From RBMs to BN2A Models: Parameter Transformation for Interpretable Educational Diagnostics

Iván Pérez[1,2,3](✉), Jiří Vomlel[1], and Patrícia Martinková[2]

[1] Institute of Information Theory and Automation, Czech Academy of Sciences, Prague, Czechia
{cabrera,vomlel}@utia.cas.cz
[2] Institute of Computer Science, Czech Academy of Sciences, Prague, Czechia
martinkova@cs.cas.cz
[3] Faculty of Mathematics and Physics, Charles University, Prague, Czechia

Abstract. Restricted Boltzmann Machines (RBMs) are bipartite graphical models with binary latent and observed variables that have shown promise for representation learning. However, their lack of interpretable parameters limits their utility in domains requiring explainability, like educational assessment. Despite extensive RBM research, non-negativity constraints on weights—essential for monotonicity in educational contexts—remain largely unexplored. To address this, we propose a method to translate RBMs into a specialized class of bipartite Bayesian networks, which we term BN2A networks, characterized by strict 2-layer separation (hidden and observed variables), Noisy-AND conditional probability tables, and directly interpretable parameters for educational models. Our work establishes a mathematical transformation from RBM weights to BN2A's interpretable parameters (leak and penalty probabilities), theoretical analysis showing BN2A's constrained connectivity is a subset of RBM architectures, and empirical evidence that the transformation preserves model fidelity under realistic conditions. By bridging these paradigms, our method leverages RBM's representational power while achieving BN2A's interpretability, opening new possibilities for adaptive learning systems and diagnostic tools.

Keywords: Restricted Boltzmann Machines · BN2A models · Educational Assessment · Interpretable machine learning

1 Introduction

Bayesian networks (BNs) [8,9,13] are a widely used framework for modeling probabilistic relationships between random variables. This paper develops an analytical framework for a specific subclass of BNs that we termed BN2A [15]: BNs with 2-layer structure and noisy-AND interactions. In this model, the top layer comprises mutually independent hidden (unobservable) binary variables, while the bottom layer consists of observed binary variables that depend exclusively on the hidden layer. These dependencies are characterized by conditional

K. Sauerwald and M. Thimm (Eds.): ECSQARU 2025, LNAI 16099, pp. 104–117, 2026.
https://doi.org/10.1007/978-3-032-05134-9_8

probability tables (CPTs) based on noisy-AND models, a well-established family of canonical CPT representations [2,6]. Figure 1 illustrates the directed bipartite graph structure of a BN2A model.

The BN2A framework is particularly relevant in psychometrics [14], where it is applied to cognitive diagnostic modeling of students. Here, hidden variables represent latent skills, while observed variables correspond to test responses. A key feature of this domain is that correct answers typically require mastery of all relevant skills, though compensatory mechanisms may exist (e.g., alternative knowledge compensating for a missing skill). The Noisy-AND model naturally captures this "conjunctive" relationship between skills and responses.

Restricted Boltzmann Machines (RBMs) [7,18] represent another two-layer architecture widely used for unsupervised learning. Like BN2A models, RBMs consist of visible and hidden units with restricted connectivity, enabling them to learn complex data distributions [4]. A key distinction, however, is that RBM connections are undirected rather than directed. Recent studies have demonstrated their potential in psychometrics, particularly for modeling latent traits and response patterns [10].

Despite nearly four decades of research on RBMs, there is limited literature addressing non-negativity constraints on RBM weights. This gap is particularly relevant in educational contexts where monotonicity properties (a direct consequence of non-negativity constraints) are essential, students with higher skill levels should have higher probabilities of correct responses.

Although prior work has explored connections between BNs and neural networks [1], studies specifically linking RBMs to BNs interpretations are limited. This work demonstrates how constrained RBMs with non-negative weights can be effectively applied to educational diagnosis, and more importantly, how their parameters can be directly interpreted through transformation to BN2A parameters, bridging the gap between representational power and interpretability in cognitive diagnostic modeling.

The paper is organized as follows: Sect. 2 introduces RBMs and their mathematical formulations, while Sect. 3 formally defines BN2A models and their educational testing applications. In Sect. 4, we establish conditions linking BN2A models and RBMs, alongside a modified Contrastive Divergence (CD) algorithm for RBMs parameter learning. Experimental results are presented in Sect. 5, followed by conclusions in Sect. 6.

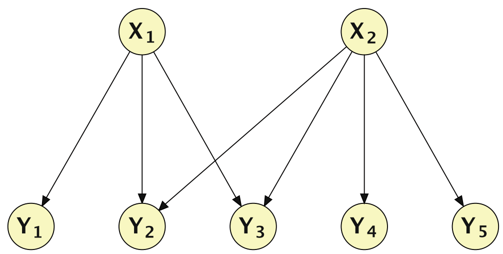

Fig. 1. An example of a BN2A model structure.

