
Factored Temporal Sigmoid Belief Networks for Sequence Learning

Jiaming Song[†]
Zhe Gan[‡]
Lawrence Carin[‡]

JIAMING.TSONG@GMAIL.COM
ZHE.GAN@DUKE.EDU
LCARIN@DUKE.EDU

[†]Department of Computer Science and Technology, Tsinghua University, Beijing, 100084, China

[‡]Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA

Abstract

Deep conditional generative models are developed to simultaneously learn the temporal dependencies of multiple sequences. The model is designed by introducing a three-way weight tensor to capture the multiplicative interactions between *side information* and sequences. The proposed model builds on the Temporal Sigmoid Belief Network (TSBN), a sequential stack of Sigmoid Belief Networks (SBNs). The transition matrices are further factored to reduce the number of parameters and improve generalization. When side information is not available, a general framework for semi-supervised learning based on the proposed model is constituted, allowing robust sequence classification. Experimental results show that the proposed approach achieves state-of-the-art predictive and classification performance on sequential data, and has the capacity to synthesize sequences, with controlled style transitioning and blending.

1. Introduction

The Restricted Boltzmann Machine (RBM) is a well-known undirected generative model with state-of-the-art performance on various problems. It serves as a building block for many deep generative models, such as the Deep Belief Network (DBN) (Hinton et al., 2006) and the Deep Boltzmann Machine (DBM) (Salakhutdinov & Hinton, 2009). Other variants of the RBM have been used for modeling discrete time series data, such as human motion (Taylor et al., 2006), videos (Sutskever et al., 2009), and weather prediction (Mittelman et al., 2014). Among these variants is the Conditional Restricted Boltzmann Machine

(CRBM) (Taylor et al., 2006), where the hidden and visible states at the current time step are dependent on directed connections from observations at the last few time steps. Most RBM-based models use Contrastive Divergence (CD) (Hinton, 2002) for efficient learning.

Recently, there has been a surging interest in deep *directed* generative models, with applications in both static and dynamic data. In particular, advances in variational methods (Mnih & Gregor, 2014; Kingma & Welling, 2014; Rezende et al., 2014) have yielded scalable and efficient learning and inference for such models, avoiding poor inference speed caused by the “explaining away” effect (Hinton et al., 2006). One directed graphical model that is related to the RBM is the Sigmoid Belief Network (SBN) (Neal, 1992).

The Temporal Sigmoid Belief Network (TSBN) (Gan et al., 2015c) is a fully directed generative model for discretely sampled time-series data, defined by a sequential stack of SBNs. The hidden state for each SBN is inherited from the states of the previous SBNs in the sequence. The TSBN can be regarded as a generalization of Hidden Markov Models (HMM), with compact hidden state representations, or as a generalization of Linear Dynamical Systems (LDS), characterized by non-linear dynamics. The model can be utilized to analyze many kinds of data, *e.g.*, binary, real-valued, and counts, and has demonstrated state-of-the-art performance in many tasks.

However, the TSBN exhibits certain limitations, in that it does not discriminate different types of sequences when training, nor does it utilize (often available) side information during generation. For example, in the case of modeling human motion, the TSBN does not regard “walking” and “running” as different motions, and the style of generated motions is not controlled. To allow for conditional generation, we first propose a straightforward modification of the TSBN, introducing three-way tensors for weight transitions, from which weight matrices for the TSBN are extracted according to the side information provided. We then adopt ideas from Taylor & Hinton (2009), where the

three-way weight tensor is factored into the multiplication of matrices, effectively reducing the model parameters. In our case, factoring is not directly imposed on the tensor parameters but over the style-dependent transition matrices, which provides a more-compact representation of the transitions, where sequences with different attributes are modeled by a shared set of parameters. The model is able to capture the subtle similarities across these sequences, while still preserving discriminative attributes. Experiments show that the factored model yields better prediction performance than its non-factored counterpart.

For cases where side information is not available for the entire training set, we propose a framework for semi-supervised learning, inspired by Kingma et al. (2014), to infer that information from observations (*e.g.*, by using a classifier to infer the motion style), while simultaneously training a generative model.

The principal contributions of this paper are as follows: (i) A fully *directed* deep generative model for sequential data is developed, that permits fast conditional generation of synthetic data, and controlled transitioning and blending of different styles. (ii) A new factoring method is proposed, that reduces the number of parameters and improves performance across multiple sequences. (iii) A general framework for semi-supervised learning is constituted, allowing robust sequence classification. (iv) A new recognition model with factored parameters is utilized to perform scalable learning and inference.

2. Model Formulation

2.1. Temporal Sigmoid Belief Network

The Temporal Sigmoid Belief Network (TSBN) (Gan et al., 2015c) models uniformly sampled time series data of length T through a sequence of SBNs, such that at any given time step SBN biases depend on the states of SBNs in the previous time steps. Assume that each observation is real-valued, and the t -th time step is denoted $\mathbf{v}_t \in \mathbb{R}^M$. The TSBN defines the joint probability of visible data \mathbf{V} and hidden states \mathbf{H} as

$$p_{\theta}(\mathbf{V}, \mathbf{H}) = \prod_{t=1}^T p(\mathbf{h}_t | \mathbf{h}_{t-1}, \mathbf{v}_{t-1}) \cdot p(\mathbf{v}_t | \mathbf{h}_t, \mathbf{v}_{t-1}) \quad (1)$$

where $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_T]$, $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_T]$, and $\mathbf{h}_t \in \{0, 1\}^J$ is the hidden state corresponding to time step t . Each conditional distribution in (1) is expressed as

$$p(h_{jt} = 1 | \mathbf{h}_{t-1}, \mathbf{v}_{t-1}) = \sigma(\tilde{h}_{jt}) \quad (2)$$

