A Survey on Bayesian Deep Learning

HAO WANG, Massachusetts Institute of Technology, USA

DIT-YAN YEUNG, Hong Kong University of Science and Technology, Hong Kong

108

A comprehensive artificial intelligence system needs to not only perceive the environment with different “senses” (e.g., seeing and hearing) but also infer the world’s conditional (or even causal) relations and corresponding uncertainty. The past decade has seen major advances in many perception tasks, such as visual object recognition and speech recognition, using deep learning models. For higher-level inference, however, probabilistic graphical models with their Bayesian nature are still more powerful and flexible. In recent years, *Bayesian deep learning* has emerged as a unified probabilistic framework to tightly integrate deep learning and Bayesian models.¹ In this general framework, the perception of text or images using deep learning can boost the performance of higher-level inference and, in turn, the feedback from the inference process is able to enhance the perception of text or images. This survey provides a comprehensive introduction to *Bayesian deep learning* and reviews its recent applications on recommender systems, topic models, control, and so on. We also discuss the relationship and differences between Bayesian deep learning and other related topics, such as Bayesian treatment of neural networks.

CCS Concepts: • **Mathematics of computing** → **Probabilistic representations**; • **Information systems** → **Data mining**; • **Computing methodologies** → **Neural networks**;

Additional Key Words and Phrases: Deep learning, Bayesian networks, probabilistic graphical models, generative models

ACM Reference format:

Hao Wang and Dit-Yan Yeung. 2020. A Survey on Bayesian Deep Learning. *ACM Comput. Surv.* 53, 5, Article 108 (September 2020), 37 pages.

<https://doi.org/10.1145/3409383>

1 INTRODUCTION

Over the past decade, deep learning has achieved significant success in many popular perception tasks, including visual object recognition, text understanding, and speech recognition. These tasks correspond to artificial intelligence (AI) systems’ ability to *see*, *read*, and *hear*, respectively, and they are undoubtedly indispensable for AI to effectively perceive the environment. However, to build a practical and comprehensive AI system, simply being able to perceive is far from sufficient. It should, above all, possess the ability of *thinking*.

¹See a curated and updated list of papers related to Bayesian deep learning at <https://github.com/js05212/BayesianDeepLearning-Survey>.

Authors’ addresses: H. Wang, Massachusetts Institute of Technology, G928, 32 Vassar Street, Cambridge, MA 02139; email: hoguwang@gmail.com; D.-Y. Yeung, Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Room 3531, Clear Water Bay, Kowloon, Hong Kong; email: dyyeung@cse.ust.hk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0360-0300/2020/09-ART108 \$15.00

<https://doi.org/10.1145/3409383>

A typical example is medical diagnosis, which goes far beyond simple perception: besides *seeing* visible symptoms (or medical images from CT) and *hearing* descriptions from patients, a doctor also has to look for relations among all the symptoms and preferably infer their corresponding etiology. Only after that can the doctor provide medical advice for the patients. In this example, although the abilities of *seeing* and *hearing* allow the doctor to acquire information from the patients, it is the *thinking* part that defines a doctor. Specifically, the ability of *thinking* here could involve identifying conditional dependencies, causal inference, logic deduction, and dealing with uncertainty, which are apparently beyond the capability of conventional deep learning methods. Fortunately, another machine learning paradigm, probabilistic graphical models (PGM), excels at probabilistic or causal inference and at dealing with uncertainty. The problem is that PGM is not as good as deep learning models at perception tasks, which usually involve large-scale and high-dimensional signals (e.g., images and videos). To address this problem, it is therefore a natural choice to unify deep learning and PGM within a principled probabilistic framework, which we call *Bayesian deep learning* (BDL) in this article.

In the example above, the *perception task* involves perceiving the patient's symptoms (e.g., by *seeing* medical images), while the *inference task* involves handling conditional dependencies, causal inference, logic deduction, and uncertainty. With the principled integration in Bayesian deep learning, the perception task and inference task are regarded as a whole and can benefit from each other. Concretely, being able to see the medical image could help with the doctor's diagnosis and inference. However, diagnosis and inference can, in turn, help understand the medical image. Suppose the doctor may not be sure about what a dark spot in a medical image is, but if she is able to *infer* the etiology of the symptoms and disease, it can help her better decide whether the dark spot is a tumor or not.

Take recommender systems [1, 70, 71, 92, 121] as another example. A highly accurate recommender system requires (1) thorough understanding of item content (e.g., content in documents and movies) [85], (2) careful analysis of users' profiles/preferences [126, 130, 134], and (3) proper evaluation of similarity among users [3, 12, 46, 109]. Deep learning with its ability to efficiently process dense high-dimensional data such as movie content is good at the first subtask, while PGM specializing in modeling conditional dependencies among users, items, and ratings (see Figure 7 as an example, where $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{R}$ are user latent vectors, item latent vectors, and ratings, respectively) excels at the other two. Hence, unifying them two in a single principled probabilistic framework gets us the best of both worlds. Such integration also comes with additional benefit that uncertainty in the recommendation process is handled elegantly. What is more, one can also derive Bayesian treatments for concrete models, leading to more robust predictions [68, 121].

As a third example, consider controlling a complex dynamical system according to the live video stream received from a camera. This problem can be transformed into iteratively performing two tasks, perception from raw images and control based on dynamic models. The perception task of processing raw images can be handled by deep learning while the control task usually needs more sophisticated models such as hidden Markov models and Kalman filters [35, 74]. The feedback loop is then completed by the fact that actions chosen by the control model can affect the received video stream in turn. To enable an effective iterative process between the perception task and the control task, we need information to flow back and forth between them. The perception component would be the basis on which the control component estimates its states and the control component with a dynamic model built in would be able to predict the future trajectory (images). Therefore Bayesian deep learning is a suitable choice [125] for this problem. Note that similar to the recommender system example, both noise from raw images and uncertainty in the control process can be naturally dealt with under such a probabilistic framework.

