

---

# Thurstonian Boltzmann Machines: Learning from Multiple Inequalities

---

Truyen Tran<sup>†‡</sup>  
Dinh Phung<sup>†</sup>  
Svetha Venkatesh<sup>†</sup>

TRUYEN.TRAN@DEAKIN.EDU.AU  
DINH.PHUNG@DEAKIN.EDU.AU  
SVETHA.VENKATESH@DEAKIN.EDU.AU

<sup>†</sup>Center for Pattern Recognition and Data Analytics (PRaDA), Deakin University, Australia

<sup>‡</sup>Institute for Multi-sensor Processing & Content Analysis (IMPCA), Curtin University, Australia

## Abstract

We introduce *Thurstonian Boltzmann Machines* (TBM), a unified architecture that can naturally incorporate a wide range of data inputs at the same time. Our motivation rests in the Thurstonian view that many discrete data types can be considered as being generated from a subset of underlying latent continuous variables, and in the observation that each realisation of a discrete type imposes certain inequalities on those variables. Thus learning and inference in TBM reduce to making sense of a set of inequalities. Our proposed TBM naturally supports the following types: Gaussian, intervals, censored, binary, categorical, muticategorical, ordinal, (in)-complete rank with and without ties. We demonstrate the versatility and capacity of the proposed model on three applications of very different natures; namely handwritten digit recognition, collaborative filtering and complex social survey analysis.

## 1. Introduction

Restricted Boltzmann machines (RBMs) have proved to be a versatile tool for a wide variety of machine learning tasks and as a building block for deep architectures (Hinton and Salakhutdinov, 2006; Salakhutdinov and Hinton, 2009a; Smolensky, 1986). The original proposals mainly handle binary visible and hidden units. Whilst binary hidden units are broadly applicable as feature detectors, non-binary visible data requires different designs. Recent extensions to other data types result in type-dependent models: the Gaussian for continuous inputs (Hinton

and Salakhutdinov, 2006), Beta for bounded continuous inputs (Le Roux et al., 2011), Poisson for count data (Gehler et al., 2006), multinomial for unordered categories (Salakhutdinov and Hinton, 2009b), and ordinal models for ordered categories (Truyen et al., 2009; Tran et al., 2012).

The Boltzmann distribution permits several types to be jointly modelled, thus making the RBM a good tool for multimodal and complex social survey analysis. The work of (Ngiam et al., 2011; Srivastava and Salakhutdinov, 2012; Xing et al., 2005) combines continuous (e.g., visual and audio) and discrete modalities (e.g., words). The work of (Tran et al., 2011) extends the idea further to incorporate ordinal and rank data. However, there are conceptual drawbacks: First, conditioned on the hidden layer, they are still separate type-specific models; second, handling ordered categories and ranks is not natural; and third, specifying direct correlation between these types remains difficult.

The main thesis of this paper is that many data types can be captured in one unified model. The key observations are that (i) type-specific properties can be modelled using one or several *underlying continuous variables*, in the spirit of Thurstonian models<sup>1</sup> (Thurstone, 1927), and (ii) evidences be expressed in the form of one or several *inequalities* of these underlying variables. For example, a binary visible unit is turned on if the underlying variable is beyond a threshold; and a category is chosen if its utility is the largest among all those of competing categories. The use of underlying variables is desirable when we want to explicitly model the generative mechanism of the data. In psychology and economics, for example, it gives much better interpretation on why a particular choice is made given

---

*Proceedings of the 30<sup>th</sup> International Conference on Machine Learning*, Atlanta, Georgia, USA, 2013. JMLR: W&CP volume 28. Copyright 2013 by the author(s).

---

<sup>1</sup>Whilst Thurstonian models often refer to human’s judgment of discrete choices, we use the term “Thurstonian” more freely without the notion of human’s decision.

the perceived utilities (Böckenholt, 2006). Further, it is natural to model the correlation among type-specific inputs using a covariance structure on the underlying variables.

The inequality observation is interesting in its own right: Instead of learning from assigned values, we learn from the *inequality expression of evidences*, which can be much more relaxed than the value assignments. This class of evidences indeed covers a wide range of practical situations, many of which have not been studied in the context of Boltzmann machines, as we shall see throughout the paper.

To this end, we propose a novel class of models called *Thurstonian Boltzmann Machine* (TBM). The TBM utilises the Gaussian restricted Boltzmann machine (GRBM): The top layer consists of binary hidden units as in standard RBMs; the bottom layer contains a collection of Gaussian variable groups, one per input type. The main difference is that TBM does not require valued assignments for the bottom layer *but a set of inequalities expressing the constraints imposed by the evidences*. Except for a limiting case of point assignments where the inequalities are strictly equalities, the Gaussian layer is never fully observed. The TBM supports more data types in a unified manner than ever before: For any combination of the point assignments, intervals, censored values, binary, unordered categories, multi-categories, ordered categories, (in)-complete ranks with and without ties, all we need to do is to supply relevant subset of inequalities.

We evaluate the proposed model on three applications of very different natures: handwritten digit recognitions, collaborative filtering and complex survey analysis. For the first two applications, the performance is competitive against methods designed for those data types. On the last application, we believe we are among the first to propose a scalable and generic machinery for handle those complex data types.

## 2. Gaussian RBM

Let  $\mathbf{x} = (x_1, x_2, \dots, x_N)^\top \in \mathbb{R}^N$  be a vector of input variables. Let  $\mathbf{h} = (h_1, h_2, \dots, h_K)^\top \in \{0, 1\}^K$  be a set of hidden factors which are designed to capture the variations in the observations. The input layer and the hidden layer form an undirected bipartite graph, i.e., only cross-layer connections are allowed. The model admits the Boltzmann distribution

$$P(\mathbf{x}, \mathbf{h}) = \frac{1}{Z} \exp\{-E(\mathbf{x}, \mathbf{h})\} \quad (1)$$

where  $Z = \sum_{\mathbf{h}} \int \exp\{-E(\mathbf{x}, \mathbf{h})\} d\mathbf{x}$  is the normalising constant and  $E(\mathbf{x}, \mathbf{h})$  is the state energy. The energy is decomposed as

$$E(\mathbf{x}, \mathbf{h}) = \sum_i \left( \frac{x_i^2}{2} - (\alpha_i + W_{i\bullet} \mathbf{h}) x_i \right) - \gamma^\top \mathbf{h}$$

where  $\{\alpha_i\}_{i=1}^N$ ,  $W = \{W_{ik}\}$ ,  $\gamma = \{\gamma_k\}$  are free parameters and  $W_{i\bullet}$  denotes the  $i$ -th row.

