

# Multi-Dimensional Recurrent Neural Networks

Alex Graves<sup>1</sup>, Santiago Fernández<sup>1</sup>, Jürgen Schmidhuber<sup>1,2</sup>

<sup>1</sup> IDSIA , Galleria 2, 6928 Manno-Lugano, Switzerland

<sup>2</sup> TU Munich, Boltzmannstr. 3, 85748 Garching, Munich, Germany  
{alex,santiago,juergen}@idsia.ch


**Abstract.** Recurrent neural networks (RNNs) have proved effective at one dimensional sequence learning tasks, such as speech and online handwriting recognition. Some of the properties that make RNNs suitable for such tasks, for example robustness to input warping, and the ability to access contextual information, are also desirable in multi-dimensional domains. However, there has so far been no direct way of applying RNNs to data with more than one spatio-temporal dimension. This paper introduces multi-dimensional recurrent neural networks, thereby extending the potential applicability of RNNs to vision, video processing, medical imaging and many other areas, while avoiding the scaling problems that have plagued other multi-dimensional models. Experimental results are provided for two image segmentation tasks.

## 1 Introduction


Recurrent neural networks (RNNs) were originally developed as a way of extending neural networks to sequential data. The addition of recurrent connections allows RNNs to make use of previous context, as well as making them more robust to warping along the time axis than non-recursive models [15]. Access to contextual information and robustness to warping are also important when dealing with multi-dimensional data. For example, a face recognition algorithm should use the entire face as context, and should be robust to changes in perspective, distance etc. It therefore seems desirable to apply RNNs to such tasks.

However, the standard RNN architectures are explicitly one dimensional, meaning that in order to use them for multi-dimensional tasks, the data must be pre-processed to one dimension, for example by presenting one vertical line of an image at a time to the network. One successful use of neural networks for multi-dimensional data has been the application of convolution networks [10] to image processing tasks such as digit recognition [14]. One disadvantage of convolution nets is that because they are not recurrent, they rely on hand specified kernel sizes to introduce context. Another disadvantage is that they do not scale well to large images. For example, sequences of handwritten digits must be pre-segmented into individual characters before they can be recognised by convolution networks [10].

Various statistical models have been proposed for multi-dimensional data, notably multi-dimensional HMMs. However, multi-dimensional HMMs suffer from two serious drawbacks: (1) the time required to run the Viterbi algorithm, and thereby calculate the optimal state sequences, grows exponentially with the size of the data exemplars; (2) the



**Fig. 1.** 2D RNN Forward pass.



**Fig. 2.** 2D RNN Backward pass.

number of transition probabilities, and hence the required memory, grows exponentially with the data dimensionality. Numerous approximate methods have been proposed to alleviate one or both of these problems, including pseudo 2D and 3D HMMs [7], isolating elements [11], approximate Viterbi algorithms [9], and dependency tree HMMs [8]. However, none of these methods exploit the full multi-dimensional structure of the data.

As we will see, multi dimensional recurrent neural networks (MDRNNs) bring the benefits of RNNs to multi-dimensional data, without suffering from the scaling problems described above.

Section 2 describes the MDRNN architecture, Section 3 presents two experiments on image segmentation, and concluding remarks are given in Section 4.

## 2 Multi-Dimensional Recurrent Neural Networks


The basic idea of MDRNNs is to replace the single recurrent connection found in standard RNNs with as many recurrent connections as there are dimensions in the data. During the forward pass, at each point in the data sequence, the hidden layer of the network receives both an external input and its own activations from one step back along all dimensions. Figure 1 illustrates the two dimensional case.

Note that, although the word *sequence* usually denotes one dimensional data, we will use it to refer to data exemplars of any dimensionality. For example, an image is a two dimensional sequence, a video is a three dimensional sequence, and a series of fMRI brain scans is a four dimensional sequence.

Clearly, the data must be processed in such a way that when the network reaches a point in an  $n$ -dimensional sequence, it has already passed through all the points from which it will receive its previous activations. This can be ensured by following a suitable ordering on the set of points  $\{(x_1, x_2, \dots, x_n)\}$ . One example of a suitable ordering is  $(x_1, \dots, x_n) < (x'_1, \dots, x'_n)$  if  $\exists m \in (1, \dots, n)$  such that  $x_m < x'_m$  and  $x_i = x'_i \forall i \in (1, \dots, m-1)$ . Note that this is not the only possible ordering, and that its realisation for a particular sequence depends on an arbitrary choice of axes. We will return to this point in Section 2.1. Figure 3 illustrates the ordering for a 2 dimensional sequence.