$$p(\mathbf{v}_t | \mathbf{h}_t, \mathbf{v}_{t-1}) = \mathcal{N}(\boldsymbol{\mu}_t, \text{diag}(\boldsymbol{\sigma}_t^2)) \quad (3)$$

$$\tilde{\mathbf{h}}_t = \mathbf{W}_1 \mathbf{h}_{t-1} + \mathbf{W}_3 \mathbf{v}_{t-1} + \mathbf{b} \quad (4)$$

$$\boldsymbol{\mu}_t = \mathbf{W}_2 \mathbf{h}_t + \mathbf{W}_4 \mathbf{v}_{t-1} + \mathbf{c} \quad (5)$$

$$\log \boldsymbol{\sigma}_t^2 = \mathbf{W}'_2 \mathbf{h}_t + \mathbf{W}'_4 \mathbf{v}_{t-1} + \mathbf{c}' \quad (6)$$

where $\sigma(x) = 1/(1 + \exp(-x))$, and $\text{diag}(\mathbf{v})$ is the diagonal matrix whose diagonal entries are \mathbf{v} . \mathbf{h}_0 and \mathbf{v}_0 are defined as zero vectors. The model parameters, θ , are specified as $\mathbf{W}_1 \in \mathbb{R}^{J \times J}$, $\mathbf{W}_3 \in \mathbb{R}^{J \times M}$, $\{\mathbf{W}_2, \mathbf{W}'_2\} \in \mathbb{R}^{M \times J}$, $\{\mathbf{W}_4, \mathbf{W}'_4\} \in \mathbb{R}^{M \times M}$, $\mathbf{b} \in \mathbb{R}^J$ and $\{\mathbf{c}, \mathbf{c}'\} \in \mathbb{R}^M$.

The TSBN with $\mathbf{W}_3 = \mathbf{0}$ and $\{\mathbf{W}_4, \mathbf{W}'_4\} = \mathbf{0}$ is an HMM with an exponentially large state space. Specifically, in the HMM, each hidden state is represented as a one-hot-encoded vector of length J , while the hidden states of the TSBN are encoded by length- J binary vectors, which can express 2^J different states. Compared with the Temporal Restricted Boltzmann Machine (TRBM) (Sutskever & Hinton, 2007), the TSBN is fully directed, allowing fast sequence generation from the inferred model.

2.2. Conditional Temporal SBN

We consider modeling *multiple* styles of time series data by exploiting additional side information, which can appear in the form of “one-hot” encoded vectors in the case of style labels, or real-valued vectors in other situations (*e.g.*, the longitude and latitude information in the weather prediction task considered in Section 5.5). Let $\mathbf{y}_t \in \mathbb{R}^S$ denote the side information at time step t . The joint probability is now described as

$$p_{\theta}(\mathbf{V}, \mathbf{H} | \mathbf{Y}) = \prod_{t=1}^T p(\mathbf{h}_t | \mathbf{h}_{t-1}, \mathbf{v}_{t-1}, \mathbf{y}_t) \cdot p(\mathbf{v}_t | \mathbf{h}_t, \mathbf{v}_{t-1}, \mathbf{y}_t) \quad (7)$$

where $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_T]$. A straightforward method to incorporate side information is by modifying (4)-(6) to allow the weight parameters to be dependent on \mathbf{y}_t . Specifically,

$$\tilde{\mathbf{h}}_t = \mathbf{W}_1^{(y)} \mathbf{h}_{t-1} + \mathbf{W}_3^{(y)} \mathbf{v}_{t-1} + \mathbf{b}^{(y)} \quad (8)$$

$$\boldsymbol{\mu}_t = \mathbf{W}_2^{(y)} \mathbf{h}_t + \mathbf{W}_4^{(y)} \mathbf{v}_{t-1} + \mathbf{c}^{(y)} \quad (9)$$

$$\log \boldsymbol{\sigma}_t^2 = \mathbf{W}'_2^{(y)} \mathbf{h}_t + \mathbf{W}'_4^{(y)} \mathbf{v}_{t-1} + \mathbf{c}'^{(y)} \quad (10)$$

where $\mathbf{b}^{(y)} = \mathbf{B}\mathbf{y}_t$, $\mathbf{c}^{(y)} = \mathbf{C}\mathbf{y}_t$; for $i \in \{1, 2, 3, 4\}$, $\mathbf{W}_{ijk}^{(y)} = \sum_{s=1}^S \hat{\mathbf{W}}_{ijks} y_{st}$, and $\hat{\mathbf{W}}_i$ is a three-way tensor, shown in Figure 1 (left).¹ The model parameters are specified as $\hat{\mathbf{W}}_1 \in \mathbb{R}^{J \times J \times S}$, $\hat{\mathbf{W}}_2 \in \mathbb{R}^{M \times J \times S}$, $\hat{\mathbf{W}}_3 \in \mathbb{R}^{J \times M \times S}$, $\hat{\mathbf{W}}_4 \in \mathbb{R}^{M \times M \times S}$, $\mathbf{B} \in \mathbb{R}^{J \times S}$, and $\mathbf{C} \in \mathbb{R}^{M \times S}$. When the side information are one-hot encoded vectors, this is equivalent to training one TSBN for each style of sequences. We name this model the Conditional Temporal Sigmoid Belief Network (CTSBN).

¹ $\mathbf{W}'_2^{(y)}$ and $\mathbf{W}'_4^{(y)}$ are also similarly defined, and we omit them when discussing $\mathbf{W}_2^{(y)}$ and $\mathbf{W}_4^{(y)}$ in the following sections. $\mathbf{b}^{(y)}$, $\mathbf{c}^{(y)}$ and $\mathbf{W}_{ijk}^{(y)}$ are dependent on t , but we choose to omit t for simplicity.




Figure 1: Graphical illustration for non-factored weights (left) and factored weights (right) from h_{t-1} to h_t .