2 RBM Model

RBMs have emerged as powerful generative neural networks with applications spanning computer vision, recommendation systems, and educational assessment [7].

2.1 Definition and Structure

An RBM is a bipartite undirected graphical model consisting of two distinct layers. The visible layer ($\mathbf{v}$) contains binary observed variables, where $v_i \in \{0, 1\}$. The hidden layer ($\mathbf{h}$) comprises binary latent variables ($h_j \in \{0, 1\}$) that capture underlying latent features. The model exhibits *restricted connectivity*, meaning there are no intra-layer connections between visible-visible or hidden-hidden units (Fig. 2).

This architecture enables the model to learn latent representations through symmetric connections between layers. The bipartite structure ensures that units within the same layer are conditionally independent given the state of the other layer, simplifying inference procedures.

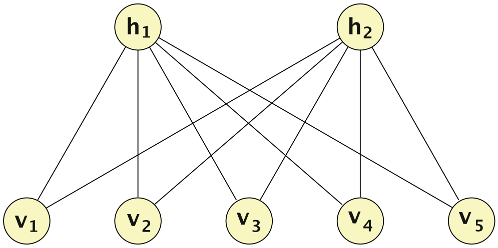

Fig. 2. Architecture of an RBM with 5 visible and 2 hidden units.

2.2 Mathematical Formulation

The model defines a probability distribution through an energy function that captures the compatibility between visible and hidden unit configurations:

$$E(\mathbf{v}, \mathbf{h}) = -\sum_i b_i v_i - \sum_j c_j h_j - \sum_{i,j} v_i W_{ij} h_j, \tag{1}$$

where b_i and c_j are biases for visible and hidden units respectively, and W_{ij} represents the weight connecting visible unit i to hidden unit j.

Lower energy configurations correspond to higher probability states. The joint probability distribution follows a Boltzmann distribution:

$$P(\mathbf{v}, \mathbf{h}) = \frac{1}{Z} e^{-E(\mathbf{v}, \mathbf{h})}, \quad Z = \sum_{\mathbf{v}, \mathbf{h}} e^{-E(\mathbf{v}, \mathbf{h})}. \tag{2}$$

The partition function Z ensures normalization but is generally intractable to compute. Due to the bipartite structure, conditional probabilities take sigmoid forms:

$$P(h_j = 1|\mathbf{v}) = \sigma\left(c_j + \sum_i v_i W_{ij}\right), \tag{3}$$

$$P(v_i = 1|\mathbf{h}) = \sigma\left(b_i + \sum_j h_j W_{ij}\right), \tag{4}$$

where $\sigma(x) = 1/(1+e^{-x})$ is the logistic function. These conditional independence properties enable efficient sampling through alternating Gibbs sampling between layers.

3 BN2A Model

Let $\mathbf{X}$ denote the vector $(X_1, \ldots, X_K)$ of K hidden variables, and similarly let $\mathbf{Y}$ denote the vector $(Y_1, \ldots, Y_L)$ of L observed dependent variables. The hidden variables are also called attributes or skills in the context of cognitive diagnostic models (CDMs) [19], and observed dependent variables are usually called items or questions in the same context. All variables are assumed to be binary, taking states from $\{0, 1\}$. The state space of the multidimensional variable $\mathbf{X}$ is denoted $\mathbb{X}$ and is equal to the Cartesian product of the state spaces of $X_k, k = 1, \ldots, K$:

$$\mathbb{X} = \times_{k=1}^{K} \mathbb{X}_k = \{0, 1\}^K. \tag{5}$$

Similarly, the state space of multidimensional variable $\mathbf{Y}$ is denoted $\mathbb{Y}$ and is equal to the Cartesian product of state spaces of $Y_\ell, \ell = 1, \ldots, L$:

$$\mathbb{Y} = \times_{\ell=1}^{L} \mathbb{Y}_\ell = \{0, 1\}^L. \tag{6}$$

The basic building blocks of a BN2A model are conditional probability tables (CPTs) specified in the form of a Noisy-AND model. Let Y_ℓ be an observed dependent variable and $pa(Y_\ell)$ be the subset of indices of related variables from $\mathbf{X}$. They are referred to as the parents of Y_ℓ.