The above examples demonstrate BDL's major advantages as a principled way of unifying deep learning and PGM: information exchange between the *perception task* and the *inference task*, conditional dependencies on high-dimensional data, and effective modeling of uncertainty. In terms of uncertainty, it is worth noting that when BDL is applied to complex tasks, there are *three kinds of parameter uncertainty* that need to be taken into account:

- (1) Uncertainty on the neural network parameters.
- (2) Uncertainty on the task-specific parameters.
- (3) Uncertainty of exchanging information between the perception component and the task-specific component.

By representing the unknown parameters using distributions instead of point estimates, BDL offers a promising framework to handle these three kinds of uncertainty in a unified way. It is worth noting that the third uncertainty could only be handled under a unified framework like BDL; training the perception component and the task-specific component separately is equivalent to assuming no uncertainty when *exchanging information* between them two. Note that neural networks are usually over-parameterized and therefore pose additional challenges in efficiently handling the uncertainty in such a large parameter space. However, graphical models are often more concise and have smaller parameter space, providing better interpretability.

Besides the advantages above, another benefit comes from the implicit regularization built in BDL. By imposing a prior on hidden units, parameters defining a neural network, or the model parameters specifying the conditional dependencies, BDL can to some degree avoid overfitting, especially when we have insufficient data. Usually, a BDL model consists of two components, a *perception component* that is a Bayesian formulation of a certain type of neural networks and a *task-specific component* that describes the relationship among different hidden or observed variables using PGM. Regularization is crucial for them both. Neural networks are usually heavily over-parameterized and therefore needs to be regularized properly. Regularization techniques such as weight decay and dropout [103] are shown to be effective in improving performance of neural networks and they both have Bayesian interpretations [22]. In terms of the task-specific component, expert knowledge or prior information, as a kind of regularization, can be incorporated into the model through the prior we imposed to guide the model when data are scarce.

There are also challenges when applying BDL to real-world tasks. (1) First, it is nontrivial to design an efficient Bayesian formulation of neural networks with reasonable time complexity. This line of work is pioneered by References [42, 72, 80], but it has not been widely adopted due to its lack of scalability. Fortunately, some recent advances in this direction [2, 9, 31, 39, 58, 119, 121] seem to shed light² on the practical adoption of Bayesian neural network.³ (2) The second challenge is to ensure efficient and effective information exchange between the perception component and the task-specific component. Ideally, both the first-order and second-order information (e.g., the mean and the variance) should be able to flow back and forth between the two components. A natural way is to represent the perception component as a PGM and seamlessly connect it to the task-specific PGM, as done in References [24, 118, 121].

²In summary, reduction in time complexity can be achieved via expectation propagation [39], the reparameterization trick [9, 58], probabilistic formulation of neural networks with maximum *a posteriori* estimates [121], approximate variational inference with natural-parameter networks [119], knowledge distillation [2], and so on. We refer readers to Reference [119] for a detailed overview.

³Here, we refer to the Bayesian treatment of neural networks as Bayesian neural networks. The other term, Bayesian deep learning, is retained to refer to complex Bayesian models with both a perception component and a task-specific component. See Section 4.1 for a detailed discussion.

This survey provides a comprehensive overview of BDL with concrete models for various applications. The rest of the survey is organized as follows: In Section 2, we provide a review of some basic deep learning models. Section 3 covers the main concepts and techniques for PGM. These two sections serve as the preliminaries for BDL, and the next section, Section 4, demonstrates the rationale for the unified BDL framework and details various choices for implementing its *perception component* and *task-specific component*. Section 5 reviews the BDL models applied to various areas such as recommender systems, topic models, and control, showcasing how BDL works in supervised learning, unsupervised learning, and general representation learning, respectively. Section 6 discusses some future research issues and concludes the article.

2 DEEP LEARNING

Deep learning normally refers to neural networks with more than two layers. To better understand deep learning, here we start with the simplest type of neural networks, multilayer perceptrons (MLP), as an example to show how conventional deep learning works. After that, we will review several other types of deep learning models based on MLP.

2.1 Multilayer Perceptrons

Essentially, a multilayer perceptron is a sequence of parametric nonlinear transformations. Suppose we want to train a multilayer perceptron to perform a regression task that maps a vector of M dimensions to a vector of D dimensions. We denote the input as a matrix $\mathbf{X}_0$ (0 means it is the 0th layer of the perceptron). The j th row of $\mathbf{X}_0$, denoted as $\mathbf{X}_{0,j*}$, is an M -dimensional vector representing one data point. The target (the output we want to fit) is denoted as $\mathbf{Y}$. Similarly, $\mathbf{Y}_{j*}$ denotes a D -dimensional row vector. The problem of learning an L -layer multilayer perceptron can be formulated as the following optimization problem:

$$\begin{aligned} \min_{\{\mathbf{W}_l\}, \{\mathbf{b}_l\}} \quad & \|\mathbf{X}_L - \mathbf{Y}\|_F + \lambda \sum_l \|\mathbf{W}_l\|_F^2, \\ \text{subject to} \quad & \mathbf{X}_l = \sigma(\mathbf{X}_{l-1}\mathbf{W}_l + \mathbf{b}_l), l = 1, \dots, L-1, \\ & \mathbf{X}_L = \mathbf{X}_{L-1}\mathbf{W}_L + \mathbf{b}_L, \end{aligned}$$

where $\sigma(\cdot)$ is an element-wise sigmoid function for a matrix and $\sigma(x) = \frac{1}{1+\exp(-x)}$. $\|\cdot\|_F$ denotes the Frobenius norm. The purpose of imposing $\sigma(\cdot)$ is to allow nonlinear transformation. Normally other transformations like $\tanh(x)$ and $\max(0, x)$ can be used as alternatives of the sigmoid function.