The forward pass of an MDRNN can then be carried out by feeding forward the input and the  $n$  previous hidden layer activations at each point in the ordered input sequence, and storing the resulting hidden layer activations. Care must be taken at the sequence boundaries not to feed forward activations from points outside the sequence.

Note that the ‘points’ in the input sequence will in general be multivalued vectors. For example, in a two dimensional colour image, the inputs could be single pixels rep-



**Fig. 3.** 2D sequence ordering. The MDRNN forward pass starts at the origin and follows the direction of the arrows. The point  $(i,j)$  is never reached before both  $(i-1,j)$  and  $(i,j-1)$ .

resented by RGB triples, or blocks of pixels, or the outputs of a preprocessing method such as a discrete cosine transform.

The error gradient of an MDRNN (that is, the derivative of some objective function with respect to the network weights) can be calculated with an  $n$ -dimensional extension of backpropagation through time (BPTT) [15]. As with one dimensional BPTT, the sequence is processed in the reverse order of the forward pass. At each timestep, the hidden layer receives both the output error derivatives and its own  $n$  ‘future’ derivatives. Figure 2 illustrates the BPTT backward pass for two dimensions. Again, care must be taken at the sequence boundaries.

At a point  $\mathbf{x} = (x_1, \dots, x_n)$  in an  $n$ -dimensional sequence, define  $i_j^{\mathbf{x}}$  and  $h_k^{\mathbf{x}}$  respectively as the activations of the  $j^{th}$  input unit and the  $k^{th}$  hidden unit. Define  $w_{kj}$  as the weight of the connection going from unit  $j$  to unit  $k$ . Then for an  $n$ -dimensional MDRNN whose hidden layer consists of  $H$  summation units, where  $\theta_k$  is the activation function of hidden unit  $k$ , the forward pass for a sequence with dimensions  $(X_1, X_2, \dots, X_n)$  can be summarised as follows:

```

for  $x_1 = 0$  to  $X_1 - 1$  do
  for  $x_2 = 0$  to  $X_2 - 1$  do
    ...
    for  $x_n = 0$  to  $X_n - 1$  do
      for  $k = 1$  to  $H$  do
         $a_k = \sum_j i n_j^{(x_1, \dots, x_n)} w_{kj}$ 
        for  $i = 1$  to  $n$  do
          if  $x_i > 0$  then
             $a_k += \sum_j h_j^{(x_1, \dots, x_{i-1}, \dots, x_n)} w_{kj}$ 
         $h_k^{\mathbf{x}} = \theta_k(a_k)$ 

```

**Algorithm 1:** MDRNN Forward Pass


Defining  $\hat{o}_j^{\mathbf{x}}$  and  $\hat{h}_k^{\mathbf{x}}$  respectively as the derivatives of the objective function with respect to the activations of the  $j^{th}$  output unit and the  $k^{th}$  hidden unit at point  $\mathbf{x}$ , the backward pass is:

```


for  $x_1 = X_1 - 1$  to 0 do
  for  $x_2 = X_2 - 1$  to 0 do
    ...
    for  $x_n = X_n - 1$  to 0 do
      for  $k = 1$  to  $H$  do
         $e_k \leftarrow \sum_j \hat{o}_j^{(x_1, \dots, x_n)} w_{jk}$ 
        for  $i = 1$  to  $n$  do
          if  $x_i < X_i - 1$  then
             $e_k += \sum_j \hat{h}_j^{(x_1, \dots, x_i+1, \dots, x_n)} w_{jk}$ 
         $\hat{h}_k^{\mathbf{x}} \leftarrow \theta'_k(e_k)$ 

```

**Algorithm 2:** MDRNN Backward Pass



**Fig. 4.** Context available at (i,j) to a 2D RNN with a single hidden layer.



**Fig. 5.** Context available at (i,j) to a multi-directional 2D RNN.


Since the forward and backward pass require one pass each through the data sequence, the overall complexity of MDRNN training is linear in the number of data points and the number of network weights.