Definition 1 (Noisy-AND model).
A conditional probability table $P(Y_\ell|\mathbf{X}_{pa(Y_\ell)})$ represents a Noisy-AND model if

$$P(Y_\ell = y_\ell|\mathbf{X}_{pa(Y_\ell)} = \mathbf{x}_{pa(Y_\ell)}) = \begin{cases} q_{\ell,0} \cdot \prod_{i \in pa(Y_\ell)} (q_{\ell,i})^{(1-x_i)} & \text{if } y_\ell = 1 \\ 1 - q_{\ell,0} \cdot \prod_{i \in pa(Y_\ell)} (q_{\ell,i})^{(1-x_i)} & \text{if } y_\ell = 0. \end{cases} \tag{7}$$

Note that if $x_i = 1$ then $(q_{\ell,i})^{(1-x_i)} = 1$ and if $x_i = 0$ then $(q_{\ell,i})^{(1-x_i)} = q_{\ell,i}$. The interpretation is that if $X_i = 1$, then this variable definitely enters the AND relation with the value 1. If $X_i = 0$, then there is still a probability $q_{\ell,i}$ that it enters the AND relation with value 1. The model also contains an auxiliary parent X_0 which is always 0 and thus enters the AND relation with probability $q_{\ell,0}$ for the value 1. This probability is traditionally called *leak* probability and allows non-zero probability of $Y_\ell = 0$ even if all parents of Y_ℓ have value 1. In educational diagnosis, the leak parameters $q_{\ell,0}$ represent the probability of answering question ℓ correctly when the student masters all required skills, while the parameters $q_{\ell,k}$ represent penalty factors for lacking the k-th skill when answering question ℓ. In the area of psychometrics, this model belongs to CDMs and is known as the Reduced Reparametrized Unified Model (RRUM) [5].

The prior probability of the hidden skill for $k = 1, \dots, K$ is defined as

$$P(X_k = x_k) = (p_k)^{x_k}(1 - p_k)^{(1-x_k)}, \tag{8}$$

which means that if $x_k = 1$ then it is p_k and if $x_k = 0$ then it equals $1 - p_k$.

Now we are ready to define a special class of BNs models with hidden variables, called BN2A model.

Definition 2 (BN2A model).
A BN2A model is a pair (G, P), where G is a directed bipartite graph with its nodes divided into two layers. The nodes of the top layer correspond to the hidden variables $X_1, \dots, X_K$ and the nodes of the bottom layer correspond to the observed variables $Y_1, \dots, Y_L$. All edges are directed from a node of the top layer to a node of the bottom layer. The symbol P refers to the joint probability distribution over the variables corresponding to the nodes of the graph G. The probability distribution is parameterized by a vector of model parameters $(\mathbf{p}, \mathbf{q})$:

$$(\mathbf{p}, \mathbf{q}) = \left((p_k)_{k\in\{1,\dots,K\}}, (q_{\ell,k})_{\ell\in\{1,\dots,L\}, k\in\{0\}\cup pa(Y_\ell)}\right). \tag{9}$$

The bipartite graph G of a BN2A model can also be specified by an incidence matrix, in the context of CDM, traditionally denoted by $\mathbf{Q}$. A $\mathbf{Q}$-matrix is an $L \times K$ binary matrix, with entries $\mathbf{Q}_{\ell,k} \in \{0, 1\}$ that indicate whether or not the ℓ^{th} observed dependent variable is linked to the k^{th} hidden variable:

$$\mathbf{Q}[\ell, k] = \begin{cases} 1 \text{ if } X_k \in pa(Y_\ell) \\ 0 \text{ otherwise.} \end{cases}$$

Monotonicity in BN2A

In the context of BNs, the concept of monotonicity constraint has been discussed in the literature for a long time [3,20]. More recent papers in this topic are [17], [12] and [16].

BNs model the probabilistic influences between its variables. Considering binary variables X and Y, a positive qualitative influence of a variable X on a variable Y along an arc $X \to Y$ in the network means that the occurrence of X increases the probability of Y occurs, assuming that the values of the other parents of Y remain the same. It means that

$$P(Y = 1|X = 1, \mathbf{z}) \geq P(Y = 1|X = 0, \mathbf{z}),$$

for any combination of values $\mathbf{z}$ for the set of parents of Y other than X [11,12]. In the context of educational testing, a positive influence is commonly assumed since mastering the skill X increases the probability of answering correctly the question Y.

BN2A models inherently satisfy the monotonicity property through their positive qualitative influences. Consider a variable Y with parent set $pa(Y)$ and two parent configurations $\mathbf{x}_{pa(Y)}$ and $\mathbf{x}'_{pa(Y)}$ that differ only in one skill: $x_i = 1$ in $\mathbf{x}_{pa(Y)}$ and $x'_i = 0$ in $\mathbf{x}'_{pa(Y)}$, while all other entries remain identical.

