



City Research Online

City St George's, University of London

Citation: Tran, S. N., Cherla, S., Garcez, A. & Weyde, T. (2017). The Recurrent Temporal Discriminative Restricted Boltzmann Machines. CoRR,

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/23671/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Linear-Time Sequence Classification using Restricted Boltzmann Machines

Son N. Tran

The Australian E-Health Research Centre, CSIRO
Brisbane, QLD 2026, Australia
son.tran@csiro.au

Srikanth Cherla, Artur d'Avila Garcez, Tillman Weyde

City University London
Northampton Square, London, EC1V 0HB, UK
{srikanth.cherla.1,a.garcez,t.e.veyde}@city.ac.uk

Abstract

Classification of sequence data is the topic of interest for dynamic Bayesian models and Recurrent Neural Networks (RNNs). While the former can explicitly model the temporal dependencies between class variables, the latter have a capability of learning representations. Several attempts have been made to improve performance by combining these two approaches or increasing the processing capability of the hidden units in RNNs. This often results in complex models with a large number of learning parameters. In this paper, a compact model is proposed which offers both representation learning and temporal inference of class variables by rolling Restricted Boltzmann Machines (RBMs) and class variables over time. We address the key issue of intractability in this variant of RBMs by optimising a conditional distribution, instead of a joint distribution. Experiments reported in the paper on melody modelling and optical character recognition show that the proposed model can outperform the state-of-the-art. Also, the experimental results on optical character recognition, part-of-speech tagging and text chunking demonstrate that our model is comparable to recurrent neural networks with complex memory gates while requiring far fewer parameters.

1 Introduction

Modelling sequences is an important research topic with various applications in audio/music informatics, natural language processing and computer vision. While some work focus on synthesising time-series events [22, 19, 20], classification with sequential data also receives much attention [1, 15, 3]. For such classification problems, there are two main areas of interest. The first is known as *sequence modelling* which is concerned with predicting the next event in time given previous events. The second is *sequence labelling*, i.e. how to label a sequence of events, with a wide range of applications such as optical character recognition (OCR) and part-of-speech (POS) tagging.

The classification problems mentioned above are very attractive to the use of dynamic Bayesian models such as Hidden Markov Models (HMMs) [17] and Conditional Random Fields (CRFs) [12]. An advantage of such models is the ability to learn relationships between sequence labels, which is useful for temporal reasoning. Recent research has seen an increasing interest in recurrent neural networks (RNNs) for sequence modelling. Differently from dynamic Bayesian models, RNNs assume that the class labels in a sequence are independent, given the sequence inputs. This makes inference easier, but with the sacrifice of valuable information: the temporal dependencies between sequence

labels. A main advantage of RNNs is the ability to learn temporal representations from data using recurrent hidden layers. A major difficulty in this process is learning from long sequences using back-propagation through time [10]. This has been addressed mainly by the adoption of complex gates such as in the Long Short Term Memory (LSTM) networks and in networks with Gated Recurrent Units (GRUs) [6]. There are also many attempts to combine the advantages of representation learning with dynamic inference, mostly by integrating a dynamic Bayesian model on top of a deep neural network [9].

In this paper, we propose a novel and compact model based on Restricted Boltzmann Machines (RBMs), which we call Sequence Classification RBMs (SCRBM), to support representation learning and dynamic inference on the classification of sequences. SCRBM is constructed by rolling RBMs with their class nodes over time. Each RBM at time t has a layer of visible units (X^t) and a layer of hidden units (H^t). Together with class nodes Y^t denoting the labels at time t , they form a model representing a distribution: $p(y^{1:T}, \mathbf{x}^{1:T}, \mathbf{h}^{1:T})$ (shorthand for $p(Y^1 = y^1, \dots, Y^T = y^T, X^1 = \mathbf{x}^1, \dots, X^T = \mathbf{x}^T, H^1 = \mathbf{h}^1, \dots, H^T = \mathbf{h}^T)$). When it comes to inference, there are two questions to answer: first, in order to compute gradients, one needs to infer the hidden states given the state of the input layer and the class labels from the distribution $p(H^{1:T} | \mathbf{x}^{1:T}, y^{1:T})$. This can be done using variational methods, i.e. treating a hidden unit as a mean-field, similarly to [20]. Second, to predict the state of the class labels given the state of the input layer, in the best case, one can search for the most probable labels from the conditional distribution, i.e. solving $\arg \max_{Y^{1:T}} p(Y^{1:T} | \mathbf{x}^{1:T})$. This may take polynomial time similarly to HMMs and CRFs. Instead, in this paper, we introduce an inference algorithm which only takes linear time. We employ the forward inference in RNNs with the modification that the hidden state at time t is dependent on the previous prediction (class labels at $t - 1$). By association, the state of the class labels at any time point is dependent on the prediction of the labels at the previous time point. For learning, we maximise the log-likelihood of the training data. Although it has been shown that learning a local RBM with labels is tractable [13, 18], learning the entire sequence is not. This is because it is not possible to marginalise out a sequence of hidden units, which is easy in the case of a single discriminative RBM. Also, computation becomes expensive due to the exponential growth of possible assignments to the sequence of labels. To solve this problem, we use the mean-field technique mentioned above to factorise the conditional probability of a sequence into the product of probabilities of local discriminative RBMs, as detailed in the next section.