Figure 2: Generative model (left) and recognition model (right) of the FCTSBN. Red, dashed arrows indicate factored weights described in Figure 1 (right).

2.3. Factoring Weight Parameters

While the CTSBN enables conditional generation, it has several disadvantages: (i) the number of parameters is proportional to S (recall that this is the number of different styles for the time-series data), which is prohibitive for large S ; (ii) parameters for different attributes/styles are not shared (no sharing of “statistical strength”), therefore the model fails to capture the underlying regularities between different data, resulting in poor generalization.

To remedy these drawbacks, we factor the weight matrices $\mathbf{W}_i^{(y)}$ defined in (8)-(10), for $i \in \{1, 2, 3, 4\}$, as

$$\mathbf{W}^{(y)} = \mathbf{W}_a \cdot \text{diag}(\mathbf{W}_b \mathbf{y}_t) \cdot \mathbf{W}_c \quad (11)$$

Suppose $\mathbf{W}^{(y)} \in \mathbb{R}^{J \times M}$, then $\mathbf{W}_a \in \mathbb{R}^{J \times F}$, $\mathbf{W}_b \in \mathbb{R}^{F \times S}$ and $\mathbf{W}_c \in \mathbb{R}^{F \times M}$, where F is the number of factors. \mathbf{W}_a and \mathbf{W}_c are shared among different styles, which capture the input-to-factor and factor-to-output relationships, respectively; the diagonal term, $\text{diag}(\mathbf{W}_b \mathbf{y}_t)$, models the unique factor-to-factor relationship for each style. If the number of factors is comparable to that of other terms (*i.e.*, J and M), this reduces the number of parameters needed for all the $\mathbf{W}^{(y)}$ from $J \cdot M \cdot S$ to $(J + M + S) \cdot F$, or equivalently, $O(N^3)$ to $O(N^2)$. In practice, we factor all parameters except for \mathbf{W}_2 to improve generation by allowing a more flexible transition from h_t to v_t across multiple styles.² We call this model the Factored Conditional Temporal Sigmoid Belief Network (FCTSBN).

2.4. Deep Architecture with FCTSBN

The shallow model described above may be restrictive in terms of representational power. Therefore, we propose a

² \mathbf{W}_2 is independent to the order n introduced in Section 5; when n is large, the number of parameters in \mathbf{W}_2 is relatively small, thus factoring \mathbf{W}_2 does not provide much advantage for reducing parameters.




Figure 3: Generative model (left) and recognition model (right) of a deep FCTSBN with two layers. Red, dashed arrows indicate factored weights described in Figure 1 (right).

deep architecture by adding stochastic hidden layers. We consider a deep FCTSBN with hidden layers $h_t^{(\ell)}$ for $t = 1 \dots T$ and $\ell = 1 \dots L$, where we denote $h_t^{(0)} = v_t$ and $h_t^{(L+1)} = \mathbf{0}$ for conciseness. Each of the hidden layers $h_t^{(\ell)}$ contains stochastic binary hidden variables, which is generated by $p(h_t^{(\ell)}) = \prod_{j=1}^{J^{(\ell)}} p(h_{jt}^{(\ell)} | h_t^{(\ell+1)}, h_{t-1}^{(\ell)}, h_{t-1}^{(\ell-1)}, \mathbf{y}_t)$, where each conditional distribution is parameterized by a sigmoid function, as in (2). Learning and inference for the model with stochastic hidden layers is provided in Section 3.3. We choose not to consider deterministic layers here, as was considered in Gan et al. (2015c), since such complicates the gradient computation in the FCTSBN, while having similar empirical results compared with the usage of stochastic layers (we tried both in practice).

2.5. Semi-supervised Learning with FCTSBN

As expressed thus far, training the FCTSBN requires side information for each time step. In many applications (Le & Mikolov, 2014; Srivastava et al., 2015), however, unlabeled sequential data might be abundant, while obtaining the side information for all the data is expensive or sometimes impossible. We propose a framework for semi-supervised learning based on FCTSBN, with the capacity to simultaneously train a generative model and a classifier from labeled and unlabeled data.

Specifically, assume that the side information is generated from the prior distribution $p_\theta(\mathbf{Y}; \pi)$ parametrized by π (which can be a multinomial distribution if \mathbf{y}_t are labels), then the generative model can be described as

$$p_\theta(\mathbf{V}, \mathbf{H}, \mathbf{Y}) = p_\theta(\mathbf{Y}; \pi) \cdot p_\theta(\mathbf{V}, \mathbf{H} | \mathbf{Y}) \quad (12)$$

where $p_\theta(\mathbf{V}, \mathbf{H} | \mathbf{Y})$ is the conditional generative model described in (7). Details on training $p_\theta(\mathbf{V}, \mathbf{H}, \mathbf{Y})$ are discussed in Section 3.4.

3. Scalable Learning and Inference

The exact posterior over the hidden variables in (7) is intractable, and methods like Gibbs sampling and mean-field

variational Bayes (VB) inference, can be inefficient (particularly for inference at test time). We follow Gan et al. (2015c), and apply the Neural Variational Inference and Learning (NVIL) algorithm described in Mnih & Gregor (2014), allowing tractable and scalable parameter learning and inference by introducing a new recognition model.