Using the Noisy-AND conditional probability from Eq. 7, for configuration $\mathbf{x}'_{pa(Y)}$ where $x'_i = 0$, the probability becomes:

$$P(Y_\ell = 1|\mathbf{x}'_{pa(Y_\ell)}) = q_{\ell,0} \cdot q_{\ell,i} \cdot \prod_{j \in pa(Y_\ell), j \neq i} q_{\ell,j}^{(1-x'_j)}. \tag{10}$$

The key difference is the additional penalty factor $q_{\ell,i}$ that appears when $x'_i = 0$. Since $0 < q_{\ell,i} \leq 1$ by model constraints, this factor can only reduce the probability. Therefore:

$$P(Y_\ell = 1|\mathbf{x}_{pa(Y_\ell)}) \geq P(Y_\ell = 1|\mathbf{x}'_{pa(Y_\ell)}). \tag{11}$$

4 Relation Between RBM and BN2A Models

While RBMs and BN2A models arise from different theoretical frameworks, neural networks and BNs, their two-layer architectures reveal intriguing parallels. Both models use hidden variables to influence observed responses, yet differ substantially in interpretability and parameter constraints.

The structural similarity becomes apparent when examining core components. BN2A models explicitly define hidden-observed relationships through prior probabilities $\boldsymbol{p}$ (analogous to RBM hidden activations $\mathbf{h}$) and structured parameters: leak probabilities and penalty parameters that, in educational contexts, determine how missing skills affect response probabilities. RBMs capture these relationships through a learned weight matrix $\mathbf{W}$ and bias terms, but without explicit cognitive interpretation.

The key difference lies in their design approaches. BN2A models prioritize interpretability and measurement consistency, constraining parameters to ensure meaningful relationships. In educational contexts, each penalty parameter $q_{j,k}$ directly represents the probability reduction when skill k is absent for question j, while $q_{j,0}$ represents success probability for question j when all required skills are present. RBMs optimize for representational flexibility, allowing weights to take any values that minimize reconstruction error, potentially sacrificing cognitive interpretability. This distinction leads to critical differences in monotonicity satisfaction. While BN2A models inherently satisfy monotonicity through constrained parameter spaces, RBMs require explicit modifications to ensure this fundamental property.

4.1 Lack of Monotonicity in Standard Binary RBMs

RBMs with binary units do not inherently satisfy the monotonicity property when interpreted in a cognitive diagnostic context, where hidden units represent latent skills and visible units correspond to test questions. In a binary RBM, the conditional probability of a visible unit v_i being active given the hidden layer configuration $\mathbf{h}$ is defined in Eq. 4.

The fundamental issue arises from the fact that weights W_{ij} can take negative values in standard RBM formulations. When $W_{ij} < 0$, activating hidden unit j (i.e., mastering skill j) actually decreases the logit and consequently reduces the probability of correctly answering question i. This counterintuitive behavior violates the monotonicity assumption that acquiring additional skills should not decrease performance. Consider two skill mastery patterns $\mathbf{h}$ and $\mathbf{h}'$ where $\mathbf{h} \preceq \mathbf{h}'$ (meaning $h_j \leq h'_j$ for all j). If there exists at least one negative weight $W_{ij} < 0$ for a skill j where $h_j < h'_j$, then the contribution $W_{ij}(h'_j - h_j) < 0$ to the logit sum, potentially making $P(v_i = 1|\mathbf{h}') < P(v_i = 1|\mathbf{h})$, thus violating monotonicity.

4.2 Ensuring Monotonicity Through Non-negative Weight Constraints

Monotonicity can be guaranteed in binary RBMs by imposing non-negative constraints on the weight matrix, i.e., $W_{ij} \geq 0$ for all i, j. Under this restriction, the model becomes monotonic with respect to hidden unit activations. Formally, for any two hidden layer configurations $\mathbf{h}$ and $\mathbf{h}'$ such that $\mathbf{h} \preceq \mathbf{h}'$, we have:

$$\sum_j W_{ij} h'_j + b_i \geq \sum_j W_{ij} h_j + b_i.$$

This inequality holds because $W_{ij} \geq 0$ and $h'_j \geq h_j$ for all j, ensuring that $\sum_j W_{ij}(h'_j - h_j) \geq 0$. Since the sigmoid function $\sigma(\cdot)$ is strictly increasing, this logit ordering directly translates to probability ordering: $P(v_i = 1|\mathbf{h}') \geq P(v_i = 1|\mathbf{h})$. This constrained RBM formulation aligns with the cognitive interpretation where mastering additional skills should either maintain or improve test performance.

The non-negative weight constraint can be implemented during training through various approaches, including projected gradient methods, penalty functions, or weight reparameterization techniques. Among these alternatives, projection based methods offer computational efficiency and direct interpretability. In our approach, we implement this constraint within an algorithm by applying weight projection after each parameter update. Specifically, we enforce non-negativity through element-wise projection: $W \leftarrow \max(W, 0)$, ensuring that all weights remain non-negative throughout the training process.

4.3 From RBM to BN2A Parameters

Once we restrict the RBM model to non-negative weights, we can interpret the leak and penalty parameters of the BN2A model in terms of the biases of the observed variables and the weights connecting the hidden and observed layers in the RBM.