## 2.1 Multi-Directional MDRNNs

At a point  $(x_1, \dots, x_n)$  in the input sequence, the network described above has access to all points  $(x'_1, \dots, x'_n)$  such that  $x'_i \leq x_i \forall i \in (1, \dots, n)$ . This defines an n-dimensional ‘context region’ of the full sequence, as illustrated in Figure 4. For some tasks, such as object recognition, this would in principle be sufficient. The network could process the image as usual, and output the object label at a point when the object to be recognized is entirely contained in the context region.

Intuitively however, we would prefer the network to have access to the surrounding context in all directions. This is particularly true for tasks where precise localization is required, such as image segmentation. For one dimensional RNNs, the problem of multi-directional context was solved by the introduction of bidirectional recurrent neural networks (BRNNs) [13]. BRNNs contain two separate hidden layers that process the input sequence in the forward and reverse directions. The two hidden layers are connected to a single output layer, thereby providing the network with access to both past and future context.

BRNNs can be extended to n-dimensional data by using  $2^n$  separate hidden layers, each of which processes the sequence using the ordering defined above, but with a



**Fig. 6.** Axes used by the 4 hidden layers in a multi-directional 2D network. The arrows inside the rectangle indicate the direction of propagation during the forward pass.

different choice of axes. The axes are chosen so that each one has its origin on a distinct vertex of the sequence. The 2 dimensional case is illustrated in Figure 6. As before, the hidden layers are connected to a single output layer, which now has access to all surrounding context (see Figure 5).

If the size of the hidden layers is held constant, the multi-directional MDRNN architecture scales as  $O(2^n)$  for  $n$ -dimensional data. In practice however, the computing power of the network tends to be governed by the overall number of weights, rather than the size of the hidden layers. This is presumably because the data processing is shared between the layers. Since the complexity of the algorithm is linear in the number of parameters, the  $O(2^n)$  scaling factor can be removed by simply using smaller hidden layers for higher dimensions. Furthermore, the complexity of a task, and therefore the number of weights likely to be needed for it, does not necessarily increase with the dimensionality of the data. For example, both the networks described in this paper have less than half the weights than the one dimensional networks we have previously applied to speech recognition [4,3]. For a given task, we have also found that using a multi-directional MDRNN gives better results than a uni-directional MDRNN with the same overall number of weights, as previously demonstrated in one dimension [4].

For a multi-directional MDRNN, the forward and backward passes through an  $n$ -dimensional sequence can be summarised as follows:

- 1: For each of the  $2^n$  hidden layers choose a distinct vertex of the sequence, then define a set of axes such that the vertex is the origin and all sequence co-ordinates are  $\geq 0$
- 2: Repeat Algorithm 1 for each hidden layer
- 3: At each point in the sequence, feed forward all hidden layers to the output layer

**Algorithm 3:** Multi-Directional MDRNN Forward Pass

- 1: At each point in the sequence, calculate the derivative of the objective function with respect to the activations of output layer
- 2: With the same axes as above, repeat Algorithm 2 for each hidden layer

**Algorithm 4:** Multi-Directional MDRNN Backward Pass

## 2.2 Multi-Dimensional Long Short-Term Memory

So far we have implicitly assumed that the network can make use of all context to which it has access. For standard RNN architectures however, the range of context that can practically be used is limited. The problem is that the influence of a given input on the hidden layer, and therefore on the network output, either decays or blows up exponentially as it cycles around the network’s recurrent connections. This is usually referred to as the *vanishing gradient problem* [5].

Long Short-Term Memory (LSTM) [6,2] is an RNN architecture specifically designed to address the vanishing gradient problem. An LSTM hidden layer consists of multiple recurrently connected subnets, known as memory blocks. Each block contains a set of internal units, known as cells, whose activation is controlled by three multiplicative units: the input gate, forget gate and output gate. The effect of the gates is to allow the cells to store and access information over long periods of time, thereby avoiding the vanishing gradient problem.

The standard formulation of LSTM is explicitly one-dimensional, since the cell contains a single self connection, whose activation is controlled by a single forget gate. However we can easily extend this to  $n$  dimensions by using instead  $n$  self connections (one for each of the cell’s previous states along every dimension) with  $n$  forget gates.