Experiments have shown that SCRBM outperforms the state-of-the-art on two benchmark datasets in melody modelling and optical character recognition. Furthermore, we perform extensive evaluations of SCRBM against standard sequence models such as HMMs, CRFs and RNNs. Finally, we also compare SCRBM with RNNs with complex memory gates including GRUs [6] and LSTMs [10] on the MIT OCR, Penn-Tree bank POS tagging, and Conll2000 Chunking problems. The results show that the performance of the SCRBM model is better than that of standard HMMs and CRFs, and is comparable to that of GRUs and LSTMs while at the same time being considerably more compact with far fewer parameters.

The remainder of the paper is organised as follows. In the next section, we present the theory of *sequence classification RBMs*. Then, we discuss the related literature which inspired this work. In section 4, we present the empirical evaluations showing the effectiveness of SCRBM. The last section concludes the paper and discusses future work.

2 The SCRBM

2.1 Model

The Sequence Classification Restricted Boltzmann Machine (SCRBM) is constructed by rolling RBMs with class labels over time. The model defines a probability distribution:

$$p(y^{1:T}, \mathbf{x}^{1:T}, \mathbf{h}^{1:T}) = \prod_{t=1}^T p(y^t, \mathbf{x}^t, \mathbf{h}^t | \mathbf{h}^{t-1}) \quad (1)$$

where $\mathbf{x}^{1:T}$, $\mathbf{h}^{1:T}$ are time-series of the visible and hidden states; $y^{1:T}$ is the class-label sequence; \mathbf{h}^0 are the biases of the hidden units.

The main problem of this model, as highlighted in [19], is that inference is intractable. This, however, can be solved by adding recurrent connections, as done for the Recurrent Temporal Restricted

Boltzmann Machine (RTRBM) [20]. In RTRBMs, class labels are not included. In SCRBM, the local distribution at time t : $p(y^t, \mathbf{x}^t, \mathbf{h}^t | \mathbf{h}^{t-1})$ is replaced by:

$$p(y^t, \mathbf{x}^t, \mathbf{h}^t | \hat{\mathbf{h}}^{t-1}) = \frac{\exp(-E_\theta(y^t, \mathbf{x}^t, \mathbf{h}^t; \hat{\mathbf{h}}^{t-1}))}{\sum_{y', \mathbf{x}', \mathbf{h}'} \exp(-E_\theta(y', \mathbf{x}', \mathbf{h}'; \hat{\mathbf{h}}^{t-1}))} \quad (2)$$

where $\hat{\mathbf{h}}^{t-1}$ is the expected values of the hidden units at $t - 1$:

$$\hat{\mathbf{h}}^{t-1} = \mathbb{E}[\mathbf{H}^{t-1} | \mathbf{x}^{1:t-1}, y^{1:t-1}] \quad (3)$$

with local energy function:

$$E_\theta(y^t, \mathbf{x}^t, \mathbf{h}^t; \hat{\mathbf{h}}^{t-1}) = -[(\mathbf{x}^t)^\top \mathbf{W} + \mathbf{u}_{y^t}^\top + (\hat{\mathbf{h}}^{t-1})^\top \mathbf{W}_{hh}] \mathbf{h}^t - \mathbf{a}^\top \mathbf{x}^t - b_{y^t} - \mathbf{c}^\top \mathbf{h}^t \quad (4)$$

which is characterised by the parameters: $\theta = \{\mathbf{W}, \mathbf{W}_{hh}, \mathbf{U}, \mathbf{a}, \mathbf{b}, \mathbf{c}\}$. Here, \mathbf{W} is the weight matrix between visible units and hidden units; \mathbf{W}_{hh} is the recurrent/temporal connection weights of the hidden units; \mathbf{U} is the weight matrix between the class units (represented by one-hot vectors); \mathbf{a}, \mathbf{b} and \mathbf{c} are the biases of the visible units, hidden units and class units respectively; \mathbf{u}_{y^t} is the y^{th} column vector of \mathbf{U} . This set of parameters can be reduced by omitting the biases of the visible units because it will be cancelled out in the conditional distribution calculation, as shown in the next section.