Here, $\mathbf{X}_l$ ($l = 1, 2, \dots, L-1$) is the hidden units. As we can see, $\mathbf{X}_L$ can be easily computed once $\mathbf{X}_0$, $\mathbf{W}_l$, and $\mathbf{b}_l$ are given. Since $\mathbf{X}_0$ is given as input, one only needs to learn $\mathbf{W}_l$ and $\mathbf{b}_l$ here. Usually this is done using backpropagation and stochastic gradient descent (SGD). The key is to compute the gradients of the objective function with respect to $\mathbf{W}_l$ and $\mathbf{b}_l$. Denoting the value of the objective function as E , one can compute the gradients using the chain rule as

$$\begin{aligned} \frac{\partial E}{\partial \mathbf{X}_L} &= 2(\mathbf{X}_L - \mathbf{Y}), \quad \frac{\partial E}{\partial \mathbf{X}_l} = \left(\frac{\partial E}{\partial \mathbf{X}_{l+1}} \circ \mathbf{X}_{l+1} \circ (1 - \mathbf{X}_{l+1}) \right) \mathbf{W}_{l+1}, \\ \frac{\partial E}{\partial \mathbf{W}_l} &= \mathbf{X}_{l-1}^T \left(\frac{\partial E}{\partial \mathbf{X}_l} \circ \mathbf{X}_l \circ (1 - \mathbf{X}_l) \right), \quad \frac{\partial E}{\partial \mathbf{b}_l} = \text{mean} \left(\frac{\partial E}{\partial \mathbf{X}_l} \circ \mathbf{X}_l \circ (1 - \mathbf{X}_l), 1 \right), \end{aligned}$$

where $l = 1, \dots, L$ and the regularization terms are omitted. $\circ$ denotes the element-wise product and $\text{mean}(\cdot, 1)$ is the matlab operation on matrices. In practice, we only use a small part of the data (e.g., 128 data points) to compute the gradients for each update. This is called stochastic gradient descent.

Fig. 1. Left: A two-layer SDAE with $L = 4$. Right: A convolutional layer with four input feature maps and two output feature maps.

As we can see, in conventional deep learning models, only $\mathbf{W}_l$ and $\mathbf{b}_l$ are free parameters, which we will update in each iteration of the optimization. $\mathbf{X}_l$ is not a free parameter, since it can be computed exactly if $\mathbf{W}_l$ and $\mathbf{b}_l$ are given.

2.2 Autoencoders

An autoencoder (AE) is a feedforward neural network to encode the input into a more compact representation and reconstruct the input with the learned representation. In its simplest form, an autoencoder is no more than a multilayer perceptron with a bottleneck layer (a layer with a small number of hidden units) in the middle. The idea of autoencoders has been around for decades [10, 29, 43, 63] and abundant variants of autoencoders have been proposed to enhance representation learning, including sparse AE [88], contrastive AE [93], and denoising AE [111]. For more details, please refer to a recent book on deep learning [29]. Here, we introduce a kind of multilayer denoising AE, known as stacked denoising autoencoders (SDAE), both as an example of AE variants and as background for its applications on BDL-based recommender systems in Section 4.

SDAE [111] is a feedforward neural network for learning representations (encoding) of the input data by learning to predict the clean input itself in the output, as shown in Figure 1 (left). The hidden layer in the middle, i.e., $\mathbf{X}_2$ in the figure, can be constrained to be a bottleneck to learn compact representations. The difference between traditional AE and SDAE is that the input layer $\mathbf{X}_0$ is a *corrupted* version of the *clean* input data $\mathbf{X}_c$. Essentially, an SDAE solves the following optimization problem:

$$\begin{aligned} \min_{\{\mathbf{W}_l\}, \{\mathbf{b}_l\}} \|\mathbf{X}_c - \mathbf{X}_L\|_F^2 + \lambda \sum_l \|\mathbf{W}_l\|_F^2, \\ \text{subject to } \mathbf{X}_l = \sigma(\mathbf{X}_{l-1}\mathbf{W}_l + \mathbf{b}_l), l = 1, \dots, L-1, \\ \mathbf{X}_L = \mathbf{X}_{L-1}\mathbf{W}_L + \mathbf{b}_L, \end{aligned}$$

where λ is a regularization parameter. Here, SDAE can be regarded as a multilayer perceptron for regression tasks described in the previous section. The input $\mathbf{X}_0$ of the MLP is the corrupted version of the data and the target $\mathbf{Y}$ is the clean version of the data $\mathbf{X}_c$. For example, $\mathbf{X}_c$ can be the raw data matrix, and we can randomly set 30% of the entries in $\mathbf{X}_c$ to 0 and get $\mathbf{X}_0$. In a nutshell, SDAE learns a neural network that takes the noisy data as input and recovers the clean data in the last layer. This is what “denoising” in the name means. Normally, the output of the middle layer, i.e., $\mathbf{X}_2$ in Figure 1(left), would be used to compactly represent the data.

2.3 Convolutional Neural Networks

Convolutional neural networks (CNN) can be viewed as another variant of MLP. Different from AE, which is initially designed to perform dimensionality reduction, CNN is biologically inspired.