Given the input  $\mathbf{x}$ , the posterior has a simple form

$$\begin{aligned} P(\mathbf{h} | \mathbf{x}) &= \prod_k P(h_k | \mathbf{x}) \quad (2) \\ P(h_k = 1 | \mathbf{x}) &= \frac{1}{1 + e^{-\gamma_k - W'_{\bullet k} \mathbf{x}}} \end{aligned}$$

where  $W_{\bullet k}$  denotes the  $k$ -th column. Similarly, the generative process given the binary factor  $\mathbf{h}$  is also factorisable

$$\begin{aligned} P(\mathbf{x} | \mathbf{h}) &= \prod_i P(x_i | \mathbf{h}) \quad (3) \\ P(x_i | \mathbf{h}) &= \mathcal{N}(\alpha_i + W_{i\bullet} \mathbf{h}, 1) \end{aligned}$$

where  $\mathcal{N}(\mu, 1)$  is the normal distribution of mean  $\mu$  and unit deviation.

## 3. Thurstonian Boltzmann Machines

We now generalise the Gaussian RBM into the Thurstonian Boltzmann Machine (TBM). Denote by  $\mathbf{e}$  an observed *evidence* of  $\mathbf{x}$ . Standard evidences are the point assignment of  $\mathbf{x}$  to some specific real-valued vector, i.e.,  $\mathbf{x} = \mathbf{v}$ . Generalised evidences can be expressed using *inequality constraints*

$$\mathbf{b} \leq_{\odot} A \mathbf{x} \leq_{\odot} \mathbf{c} \quad (4)$$

for some transform matrix  $A \in \mathbb{R}^{M \times N}$  and vectors  $\mathbf{b}, \mathbf{c} \in \mathbb{R}^M$ , where  $\leq_{\odot}$  denotes element-wise inequalities. Thus an evidence can be completely realised by specifying the triple  $\langle A, \mathbf{b}, \mathbf{c} \rangle$ . For example, for the point assignment,  $\langle A = \mathbf{I}, \mathbf{b} = \mathbf{v}, \mathbf{c} = \mathbf{v} \rangle$ , where  $\mathbf{I}$  is the identity matrix. In what follows, we will detail other useful popular realisations of these quantities.

### 3.1. Boxed Constraints

This refers to the case where input variables are independently constrained, i.e.,  $A = \mathbf{I}$ , and thus we need only to specify the pair  $\langle \mathbf{b}, \mathbf{c} \rangle$ .

**Censored observations.** This refers to situation where we only know the continuous observation beyond a certain point, i.e.,  $\mathbf{b} = \mathbf{e}$  and  $\mathbf{c} = +\infty$ . For example, in survival analysis, the life expectancy of a person might be observed up to a certain age, and we have no further information afterward.

**Interval observations.** When the measurements are imprecise, it may be better to specify the range of possible observations with greater confidence rather

than a single point, i.e.,  $\mathbf{b} = \mathbf{e} - \boldsymbol{\delta}$  and  $\mathbf{c} = \mathbf{e} + \boldsymbol{\delta}$  for some pair  $(\mathbf{e}, \boldsymbol{\delta})$ . For instance, missile tracking may estimate the position of the target with certain precision.

**Binary observations.** A binary observation  $e_i$  can be thought as a result of clipping  $x_i$  by a threshold  $\theta_i$ , that is  $e_i = 1$  if  $x_i \geq \theta_i$  and  $e_i = 0$  otherwise. The boundaries in Eq. (3.2) become:

$$\langle b_i, c_i \rangle = \begin{cases} \langle -\infty, \theta_i \rangle & e_i = 0 \\ \langle \theta_i, +\infty \rangle & e_i = 1 \end{cases} \quad (5)$$

Thus, this model offers an alternative<sup>2</sup> to standard binary RBMs of (Smolensky, 1986; Freund and Haussler, 1993).

**Ordinal observations.** Denote by  $\mathbf{e} = (e_1, e_2, \dots, e_N)$  the set of ordinal observations, where each  $e_i$  is drawn from an ordered set  $\{1, 2, \dots, L\}$ . The common assumption is that the ordinal level  $e_i = l$  is observed given  $x_i \in [\theta_{l-1}, \theta_l]$  for some thresholds  $\theta_1 \leq \theta_2 \leq \dots \leq \theta_{L-1}$ . The boundaries thus read

$$\langle b_i, c_i \rangle = \begin{cases} \langle -\infty, \theta_1 \rangle & l = 1 \\ \langle \theta_{l-1}, \theta_l \rangle & l = 2, 3, \dots, L-1 \\ \langle \theta_{L-1}, +\infty \rangle & \text{otherwise} \end{cases} \quad (6)$$

This offers an alternative<sup>3</sup> to the ordinal RBMs of (Truyen et al., 2009).

### 3.2. Inequality Constraints

**Categorical observations.** This refers to the situation where out of an *unordered* set of categories, we observe only one category at a time. This can be formulated as follows. Each category is associated with a ‘‘utility’’ variable. The category  $l$  is observed (i.e.,  $e_i = m$ ) if it has the largest utility, that is  $x_{il} \geq \max_{m \neq l} x_{im}$ . Thus,  $x_{il}$  is the upper-threshold for all other utilities. On the other hand,  $\max_{m \neq l} x_{im}$  is the lower-threshold for  $x_{il}$ . This suggests an *EM-style* procedure: (i) fix  $x_{il}$  (or treat it as a threshold) and learn the model under the intervals  $x_{im} \leq x_{il}$  for all  $m \neq l$ , and (ii) fix all categories other than  $l$ , learn the model under the interval  $x_{il} \geq \max_{m \neq l} x_{im}$ . This offers an alternative<sup>4</sup> to the multinomial logit treatment in (Salakhutdinov et al., 2007).