3.1. Lower Bound Objective

We are interested in training the CTSBN and FCTSBN models, $p_{\theta}(\mathbf{V}, \mathbf{H}|\mathbf{Y})$, both of which come in the form of (7). Given an observation \mathbf{V} , we introduce a fixed-form distribution $q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y})$ with parameters ϕ , to approximate the true posterior $p(\mathbf{H}|\mathbf{V}, \mathbf{Y})$. According to the variational principle (Jordan et al., 1999), we construct a lower bound on the marginal log-likelihood with the following form

$$\mathcal{L}(\mathbf{V}|\mathbf{Y}, \theta, \phi) = \mathcal{J}(q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y}), p_{\theta}(\mathbf{V}, \mathbf{H}|\mathbf{Y})) \quad (13)$$

where $\mathcal{J}(q, p) = \mathbb{E}_q[\log p - \log q]$ is used to save space, and $q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y})$ is the recognition model. For both models, the recognition model is expressed as

$$q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y}) = \prod_{t=1}^T q(\mathbf{h}_t | \mathbf{h}_{t-1}, \mathbf{v}_t, \mathbf{v}_{t-1}, \mathbf{y}_t) \quad (14)$$

and each conditional distribution is specified as

$$q(h_{jt} = 1 | \mathbf{h}_{t-1}, \mathbf{v}_t, \mathbf{v}_{t-1}, \mathbf{y}_t) = \sigma(\hat{h}_{jt}) \quad (15)$$

$$\hat{\mathbf{h}}_t = \mathbf{U}_1^{(y)} \mathbf{h}_{t-1} + \mathbf{U}_2^{(y)} \mathbf{v}_t + \mathbf{U}_3^{(y)} \mathbf{v}_{t-1} + \mathbf{d}^{(y)} \quad (16)$$

For CTSBN, $\mathbf{U}_{ijk}^{(y)} = \sum_{l=1}^L \hat{\mathbf{U}}_{ijkl} y_{lt}$, where $\hat{\mathbf{U}}_i$ is a three-way tensor similar to $\hat{\mathbf{W}}_i$; for FCTSBN, the weight matrix $\mathbf{U}_i^{(y)}$ is factored as in (11). The recognition model introduced in (14), allows fast and scalable inference.

3.2. Parameter Learning

To optimize the lower bound described in (13), we apply Monte Carlo integration to approximate the expectations over q_{ϕ} , and stochastic gradient descent (SGD) for parameter optimization. The gradients of \mathcal{L} with respect to θ and ϕ are expressed as

$$\nabla_{\theta} \mathcal{L} = \mathbb{E}_{q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y})} [\nabla_{\theta} \log p_{\theta}(\mathbf{V}, \mathbf{H}|\mathbf{Y})] \quad (17)$$

$$\nabla_{\phi} \mathcal{L} = \mathbb{E}_{q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y})} [l_{\phi}(\mathbf{V}, \mathbf{H}, \mathbf{Y}) \cdot \nabla q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y})] \quad (18)$$

where $l_{\phi} = \log p_{\theta}(\mathbf{V}, \mathbf{H}|\mathbf{Y}) - \log q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y})$, termed the *learning signal* for the recognition parameters ϕ .

The expectation of l_{ϕ} is exactly the lower bound (13). The recognition model gradient estimator, however, can be very noisy, since the estimated learning signal is potentially large. According to Mnih & Gregor (2014), we apply two variance-reduction techniques. The first technique

is applied by centering the learning signal by subtracting the data-dependent baseline and data-independent baseline learned during training. The second technique is variance normalization, which is normalizing the learning signal by a running estimate of its standard deviation. All learning details and evaluation of gradients are provided in Supplementary Sections B.1-B.3.

3.3. Extension for deep models

The recognition model with respect to the deep FCTSBN is shown in Figure 3 (right). Since the middle layers are also stochastic, the calculation of the lower bound includes more terms, in the form of

$$\mathcal{L} = \sum_{\ell=0}^L \mathcal{J}(q_{\phi}(\mathbf{H}^{(\ell+1)} | \mathbf{H}^{(\ell)}, \mathbf{Y}), p_{\theta}(\mathbf{H}^{(\ell+1)}, \mathbf{H}^{(\ell)} | \mathbf{Y}))$$

Each layer has a unique set of parameters whose gradient is zero for the lower bound of the other layers. Therefore, we can calculate the gradients separately for each layer of parameters, in a similar fashion to the single-layer FCTSBN. All details are provided in Supplementary Section B.4.

3.4. Extension for Semi-supervised Learning

The recognition model in Section 3.1 provides a fast bottom-up probabilistic mapping from observations to hidden variables. Following Kingma et al. (2014), for the semi-supervised model (12), we introduce a recognition model for both \mathbf{H} and \mathbf{Y} , with the factorized form $q_{\phi}(\mathbf{H}, \mathbf{Y}|\mathbf{V}) = q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y}) \cdot q_{\phi}(\mathbf{Y}|\mathbf{V})$. The first term, $q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y})$, is the same as in (14). When \mathbf{y}_t is a one-hot encoded vector, we assume a discriminative classifier $q_{\phi}(\mathbf{y}_t | \mathbf{v}_t) = \text{Cat}(\mathbf{v}_t; \pi_{\phi}(\mathbf{v}_t))$, where $\text{Cat}(\mathbf{x}; \pi)$ denotes the categorical distribution.

When the label corresponding to a sequence is missing, it is treated as a latent variable over which we perform posterior inference. The variational lower bound for the unlabeled data is

$$\mathcal{L}_u = \mathcal{J}(q_{\phi}(\mathbf{H}, \mathbf{Y}|\mathbf{V}), p_{\theta}(\mathbf{H}, \mathbf{V}, \mathbf{Y})) \quad (19)$$

When the label is observed, the variational bound is a simple extension of (13), expressed as

$$\mathcal{L} = \mathcal{J}(q_{\phi}(\mathbf{H}|\mathbf{V}, \mathbf{Y}), p_{\theta}(\mathbf{V}, \mathbf{H}|\mathbf{Y})) + \log p_{\theta}(\mathbf{Y}; \rho) \quad (20)$$

where $p_{\theta}(\mathbf{Y}; \rho)$ is the prior distribution for \mathbf{Y} , a constant term that can be omitted. To exploit the discriminative power of the labeled data, we further add a classification loss to (20). The resulting objective function is

$$\mathcal{L}_l = \mathcal{L} + \alpha \cdot \mathbb{E}_{\hat{p}_l(\mathbf{V}, \mathbf{Y})} [\log q_{\theta}(\mathbf{Y}|\mathbf{V})] \quad (21)$$

where \mathcal{L} is the lower bound for the generative model, defined in (20), and $\hat{p}_l(\mathbf{V}, \mathbf{Y})$ denotes the empirical distribution. Parameter α controls the relative weight between

the generative and discriminative components within the semi-supervised learning framework³, where the extended objective function $\mathcal{L}_s = \mathcal{L}_l + \mathcal{L}_u$ takes both labeled and unlabeled data into account. Details for optimizing \mathcal{L}_s are included in Supplementary Section B.5.