2.2 Inference

As mentioned earlier, inference of units in the hidden layer given the inputs and labels is easy, as in [20], where each hidden unit can be treated as a mean-field, as follows:

$$\hat{\mathbf{h}}^t = \sigma(\mathbf{W}^\top \mathbf{x}^t + \mathbf{u}_{y^t} + \mathbf{W}_{hh}^\top \mathbf{h}^{t-1} + \mathbf{c}^t) \quad (5)$$

For classification with SCRBM one would like to search for the most probable assignment of $Y^{1:T}$ given the inputs $x^{1:T}$. Technically, this should be done in a similar way as in HMMs and CRFs. However, in the best case this will take polynomial time. In this paper, we follow the inference mechanism of RNNs which only needs linear time. Note that differently from RNNs, in SCRBM one needs to infer the state of the hidden layer while, at the same time, performing prediction. As discussed earlier, once the class labels are known, one can infer the state of the hidden layer straightforwardly. We now show that inference of class labels is also easy given the state of the hidden layer. In particular, from the mean-field values of the hidden layer in the previous time step we can infer the class labels using the following conditional distribution:

$$p(y^t | \mathbf{x}^t, \hat{\mathbf{h}}^{t-1}) = \frac{\exp(-\mathcal{F}(\mathbf{x}^t, y^t, \hat{\mathbf{h}}^{t-1}))}{\sum_{y'} \exp(-\mathcal{F}(\mathbf{x}^t, y', \hat{\mathbf{h}}^{t-1}))} \quad (6)$$

with free energy:

$$\mathcal{F}(\mathbf{x}^t, y, \hat{\mathbf{h}}^t) = -b_{y^t} - \sum_j \log(1 + \exp(\mathbf{w}_j^\top \mathbf{x}^t + u_{yj} + c_j)) \quad (7)$$

where \mathbf{w}_j is the j^{th} column of the weight matrix \mathbf{W} between the hidden units and the units of the visible layer corresponding to the inputs, and u_{yj} is an element of the weight matrix U between the hidden units and the units in the visible layer corresponding to the class labels. Here, the visible biases \mathbf{a} have been cancelled out which makes the number of parameters equivalent to that of a standard RNN with the same number of hidden units, i.e. $\theta = \{\mathbf{W}, \mathbf{W}_{hh}, \mathbf{U}, \mathbf{b}, \mathbf{c}\}$.

From (6) we can predict the value of the class labels. Once \mathbf{y}^t is known we use it to infer the mean-field values $\hat{\mathbf{h}}^t$ as in (5). Let us put this in a specific context, starting from $t = 1$: the conditional distribution $p(y^1 | \mathbf{x}^1, \hat{\mathbf{h}}^0)$ can be computed exactly by marginalising out the hidden variable \mathbf{h}^1 while having $\hat{\mathbf{h}}^0$ as parameters. In order to calculate from the prediction of the previous step, the values of the current step we do not use the predicted value of y^1 , instead we use the distribution to infer the

Algorithm 1 Inference with SCRBM

Data: Input: $\mathbf{x}^{(1:T)}$ **Result:** Output: $\mathbf{y}^{(1:T)}$ **for** $t = 1 : T$ **do** set $\hat{\mathbf{y}}^{(t)} = p(\mathbf{y}^{(t)} | \mathbf{x}^{(t)}, \hat{\mathbf{h}}^{(t-1)})$ set $y^{(t)} = \arg \max_k \hat{y}_k^{(t)}$ set $\hat{\mathbf{h}}^{(t)} = \sigma(\mathbf{W}\mathbf{x}^{(t)} + \mathbf{U}\hat{\mathbf{y}}^{(t)} + \mathbf{W}_{hh}\mathbf{h}^{t-1} + \mathbf{c}^{(t)})$ **end**

mean-field value $\hat{\mathbf{h}}^1$. This value is then passed to the next prediction step, and so on. The details of the SCRBM's inference algorithm are given in Algorithm 1.

Although this algorithm does not use dynamic programming as in the Viterbi algorithms for HMMs and CRFs, it should still capture the dependencies between class variables over time through the inference of hidden units using the expected values of the class units.