Fig. 2. Left: A conventional feedforward neural network with one hidden layer, where x is the input, z is the hidden layer, and o is the output, W and V are the corresponding weights (biases are omitted here). Middle: A recurrent neural network with input $\{x_t\}_{t=1}^T$, hidden states $\{h_t\}_{t=1}^T$, and output $\{o_t\}_{t=1}^T$. Right: An unrolled RNN that is equivalent to the one in Figure 2 (middle). Here, each node (e.g., x_1 , h_1 , or o_1) is associated with one particular time step.

According to Reference [53], two types of cells have been identified in the cat's visual cortex. One is simple cells that respond maximally to specific patterns within their receptive field, and the other is complex cells with larger receptive field that are considered locally invariant to positions of patterns. Inspired by these findings, the two key concepts in CNN are then developed: convolution and max-pooling.

Convolution: In CNN, a feature map is the result of the convolution of the input and a linear filter, followed by some element-wise nonlinear transformation. The *input* here can be the raw image or the feature map from the previous layer. Specifically, with input X , weights W^k , bias b^k , the k th feature map H^k can be obtained as follows:

$$H_{ij}^k = \tanh((W^k * X)_{ij} + b^k).$$

Note that in the equation above, we assume one single input feature map and multiple output feature maps. In practice, CNN often has multiple input feature maps as well due to its deep structure. A convolutional layer with four input feature maps and two output feature maps is shown in Figure 1 (right).

Max-Pooling: Traditionally, a convolutional layer in CNN is followed by a max-pooling layer, which can be seen as a type of nonlinear downsampling. The operation of max-pooling is simple. For example, if we have a feature map of size 6×9 , then the result of max-pooling with a 3×3 region would be a downsampled feature map of size 2×3 . Each entry of the downsampled feature map is the maximum value of the corresponding 3×3 region in the 6×9 feature map. Max-pooling layers can not only reduce computational cost by ignoring the non-maximal entries but also provide local translation invariance.

Putting it all together: Usually to form a complete and working CNN, the input would alternate between convolutional layers and max-pooling layers before going into an MLP for tasks such as classification or regression. One classic example is the LeNet-5 [64], which alternates between two convolutional layers and two max-pooling layers before going into a fully connected MLP for target tasks.

2.4 Recurrent Neural Network

When reading an article, one normally takes in one word at a time and try to understand the current word based on previous words. This is a recurrent process that needs short-term memory. Unfortunately conventional feedforward neural networks like the one shown in Figure 2 (left) fail to do so. For example, imagine we want to constantly predict the next word as we read an article. Since the feedforward network only computes the output o as $Vq(Wx)$, where the function $q(\cdot)$

denotes element-wise nonlinear transformation, it is unclear how the network could naturally model the sequence of words to predict the next word.

2.4.1 Vanilla Recurrent Neural Network. To solve the problem, we need a recurrent neural network [29] instead of a feedforward one. As shown in Figure 2 (middle), the computation of the current hidden states $\mathbf{h}_t$ depends on the current input $\mathbf{x}_t$ (e.g., the t th word) and the previous hidden states $\mathbf{h}_{t-1}$. This is why there is a loop in the RNN. It is this loop that enables short-term memory in RNNs. The $\mathbf{h}_t$ in the RNN represents what the network knows so far at the t th time step. To see the computation more clearly, we can unroll the loop and represent the RNN as in Figure 2 (right). If we use hyperbolic tangent nonlinearity ($\tanh$), then the computation of output $\mathbf{o}_t$ will be as follows:

$$\mathbf{a}_t = \mathbf{W}\mathbf{h}_{t-1} + \mathbf{Y}\mathbf{x}_t + \mathbf{b}, \quad \mathbf{h}_t = \tanh(\mathbf{a}_t), \quad \mathbf{o}_t = \mathbf{V}\mathbf{h}_t + \mathbf{c},$$

where $\mathbf{Y}$, $\mathbf{W}$, and $\mathbf{V}$ denote the weight matrices for input-to-hidden, hidden-to-hidden, and hidden-to-output connections, respectively, and $\mathbf{b}$ and $\mathbf{c}$ are the corresponding biases. If the task is to classify the input data at each time step, then we can compute the classification probability as $\mathbf{p}_t = \text{softmax}(\mathbf{o}_t)$, where

$$\text{softmax}(\mathbf{q}) = \frac{\exp(\mathbf{q})}{\sum_i \exp(\mathbf{q}_i)}.$$

Similar to feedforward networks, an RNN is trained with a generalized back-propagation algorithm called *back-propagation through time* (BPTT) [29]. Essentially, the gradients are computed through the unrolled network as shown in Figure 2 (right) with shared weights and biases for all time steps.

2.4.2 Gated Recurrent Neural Network. The problem with the vanilla RNN above is that the gradients propagated over many time steps are prone to vanish or explode, making the optimization notoriously difficult. In addition, the signal passing through the RNN decays exponentially, making it impossible to model long-term dependencies in long sequences. Imagine we want to predict the last word in the paragraph “I have many books... I like *reading*.” To get the answer, we need “long-term memory” to retrieve information (the word “books”) at the start of the text. To address this problem, the long short-term memory model (LSTM) is designed as a type of gated RNN to model and accumulate information over a relatively long duration. The intuition behind LSTM is that when processing a sequence consisting of several subsequences, it is sometimes useful for the neural network to summarize or forget the old states before moving on to process the next subsequence [29]. Using $t = 1 \dots T_j$ to index the words in the sequence, the formulation of LSTM is as follows (we drop the item index j for notational simplicity):