The relationship between both parameter sets can be derived from Eqs. 7 and 4.

$$q_{\ell,0} = \sigma\left(b_j + \sum_k W_{jk}\right), \tag{12}$$

$$q_{\ell,i} = \sigma\left(b_j + \sum_{k \neq i} W_{jk}\right) \cdot \frac{1}{q_{\ell,0}}, \tag{13}$$

where $\sigma(x) = 1/(1 + e^{-x})$ is the logistic function.

4.4 Algorithm

A key motivation for this research was the observation that RBM and BN2A models share very similar structures. Initially, we applied the well-established Contrastive Divergence (CD) [4,7] algorithm to data generated from a BN2A model (Sect. 5), but the resulting RBM exhibited some negative weights. We were unable to find a meaningful interpretation connecting these negative weights, or combinations thereof, to the parameters of the original BN2A model. This limitation motivated us to propose a modification to the CD algorithm.

The proposed approach extends the original CD algorithm for RBMs by incorporating two domain-specific constraints critical for educational testing: structural correspondence and non-negativity of weights.

Model structure: The algorithm enforces a binary mask (Q-matrix) that defines permissible connections between observed questions and hidden skills. During initialization, weights $\mathbf{W}$ are element-wise multiplied ($\odot$) by the Q-matrix $\mathbf{Q}$, ensuring zero weights for prohibited connections. This masking operation is reapplied after each weight update, maintaining the predefined structure throughout training.

Non-Negativity: Weights are constrained via $\max(\mathbf{W}, 0)$ to guarantee non-negative values, ensuring monotonicity between skill mastery and correct response probability, a cognitive assumption stating that higher skill proficiency cannot decrease the probability of correct responses.

The training process follows these key phases for each mini-batch:

Positive Phase: Computes hidden unit activations $\mathbf{h}_{\text{prob}}$ from observed data using the sigmoidal function σ, followed by stochastic binarization.

Negative Phase: Reconstructs visible units $\mathbf{V}_{\text{recon}}$ from sampled hidden states, then recomputes hidden activations.

Gradient Update: Adjusts parameters using the difference between positive and negative phase statistics, scaled by the learning rate η. Biases $\mathbf{b}$ and $\mathbf{c}$ are updated via gradient estimates [4].

This constrained approach maintains the computational efficiency of standard RBMs ($O(NLK)$ per epoch for N samples, L questions, K skills) while

Input : $\mathbf{V} \in \{0,1\}^{N \times L}$ – Binary observed data
K – Number of hidden units
$\mathbf{Q} \in \{0,1\}^{L \times K}$ – Mask matrix (Q-matrix)
T – Number of epochs
B – Batch size
η – Learning rate
Output: $\mathbf{W} \in \mathbb{R}_+^{L \times K}$ – Non-negative weights
$\mathbf{b} \in \mathbb{R}^L$ – Visible layer biases
$\mathbf{c} \in \mathbb{R}^K$ – Hidden layer biases

Initialize $\mathbf{W} \leftarrow |\mathcal{N}(0, 0.01)| \odot \mathbf{Q}$;
Initialize $\mathbf{b} \leftarrow \mathbf{0}$, $\mathbf{c} \leftarrow \mathbf{0}$;
for *epoch* $t \leftarrow 1$ **to** T **do**
  for *each batch* $\mathbf{V}_{batch} \subset \mathbf{V}$ *(size B)* **do**
    Positive phase:;
      $\mathbf{h}_{\text{prob}} \leftarrow \sigma(\mathbf{V}_{\text{batch}}\mathbf{W} + \mathbf{1c}^T)$;
      $\mathbf{h}_{\text{states}} \leftarrow \text{Bernoulli}(\mathbf{h}_{\text{prob}})$;
    Negative phase:;
      $\mathbf{V}_{\text{recon}} \leftarrow \sigma(\mathbf{h}_{\text{states}}\mathbf{W}^T + \mathbf{1b}^T)$;
      $\mathbf{h}^{\text{neg}}_{\text{prob}} \leftarrow \sigma(\mathbf{V}_{\text{recon}}\mathbf{W} + \mathbf{1c}^T)$;
    Parameter update:;
      $\Delta\mathbf{W} \leftarrow \eta(\mathbf{V}^T_{\text{batch}}\mathbf{h}_{\text{prob}} - \mathbf{V}^T_{\text{recon}}\mathbf{h}^{\text{neg}}_{\text{prob}})/B$;
      $\mathbf{W} \leftarrow (\mathbf{W} + \Delta\mathbf{W}) \odot \mathbf{Q}$; `// Restrict structure`
      $\mathbf{W} \leftarrow \max(\mathbf{W}, 0)$; `// Enforce non-negativity`
      $\mathbf{b} \leftarrow \mathbf{b} + \eta \cdot \text{mean}(\mathbf{V}_{\text{batch}} - \mathbf{V}_{\text{recon}})$;
      $\mathbf{c} \leftarrow \mathbf{c} + \eta \cdot \text{mean}(\mathbf{h}_{\text{prob}} - \mathbf{h}^{\text{neg}}_{\text{prob}})$;
  end
end

Algorithm 1: Modified CD algorithm with Non-Negativity and Structure Constraints.

producing interpretable parameters that respect cognitive diagnostic requirements. The non-negativity constraint particularly differentiates it from conventional RBMs, guaranteeing plausible monotonic relationships.