2.3 Learning

In SCRBM, we are usually interested in learning the conditional distribution:

$$p(y^{1:T} | \mathbf{x}^{1:T}) = \frac{p(\mathbf{x}^{1:T}, y^{1:T})}{\sum_{y^{1:T}} p(\mathbf{x}^{1:T}, y^{1:T})} \quad (8)$$

However, it is difficult to marginalise out all hidden variables $\mathbf{h}^{1:T}$ to compute exactly this distribution. Moreover, the complexity of learning our model would increase exponentially with the length of the sequence, due to the need to sum over all possible classes at every time step. So, instead of computing the distribution directly we simplify it by marginalising out the hidden variable at each time t using the expectation of the hidden state at the previous time $t - 1$. Let us consider:

$$\begin{aligned} p(\mathbf{x}^{1:T}, y^{1:T}) &= \sum_{\mathbf{h}^{1:T}} p(y^{1:T}, \mathbf{x}^{1:T}, \mathbf{h}^{1:T}) \\ &= \sum_{\mathbf{h}^{1:T}} \prod_{t=1}^T p(y^t, \mathbf{x}^t, \mathbf{h}^t | \mathbf{h}^{t-1}) \end{aligned} \quad (9)$$

If we first compute the expectation of \mathbf{h}^{t-1} given the previous input states $\mathbf{x}^{1:t-1}$ and $y^{1:t-1}$, which is equivalent to minimising the total energy function of the SCRBM, then we have:

$$q(\mathbf{x}^{1:T}, y^{1:T}) = \prod_{t=1}^T p(y^t, \mathbf{x}^t | \hat{\mathbf{h}}^{t-1}) \quad (10)$$

One can see this as an expectation step to be followed by an optimisation step which maximises the log-likelihood of this simplified distribution:

$$\begin{aligned} q(y^{1:T} | \mathbf{x}^{1:T}) &= \frac{\prod_{t=1}^T p(y^t, \mathbf{x}^t | \hat{\mathbf{h}}^{t-1})}{\sum_{y^{1:T}} \prod_{t=1}^T p(y^t, \mathbf{x}^t | \hat{\mathbf{h}}^{t-1})} \\ &= \frac{\prod_{t=1}^T p(y^t, \mathbf{x}^t | \hat{\mathbf{h}}^{t-1})}{\prod_{t=1}^T \sum_{y^t} p(y^t, \mathbf{x}^t | \hat{\mathbf{h}}^{t-1})} \\ &= \prod_{t=1}^T p(y^t | \mathbf{x}^t, \hat{\mathbf{h}}^{t-1}) \end{aligned} \quad (11)$$

Since $p(y^t | \mathbf{x}^t, \hat{\mathbf{h}}^{t-1})$ is tractable as shown in (6), we can compute the above distribution exactly. Now, one can learn the model by maximising the log-likelihood function:

$$\ell = \sum_{y^{1:T}, \mathbf{x}^{1:T}} \sum_{t=1}^T \log p(y^t | \mathbf{x}^t, \hat{\mathbf{h}}^{t-1}). \quad (12)$$

Similarly to other time-series connectionist models, such as standard RNNs and RTRBMs, we train the model using back-propagation through time. The update of the model’s set of parameters, denoted by θ , is shown below.

$$\nabla\theta = \sum_{t=1}^T \left(\frac{\partial_{\theta} \log p(y^t | \mathbf{x}^t, \hat{\mathbf{h}}^{t-1})}{\partial\theta} + \frac{\partial_{\theta} \hat{\mathbf{h}}^t}{\partial\theta} \mathcal{O}^t \right) \quad (13)$$

where $\frac{\partial_{\theta} \hat{\mathbf{h}}^t}{\partial\theta} = \frac{\partial_{\theta} \sigma(\mathbf{W}^{\top} \mathbf{x}^t + \mathbf{u}_{y^t} + \mathbf{W}_{hh}^{\top} \hat{\mathbf{h}}^{t-1} + \mathbf{c}^t)}{\partial\theta}$ and $\frac{\partial_{\theta} \log p(y^t | \mathbf{x}^t, \hat{\mathbf{h}}^{t-1})}{\partial\theta}$ are local derivatives, where $\hat{\mathbf{h}}^{t-1}$ is a value, not a function of θ ; and for mathematical convenience,

$$\mathcal{O}^t = \mathbf{W}_{hh}^{\top} \hat{\mathbf{h}}^{t+1} (1 - \hat{\mathbf{h}}^{t+1}) \mathcal{O}^{t+1} + \frac{\partial \log p(y^{t+1} | \mathbf{x}^{t+1}, \hat{\mathbf{h}}^t)}{\partial \hat{\mathbf{h}}^t} \quad (14)$$