4. Related Work

Probabilistic models for sequential data in the deep learning literature can be roughly divided into two categories. The first category includes generative models *without* latent variables, which rely on use of Recurrent Neural Networks (RNNs) (Sutskever et al., 2011; Graves, 2013; Chung et al., 2015a). The second category contains generative models *with* latent variables, which can be further divided into two subcategories: (i) *undirected* latent variable models, utilizing the RBM as the building block (Taylor et al., 2006; Sutskever & Hinton, 2007; Sutskever et al., 2009; Boulanger-Lewandowski et al., 2012; Mittelman et al., 2014); (ii) *directed* latent variable models, *e.g.*, extending the variational auto-encoders (Bayer & Osendorfer, 2014; Fabius et al., 2014; Chung et al., 2015b), or utilizing the SBN as the building block (Gan et al., 2015a;b;c).

Among this work, the Factored Conditional Restricted Boltzmann Machine (FCRBM) (Taylor & Hinton, 2009) and Temporal Sigmoid Belief Network (TSBN) (Gan et al., 2015c) are most related to the work reported here. However, there exist several key differences. Compared with FCRBM, our proposed model is fully *directed*, allowing fast generation through ancestral sampling; while in the FCRBM, alternating Gibbs sampling is required to obtain a sample at each time-step. Compared with TSBN, where generation of different styles was purely based on initialization, our model utilizes side information to gate the connections of a TSBN, which makes the model context-sensitive, and permits controlled transitioning and blending (detailed when presenting experiments in Section 5).

Another novelty of our model lies in the utilization of a factored three-way tensor to model the multiplicative interactions between side information and sequences. Similar ideas have also been exploited in (Memisevic & Hinton, 2007; Taylor & Hinton, 2009; Sutskever et al., 2011; Kiros et al., 2014a;b). We note that our factoring method is different from the one used in FCRBM. Specifically, in FCRBM, the energy function was factored, while our model factors over the transition parameters, by which we capture the underlying similarities within different sequences, boosting performance.

Most deep generative models focus on exploring the generative ability (Kingma & Welling, 2014), and little work

³We use $\alpha = 2 \cdot T$ throughout the experiments, but obtain similar performance for other values of α within $[0.1 \cdot T, 4 \cdot T]$.

has been done on examining the discriminative ability of deep generative models, except Kingma et al. (2014); Li et al. (2015). However, both works are restricted to the application of static data. Our paper is the first to develop a semi-supervised sequence classification method with deep generative models.

In terms of inference, the wake-sleep algorithm (Hinton et al., 1995), Stochastic Gradient Variational Bayes (SGVB) (Kingma & Welling, 2014) and Neural Variational Inference and Learning (NVIL) (Mnih & Gregor, 2014) are widely studied for training recognition models. We utilize NVIL for scalable inference, and our method is novel in designing a new factored multiplicative recognition model.

5. Experiments

Sections 5.1-5.4 report the results of training several models with data from the CMU Motion Capture Database. Specifically, we consider two datasets: (i) motion sequences performed by subject 35 (Taylor et al., 2006)(**mocap2**), which contains two types of motions, *i.e.*, walking and running; (ii) motion sequences performed by subject 137 (Taylor & Hinton, 2009)(**mocap10**), which contains 10 different styles of walking. Both datasets are preprocessed using downsampling and scaling to have zero mean and unit variance. In Section 5.5 and 5.6, we present additional experiments on weather prediction and conditional text generation to further demonstrate the versatility of the proposed model. The weather prediction dataset (Liu et al., 2010) contains monthly observations of time series data of 18 climate agents (data types) over different places in North America.

The FCTSBN model with $\mathbf{W}_3^{(y)} \equiv \mathbf{0}$ and $\mathbf{W}_4^{(y)} \equiv \mathbf{0}$ is denoted Hidden Markov FCTSBN. The deep extension to FCTSBN is abbreviated as dFCTSBN. We use this abbreviation to denote the conditional generative model instead of the semi-supervised model, unless explicitly stated. Furthermore, we allow each observation to be dependent on the hidden and visible states of the previous n time steps, instead of $n = 1$. We refer to n as the *order* of the model. The choice of n can be different according to the specific scenario. For mocap2 prediction and mocap10 generation, n is selected to align with previous methods on these tasks; for other prediction tasks, n is chosen to balance between performance and model complexity; for semi-supervised tasks, $n = 6$ is considered to allow consecutive frames of data to be utilized effectively by the classifier.

For CTsBN, the model parameters for weights are initialized by sampling from $\mathcal{N}(\mathbf{0}, 0.001^2 \mathbf{I})$, whereas the bias parameters are initialized as zero. For FCTSBN, the parameters are initialized differently, since the actual initialization value of the weight parameters depends on the product

of factors. To ensure faster convergence, the initial values \mathbf{W}_a , \mathbf{W}_b and \mathbf{W}_c are sampled from $\mathcal{N}(\mathbf{0}, 0.01^2 \mathbf{I})$. We use RMSprop (Tieleman & Hinton, 2012) throughout all the experiments. This is a stochastic gradient descent (SGD) method that allows the gradients to be adaptively rescaled by a running average of their recent magnitude. The data-dependent baseline is implemented with a single-hidden-layer neural network with 100 tanh units. We update the estimated learning signal with a momentum of 0.9.