$$\mathbf{x}_t = \mathbf{W}_w \mathbf{e}_t, \quad \mathbf{s}_t = \mathbf{h}_{t-1}^f \odot \mathbf{s}_{t-1} + \mathbf{h}_{t-1}^i \odot \sigma(\mathbf{Y}\mathbf{x}_{t-1} + \mathbf{W}\mathbf{h}_{t-1} + \mathbf{b}), \quad (1)$$

where $\mathbf{x}_t$ is the word embedding of the t th word, $\mathbf{W}_w$ is a K_w -by- S word embedding matrix, and $\mathbf{e}_t$ is the 1-of- S representation, $\odot$ stands for the element-wise product operation between two vectors, $\sigma(\cdot)$ denotes the sigmoid function, $\mathbf{s}_t$ is the cell state of the t th word, and $\mathbf{b}$, $\mathbf{Y}$, and $\mathbf{W}$ denote the biases, input weights, and recurrent weights, respectively. The forget gate units $\mathbf{h}_t^f$ and the input gate units $\mathbf{h}_t^i$ in Equation (1) can be computed using their corresponding weights and biases $\mathbf{Y}^f$, $\mathbf{W}^f$, $\mathbf{Y}^i$, $\mathbf{W}^i$, $\mathbf{b}^f$, and $\mathbf{b}^i$:

$$\mathbf{h}_t^f = \sigma(\mathbf{Y}^f \mathbf{x}_t + \mathbf{W}^f \mathbf{h}_t + \mathbf{b}^f), \quad \mathbf{h}_t^i = \sigma(\mathbf{Y}^i \mathbf{x}_t + \mathbf{W}^i \mathbf{h}_t + \mathbf{b}^i).$$

The output depends on the output gate $\mathbf{h}_t^o$, which has its own weights and biases $\mathbf{Y}^o$, $\mathbf{W}^o$, and $\mathbf{b}^o$:

$$\mathbf{h}_t = \tanh(\mathbf{s}_t) \odot \mathbf{h}_{t-1}^o, \quad \mathbf{h}_t^o = \sigma(\mathbf{Y}^o \mathbf{x}_t + \mathbf{W}^o \mathbf{h}_t + \mathbf{b}^o).$$

Fig. 3. The encoder-decoder architecture involving two LSTMs. The encoder LSTM (in the left rectangle) encodes the sequence “ABC” into a representation and the decoder LSTM (in the right rectangle) recovers the sequence from the representation. “\$” marks the end of a sentence.

Fig. 4. The probabilistic graphical model for LDA, J is the number of documents, D is the number of words in a document, and K is the number of topics.

Note that in the LSTM, information of the processed sequence is contained in the cell states $\mathbf{s}_t$ and the output states $\mathbf{h}_t$, both of which are column vectors of length K_W .

Similar to References [16, 108], we can use the output state and cell state at the last time step ($\mathbf{h}_{T_j}$ and $\mathbf{s}_{T_j}$) of the first LSTM as the initial output state and cell state of the second LSTM. This way the two LSTMs can be concatenated to form an encoder-decoder architecture, as shown in Figure 3.

Note that there is a vast literature on deep learning and neural networks. The introduction in this section intends to serve only as the background of Bayesian deep learning. Readers are referred to Reference [29] for a comprehensive survey and more details.

3 PROBABILISTIC GRAPHICAL MODELS

Probabilistic Graphical Models (PGM) use diagrammatic representations to describe random variables and relationships among them. Similar to a graph that contains nodes (vertices) and links (edges), PGM has nodes to represent random variables and links to indicate probabilistic relationships among them.

3.1 Models

There are essentially two types of PGM, directed PGM (also known as Bayesian networks) and undirected PGM (also known as Markov random fields) [5]. In this survey, we mainly focus on directed PGM.⁴ For details on undirected PGM, readers are referred to Reference [5].

A classic example of PGM would be latent Dirichlet allocation (LDA), which is used as a topic model to analyze the generation of words and topics in documents [8]. Usually PGM comes with a graphical representation of the model and a generative process to depict the story of how the random variables are generated step by step. Figure 4 shows the graphical model for LDA and the corresponding generative process is as follows:

- For each document j ($j = 1, 2, \dots, J$),
 - (1) Draw topic proportions $\theta_j \sim \text{Dirichlet}(\alpha)$.
 - (2) For each word w_{jn} of item (document) $\mathbf{w}_j$,
 - (a) Draw topic assignment $z_{jn} \sim \text{Mult}(\theta_j)$,
 - (b) Draw word $w_{jn} \sim \text{Mult}(\beta_{z_{jn}})$.

⁴For convenience, PGM stands for directed PGM in this survey unless specified otherwise.

The generative process above provides the story of how the random variables are generated. In the graphical model in Figure 4, the shaded node denotes observed variables while the others are latent variables (θ and z) or parameters (α and β). Once the model is defined, learning algorithms can be applied to automatically learn the latent variables and parameters.

Due to its Bayesian nature, PGM such as LDA is easy to extend to incorporate other information or to perform other tasks. For example, following LDA, different variants of topic models have been proposed. References [7, 113] are proposed to incorporate temporal information, and Reference [6] extends LDA by assuming correlations among topics. Reference [44] extends LDA from the batch mode to the online setting, making it possible to process large datasets. On recommender systems, collaborative topic regression (CTR) [112] extends LDA to incorporate rating information and make recommendations. This model is then further extended to incorporate social information [89, 115, 116].