To illustrate the point, suppose there are only four variables  $z_1, z_2, z_3, z_4$ , and  $z_1$  is observed, then we have  $z_1 \geq \max\{z_2, z_3, z_4\}$ . This can be expressed as  $z_1 - z_2 \geq 0$ ;  $z_1 - z_3 \geq 0$  and  $z_1 - z_4 \geq 0$ . These are

<sup>2</sup>To be consistent with the statistical literature, we can call it the *probit RBM*, which we will study in Section 6.1.

<sup>3</sup>This can be called *ordered probit RBM*.

<sup>4</sup>We can call this model *multinomial probit RBM*.

equivalent to

$$\left\langle A = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 1 & 0 & -1 & 0 \\ 1 & 0 & 0 & -1 \end{bmatrix}; \quad \mathbf{b} = \mathbf{0}; \quad \mathbf{c} = +\infty \right\rangle$$

**Imprecise categorical observations.** *This generalises the categorical case:* The observation is a subset of a set, where *any* member of the subset can be a possible observation<sup>5</sup>. For example, when asked to choose the best sport team of interest, a person may pick two teams without saying which is the best. For instance, suppose the subset is  $\{z_1, z_2\}$ , then  $\min\{z_1, z_2\} \geq \max\{z_3, z_4\}$ , which can be expressed as  $z_1 - z_3 \geq 0$ ;  $z_2 - z_3 \geq 0$ ,  $z_1 - z_4 \geq 0$  and  $z_2 - z_4 \geq 0$ . This translates to the following triple

$$\left\langle A = \begin{bmatrix} 1 & 0 & -1 & 0 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{bmatrix}; \quad \mathbf{b} = \mathbf{0}; \quad \mathbf{c} = +\infty \right\rangle$$

**Rank (with Ties) observations.** *This generalises the imprecise categorical cases:* Here we have a (partially) ranked set of categories. Assume that the rank is produced in a stagewise manner as follows: The best category subset is selected out of all categories, the second best is selected out of all categories except for the best one, and so on. Thus, at each stage we have an imprecise categorical setting, but now the utilities of middle categories are constrained from both sides – the previous utilities as the upper-bound, and the next utilities as the lower-bound.

As an illustration, suppose there are four variables  $z_1, z_2, z_3, z_4$  and a particular rank (with ties) imposes that  $\min\{z_1, z_2\} \geq z_3 \geq z_4$ . This be rewritten as  $z_1 \geq z_3$ ;  $z_2 \geq z_3$ ;  $z_3 \geq z_4$ , which is equivalent to

$$\left\langle A = \begin{bmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix}; \quad \mathbf{b} = \mathbf{0}; \quad \mathbf{c} = +\infty \right\rangle$$

## 4. Inference Under Linear Constraints

Under the TBM, MCMC-based inference without evidences is simple: we alternate between  $P(\mathbf{h} | \mathbf{x})$  and  $P(\mathbf{x} | \mathbf{h})$ . This is efficient because of the factorisations in Eqs. (2,3). Inference with inequality-based evidence  $\mathbf{e}$  is, however, much more involved except for the limiting case of point assignments.

Denote by  $\Omega(\mathbf{e}) = \{\mathbf{x} | \mathbf{b} \leq_{\odot} \mathbf{A}\mathbf{x} \leq_{\odot} \mathbf{c}\}$  the constrained domain of  $\mathbf{x}$  defined by the evidence  $\mathbf{e}$ . Now we need to specify and sample from the constrained distribution  $P(\mathbf{x}, \mathbf{h} | \mathbf{e})$  defined on  $\Omega(\mathbf{e})$ . Sampling  $P(\mathbf{h} | \mathbf{x})$  remains unchanged, and in what follows we focus on sampling from  $P(\mathbf{x} | \mathbf{h}, \mathbf{e})$ .

<sup>5</sup>This is different from saying that all the members of the subset must be observed.

#### 4.1. Inference under Boxed Constraints

For *boxed constraints* (Section 3.1), due to the conditional independence, we still enjoy the factorisation  $P(\mathbf{x} | \mathbf{h}, \mathbf{e}) = \prod_i P(x_i | \mathbf{h}, \mathbf{e})$ . We further have

$$P(x_i | \mathbf{h}, \mathbf{e}) = \frac{P(x_i | \mathbf{h})}{\Phi(c_i | \mathbf{h}) - \Phi(b_i | \mathbf{h})}$$

where  $\Phi(\cdot | \mathbf{h})$  is the normal cumulative distribution function of  $P(x_i | \mathbf{h})$ . Now  $P(x_i | \mathbf{h}, \mathbf{e})$  is a truncated normal distribution, from which we can sample using the simple rejection method, or more advanced methods such as those in (Robert, 1995).

#### 4.2. Inference under Inequality Constraints

For general *inequality constraints* (Section 3.2), the input variables are interdependent due to the linear transform  $A$ . However, we can specify the conditional distribution  $P(x_i | \mathbf{x}_{-i}, \mathbf{h}, \mathbf{e})$  (here  $\mathbf{x}_{-i} = \mathbf{x} \setminus x_i$ ) by realising that

$$b_m - \sum_{j \neq i} A_{mj} x_j \leq A_{mi} x_i \leq c_m - \sum_{j \neq i} A_{mj} x_j$$

where  $A_{mi} \neq 0$  for  $m = 1, 2, \dots, M$ . In other words,  $x_i$  is conditionally box-constrained given other variables.