The prediction of \mathbf{v}_t given $\mathbf{v}_{1:t-1}$ requires first sampling $\mathbf{h}_{1:t-1}$ from $q_\phi(\mathbf{h}_{1:t-1} | \mathbf{v}_{1:t-1}, \mathbf{y}_{1:t-1})$, then calculating the posterior $p_\theta(\mathbf{h}_t | \mathbf{h}_{1:t-1}, \mathbf{v}_{1:t-1}, \mathbf{y}_t)$ for \mathbf{h}_t . Given \mathbf{h}_t , a prediction can be made for \mathbf{v}_t by sampling from $p_\theta(\mathbf{v}_t | \mathbf{h}_{1:t}, \mathbf{v}_{1:t-1}, \mathbf{y}_t)$, or by using the expectation of sampled variables. Generating samples, on the other hand, is a similar but relatively easier task, where sequences can be generated through ancestral sampling. A special advantage of the conditional generative model appears when the side information \mathbf{y}_t changes over time, so that we can generate sequences with style transitions that do not appear in the training data. More details on generating such sequences are discussed in Section 5.4.

5.1. Mocap2 Prediction

For the first experiment, we perform the prediction task using the mocap2 dataset. We used 33 running and walking sequences, partitioned them into 31 training sequences and 2 test sequences, as in Gan et al. (2015c). Our models have 100 hidden units (and 50 factors for factored models) in each layer and the order of $n = 1$, according to the settings in Gan et al. (2015c). Deep models, including deep FCTSBN and deep TSBN, have two hidden layers of 100 units each. The side information, \mathbf{y}_t , is a one-hot encoded vector of length 2, which indicates “running” or “walking”.

From Table 1, it is observed that the one-layered conditional models have a significant improvement over TSBN, ss(spike-slab)-SRTRBM (Mittelman et al., 2014) and g(Gaussian)-RTRBM (Sutskever et al., 2009), whereas FCTSBN has better performance than CTsBN, due to the factoring mechanism. Results of deep FCTSBN have comparable performances to that of TSBN with stochastic hidden layers; this may be because only two styles of time-series data are considered.

5.2. Mocap10 Generation and Prediction

In order to demonstrate the conditional generative capacity of our model, we choose the mocap10 dataset, which contains 10 different walking styles, namely *cat*, *chicken*, *dinosaur*, *drunk*, *gangly*, *graceful*, *normal*, *old-man*, *sexy* and *strong*. A single-layer Hidden Markov FCTSBN with 100 hidden units and order of $n = 12$ is trained on the dataset for 400 epochs, whereas the parameters are updated

Method	Walking	Running
FCTSBN	4.59 \pm 0.35	2.86 \pm 0.23
CTSBN	4.67 \pm 0.22	3.41 \pm 0.65
TSBN [◊]	5.12 \pm 0.50	4.85 \pm 1.26
dFCTSBN	4.31 \pm 0.13	2.58 \pm 0.21
DTsBN-S [◊]	4.40 \pm 0.28	2.56 \pm 0.40
DTsBN-D [◊]	4.62 \pm 0.01	2.84 \pm 0.01
ss-SRTRBM [◊]	8.13 \pm 0.06	5.88 \pm 0.05
g-RTRBM [◊]	14.41 \pm 0.38	10.91 \pm 0.27

Table 1: Prediction error obtained for the mocap2 dataset. (◊) taken from Gan et al. (2015c); (◊) taken from Mittelman et al. (2014). Bold indicate the best results for shallow deep models, respectively.

10 times for each epoch. The side information \mathbf{y}_t is provided as a one-hot encoded vector of length 10. We set a fixed learning rate of 3×10^{-3} , and a decay rate of 0.9. Motions of a particular style are generated by initializing with 12 frames of the training data for that style and fixing the side information during generation.

Our one-layer model can generate walking sequences of 9 styles out of 10, which are included in the supplementary videos. The only exception is the *old-man* style, where our model captures the general posture but fails to synthesize the subtle frame-to-frame changes in the footsteps.

We also present prediction results. For each style, we select 90% of the sequences as training data, and use the rest as testing data. We consider FCTSBN and a two-layer deep FCTSBN with order $n = 8$, 100 hidden variables and 50 factors on each layer. We also consider FCRBM with 600 hidden variables, 200 factors, 200 features and order of $n = 12$, as in Taylor & Hinton (2009). The average prediction error over 10 styles for FCRBM, FCTSBN and dFCTSBN are 0.4355, 0.4832, and 0.3599, respectively, but the performances for different styles vary significantly. Detailed prediction error results for each style is provided in Supplementary Section C.

5.3. Semi-supervised learning with mocap2 and mocap10

Our semi-supervised learning framework can train a classifier and a generative model simultaneously. In this section we present semi-supervised classification results for mocap10, and generated sequences for mocap2.⁴

For mocap10, we perform 10-class classification over the motion style with windows of 7 consecutive frames from

⁴We cannot perform both tasks on a single dataset due to the limitations of the mocap datasets: for mocap2, the 2-class classification task is so simple that the test accuracy is almost 1; for mocap10, the number of frames is too small to train a generative model with semi-supervised learning.




Figure 4: Test results on mocap10 classification. Each result is the average over 5 independent trials.

the sequences. We use FCTSBN and dFCTSBN of order $n = 6$ with 100 hidden variables on each layer, set $q_\phi(\mathbf{Y}|\mathbf{V})$ as a one-layer softmax classifier, and use $\alpha = 2 \cdot T$ throughout our experiments. Our baseline models include K-nearest neighbors (K-NN), the one-layer softmax classifier with regularization parameter λ (Softmax- λ), and transductive support vector machines (TSVM) (Gammerman et al., 1998). Sequences of certain styles are truncated such that each style has 500 frames of actual motions. We hold out 100 frames for testing, and the rest as training data, a consecutive proportion of which is provided with \mathbf{y}_t .