## 3 Experiments


### 3.1 Air Freight Data

The Air Freight database [12] is a ray-traced colour image sequence that comes with a ground truth segmentation based on textural characteristics (Figure 7). The sequence is 455 frames long and contains 155 distinct textures. Each frame is 120 pixels high and 160 pixels wide.

The advantage of ray-traced data is the true segmentation can be defined directly from the 3D models. Although the images are not real, they are realistic in the sense that they have significant lighting, specular effects etc.

We used the sequence to define a 2D image segmentation task, where the aim was to assign each pixel in the input data to the correct texture class. We divided the data at random into a 250 frame train set, a 150 frame test set and a 55 frame validation set. Note that we could have instead defined a 3D task where the network processed the entire video as one sequence. However, this would have left us with only one exemplar.

For this task we used a multi-directional MDRNN with 4 LSTM hidden layers. Each layer consisted of 25 memory blocks, each containing 1 cell, 2 forget gates, 1 input gate, 1 output gate and 5 peephole weights. This gave a total 600 hidden units. The input and output activation functions of the cells were both tanh, and the activation



**Fig. 7.** Frame from the Air Freight database, showing the original image (left) and the colour-coded texture segmentation (right).

function for the gates was the logistic sigmoid in the range  $[0, 1]$ . The input layer was size 3 (one each for the red, green and blue components of the pixels) and the output layer was size 155 (one unit for each textural class). The network contained 43,257 trainable weights in total. As is standard for classification tasks, the softmax activation function was used at the output layer, with the cross-entropy objective function [1]. The network was trained using online gradient descent (weight updates after every training sequence) with a learning rate of  $10^{-6}$  and a momentum of 0.9.

The final pixel classification error rate was 7.1 % on the test set.

### 3.2 MNIST Data

The MNIST database [10] of isolated handwritten digits is a subset of a larger database available from NIST. It consists of size-normalized, centered images, each of which is 28 pixels high and 28 pixels wide and contains a single handwritten digit. The data comes divided into a training set with 60,000 images and a test set with 10,000 images. We used 10,000 of the training images for validation, leaving 50,000 for training.


The usual task on MNIST is to label the images with the corresponding digits. This is benchmark task for which many algorithms have been evaluated.

We carried out a slightly modified task where each pixel was classified according to the digit it belonged to, with an additional class for background pixels. However, the original task can be recovered by simply choosing the digit whose corresponding output unit had the highest cumulative activation for the entire sequence.

To test the network's robustness to input warping, we also evaluated it on an altered version of the MNIST test set, where elastic deformations had been applied to every image (see Figure 8).

We compared our results with the convolution neural network that has achieved the best results so far on MNIST [14]. Note that we re-implemented the convolution network ourselves, and we did not augment the training set with elastic distortions, which gives a substantial improvement in performance.

The MDRNN for this task was identical to that for the Air Freight task with the following exceptions: the sizes of the input and output layers were now 1 (for grayscale pixels) and 11 (one for each digit, plus background) respectively, giving 27,511 weights in total, and the gradient descent learning rate was  $10^{-5}$ .



**Fig. 8.** MNIST image before and after deformation.

For the distorted test set, we used the same degree of elastic deformation used by Simard [14] to augment the training set ( $\sigma = 4.0$ ,  $\alpha = 34.0$ ), with a different initial random field for every sample image.

Table 1 shows that the MDRNN matched the convolution net on the clean test set, and was considerably better on the warped test set. This suggests that MDRNNs are more robust to input warping than convolution networks.

**Table 1.** Image error rates on MNIST (pixel error rates in brackets)


Algorithm	Clean Test Set	Warped Test Set
MDRNN	0.9 % (0.4 %)	7.1 % (3.8 %)
Convolution	0.9 %	11.3 %

### 3.3 Analysis

One benefit of two dimensional tasks is that the operation of the network can be easily visualised. Figure 9 shows the network activations during a frames from the Air Freight database. As can be seen, the network segments this image almost perfectly, in spite of difficult, reflective surfaces such as the glass and metal tube running from left to right. Clearly, classifying individual pixels in such a surface requires the use of context.