This suggests a Gibbs procedure by looping through  $x_1, x_2, \dots, x_N$ . With some abuse of notation, let  $\tilde{b}_{mi} = (b_m - \sum_{j \neq i} A_{mj} x_j) / A_{mi}$  and  $\tilde{c}_{mi} = (c_m - \sum_{j \neq i} A_{mj} x_j) / A_{mi}$ . The constraints can be summarised as

$$\begin{aligned} x_i &\in \bigcap_{m=1}^M \left[ \min \{ \tilde{b}_{mi}, \tilde{c}_{mi} \}, \max \{ \tilde{b}_{mi}, \tilde{c}_{mi} \} \right] \\ &= \left[ \max_m \min \{ \tilde{b}_{mi}, \tilde{c}_{mi} \}, \min_m \max \{ \tilde{b}_{mi}, \tilde{c}_{mi} \} \right] \end{aligned}$$

The min and max operators are needed to handle change in inequality direction with the sign of  $A_{mi}$ , and the join operator is due to multiple constraints.

For more sophisticated Gibbs procedures, we refer to the work in (Geweke, 1991).

#### 4.3. Estimating the Binary Posteriors

We are often interested in the posteriors  $P(\mathbf{h} | \mathbf{e})$ , e.g., for further processing. Unlike the standard RBMs, the binary latent variables here are coupled through the unknown Gaussians and thus there are no exact solutions unless the evidences are all point assignments. The MCMC-based techniques described above offer an approximate estimation by averaging the samples  $\left\{ \mathbf{h}^{(s)} \right\}_{s=1}^S$ . For the case of boxed constraints, mean-field offers an alternative approach which may be numerically faster. In particular, the mean-field updates are recursive:

$$\begin{aligned} Q_k &\leftarrow \frac{1}{1 + \exp \{ -\gamma_k - \sum_i W_{ik} \hat{\mu}_i \}} \\ \mu_i &\leftarrow \alpha_i + \sum_k W_{ik} Q_k \\ \hat{\mu}_i &= \mu_i + \frac{\phi(b_i - \mu_i) - \phi(c_i - \mu_i)}{\Phi(c_i - \mu_i) - \Phi(b_i - \mu_i)} \end{aligned}$$

where  $Q_k$  is the probability of the unit  $k$  being activated,  $\hat{\mu}_i$  is the mean of the normal distribution truncated in the interval  $[b_i, c_i]$ ,  $\phi(z)$  is the probability density function, and  $\Phi(z)$  is normal cumulative distribution function. Interested readers are referred to the Supplement<sup>6</sup> for more details.

#### 4.4. Estimating Probability of Evidence Generation

Given the hidden states  $\mathbf{h}$  we want to estimate the probability that hidden states generate a particular evidence  $\mathbf{e}$

$$P(\mathbf{e} | \mathbf{h}) = \int_{\Omega(\mathbf{e})} P(\mathbf{x} | \mathbf{h}) d\mathbf{x}$$

For boxed constraints, analytic solution is available since the Gaussian variables are decoupled, i.e.,  $P(e_i | \mathbf{h}) = \Phi(c_i - \mu_i) - \Phi(b_i - \mu_i)$ , where  $\mu_i = \alpha_i + \sum_k W_{ik} h_k$ . For general inequality constraints, however, these variables are coupled by the inequalities. The general strategy is to sample from  $P(\mathbf{x} | \mathbf{h})$  and compute the portion of samples falling into the constrained domain  $\Omega(\mathbf{e})$ . For certain classes of inequalities we can approximate the Gaussian by appropriate distributions from which the integration has the closed form. In particular, those inequalities imposed by the categorical and rank evidences can be dealt with by using the *extreme value distributions*. The integration will give the logit form on distribution of categories and Plackett-Luce distribution of ranks. For details, we refer to the Supplement.

### 5. Stochastic Gradient Learning with Persistent Markov Chains

Learning is based on maximising the evidence likelihood

$$\mathcal{L} = \log P(\mathbf{e}) = \log \sum_{\mathbf{h}} \int_{\Omega(\mathbf{e})} P(\mathbf{h}, \mathbf{x}) d\mathbf{x}$$

where  $P(\mathbf{h}, \mathbf{x})$  is defined in Eq. (1). Let  $Z(\mathbf{e}) = \sum_{\mathbf{h}} \int_{\Omega(\mathbf{e})} \exp \{ -E(\mathbf{x}, \mathbf{h}) \} d\mathbf{x}$ , then  $\mathcal{L} = \log Z(\mathbf{e}) - \log Z$ . The gradient w.r.t. the mapping parameter reads

$$\partial_{W_{ik}} \mathcal{L} = \mathbb{E}_{P(x_i, h_k | \mathbf{e})} [x_i h_k] - \mathbb{E}_{P(x_i, h_k)} [x_i h_k] \quad (7)$$

<sup>6</sup>Supplement material will be available at: <http://truyen.vietlabs.com>

The derivation is left to the Supplement.

### 5.1. Estimating Data Statistics

The *data-dependent* statistics  $\mathbb{E}_{P(x_i, h_k | \mathbf{e})} [x_i h_k]$  and the *data-independent* statistics  $\mathbb{E}_{P(x_i, h_k)} [x_i h_k]$  are not tractable to compute in general, and thus approximations are needed.

**Data-dependent statistics.** Under the box constraints, the mean-field technique (Section 4.3) can be employed as follows

$$\mathbb{E}_{P(x_i, h_k | \mathbf{e})} [x_i h_k] \approx \hat{\mu}_i Q_k$$

For general cases, sampling methods are applicable. In particular, we maintain one persistent Markov chain (Younes, 1989; Tieleman, 2008) per data instance and estimate the statistics after a very short run. This would explore the space of the data-dependent distribution  $P(\mathbf{x}, \mathbf{h} | \mathbf{e})$  by alternating between  $P(\mathbf{h} | \mathbf{x})$  and  $P(\mathbf{x} | \mathbf{h}, \mathbf{e})$  using techniques described in Section 4.