Figure 4 shows the average test accuracy of different methods under various percentage of labeled training data, where our models consistently have better performance than other methods, even with a simple, one-layer classifier q_ϕ . To ensure fair comparison, the other parametric methods have roughly the same number of parameters. This would demonstrate that the \mathcal{L}_u term in (19) serves as an effective regularizer that helps prevent overfitting. Our deep FCTSBN model has better performance than the shallow model, which may be attributed to a more flexible regularizer due to more parameters in the generative model objective.

For mocap2 generation, the training data are processed such that one sequence for each style is provided with \mathbf{y}_t , whereas the remaining 31 sequences are unlabeled. Then the data are used to train a Hidden Markov FCTSBN of order $n = 6$ with 100 hidden variables and 50 factors, and $q_\phi(\mathbf{Y}|\mathbf{V})$ as a one-layer softmax classifier. Although the proportion of labeled data is very small, we are able to synthesize smooth, high-quality motion sequences for both styles. Related videos are presented in the Supplementary Material.




Figure 5: Intensity of \mathbf{y}_t for different styles during transition. Pre-transition style is denoted as “prev”; post-transition style as “post”.




Figure 6: **Left:** *Drunk* to *normal* (6 skeletons for each style). **Right:** *Graceful* to *gangly* (5 skeletons for each style). Both figures are generated within 800 frames, with the side information given according to Figure 5.

5.4. Generating sequences with style transitions and blending

One particular advantage of modeling multiple time-series data with one FCTSBN over training several TSBNs for each sequence is that the FCTSBN has a shared hidden representation for all styles, so \mathbf{y}_t are no longer restricted to a constant value during generation. This flexibility allows generation of sequences that do not appear in the training set, such as transitions between styles, or combinations of multiple styles.

Videos displaying smooth and efficient style transitions and blending are provided in the supplementary material. Two examples are provided in Figure 6. We also include transitions between walking and running with the generative model trained in Section 5.3. To enable such transition, we transform the side information from one to another using a sigmoid-shaped form as in Figure 5. The entire transition process happens over around 60 frames, which is much faster than the 200-frame transition proposed for FCRBM (Taylor & Hinton, 2009)⁵. For style blending, the side information is provided by a convex combination of several one-hot-encoded side information vectors.

5.5. Weather prediction

We demonstrate weather prediction in multiple places, and discuss how different representations of side information can affect performance. We select 25 places in the United States, which are located in a 5×5 grid spanning across

⁵In practice, a sudden change of labels will result in “jerky” transitions (Taylor & Hinton, 2009), but we aim to use as few frames as possible during transition.

	2d	10d	25d
CTSBN	6.45 \pm 0.11	3.83 \pm 0.20	3.78 \pm 0.01
FCTSBN	5.46 \pm 0.06	3.43 \pm 0.03	3.37 \pm 0.02
dFCTSBN	5.09 \pm 0.11	3.37 \pm 0.01	3.35 \pm 0.02
FCRBM	5.62 \pm 0.35	3.77 \pm 0.38	3.75 \pm 0.08

Table 2: Average prediction error of 25 locations on the weather dataset.

multiple climate regions. The weather data are processed so that for every month in 1990-2002, each place has 17 climate indicator values (*e.g.*, temperature). We use only longitude and latitude as side information.

One way of providing \mathbf{y}_t is by using the geographic information directly, such that \mathbf{y}_t would be a real-valued vector of length 2 (we call this setting 2d for simplicity). Such representation is oversimplified, since the transition parameters \mathbf{W} of a certain location cannot be described as a linear combination of longitude and latitude. Therefore, we consider two other representations of side information. One uses a vector with 10 binary values (10d), which is the concatenation of two one-hot encoded vector of length 5 for longitude and latitude, respectively; the other assumes a one-hot encoded vector of 25, one for each location (25d).

We use the monthly climate data from 1990 to 2001 as training set, and the data of 2002 as test set. We consider CTSBN, FCTSBN, and two-layer deep FCTSBN of order $n = 6$, with 100 hidden variables and 50 factors on each hidden layer. During test time, we predict the weather of every month in 2002 given the information of previous 6 months.

Average prediction error of 5 trials with random initializations are provided in Table 2. We observe that the factored model has significantly higher performance than the non-factored CTSBN, since the factored model has shared input-to-factor and factor-to-output parameters, and the factor-to-factor parameters have the potential to capture the correlations among side information at a higher level. Moreover, the performance improves as we increase the dimension of side information, which might be due to the increase of the number of parameters, along with the one-hot embeddings of side information. The 10d model has a significant performance boost compared with the 2d model, while using the 25d model further improves the performance slightly. We also compare our method with the one layer FCRBM model, which has 100 hidden variables and 50 variables for each of the three style factors and features, and achieved better results on the one-layer setting.

5.6. Conditional Text Generation

Finally, we use a simple experiment to demonstrate the versatility of our model over the task of conditional text gen-

B	<#> and the evening and the morning were the fourth day . <#> and god said , let the waters the heaven after his kind : <#> god blessed their thing months...
C	we shall hold modern into tends ; and circumstance between seem understood retained defendand'ts to has that belief are not the recalled and will be led constituent...
BC	<#> and unto the rendered fair violence , morning turn the human whole been so eyes . <#> that god of air of the mountain show <#>the waters of fish and him would expect application : are gradual obliged that...

Table 3: Generated text. <#> denote markers for verses.

eration. We select 4 books (*i.e.*, *Napoleon the Little*, *The Common Law*, *Mysticism and Logic* and *The Bible*) in the Gutenberg corpus, and use *word2vec* (Mikolov et al., 2013) to extract word embeddings, where words with few occurrences are omitted. This preprocessing step provides us with real-valued vectors, which are treated as observations and fit into our current model.