3.2 Inference and Learning

Strictly speaking, the process of finding the parameters (e.g., α and β in Figure 4) is called learning and the process of finding the latent variables (e.g., θ and z in Figure 4) given the parameters is called inference. However, given only the observed variables (e.g., w in Figure 4), learning and inference are often intertwined. Usually the learning and inference of LDA would alternate between the updates of latent variables (which correspond to inference) and the updates of the parameters (which correspond to learning). Once the learning and inference of LDA is completed, one could obtain the learned parameters α and β . If a new document comes, then one can now fix the learned α and β and then perform inference alone to find the topic proportions θ_j of the new document.⁵

Similar to LDA, various learning and inference algorithms are available for each PGM. Among them, the most cost-effective one is probably maximum *a posteriori* (MAP), which amounts to maximizing the posterior probability of the latent variable. Using MAP, the learning process is equivalent to minimizing (or maximizing) an objective function with regularization. One famous example is the probabilistic matrix factorization (PMF) [96], where the learning of the graphical model is equivalent to factorizing a large matrix into two low-rank matrices with L2 regularization.

MAP, as efficient as it is, gives us only *point estimates* of latent variables (and parameters). To take the uncertainty into account and harness the full power of Bayesian models, one would have to resort to Bayesian treatments such as variational inference and Markov chain Monte Carlo (MCMC). For example, the original LDA uses variational inference to approximate the true posterior with factorized variational distributions [8]. Learning of the latent variables and parameters then boils down to minimizing the KL-divergence between the variational distributions and the true posterior distributions. Besides variational inference, another choice for a Bayesian treatment is MCMC. For example, MCMC algorithms such as Reference [86] have been proposed to learn the posterior distributions of LDA.

4 BAYESIAN DEEP LEARNING

With the preliminaries on deep learning and PGM, we are now ready to introduce the general framework and some concrete examples of BDL. Specifically, in this section, we will list some recent BDL models with applications on recommender systems, topic models, control, and so on. A summary of these models is shown in Table 1.

⁵For convenience, we use “learning” to represent both “learning and inference” in the following text.

Table 1. Summary of BDL Models with Different Learning Algorithms (MAP: Maximum a Posteriori, VI: Variational Inference, Hybrid MC: Hybrid Monte Carlo) and Different Variance Types (ZV: Zero-Variance, HV: Hyper-Variance, LV: Learnable-Variance)

Applications	Models	Variance of $\Omega_{\mathbf{h}}$	MAP	VI	Gibbs Sampling	Hybrid MC
Recommender Systems	Collaborative Deep Learning (CDL) [121]	HV	✓			
	Bayesian CDL [121]	HV			✓	
	Marginalized CDL [66]	LV	✓			
	Symmetric CDL [66]	LV	✓			
	Collaborative Deep Ranking [131]	HV	✓			
	Collaborative Knowledge Base Embedding [132]	HV	✓			
	Collaborative Recurrent AE [122]	HV	✓			
	Collaborative Variational Autoencoders [68]	HV		✓		
Topic Models	Relational SDAE	HV	✓			
	Deep Poisson Factor Analysis with Sigmoid Belief Networks [24]	ZV			✓	✓
	Deep Poisson Factor Analysis with Restricted Boltzmann Machine [24]	ZV			✓	✓
	Deep Latent Dirichlet Allocation [18]	LV				✓
	Dirichlet Belief Networks [133]	LV			✓	
Control	Embed to Control [125]	LV		✓		
	Deep Variational Bayes Filters [57]	LV		✓		
	Probabilistic Recurrent State-Space Models [19]	LV		✓		
	Deep Planning Networks [34]	LV		✓		
Link Prediction	Relational Deep Learning [120]	LV	✓	✓		
	Graphite [32]	LV		✓		
	Deep Generative Latent Feature Relational Model [75]	LV		✓		
NLP	Sequence to Better Sequence [77]	LV		✓		
	Quantifiable Sequence Editing [69]	LV		✓		
Computer Vision	Asynchronous Temporal Fields [102]	LV		✓		
	Attend, Infer, Repeat (AIR) [20]	LV		✓		
	Fast AIR [105]	LV		✓		
	Sequential AIR [60]	LV		✓		
Speech	Factorized Hierarchical VAE [48]	LV		✓		
	Scalable Factorized Hierarchical VAE [47]	LV		✓		
	Gaussian Mixture Variational Autoencoders [49]	LV		✓		
	Recurrent Poisson Process Units [51]	LV	✓	✓		
	Deep Graph Random Process [52]	LV	✓	✓		
Time Series Forecasting	DeepAR [21]	LV		✓		
	DeepState [90]	LV		✓		
	Spline Quantile Function RNN [27]	LV		✓		
	DeepFactor [124]	LV		✓		
Health Care	Deep Poisson Factor Models [38]	LV				✓
	Deep Markov Models [61]	LV		✓		
	Black-Box False Discovery Rate [110]	LV		✓		
	Bidirectional Inference Networks [117]	LV	✓			

4.1 A Brief History of Bayesian Neural Networks and Bayesian Deep Learning

One topic highly related to BDL is Bayesian neural networks (BNN) or Bayesian treatments of neural networks. Similar to any Bayesian treatment, BNN imposes a prior on the neural network's parameters and aims to learn a posterior distribution of these parameters. During the inference phrase, such a distribution is then marginalized out to produce final predictions. In general such a process is called Bayesian model averaging [5] and can be seen as learning an infinite number of (or a distribution over) neural networks and then aggregating the results through ensembling.