3 Related work

Although SCRBM is a sequence learning connectionist network having the same set of parameters as a recurrent neural network, they are much different in terms of structure, learning and inference, as showed above. Modelling sequence data with RBMs has been studied previously [19, 20]. However, as in generative models, they are not easy to apply to classification tasks. The key problem is that the exact gradient can not be computed analytically which requires the use of approximation algorithms. In contrast, SCRBM is a discriminative model whose log-likelihood is tractable. Besides such tractability, another motivation for having a discriminatively learned variant of RTRBMs is the desire for better classification performance. It has been shown in [14] that with sufficient training examples, discriminative learning tends to do better on the task it is optimized for than its generative counterpart on the same model. This has been confirmed by the findings in [13]. In more recent work on the Recurrent Temporal Discriminative Restricted Boltzmann Machines (RTDRBM) [4], a generalised version of the RTRBM has been tailored with discriminative inference and learning specifically for melody modelling. Here, we make that work, otherwise focused on such application, applicable to other sequence learning tasks, especially sequence labelling as seen in the following section. Finally, RTDRBMs require more parameters than SCRBM because it generalises RTRBMs by including connections from previous hidden layers to the current class labels. Another, even more recent related model is the Dynamic Boltzmann Machine [7] which supports online learning for sequence prediction, mainly applied to regression.

4 Experiments

4.1 Music Modelling

In the task of music prediction addressed here, one seeks to predict the next pitch in a sequence which is a time series representation of the music. We use the benchmark from [4] which consists of 8 folk and chorale melody datasets. For the evaluation we use cross-entropy to measure the mean divergence between the entropy calculated from the predicted distribution and the correct pitch. Given a test data \mathcal{D}_{test} , the cross-entropy can be computed over all the events belonging to different sequences as follows:

$$H_c(p_{mod}, \mathcal{D}_{test}) = \frac{-\sum_{s \in \mathcal{D}_{test}} \sum_{t=1}^{T_s} \log_2 p_{mod}(s^t | s^{1:t-1})}{\sum_{s \in \mathcal{D}_{test}} T_s} \quad (15)$$

where p_{mod} is the probability assigned by the model to the pitch of the next musical event s^t in the melody, $s \in \mathcal{D}_{test}$, given its preceding pitches, and T_s is the length of s . To apply SCRBM to this task we follow the standard encoding in [1] where $y^t = s^t$ and $x^t = s^{t-1}$.

We compare the performance of SCRBM with the state-of-the-art approaches on this dataset which include RTRBMs and five other connectionist models: standard feedforward neural networks (FNN) [2], Restricted Boltzmann Machines (RBM) [4], Discriminative RBMs (DRBM) [13], a bounded order n -gram model $C2I$, and an unbounded order n -gram model C^*I , as well as standard recurrent neural networks (RNNs) and the generalised RTRBM for melody prediction (RTD-RBM) [4].

For fairness of comparison, we partition the data exactly as done in previous work [16, 5]. In particular, a resampling method is employed for evaluation of the 8 datasets [8]. For model selection,

Model	Context	Cross Entropy (bits)
SCRBM	n/a	2.712
RTDRBM	n/a	2.741
RTRBM	n/a	2.764
RNN $_{n_h=100, h_{act}=TanSigU}$	n/a	2.787
RBM $_{n_h=100, h_{act}=LogSigU}$	8	2.799
DRBM $_{n_h=25, h_{act}=LogSigU}$	5	2.819
FNN $_{n_h=200, h_{act}=ReLU}$	7	2.819
n -gram (u)	n/a	2.878
n -gram (b)	2	2.957

Table 1: Test-set performance comparison on pitch prediction time series; n_h denotes the number of hidden neurons and h_{act} the activation function used, either standard sigmoid, hyperbolic tangent or rectified linear functions. The context is the size of the window in the time-series for the models that require this to be predefined.

we use a small part of the data (5%) in each case as validation set. In order to determine the best hyperparameters for a model a grid search is carried out. Such hyperparameters include the learning rate, regularization parameters and number of hidden units.

Table 1 shows the best predictive performance of each model averaged across all 8 datasets. It indicates a consistent improvement in best-case performance from the n -gram models, the non-recurrent neural networks, and then the recurrent neural network models, with the SCRBM outperforming all of the others. The hyperparameter with the most influence in the result was the size of the hidden layer. It was found both in the case of the RTRBM and the SCRBM that a hidden layer size of 100 units resulted in the best predictive performance.