The real-valued word vector sequences are trained on a one-layer CTSBN of order 15 and 100 hidden variables. During generation, we clamp each generated value to the nearest possible word vector using cosine similarity. We focus on modeling real-valued sequences, hence the clamping approach is employed, instead of adopting a softmax layer to predict the words directly.

Table 3 displays the generated samples from the model when provided with the side information of *The Bible* (B), *The Common Law* (C) and an even mixture of these two books (BC). As can be seen, our model learns to capture the styles associated with different books. By conditioning on the average of two styles, the model can generate reasonable samples that represent a hybrid of both styles (distinctive words such as <#> and *god* from B, and *application* and *violence* from C), even though such style combinations were not observed during training.

6. Conclusions and Future Work

We have presented the Factored Conditional Temporal Sigmoid Belief Network, which can simultaneously model temporal sequences of multiple subjects. A general framework for semi-supervised sequence classification is also provided, allowing one to train a conditional generative model along with a classifier. Experimental results on several datasets show that the proposed approaches obtain superior predictive performance, boost classification accuracy, and synthesize a large family of sequences.

While we assume side information as a simple vector, it would be interesting to incorporate more structured side information into the model, *e.g.*, utilizing separate style and content components. We are also interested in models that can extract structured properties from side information.

Acknowledgements

This research was supported in part by ARO, DARPA, DOE, NGA and ONR.

References

- Bayer, J. and Osendorfer, C. Learning stochastic recurrent networks. In *arXiv:1411.7610*, 2014.
- Boulanger-Lewandowski, N., Bengio, Y., and Vincent, P. Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription. In *ICML*, 2012.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. Gated feedback recurrent neural networks. In *ICML*, 2015a.
- Chung, J., Kastner, K., Dinh, L., Goel, K., Courville, A., and Bengio, Y. A recurrent latent variable model for sequential data. In *NIPS*, 2015b.
- Fabius, O., van Amersfoort, J. R., and Kingma, D. P. Variational recurrent auto-encoders. In *arXiv:1412.6581*, 2014.
- Gamerman, A., Vovk, V., and Vapnik, V. Learning by transduction. In *UAI*, 1998.
- Gan, Z., Chen, C., Henao, R., Carlson, D., and Carin, L. Scalable deep poisson factor analysis for topic modeling. In *ICML*, 2015a.
- Gan, Z., Henao, R., Carlson, D., and Carin, L. Learning deep sigmoid belief networks with data augmentation. In *AISTATS*, 2015b.
- Gan, Z., Li, C., Henao, R., Carlson, D., and Carin, L. Deep temporal sigmoid belief networks for sequence modeling. In *NIPS*, 2015c.
- Graves, A. Generating sequences with recurrent neural networks. In *arXiv:1308.0850*, 2013.
- Hinton, G. E. Training products of experts by minimizing contrastive divergence. In *Neural computation*, 2002.
- Hinton, G. E., Dayan, P., Frey, B. J., and Neal, R. M. The “wake-sleep” algorithm for unsupervised neural networks. In *Science*, 1995.
- Hinton, G. E., Osindero, S., and Teh, Y. W. A fast learning algorithm for deep belief nets. In *Neural Computation*, 2006.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. An introduction to variational methods for graphical models. In *Machine Learning*, 1999.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *ICLR*, 2014.
- Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M. Semi-supervised learning with deep generative models. In *NIPS*, 2014.
- Kiros, R., Salakhutdinov, R., and Zemel, R. Multimodal neural language models. In *ICML*, 2014a.
- Kiros, R., Zemel, R., and Salakhutdinov, R. R. A multiplicative model for learning distributed text-based attribute representations. In *NIPS*, 2014b.
- Le, Q. V. and Mikolov, T. Distributed representations of sentences and documents. In *ICML*, 2014.
- Li, C., Zhu, J., Shi, T., and Zhang, Bo. Max-margin deep generative models. In *NIPS*, 2015.
- Liu, Y., Niculescu-Mizil, A., Lozano, A. C., and Lu, Y. Learning temporal causal graphs for relational time-series analysis. In *ICML*, 2010.
- Memisevic, R. and Hinton, G. Unsupervised learning of image transformations. In *CVPR*, 2007.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.
- Mittelman, R., Kuipers, B., Savarese, S., and Lee, H. Structured recurrent temporal restricted boltzmann machines. In *ICML*, 2014.
- Mnih, A. and Gregor, K. Neural variational inference and learning in belief networks. In *ICML*, 2014.
- Neal, R. M. Connectionist learning of belief networks. In *Artificial intelligence*, 1992.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, 2014.
- Salakhutdinov, R. and Hinton, G. Deep boltzmann machines. In *AISTATS*, 2009.
- Srivastava, N., Mansimov, E., and Salakhutdinov, R. Unsupervised learning of video representations using lstms. In *ICML*, 2015.
- Sutskever, I. and Hinton, G. E. Learning multilevel distributed representations for high-dimensional sequences. In *AISTATS*, 2007.
- Sutskever, I., Hinton, G., and Taylor, G. The recurrent temporal restricted boltzmann machine. In *NIPS*, 2009.

Sutskever, I., Martens, J., and Hinton, G. E. Generating text with recurrent neural networks. In *ICML*, 2011.

Taylor, G., Hinton, G., and Roweis, S. Modeling human motion using binary latent variables. In *NIPS*, 2006.

Taylor, G. W. and Hinton, G. E. Factored conditional restricted boltzmann machines for modeling motion style. In *ICML*, 2009.

Tieleman, T. and Hinton, G. E. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. In *COURSERA: Neural Networks for Machine Learning*, 2012.