5 Experiment

To demonstrate how we can learn RBM model parameters using Algorithm 1 and make these parameters interpretable by showing their relationship with BN2A model parameters, we conducted an experiment using a model with three hidden variables representing skills and 12 observed variables representing test questions. We used the structure shown in Fig. 3. In this model, note that some questions require only one skill, while the other questions require two skills to be answered correctly.

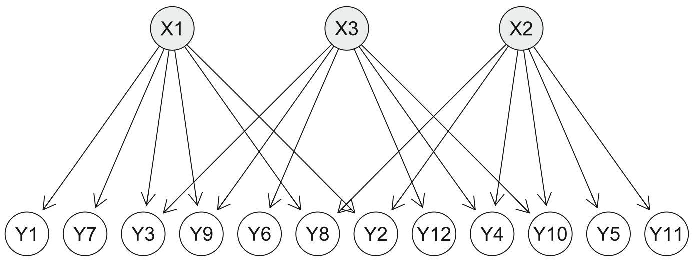

Fig. 3. BN2A model with three hidden variables and 12 observed variables.

As in any knowledge domain, some skills are easy to master and others are difficult; thus, tests typically require skills of varying mastery prevalence in the population. In our experiment, the simulated proportion of skill mastery (prior probabilities of mastering the skills) was set as $p_1 = 0.8$, $p_2 = 0.6$, and $p_3 = 0.4$. The leak parameters $q_{\ell,0}$ were selected to range from 0.6 to 0.95, and the penalty parameters $q_{\ell,k}$ were selected in the range of 0.1 to 0.4. These ranges are considered realistic for datasets where BN2A models are used in educational diagnosis [5]. The parameters of the model used in our experiment are presented in Table 1.

With the structure in Fig. 3 and the above parameters, we randomly generated a dataset $\mathbf{V}$ of size $n = 10^4$, a feasible sample size for large-scale standardized tests [14].

Table 1. True parameter values used to generate synthetic data for the BN2A model (structure shown in Fig. 3): leak and penalty parameters (left), and prior probabilities (right).

ℓ	$q_{\ell,0}$	$q_{\ell,1}$	$q_{\ell,2}$	$q_{\ell,3}$
1	0.900	0.400	-	-
2	0.800	0.200	0.100	-
3	0.700	0.100	-	0.300
4	0.600	-	0.200	0.100
5	0.700	-	0.200	-
6	0.800	-	-	0.300
7	0.950	0.300	-	-
8	0.850	0.400	0.200	-
9	0.750	0.300	-	0.200
10	0.650	-	0.200	0.400
11	0.750	-	0.100	-
12	0.850	-	-	0.200

p_1	p_2	p_3
0.800	0.600	0.400

5.1 RBM Parameters

We ran Algorithm 1 using dataset **V** in two different ways. First, using a single constraint: non-negativity on the weights **W**; and subsequently, using two constraints: non-negativity on the weights **W**, and the structure in Fig. 3 encoded in the Q-matrix **Q**.

Table 2 presents the RBM model parameters obtained with one and two constraints. Interestingly, the non-negativity constraint alone is sufficient to learn the model parameters. In Table 2(a), we highlight the **W** values that are close to zero, which indicate that those hidden variables have very little or no influence on the activation of the corresponding observed variables. Upon closer inspection, these near-zero values correspond to the 0 entries in Table 2(b). That is, Algorithm 1 is also useful for learning the model's structure—i.e., the relationships between latent and observed variables [10].

Table 2. RBM parameters learned using Algorithm 1: without Q-matrix constraint (left) vs. with Q-matrix constraint (right).

Question	b	$\mathbf{W}_1$	$\mathbf{W}_2$	$\mathbf{W}_3$
Y1	-0.518	2.782	**0.008**	**0.027**
Y2	-4.147	1.975	3.238	**0.003**
Y3	-3.553	2.225	**0.000**	1.806
Y4	-4.470	**0.005**	1.897	2.810
Y5	-1.598	**0.018**	2.523	**0.016**
Y6	-1.083	**0.033**	**0.028**	2.532
Y7	-0.763	3.777	**0.006**	**0.014**
Y8	-2.743	1.306	2.847	**0.000**
Y9	-3.068	1.350	**0.021**	2.592
Y10	-2.950	**0.013**	2.017	1.278
Y11	-2.163	**0.003**	3.232	**0.018**
Y12	-1.522	**0.018**	**0.016**	3.230