An indication of the network’s sensitivity to context can be found by analysing the derivatives of the network outputs at a particular point in the sequence with respect to the inputs at all other points. Collectively, we refer to these derivatives as the *sequential Jacobian*. Figure 10 shows the absolute value of the sequential Jacobian of an output during classification of an image from the MNIST database. It can be seen that the network responds to context from across the entire image, and seems particularly attuned to the outline of the digit.





**Fig. 9.** 2D MDRNN applied to an image from the Air Freight database. The hidden layer activations display one unit from each of the layers. A common behaviour is to ‘mask off’ parts of the image, exhibited here by layers 2 and 3.

## 4 Conclusion


We have introduced multi-dimensional recurrent neural networks (MDRNNs), thereby extending the applicability of RNNs to  $n$ -dimensional data. We have added multi-directional hidden layers that provide the network with access to all contextual information, and we have developed a multi-dimensional variant of the Long Short-Term Memory RNN architecture. We have tested MDRNNs on two image segmentation tasks, and found that it was more robust to input warping than a state-of-the-art digit recognition algorithm.

## Acknowledgments

This research was funded by SNF grant 200021-111968/1.

## References

1. J.S. Bridle. Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. In F. Fogelman-Soulie and J. Herault, editors, *Neurocomputing: Algorithms, Architectures and Applications*, pages 227–236. Springer-Verlag, 1990.
2. F. Gers, N. Schraudolph, and J. Schmidhuber. Learning precise timing with LSTM recurrent networks. *Journal of Machine Learning Research*, 3:115–143, 2002.
3. A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber. Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the International Conference on Machine Learning, ICML 2006*, Pittsburgh, USA, 2006.



**Fig. 10.** Sequential Jacobian of a 2D RNN for an image from the MNIST database. The white outputs correspond to the class ‘background’ and the light grey ones to ‘8’. The black outputs represent misclassifications. The output pixel for which the Jacobian is calculated is marked with a cross. Absolute values are plotted for the Jacobian, and lighter colours are used for higher values.

4. A. Graves and J. Schmidhuber. Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Networks*, 18(5-6):602–610, 2005.
5. S. Hochreiter, Y. Bengio, P. Frasconi, and J. Schmidhuber. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. In S. C. Kremer and J. F. Kolen, editors, *A Field Guide to Dynamical Recurrent Neural Networks*. IEEE Press, 2001.
6. S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, 1997.
7. F. Hülken, F. Wallhoff, and G. Rigoll. Facial expression recognition with pseudo-3d hidden markov models. In *Proceedings of the 23rd DAGM-Symposium on Pattern Recognition*, pages 291–297, London, UK, 2001. Springer-Verlag.
8. J. Jiten, B. Mérialdo, and B. Huet. Multi-dimensional dependency-tree hidden Markov models. In *ICASSP 2006, 31st IEEE International Conference on Acoustics, Speech, and Signal Processing, May 14-19, 2006, Toulouse, France, May 2006*.
9. D. Joshi, J. Li, and J.Z. Wang. Parameter estimation of multi-dimensional hidden markov models: A scalable approach. pages III: 149–152, 2005.
10. Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998.
11. J. Li, A. Najmi, and R. M. Gray. Image classification by a two-dimensional hidden markov model. *IEEE Transactions on Signal Processing*, 48(2):517–533, 2000.
12. G. McCarter and A. Storkey. Air Freight Image Segmentation Database. <http://homepages.inf.ed.ac.uk/amos/afreightdata.html>.
13. M. Schuster and K. K. Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45:2673–2681, November 1997.
14. Patrice Y. Simard, Dave Steinkraus, and John C. Platt. Best practices for convolutional neural networks applied to visual document analysis. In *ICDAR '03: Proceedings of the Seventh International Conference on Document Analysis and Recognition*, page 958, Washington, DC, USA, 2003. IEEE Computer Society.
15. R. J. Williams and D. Zipser. Gradient-based learning algorithms for recurrent networks and their computational complexity. In Y. Chauvin and D. E. Rumelhart, editors, *Back-propagation: Theory, Architectures and Applications*, pages 433–486. Lawrence Erlbaum Publishers, Hillsdale, N.J., 1995.