcal{O}(BLn^2 + Ln^3)$. The latter methods are less efficient in this case since $B \gg n$ is less likely to be the case compared with $BT \gg n$ when using SGD.

From a performance point of view, further experiments should be performed to better understand the difference between the unitary versus orthogonal constraint.

Acknowledgment

The authors would like to acknowledge Department of State Growth Tasmania for partially funding this work through SenseT. We would also like to thank Christfried Webers for his valuable feedback.

References

- Arjovsky, Martin, Shah, Amar, and Bengio, Yoshua. Unitary evolution recurrent neural networks. In *International Conference on Machine Learning (ICML)*, 2016.
- Chernodub, Artem and Nowicki, Dimitri. Norm-preserving orthogonal permutation linear unit activation functions (oplu). *arXiv preprint arXiv:1604.02313*, 2016.
- Gers, F. A., Eck, D., and Schmidhuber, J. Applying LSTM to time series predictable through time-window approaches. In Dorffner, Gerg (ed.), *Artificial Neural Networks – ICANN 2001 (Proceedings)*, pp. 669–676, Berlin, 2001. Springer.
- Glorot, Xavier and Bengio, Yoshua. Understanding the difficulty of training deep feedforward neural networks. In *Aistats*, volume 9, pp. 249–256, 2010.
- Henaff, Mikael, Szlam, Arthur, and LeCun, Yann. Orthogonal rnns and long-memory tasks. *arXiv preprint arXiv:1602.06662*, 2016.
- Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Hochreiter, Sepp, Bengio, Yoshua, Frasconi, Paolo, and Schmidhuber, Jürgen. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies, 2001.
- Hyland, Stephanie L. and Rätsch, Gunnar. Learning unitary operators with help from $u(n)$. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pp. 2050–2058, 2017. URL <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14930>.
- Joffrain, Thierry, Low, Tze Meng, Quintana-Ortí, Enrique S, Geijn, Robert van de, and Zee, Field G Van. Accumulating householder transformations, revisited. *ACM Transactions on Mathematical Software (TOMS)*, 32(2): 169–179, 2006.
- Kingma, Diederik P. and Ba, Jimmy. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL <http://arxiv.org/abs/1412.6980>.
- Koutník, Jan, Greff, Klaus, Gomez, Faustino J., and Schmidhuber, Jürgen. A clockwork RNN. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pp. 1863–1871, 2014. URL <http://jmlr.org/proceedings/papers/v32/koutnik14.html>.
- Le, Quoc V., Jaitly, Navdeep, and Hinton, Geoffrey E. A simple way to initialize recurrent networks of rectified linear units. *CoRR*, abs/1504.00941, 2015. URL <http://arxiv.org/abs/1504.00941>.
- Mikolov, Tomáš. *Statistical language models based on neural networks*. PhD thesis, Brno University of Technology, 2012.
- Pascanu, Razvan, Mikolov, Tomas, and Bengio, Yoshua. On the difficulty of training recurrent neural networks. *ICML (3)*, 28:1310–1318, 2013.
- Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May 2016. URL <http://arxiv.org/abs/1605.02688>.
- Vorontsov, Eugene, Trabelsi, Chiheb, Kadoury, Samuel, and Pal, Chris. On orthogonality and learning recurrent networks with long term dependencies. *arXiv preprint arXiv:1702.00071*, 2017.
- Wisdom, Scott, Powers, Thomas, Hershey, John R., Roux, Jonathan Le, and Atlas, Les E. Full-capacity unitary recurrent neural networks. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 4880–4888, 2016.