The study of BNN dates back to 1990s with notable works from References [42, 72, 80]. Over the years, a large body of works [2, 9, 31, 39, 58, 100] have emerged to enable substantially better

Fig. 5. The PGM for an example BDL. The red rectangle on the left indicates the perception component, and the blue rectangle on the right indicates the task-specific component. The hinge variable $\Omega_h = \{H\}$.

scalability and incorporate recent advancements of deep neural networks. Due to BNN's long history, the term "Bayesian deep learning" sometimes specifically refers to "Bayesian neural networks" [73, 128]. In this survey, we instead use "Bayesian deep learning" in a broader sense to refer to the probabilistic framework subsuming Bayesian neural networks. To see this, note that a BDL model with a *perception component* and an empty *task-specific component* is equivalent to a Bayesian neural network (details on these two components are discussed in Section 4.2).

Interestingly, though BNN started in 1990s, the study of BDL in a broader sense started roughly in 2014 [38, 114, 118, 121], slightly after the deep learning breakthrough in the ImageNet LSVRC contest in 2012 [62]. As we will see in later sections, BNN is usually used as a perception component in BDL models.

Today BDL is gaining more and more popularity, has found successful applications in areas such as recommender systems and computer vision, and appears as the theme of various conference workshops (e.g., the NeurIPS BDL workshop⁶).

4.2 General Framework

As mentioned in Section 1, BDL is a principled probabilistic framework with two seamlessly integrated components: a *perception component* and a *task-specific component*.

Two Components: Figure 5 shows the PGM of a simple BDL model as an example. The part inside the red rectangle on the left represents the perception component and the part inside the blue rectangle on the right is the task-specific component. Typically, the perception component would be a probabilistic formulation of a deep learning model with multiple nonlinear processing layers represented as a chain structure in the PGM. While the nodes and edges in the perception component are relatively simple, those in the task-specific component often describe more complex distributions and relationships among variables. Concretely, a task-specific component can take various forms. For example, it can be a typical Bayesian network (directed PGM) such as LDA, a deep Bayesian network [117], or a stochastic process [51, 94], all of which can be represented in the form of PGM.

Three Variable Sets: There are three sets of variables in a BDL model: perception variables, hinge variables, and task variables. In this article, we use Ω_p to denote the set of perception variables (e.g., X_0, X_1 , and W_1 in Figure 5), which are the variables in the perception component. Usually Ω_p would include the weights and neurons in the probabilistic formulation of a deep learning model. Ω_h is used to denote the set of hinge variables (e.g., H in Figure 5). These variables directly interact with the perception component from the task-specific component. The set of task variables (e.g., A, B , and C in Figure 5), i.e., variables in the task-specific component without direct relation to the perception component, is denoted as Ω_t .

⁶<http://bayesiandeeplearning.org/>.

Generative Processes for Supervised and Unsupervised Learning: If the edges between the two components *point toward* Ω_h , then the joint distribution of all variables can be written as

$$p(\Omega_p, \Omega_h, \Omega_t) = p(\Omega_p)p(\Omega_h|\Omega_p)p(\Omega_t|\Omega_h). \quad (2)$$

If the edges between the two components *originate from* Ω_h , then the joint distribution of all variables can be written as

$$p(\Omega_p, \Omega_h, \Omega_t) = p(\Omega_t)p(\Omega_h|\Omega_t)p(\Omega_p|\Omega_h). \quad (3)$$

Equations (2) and (3) assume different generative processes for the data and correspond to different learning tasks. The former is usually used for supervised learning, where the perception component serves as a probabilistic (or Bayesian) representation learner to facilitate any downstream tasks (see Section 5.1 for some examples). The latter is usually used for unsupervised learning, where the task-specific component provides structured constraints and domain knowledge to help the perception component learn stronger representations (see Section 5.2 for some examples).

Note that besides these two vanilla cases, it is possible for BDL to simultaneously have some edges between the two components pointing towards Ω_h and some originating from Ω_h , in which case the decomposition of the joint distribution would be more complex.

Independence Requirement: The introduction of hinge variables Ω_h and related conditional distributions simplifies the model (especially when Ω_h 's in-degree or out-degree is 1), facilitate learning, and provides inductive bias to concentrate information inside Ω_h . Note that hinge variables are always in the task-specific component; the connections between hinge variables Ω_h and the perception component (e.g., $X_4 \rightarrow H$ in Figure 5) should normally be independent for convenience of parallel computation in the perception component. For example, each row in H is related to only one corresponding row in X_4 . Although it is not mandatory in BDL models, meeting this requirement would significantly increase the efficiency of parallel computation in model training.

Flexibility of Variance for Ω_h : As mentioned in Section 1, one of BDL's motivations is to model the *uncertainty of exchanging information* between the perception component and the task-specific component, which boils down to modeling the uncertainty related to Ω_h . For example, such uncertainty is reflected in the variance of the conditional density $p(\Omega_h|\Omega_p)$ in Equation (2).⁷ According to the degree of flexibility, there are three types of variance for Ω_h (for simplicity, we assume the joint likelihood of BDL is Equation (2), $\Omega_p = \{p\}$, $\Omega_h = \{h\}$, and $p(\Omega_h|\Omega_p) = N(h|\mu_p, \sigma_p^2)$ in our example):

- **Zero-Variance:** Zero-Variance (ZV) assumes no uncertainty during the information exchange between the two components. In the example, zero-variance means directly setting σ_p^2 to 0.
- **Hyper-Variance:** Hyper-Variance (HV) assumes that uncertainty during the information exchange is defined through hyperparameters. In the example, HV means that σ_p^2 is a manually tuned hyperparameter.
- **Learnable Variance:** Learnable Variance (LV) uses learnable parameters to represent uncertainty during the information exchange. In the example, σ_p^2 is the learnable parameter.