4.2 Optical Character Recognition

The MIT OCR dataset¹ is a widely used benchmark for evaluating sequence labelling algorithms [21]. We use two popular partitions from [15] and [9, 3]. In the former, called here "ms" for model selection, the data is partitioned into 10 groups, each consisting of a training, validation and test set. We select models based on performance on the validation sets and report their average accuracy on the test sets. In the latter, here called "cv" for cross-validation, the data is divided into 10 folds in the usual way but without model selection.

Each model is expected to predict the correct label corresponding to the image of a character as it is drawn. All the models are evaluated using the averaged loss function per sequence $E(y, y^*)$, given by:

$$E(y, y^*) = \frac{1}{N} \sum_{i=1}^N \left[\frac{1}{L_i} \sum_{j=1}^{L_i} \mathcal{I}((y_i)_j \neq (y^*)_j) \right] \quad (16)$$

where y and y^* are the predicted and the true sequence respectively, N is the total number of test examples, L_i is the length of the i^{th} sequence, and \mathcal{I} is the 0 – 1 loss function.

4.2.1 Baselines

We compare the performance of SCRBM on the sequence labelling task with the baseline models: SVM-*multiclass*, SVM-*struct*, Max-margin Markov network (M3N), Averaged Perceptron, Search-based structured prediction model (SEARN), CRF, HMM, LogitBoost, TreeCRF, RTDRBM²; and the state-of-the-arts:

- Structured learning ensemble (SLE) [15]: An optimised ensemble of 7 effective models: SVM-*multiclass*, SVM-*struct*, M3N, Perceptron, SEARN, CRF, HMM.
- Neural CRF [9]: A combination of CRF and deep networks.
- Gradient boosting CRF (GBCRF) [3]: CRF trained by a novel gradient boosting algorithm.

¹<http://www.seas.upenn.edu/~taskar/ocr/>

²This version of RTDRBM is modified by using our inference algorithm to work with sequence labelling task.

Model	ms	cv
SCRBM	14.81	04.03
RTDRBM	15.46	-
Neural CRF ^{CML}	-	04.44
Neural CRF ^{LM}	-	04.56
SLE	20.58	-
GBCRF	-	04.64
TreeCRF	-	06.99
LogitBoost	-	09.67
RNN	22.92	13.30
M3N	25.08	13.46
Perceptron	26.40	-
SEARN	27.02	-
SVM _{multiclass}	28.54	-
SVM _{struct}	21.16	-
HMM	23.70	-
CRF	32.30	14.20

Table 2: Test errors (%) of different models in the evaluation.

4.2.2 Results

For the model selection ("ms") partition, in order to determine the best model for the task, a grid search was carried out. For fair comparison, we use the best results of the baselines reported in [15] on the same dataset. For SCRBM, RTDRBM and RNN the hyper-parameters include: the learning rate and the number of hidden unit. We also use early stopping in which the performance of the model on a validation set was determined after every 10 epochs. If the performance in one of these checks happened to be worse than the previous best one for 5 consecutive checks the training was stopped. In the sequence learning, stochastic gradient descent with backpropagation through time was used.

For the cross-validation ("cv") partition, since model selection is not applicable, we fix our SCRBM with 1000 hidden units and use adaptive learning rate (Adam [11]) started from 0.001 to get rid of the search for hyper-parameters.

Table 2 shows the comparative performance of SCRBM. It is evident that SCRBM clearly outperforms the other models. It is worth noting that the performance of our model is much better than SLE, Neural-CRF, and GBCRF.

4.3 SCRBM v.s LSTM & GRU

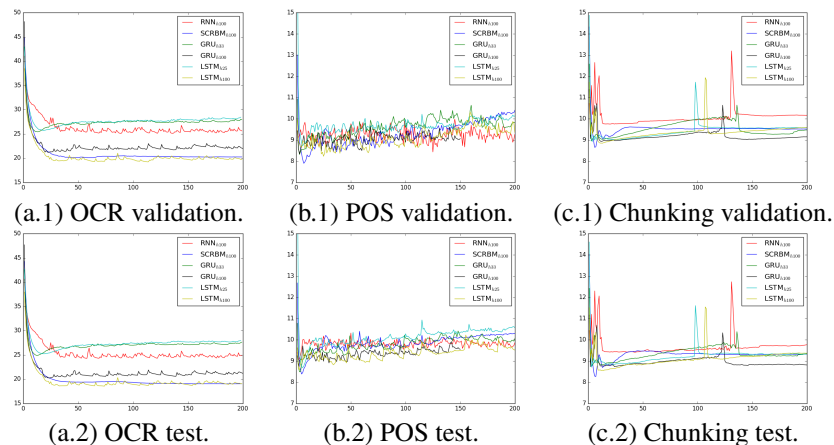


Figure 1: Prediction errors. Averaged negative log-likelihood. x-axis: number of epoch. y-axis: prediction errors.