**Data-independent statistics.** Mean-field distributions are not appropriate for exploring the entire state space because they tend to fit into one mode. One practical solution is based on the idea of Hinton’s Contrastive Divergence (CD), where we create another Markov chain on-the-fly starting from the latest state of the clamped chain. This chain will be discarded after each parameter update. This is particular useful when the models are instance-specific, e.g., in collaborative filtering, it is much cheaper to build one model per user, all share the same parameters. If it is not the case, then we can maintain a moderate set of parallel chains and collect the samples after a short run at every updating step (Younes, 1989; Tieleman, 2008).

### 5.2. Learning the Box Boundaries

In the case of boxed constraints, sometimes it is helpful to learn the boundaries  $\langle b_i, c_i \rangle$  themselves. The gradient of the log-likelihood w.r.t. the lower boundaries reads

$$\begin{aligned} \partial_{b_i} \mathcal{L} &= \frac{1}{Z(\mathbf{e})} \sum_{\mathbf{h}} \partial_{b_i} \int_{\Omega(\mathbf{e})} \exp \{-E(\mathbf{x}, \mathbf{h})\} d\mathbf{x} \\ &= \sum_{\mathbf{h}} \partial_{b_i} \int_{\Omega(\mathbf{e})} P(\mathbf{x}, \mathbf{h} | \mathbf{e}) d\mathbf{x} \\ &\approx \frac{1}{S} \partial_{b_i} \int_{b_i}^{c_i} \int_{\Omega(\mathbf{e}) \setminus [b_i, c_i]} P(\mathbf{x} | \mathbf{h}^{(s)}, \mathbf{e}) d\mathbf{x}_{-i} dx_i \\ &= -\frac{1}{S} \sum_{\mathbf{h}} P(x_i = b_i | \mathbf{h}^{(s)}, \mathbf{e}) \end{aligned}$$

where  $\{\mathbf{h}^{(s)}\}_{s=1}^S$  are samples collected during the MCMC procedure running on the data-dependent distribution  $P(\mathbf{x}, \mathbf{h} | \mathbf{e})$ . Similarly we would have the

gradient w.r.t. the upper boundaries:

$$\partial_{c_i} \mathcal{L} = \frac{1}{S} \sum_{\mathbf{h}} P(x_i = c_i | \mathbf{h}^{(s)}, \mathbf{e}).$$

## 6. Applications

In this section, we describe applications of the TBM for three realistic domains, namely handwritten digit recognition, collaborative filtering and worldwide survey analysis. Before going to the details, let us first address key implementation issues (see Supplement for more details).

One observed difficulty in training the TBM is that the hidden samples can get stuck in one of the two ends and thus learning cannot progress. The reasons might be the large mapping parameters or the unbounded nature of the underlying Gaussian variables, which can saturate the hidden units. We can control the norm of the mapping parameters, either by using the standard  $\ell_2$ -norm regularisation, or by rescaling the norm of the parameter vector for each hidden unit. To deal with the non-boundedness of the Gaussian variables, then we can restrict their range, making the model bounded.

Another effective solution is to impose a constraint on the posteriors by adding the regularising term to the log-likelihood, e.g.,

$$\lambda \left\{ \sum_k [\rho \log P(h_k^1 | \mathbf{e}) + (1 - \rho) \log (1 - P(h_k^1 | \mathbf{e}))] \right\}$$

where  $\rho \in (0, 1)$  is the expected probability that a hidden unit will turn on given the evidence and  $\lambda > 0$  is the regularisation weight. Maximising this quantity is essentially minimising the Kullback-Leibler divergence between the expected posteriors and the true posteriors. In our experiments, we found  $\rho \in (0.1, 0.3)$  and  $\lambda \in (0.1, 1)$  gave satisfying results.

The main technical issue is that  $P(h_k^1 | \mathbf{e})$  does not have a simple form due to the integration over all the constrained Gaussian variables. Approximation is thus needed. The use of mean-field methods will lead to the simple sigmoid form, but it is only applicable for boxed constraints since it breaks down deterministic constraints among variables (Section 4.3). However, we can estimate the “mean” truncated Gaussian  $\hat{\mu}_i$  by averaging the recent samples of the Gaussian variables in the data-dependent phase.

Once these safeguards are in place, learning can greatly benefit from quite large learning rate and small batches as it appears to quickly get the samples out off the local energy traps by significantly distorting the energy landscape. Depending on the problem sizes, we vary the batch sizes in the range [100, 1000].

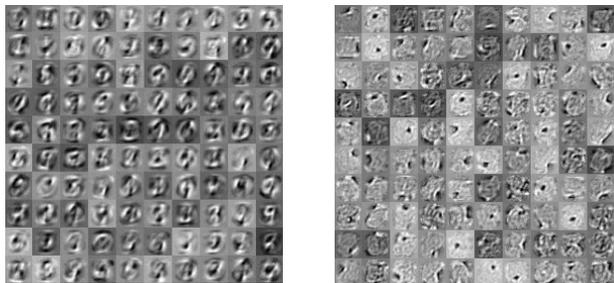


Figure 1. 100 MNIST feature weights (one image per hidden unit) learned by Probit RBM (left) and RBM with CD-1 (right).

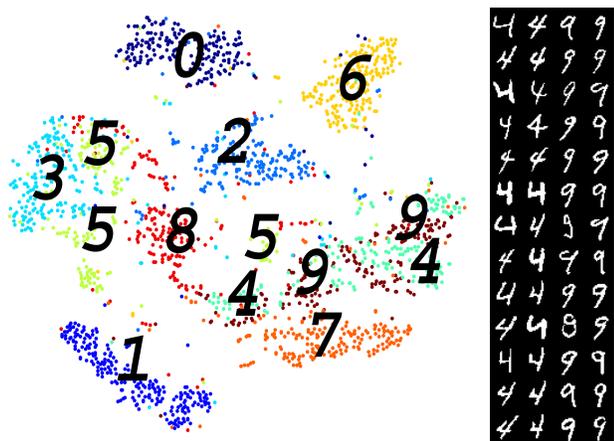


Figure 2. t-SNE visualisation of the learned representations on MNIST (left) and random samples of two confused digits (4,9) (right). Best viewed in colours.