Question	b	$\mathbf{W}_1$	$\mathbf{W}_2$	$\mathbf{W}_3$
Y1	-0.602	2.825	0	0
Y2	-4.163	1.984	3.257	0
Y3	-3.783	2.494	0	1.806
Y4	-4.532	0	1.903	2.812
Y5	-1.617	0	2.470	0
Y6	-1.113	0	0	2.516
Y7	-0.867	3.882	0	0
Y8	-2.785	1.330	2.921	0
Y9	-3.138	1.394	0	2.553
Y10	-2.945	0	2.037	1.322
Y11	-2.112	0	3.191	0
Y12	-1.447	0	0	3.272

Section 4 established that RBMs capture the relationship between hidden and observed variables through learned weight matrices **W** and bias terms, but without explicit cognitive interpretation. Below, we show how to transform the parameters to make them interpretable.

5.2 BN2A Parameters

As mentioned in Sect. 3, BN2A models can be applied in educational assessment. A particular advantage of these models is their interpretability.

Using the Eqs. 12 and 13, we can transform the parameters of the RBM model (Table 2) into the parameters of a corresponding BN2A model. The results are shown in Table 3.

Table 3. BN2A parameters derived from RBM model: without Q-matrix constraint (left) vs. with Q-matrix constraint (right).

ℓ	$q_{\ell,0}$	$q_{\ell,1}$	$q_{\ell,2}$	$q_{\ell,3}$
1	0.909	0.420	**0.999**	**0.997**
2	0.745	0.387	0.138	**0.999**
3	0.617	0.240	**1.000**	0.339
4	0.560	**0.998**	0.286	0.127
5	0.723	**0.995**	0.239	**0.996**
6	0.819	**0.994**	**0.995**	0.323
7	0.954	0.338	**1.000**	**0.999**
8	0.804	0.654	0.239	**1.000**
9	0.710	0.547	**0.994**	0.218
10	0.589	**0.995**	0.272	0.484
11	0.748	**0.999**	0.140	**0.995**
12	0.851	**0.997**	**0.998**	0.216

ℓ	$q_{\ell,0}$	$q_{\ell,1}$	$q_{\ell,2}$	$q_{\ell,3}$
1	0.902	0.392	-	-
2	0.746	0.386	0.136	-
3	0.626	0.194	-	0.345
4	0.546	-	0.278	0.123
5	0.701	-	0.236	-
6	0.803	-	-	0.308
7	0.953	0.310	-	-
8	0.812	0.657	0.232	-
9	0.692	0.516	-	0.215
10	0.602	-	0.273	0.477
11	0.746	-	0.144	-
12	0.861	-	-	0.221

Table 3(a) shows the transformation of the RBM model parameters using only the non-negativity constraint. In this table, we can observe that the weights $W_{\ell,k}$ close to zero correspond to very low penalties (i.e., $q_{\ell,k}$ values very close to one). This makes sense, as all of these highlighted relationships are absent in the original model used to generate dataset $\mathbf{V}$, which can be confirmed by comparing these values with their corresponding entries in Table 3(b). In particular, we observe that the parameters of the model with and without the Q-matrix constraint are very similar, and are also comparable to the original parameters used to generate dataset $\mathbf{V}$ (Table 1).

Finally, with the RBM model parameters and Eq. 3, it is possible to compute the prior probabilities of mastering the three skills in the model (Table 4). These probabilities are quite similar whether the model is learned with or without the Q-matrix constraint, and they are also comparable to the true prior probabilities (Table 1).

Table 4. Comparison of prior probabilities for skill mastery: without Q-matrix constraint (left) vs. with Q-matrix constraint (right).

p_1	p_2	p_3
0.780	0.577	0.394

p_1	p_2	p_3
0.778	0.576	0.392

6 Discussion

This work addresses a fundamental limitation of RBMs in educational contexts: the lack of direct interpretability of their parameters. Our main contribution is the development of a mathematical framework that transforms RBMs into BN2A models, preserving representational capacity while obtaining direct interpretability of leak and penalty parameters essential for educational diagnosis. The key modification of the CD algorithm incorporates non-negativity constraints on weights, ensuring the monotonicity property expected in educational contexts. This constraint does not compromise the model's learning capacity and facilitates subsequent transformation to interpretable BN2A parameters. The validation experiment demonstrates that transformed parameters maintain high fidelity with the originals and that model structure can be effectively recovered. The transformation is mathematically sound and practically viable under realistic conditions. Our approach enables combining RBMs' representational power with BN2A interpretability, enabling granular diagnosis, computational scalability, and dynamic adaptability. This combination is especially valuable in intelligent tutoring systems where both predictive accuracy and explainability are crucial. Main limitations include restriction to controlled synthetic data and potential expressivity limitations due to non-negativity. Future research should validate the approach with real educational data, explore extensions to more complex architectures (Deep Boltzmann Machines), and evaluate applications in specific domains such as mathematics and language.