For completeness, we now compare our model with different recurrent neural networks, especially with ones having complex memory gates in hidden units such as *gated recurrent units* (GRU) and




Figure 2: Averaged negative log-likelihood. x-axis: number of epoch,y axis: average negative log-likelihood.

long short term memory (LSTM). We use a traditional RNN with \tanh hidden units (denoted as RNN), HMMs, and CRFs as baselines.

It is worth noting that for a visible layer with M units and a hidden layer with N units, the number of weights of RNN and RTDRM are the same while GRU and LSTM have $2N(M + N + 1)$ and $3N(M + N + 1)$ more weights than the RNN and SCRBM, respectively. In general, if the number of labels is small then SCRBM is ≈ 3 times smaller than GRU and ≈ 4 times smaller than LSTM.

	OCR	POS ₅₀₀	POS ₁₀₀₀	POS ₂₀₀₀	Chunking
HMM	23.700	23.460	19.950	17.960	-
CRF	32.300	16.530	12.510	09.840	-
RNN	22.921	12.220	10.480	09.207	09.303
GRU	15.853	11.663	09.892	09.011	08.759
LSTM	13.329	12.012	10.168	09.273	08.561
SCRBM	14.808	12.119	10.061	09.178	08.286

Table 3: Test errors of HMMs, CRFs, RNNs: recurrent neural network with \tanh hidden units,GRU: gated recurrent units, LSTM: long short term memory. Model sizes and other hyperparameters are optimized in a grid search

We test both GRU and LSTM on the OCR dataset with "ms" partition above. Furthermore, we also perform experiments with POS tagging dataset from Penn Treebank³. The data is partitioned into different training sets of 500, 1000, and 2000 samples. The models are selected by using a held-out 10% samples of training sets, which are then evaluated in a test set of ~ 1600 sentences. The challenge of this data is that the lexical features are very large, around 450,000. The final dataset is used for evaluation here is Conll200 chunking⁴. Since the original data does not provide a validation set, we take 2.5% of the training samples for training, the rest is for validation. The original test set is used for testing. The results are shown in Table 3. We can see that our SCRBM outperforms the standard models: RNNs, CRFs, and HMMs. In comparison to RNNs with complex gates, SCRBM has comparable performance as GRU and LSTM. In particular, SCRBM is better than GRU in the OCR dataset and better than LSTM in the POS dataset. It is more effective than both GRU and LSTM in the Conll2000 chunking. We should note that even though the LSTM seems to beat our SCRBM in OCR dataset it needs 2000 hidden units for that best result while SCRBM only requires 100 hidden units.

Finally we compare the effectiveness of SCRBM, RNN, GRU and LSTM having the same number of parameters. We use 100 hidden units for both SCRBM and RNN, 33 hidden units for GRU and 25 hidden units for LSTM. We even include GRU and LSTM with 100 hidden units in the test to give a more comprehensive comparison. We evaluate the models on the OCR, POS tagging with 2000 training sentences, and Conll 2000 chunking. The evaluation metrics are prediction error rate and the average negative log-likelihood on validation and test sets. All models are trained using Adam

³<http://www.cis.upenn.edu/treebank>

⁴<https://www.clips.uantwerpen.be/conll2000/chunking/>

with learning rate 0.001 which, as we found, is generally good for all models on three datasets. Both Figure 1 and Figure 2 show that SCRBM generalises better than RNNs, GRU and LSTM which have the same size or same number of hidden units. For prediction errors, Figures 1b.1, 1b.2, 1c.1, 1c.2 demonstrate that in POS and Chunking datasets all models converge very quickly after few epochs with the SCRBM reaches the lowest error. We can see from Figure 1a.1 In OCR dataset, SCRBM performs similarly as LSTM with the same number of 100 hidden units. In the case of the average negative log-likelihood, its value is inversely proportional to the cost function we want to maximise in the training of all models. Therefore, as we can see in Figures 2a.1, 2a.2, 2b.1, 2b.2, 2c.1, 2c.2 that the SCRBM generalises better than other models as it achieves the lowest negative loglikelihoods on the validation and test sets. This, along with the Table 3, explains why in some cases GRUs and LSTMs may be better but only when they have more hidden units. Also, the Figures 1, 2 make the comparison more complete by answering a question on our claim about compactness, e.g what if GRU and LSTM are better with the same number of parameters as SCRBM.