### 6.1. Probit RBM for Handwritten Digits

We use the name Probit RBM to denote the special case of TBM where the observations are binary (i.e., boxed constraints, see Section 3.1). The threshold  $\theta_i$  for each visible unit  $i$  is chosen so that under the zero mean, the probability of generating a binary evidence equals the empirical probability, i.e.,  $1 - \Phi(\theta_i) = \bar{e}_i$ , and thus  $\theta_i = \Phi^{-1}(1 - \bar{e}_i)$ . Since any mismatch in thresholds can be corrected by shifting the corresponding biases, we do not need to update the thresholds further.

We report here the result of the mean-field method for computing data-dependent statistics, which are averaged over a random batch of 500 images. For the data-independent statistics, 500 persistent chains are run in parallel with samples collected after every 5 Gibbs steps. The sparsity level  $\rho$  is set to 0.3 and the sparseness weight  $\lambda$  is set to 0.5. Once the model has been learned, mean-field is used to estimate the hidden posteriors. Typically this mean-field is quite fast as it converges in a few steps.

We take the data from MNIST and binarize the im-

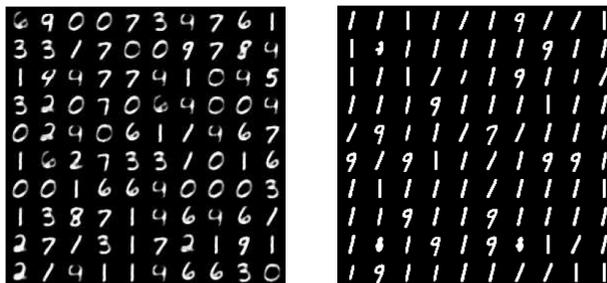


Figure 3. Samples generated from the Probit RBM (left) and RBM with CD-1 (right) by mode seeking procedures (see Supplement).

ages using a mid-intensity threshold. The learned representation is shown in Figure 2. Most digits are well separated in 2D except for digits 4 and 9. The learned representation can be used for classifications, e.g., by feeding to the multiclass logistic classifier. For 500 hidden units, the Probit RBM achieves the error rate of 3.28%, comparable with those obtained by the RBM trained with CD-1 (3.02%), and much better than the raw pixels (8.46%). The features discovered by the Probit RBM and RBM with CD-1 are very different (Figure 1), and this is expected because they operate on different input representations. The energy surface learned by the Probit RBM is smooth enough to allow efficient exploration of modes, as shown in Figure 3.

### 6.2. Rank Evidences for Collaborative Filtering

In collaborative filtering, one of the main goals is to produce a personalized ranked list of items. Until very recently, the majority in the area, on the other hand, focused on predicting the ratings, which are then used for ranking items. It can be arguably more efficient to learn a rank model directly instead of going through the intermediate steps.

We build one TBM for ranking with ties (i.e., inequality constraints, see Section 3.2) per user due to the variation in item choices but all the TBMs share the same parameter set. The handling of ties is necessary because during training, many items share the same rating. Unseen items are simply not accounted for in each model: We only need to compare the utilities between the items seen by each person. The result is that the models are very sparse and fast to run. For the data-dependent statistics, we maintain one Markov chain per user. Since there is no single model for all data instances, the data-independent statistics cannot be estimated from a small set of Markov chains. Rather we also maintain a data-independent chain per data instance, which can be persistent on their own, or restarted from the data-dependent chains after every

	<i>ERR</i>	<i>N@1</i>	<i>N@5</i>	<i>N@10</i>
Item popularity	0.587	0.560	0.680	0.835
ListRank.MF	0.653	0.673	0.751	0.873
PMOP	0.648	0.664	0.747	0.871
TBM	<b>0.678</b>	<b>0.722</b>	<b>0.792</b>	<b>0.893</b>

Table 1. Item ranking results on MovieLens – the higher the better ( $K = 50$ ). Here  $N@T$  is a shorthand for  $NDCG@T$ .

parameter updating step. The latter case, which is reported here, is in the spirit of the Hinton’s Contrastive Divergence, where the data-independent chain is just a few steps away from the data-dependent chain.

Once the model has been trained, the hidden posterior vector  $\hat{\mathbf{h}} = (\hat{h}_1, \hat{h}_2, \dots, \hat{h}_K)$ , where  $\hat{h}_k = P(h_k = 1 | \mathbf{e}_u)$ , is used as the new representation of the tastes of user  $u$ . The rank of unseen movies is the mode of the distribution  $P(\mathbf{e}^* | \hat{\mathbf{h}})$ , where  $\mathbf{e}^*$  are the rank-based evidences (see Section 4.4). For fast computation of  $P(\mathbf{e}^* | \hat{\mathbf{h}})$ , we approximate the Gaussian by a Gumbel distribution, which leads to a simple way of ranking movies using the mean “utility”  $\mu_{ui} = \alpha_i + W_i \cdot \hat{\mathbf{h}}_u$  for user  $u$  (see the Supplement for more details).

The data used in this experiment is the MovieLens, which contains 1M ratings by approximately 6K users on 4K movies. To encourage diversity in the rank lists, we remove the top 10% most popular movies. We then remove users with less than 30 ratings on the remaining movies. The most recently rated 10 movies per user are held out for testing, the next most recent 5 movies are used for tuning hyper-parameters, and the rest for training.

For comparison, we implement a simple baseline using item popularity for ranking, and thus offering a naive non-personalized solution. For personalized alternatives, we implement two recent rank-based matrix factorisation methods, namely ListRank.MF (Shi et al., 2010) and PMOP (Truyen et al., 2011). Two ranking metrics from the information retrieval literature are used: the *ERR* (Chapelle et al., 2009) and the *NDCG@T* (Järvelin and Kekäläinen, 2002). These metrics place more emphasis on the top-ranked items. Table 1 reports the movie ranking results on test subset (each user is presented with a ranked list of unseen movies), demonstrating that the TBM is a clear winner in all metrics.