Acknowledgments. The authors acknowledge financial support for this work: the first author was supported by the Czech Science Foundation through the project No. 25-18070S, and the third author was supported by the project "Research of Excellence on Digital Technologies and Wellbeing CZ.02.01.01/00/22_008/0004583" which is co-financed by the European Union.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Choi, A., Wang, R., Darwiche, A.: On the relative expressiveness of Bayesian and neural networks. Int. J. Approximate Reasoning **113**, 303–323 (2019). https://doi.org/10.1016/j.ijar.2019.07.008
2. Díez, F.J., Druzdzel, M.J.: Canonical probabilistic models for knowledge engineering. Technical report CISIAD-06-01, UNED, Madrid, Spain (2006)
3. Druzdzel, J., Henrion, M.: Efficient reasoning in qualitative probabilistic networks. In: Proceedings of the Eleventh National Conference on Artificial Intelligence, pp. 548–553 (1993)
4. Fischer, A., Igel, C.: Training restricted Boltzmann machines: an introduction. Pattern Recogn. **47**(1), 25–39 (2014). https://doi.org/10.1016/j.patcog.2013.05.025
5. Hartz, S., Roussos, L.: The fusion model for skills diagnosis: blending theory with practicality. ETS Res. Rep. Ser. **2008**(2), i–57 (2008). https://doi.org/10.1002/j.2333-8504.2008.tb02157.x

6. Henrion, M.: Some practical issues in constructing belief networks. In: Proceedings of the Third Conference on Uncertainty in Artificial Intelligence (UAI-87), pp. 161–173. Elsevier Science Publishers B.V. (North Holland) (1987)
7. Hinton, G.E.: A practical guide to training restricted Boltzmann machines. In: Montavon, G., Orr, G.B., Müller, K.-R. (eds.) Neural Networks: Tricks of the Trade. LNCS, vol. 7700, pp. 599–619. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35289-8_32
8. Jensen, F.V., Nielsen, T.D.: Bayesian Networks and Decision Graphs. Information Science and Statistics, 2nd edn. Springer New York (2007). https://doi.org/10.1007/978-0-387-68282-2
9. Koller, D., Friedman, N.: Probabilistic Graphical Models: Principles and Techniques. The MIT Press (2009)
10. Li, C., Ma, C., Xu, G.: Learning large Q-matrix by restricted Boltzmann machines. Psychometrika **87**(3), 1010–1041 (2022). https://doi.org/10.1007/s11336-021-09828-4
11. Lucas, P.J.: Bayesian network modelling through qualitative patterns. Artif. Intell. **163**(2), 233–263 (2005). https://doi.org/10.1016/j.artint.2004.10.011
12. Masegosa, A.R., Feelders, A.J., Gaag, L.C.: Learning from incomplete data in Bayesian networks with qualitative influences. Int. J. Approximate Reasoning **69**, 18–34 (2016)
13. Pearl, J.: Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann Publishers Inc., San Francisco (1988)
14. Pérez, I., Vomlel, J.: Enhancing Bayesian networks with psychometric models. In: Kwisthout, J., Renooij, S. (eds.) Proceedings of The 12th International Conference on Probabilistic Graphical Models. Proceedings of Machine Learning Research, vol. 246, pp. 401–414. PMLR (2024). https://proceedings.mlr.press/v246/perez24a.html
15. Pérez, I., Vomlel, J.: On identifiability of BN2A networks. In: Bouraoui, Z., Vesic, S. (eds.) ECSQARU 2023. LNCS, vol. 14294, pp. 136–148. Springer, Heidelberg (2023). https://doi.org/10.1007/978-3-031-45608-4_11
16. Plajner, M., Vomlel, J.: Learning bipartite Bayesian networks under monotonicity restrictions. Int. J. Gen Syst **49**(1), 88–111 (2020). https://doi.org/10.1080/03081079.2019.1692004
17. Restificar, A.C., Dietterich, T.G.: Exploiting monotonicity via logistic regression in Bayesian network learning. Technical report, Corvallis (2013)
18. Smolensky, P.: Information processing in dynamical systems: foundations of harmony theory. In: Parallel Distributed Processing, Volume 1: Explorations in the Microstructure of Cognition: Foundations. The MIT Press (1986). https://doi.org/10.7551/mitpress/5236.003.0009
19. de la Torre, J., Minchen, N.: Cognitively diagnostic assessments and the cognitive diagnosis model framework. Psicología Educativa. Rev. los Psicól. Educ. (2014). https://www.redalyc.org/articulo.oa?id=613765435004
20. Wellman, M.P.: Fundamental concepts of qualitative probabilistic networks. Artif. Intell. **44**(3), 257–303 (1990)