5 Conclusions

In this paper, we propose a novel model for classification with sequences by rolling restricted Boltzmann machines with class layer over time. We name it as *Sequence classification RBMs*. The key characteristic of this SCRBM is that it can support both representation learning and temporal reasoning. Also, the proposed model is very compact, its number of parameters is equivalent to a basic recurrent neural network having the same number of hidden units. Inference with SCRBM is done by performing prediction of the class and computing mean-field values of hidden layer at each time t , consecutively. We train SCRBM by computing the expectation of hidden units at time $t - 1$ for each element of the conditional distribution which makes learning tractable. In the experiments, we evaluate the effectiveness of SCRBM in sequence modelling with a benchmark of 8 melody datasets and in sequence labelling with two different partitions of OCR dataset. In both cases, SCRBM outperforms the state-of-the-arts. We also show the simplicity advantage of SCRBM over GRU and LSTM, two types of recurrent neural networks with complex memory gates in hidden units.

References

- [1] Yoshua Bengio, Rejean Ducharme, Pascal Vincent, and Christian Jauvin. A Neural Probabilistic Language Model. *Journal of Machine Learning Research*, 3:1137–1155, 2003.
- [2] John S Bridle. Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. In *Neurocomputing*, pages 227–236. 1990.
- [3] Tianqi Chen, Sameer Singh, Ben Taskar, and Carlos Guestrin. Efficient Second-Order Gradient Boosting for Conditional Random Fields. In *International Conference on Artificial Intelligence and Statistics*, pages 147–155, 2015.
- [4] S. Cherla, S. N. Tran, A. d. Garcez, and T. Weyde. Discriminative learning and inference in the recurrent temporal rbm for melody modelling. In *2015 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2015.
- [5] Srikanth Cherla, Tillman Weyde, Artur S d’Avila Garcez, and Marcus Pearce. A Distributed Model for Multiple-Viewpoint Melodic Prediction. In *International Society for Music Information Retrieval Conference*, pages 15–20, 2013.
- [6] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734, 2014.
- [7] Sakyasingha Dasgupta and Takayuki Osogami. Nonlinear dynamic boltzmann machines for time-series prediction. In *AAAI*, 2017.
- [8] Thomas G Dietterich. Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms. *Neural computation*, 10(7):1895–1923, 1998.
- [9] Trinh-Minh-Tri Do and Thierry Artieres. Neural conditional random fields. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 177–184, 2010.

- [10] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997.
- [11] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- [12] John Lafferty, Andrew McCallum, and Fernando CN Pereira. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *International Conference on Machine Learning*, pages 282–289, 2001.
- [13] Hugo Larochelle and Yoshua Bengio. Classification using Discriminative Restricted Boltzmann Machines. In *International Conference on Machine Learning*, pages 536–543, 2008.
- [14] Andrew Y Ng and Michael I Jordan. On Discriminative vs. Generative Classifiers: A Comparison of Logistic Regression and Naive Bayes. In *Advances in Neural Information Processing Systems*, pages 841–848, 2001.
- [15] Nam Nguyen and Yunsong Guo. Comparisons of Sequence Labeling Algorithms and Extensions. In *International Conference on Machine Learning*, pages 681–688. ACM, 2007.
- [16] Marcus Pearce and Geraint Wiggins. Improved Methods for Statistical Modelling of Monophonic Music. *Journal of New Music Research*, 33(4):367–385, 2004.
- [17] Lawrence R. Rabiner. Readings in speech recognition. chapter A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition, pages 267–296. San Francisco, USA, 1990.
- [18] Cherla Srikanth, Tran Son, N., Weyde Tillman, and Garcez Artur. Generalising the discriminative restricted boltzmann machines. In *International Conference on Artificial Neural Networks*, 2017.
- [19] Ilya Sutskever and Geoffrey E. Hinton. Learning multilevel distributed representations for high-dimensional sequences. In *AISTATS*, volume 2 of *JMLR Proceedings*, pages 548–555, 2007.
- [20] Ilya Sutskever, Geoffrey E Hinton, and Graham W Taylor. The Recurrent Temporal Restricted Boltzmann Machine. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1601–1608, 2009.
- [21] Ben Taskar, Carlos Guestrin, and Daphne Koller. Max-margin Markov Networks. *Advances in neural information processing systems*, 16:25, 2004.
- [22] Graham W. Taylor, Geoffrey E Hinton, and Sam T. Roweis. Modeling human motion using binary latent variables. In *Advances in Neural Information Processing Systems 19*, pages 1345–1352. 2007.