### 6.3. Mixed Evidences for World Attitude Analysis

Finally, we demonstrate the TBM on mixed evidences. The data is from the survey analysis domain, which mostly consists of multiple questions of different natures such as basic facts (e.g., ages and genders) and

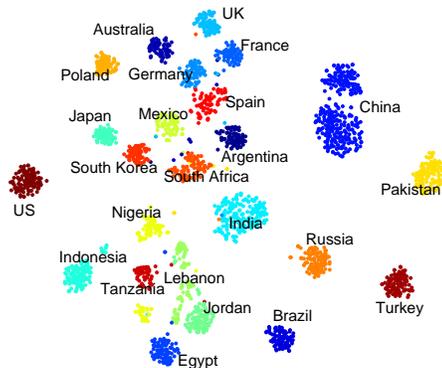


Figure 4. Distribution of global attitudes obtained by projecting the hidden posteriors (100 hidden units) on 2D using t-SNE. A dot represents one respondent. Best viewed in colours.

opinions (e.g., binary choices, single choices, multiple choices, ordinal judgments, preferences and ranks). The standard approach to deal with such heterogeneity is to perform the so-called “coding”, which converts types into some numerical representations (e.g., ordinal scales into stars, ranks into multiple pairwise comparisons) so that standard processing tools can handle. However, this coding process breaks the structure in the data and thus significant information will be lost. Thus our TBM offers a scalable and generic machinery to process the data in its native format and then convert the mixed types into a more homogeneous posterior vector.

We use the global attitude survey dataset collected by the PewResearch Centre<sup>7</sup>. The survey was conducted on 24,717 people from 24 countries during the period of March 17 – April 21, 2008 on a variety of topics concerning people’s life, opinions on issues in their countries and around the world as well as future expectations. There are 52 binary, 124 categorical (of variable category sizes), 3 continuous, 165 ordinal (of variable level sizes) question types.

Like the case of collaborative filtering, we build one TBM per respondent due to the variation in questions and answers but all the TBMs share the same parameter set. Unanswered/inappropriate questions are ignored. For each respondent, we maintain 2 persistent and non-interacting Markov chains for the data-dependent statistics and the data-independent statistics, respectively.

Figure 4 shows the 2D distribution of respondents from 24 countries obtained by feeding the posteriors to

<sup>7</sup>The datasets are publicly available from <http://pewresearch.org/>

the t-SNE (van der Maaten and Hinton, 2008) (here no explicit information of countries is used). It is interesting to see the cultural/social clustering and gaps between countries as opposed to the geographical distribution (e.g., between Indonesia and Egypt, Australia and UK and the relative separation of the China, Pakistan, Turkey and the US from the rest). To predict the 24 countries, we feed the posteriors into the standard multiclass logistic regression and achieve an error rate of 0.49%, suggesting that the TBM has captured the intrastate regularities and separated the interstate variations well.

## 7. Related Work

Latent multivariate Gaussian variables have been widely studied in statistical analysis, initially to model correlated binary data<sup>8</sup> (Ashford and Sowden, 1970; Chib and Greenberg, 1998) then now used for a variety of data types such as ordered categories (Kottas et al., 2005), unordered categories (Zhang et al., 2008), and the mixture of types (Dunson and Herring, 2005). Learning with the underlying Gaussian model is notoriously difficult for large-scale setting: independent sampling costs cubic time due to the need of inverting the covariance matrix, while MCMC techniques such as Gibbs sampling can be very slow if the graph is dense and the interactions between variables are strong. This can be partly overcome by adding one more layer of latent variables as in factor analysis (Wedel and Kamakura, 2001; Khan et al., 2010) and probabilistic principle component analysis (Tipping and Bishop, 1999). The main difference from our TBM is that those models are directed with continuous factors while ours is undirected with binary factors.

Gaussian RBMs have been used for modelling continuous data such as visual features (Hinton and Salakhutdinov, 2006), where the evidences are the value assignments, and thus a limiting case of our evidence system. Some restrictions to the continuous Boltzmann machines have been studied: In (Downs et al., 2000), Gaussian variables are assumed to be non-negative, and in (Yasuda and Tanaka, 2007), continuous variables are bounded. However, we do not make these restrictions on the model but rather placing restrictions during the training phase only. GRBMs that handle ordinal evidences have been studied in (Tran et al., 2012), which is an instance of the boxed-constraints in our TBM.

## 8. Discussion and Conclusion

Since the underlying variables of the TBM are Gaussian, various extensions can be made without much difficulty. For example, direct correlations among vari-

ables, *regardless of their types*, can be readily modelled by introducing the non-identity covariance matrix (Ranzato and Hinton, 2010). This is clearly a good choice for image modelling since nearby pixels are strongly correlated. Another situation is when the input units are associated with their own attributes. Each unit can be extended naturally by adding a linear combination of attributes to the mean structure of the Gaussian.

The additive nature of the mean-structure allows the natural extension to matrix modelling (e.g., see (Truyen et al., 2009; Tran et al., 2012)). That is, we do not distinguish the role of rows and columns, and thus each row and column can be modelled using their own hidden units (the row parameters and columns parameters are different). Conditioned on the row-based hidden units, we return to the standard TBM for column vectors. Inversely, conditioned on the column-based hidden units, we have the TBM for row vectors.

To sum up, we have proposed a generic class of models called Thurstonian Boltzmann machine (TBM) to unify many type-specific modelling problems and generalise them to the general problem of learning from multiple groups of inequalities. Our framework utilises the Gaussian restricted Boltzmann machines, but the Gaussian variables are never observed except for one limiting case. Rather, those variables are subject to inequality constraints whenever an evidence is observed. Under this representation, the TBM supports a very wide range of evidences, many of which were not possible before in the Boltzmann machine literature, *without* the need to specify type-specific models. In particular, the TBM supports any combination of the point assignments, intervals, censored values, binary, unordered categories, multi-categories, ordered categories, (in)-complete ranks with and without ties.

We demonstrated the TBM on three applications of very different natures, namely handwritten digit recognition, collaborative filtering and complex survey analysis. The results are satisfying and the performance is competitive with those obtained by type-specific models.

## References

- JR Ashford and RR Sowden. Multi-variate probit analysis. *Biometrics*, pages 535–546, 1970.
- U. Böckenholt. Thurstonian-based analyses: past, present, and future utilities. *Psychometrika*, 71(4):615–629, 2006.
- O. Chapelle, D. Metzler, Y. Zhang, and P. Grinspan. Expected reciprocal rank for graded relevance. In *CIKM*, pages 621–630. ACM, 2009.
- S. Chib and E. Greenberg. Analysis of multivariate probit models. *Biometrika*, 85(2):347–361, 1998.

<sup>8</sup>This is often known as *multivariate probit models*.

- O.B. Downs, D.J.C. MacKay, and D.D. Lee. The nonnegative Boltzmann machine. *Advances in Neural Information Processing Systems*, 12:428–434, 2000.
- D.B. Dunson and A.H. Herring. Bayesian latent variable models for mixed discrete outcomes. *Biostatistics*, 6(1):11, 2005.
- Y. Freund and D. Haussler. Unsupervised learning of distributions on binary vectors using two layer networks. *Advances in Neural Information Processing Systems*, pages 912–919, 1993.
- P.V. Gehler, A.D. Holub, and M. Welling. The rate adapting Poisson model for information retrieval and object recognition. In *Proceedings of the ICML*, pages 337–344, 2006.
- J. Geweke. Efficient simulation from the multivariate normal and student-t distributions subject to linear constraints and the evaluation of constraint probabilities. In *Computing science and statistics: Proceedings of the 23rd symposium on the interface*, pages 571–578, 1991.
- G.E. Hinton and R.R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- K. Järvelin and J. Kekäläinen. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4):446, 2002.
- E. Khan, B. Marlin, and K. Murphy. Variational bounds for mixed-data factor analysis. In *Proc. of Neural Information Processing Systems*, 2010.
- A. Kottas, P. Müller, and F. Quintana. Nonparametric Bayesian modeling for multivariate ordinal data. *Journal of Computational and Graphical Statistics*, 14(3):610–625, 2005.
- N. Le Roux, N. Heess, J. Shotton, and J. Winn. Learning a generative model of images by factoring appearance and shape. *Neural Computation*, 23(3):593–650, 2011.
- J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A.Y. Ng. Multimodal deep learning. In *ICML*, 2011.
- M.A. Ranzato and G.E. Hinton. Modeling pixel means and covariances using factorized third-order Boltzmann machines. In *CVPR*, pages 2551–2558. IEEE, 2010.
- C.P. Robert. Simulation of truncated normal variables. *Statistics and computing*, 5(2):121–125, 1995.
- R. Salakhutdinov and G. Hinton. Deep Boltzmann Machines. In *Proceedings of 20th AISTATS*, volume 5, pages 448–455, 2009a.
- R. Salakhutdinov and G. Hinton. Replicated softmax: an undirected topic model. *Advances in Neural Information Processing Systems*, 22, 2009b.
- R. Salakhutdinov, A. Mnih, and G. Hinton. Restricted Boltzmann machines for collaborative filtering. In *Proceedings of the 24th ICML*, pages 791–798, 2007.
- Y. Shi, M. Larson, and A. Hanjalic. List-wise learning to rank with matrix factorization for collaborative filtering. In *ACM RecSys*, pages 269–272. ACM, 2010.
- P. Smolensky. Information processing in dynamical systems: Foundations of harmony theory. *Parallel distributed processing: Explorations in the microstructure of cognition*, 1:194–281, 1986.
- N. Srivastava and R. Salakhutdinov. Multimodal learning with deep Boltzmann machines. In *NIPS*, pages 2231–2239, 2012.
- L.L. Thurstone. A law of comparative judgment. *Psychological review*, 34(4):273, 1927.
- T. Tieleman. Training restricted Boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th ICML*, pages 1064–1071, 2008.
- M.E. Tipping and C.M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B*, 61(3):611–622, 1999.
- T. Tran, D.Q. Phung, and S. Venkatesh. Mixed-variate restricted Boltzmann machines. In *Proc. of 3rd Asian Conference on Machine Learning (ACML)*, Taoyuan, Taiwan, 2011.
- T. Tran, D.Q. Phung, and S. Venkatesh. Cumulative restricted Boltzmann machines for ordinal matrix data analysis. In *Proc. of 4th Asian Conference on Machine Learning (ACML)*, Singapore, 2012.
- T. Truyen, D.Q. Phung, and S. Venkatesh. Probabilistic models over ordered partitions with applications in document ranking and collaborative filtering. In *Proc. of SIAM Conference on Data Mining (SDM)*, Mesa, Arizona, USA, 2011. SIAM.
- T.T. Truyen, D.Q. Phung, and S. Venkatesh. Ordinal Boltzmann machines for collaborative filtering. In *Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI)*, Montreal, Canada, June 2009.
- L. van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(2579-2605):85, 2008.
- M. Wedel and W.A. Kamakura. Factor analysis with (mixed) observed and latent variables in the exponential family. *Psychometrika*, 66(4):515–530, 2001.
- E. Xing, R. Yan, and A.G. Hauptmann. Mining associated text and images with dual-wing harmoniums. In *Proceedings of the 21st UAI*, 2005.
- M. Yasuda and K. Tanaka. Boltzmann machines with bounded continuous random variables. *Interdisciplinary Information Sciences*, 13(1):25–31, 2007.
- L. Younes. Parametric inference for imperfectly observed Gibbsian fields. *Probability Theory and Related Fields*, 82(4):625–645, 1989.
- X. Zhang, W.J. Boscardin, and T.R. Belin. Bayesian analysis of multivariate nominal measures using multivariate multinomial probit models. *Computational statistics & data analysis*, 52(7):3697–3708, 2008.