As shown above, we can see that in terms of model flexibility, $LV > HV > ZV$. Normally, if properly regularized, an LV model outperforms an HV model, which is superior to a ZV model. In Table 1, we show the types of variance for Ω_h in different BDL models. Note that although each

⁷For models with the joint likelihood decomposed as in Equation (3), the uncertainty is reflected in the variance of $p(\Omega_p|\Omega_h)$.

model in the table has a specific type, one can always adjust the models to devise their counterparts of other types. For example, while CDL in the table is an HV model, we can easily adjust $p(\Omega_h | \Omega_p)$ in CDL to devise its ZV and LV counterparts. In Reference [121], the authors compare the performance of an HV CDL and a ZV CDL and find that the former performs significantly better, meaning that sophisticatedly modeling uncertainty between two components is essential for performance.

Learning Algorithms: Due to the nature of BDL, practical learning algorithms need to meet the following criteria:

- (1) They should be online algorithms to scale well for large datasets.
- (2) They should be efficient enough to scale linearly with the number of free parameters in the perception component.

Criterion (1) implies that conventional variational inference or MCMC methods are not applicable. Usually an online version of them is needed [45]. Most SGD-based methods do not work either unless only MAP inference (as opposed to Bayesian treatments) is performed. Criterion (2) is needed because there are typically a large number of free parameters in the perception component. This means methods based on Laplace approximation [72] are not realistic, since they involve the computation of a Hessian matrix that scales quadratically with the number of free parameters.

4.3 Perception Component

Ideally, the perception component should be a probabilistic or Bayesian neural network, to be compatible with the task-specific component, which is probabilistic in nature. This is to ensure the perception component's built-in capability to handle uncertainty of parameters and its output.

As mentioned in Section 4.1, the study of Bayesian neural networks dates back to 1990s [31, 42, 72, 80]. However, pioneering work at that time was not widely adopted due to its lack of scalability. To address this issue, there has been recent development such as restricted Boltzmann machine (RBM) [40, 41], probabilistic generalized stacked denoising autoencoders (pSDAE) [118, 121], variational autoencoders (VAE) [58], probabilistic back-propagation (PBP) [39], Bayes by Backprop (BBB) [9], Bayesian dark knowledge (BDK) [2], and natural-parameter networks (NPN) [119].

More recently, generative adversarial networks (GAN) [30] prevail as a new training scheme for training neural networks and have shown promise in generating photo-realistic images. Later on, Bayesian formulations (as well as related theoretical results) for GAN have also been proposed [30, 37]. These models are also potential building blocks as the BDL framework's perception component.

In this subsection, we mainly focus on the introduction of recent Bayesian neural networks such as RBM, pSDAE, VAE, and NPN. We refer the readers to Reference [29] for earlier work in this direction.

4.3.1 Restricted Boltzmann Machine. Restricted Boltzmann Machine (RBM) is a special kind of BNN in that (1) it is not trained with back-propagation (BP) and that (2) its hidden neurons are binary. Specifically, RBM defines the following energy:

$$E(\mathbf{v}, \mathbf{h}) = -\mathbf{v}^T \mathbf{W} \mathbf{h} - \mathbf{v}^T \mathbf{b} - \mathbf{h}^T \mathbf{a},$$

where $\mathbf{v}$ denotes visible (observed) neurons, and $\mathbf{h}$ denotes binary hidden neurons. $\mathbf{W}$, $\mathbf{a}$, and $\mathbf{b}$ are learnable weights. The energy function leads to the following conditional distributions:

$$p(\mathbf{v}|\mathbf{h}) = \frac{\exp(-E(\mathbf{v}, \mathbf{h}))}{\sum_{\mathbf{v}} \exp(-E(\mathbf{v}, \mathbf{h}))}, \quad p(\mathbf{h}|\mathbf{v}) = \frac{\exp(-E(\mathbf{v}, \mathbf{h}))}{\sum_{\mathbf{h}} \exp(-E(\mathbf{v}, \mathbf{h}))}. \quad (4)$$

RBM is trained using “Contrastive Divergence” [40] rather than BP. Once trained, RBM can infer $\mathbf{v}$ or $\mathbf{h}$ by marginalizing out other neurons. One can also stack layers of RBM to form a deep belief network (DBN) [76], use multiple branches of deep RBN for multimodal learning [104], or combine DBN with convolutional layers to form a convolutional DBN [65].

4.3.2 Probabilistic Generalized SDAE. Following the introduction of SDAE in Section 2.2, if we assume that both the clean input $\mathbf{X}_c$ and the corrupted input $\mathbf{X}_0$ are observed, similar to References [4, 5, 13, 72], then we can define the following generative process of the probabilistic SDAE:

- (1) For each layer l of the SDAE network,
 - (a) For each column n of the weight matrix $\mathbf{W}_l$, draw $\mathbf{W}_{l,*n} \sim \mathcal{N}(\mathbf{0}, \lambda_w^{-1} \mathbf{I}_{K_l})$,
 - (b) Draw the bias vector $\mathbf{b}_l \sim \mathcal{N}(\mathbf{0}, \lambda_w^{-1} \mathbf{I}_{K_l})$,
 - (c) For each row j of $\mathbf{X}_l$, draw

$$\mathbf{X}_{l,j*} \sim \mathcal{N}(\sigma(\mathbf{X}_{l-1,j*} \mathbf{W}_l + \mathbf{b}_l), \lambda_s^{-1} \mathbf{I}_{K_l}). \quad (5)$$

- (2) For each item j , draw a clean input⁸ $\mathbf{X}_{c,j*} \sim \mathcal{N}(\mathbf{X}_{L,j*}, \lambda_n^{-1} \mathbf{I}_B)$.

Note that if λ_s goes to infinity, the Gaussian distribution in Equation (5) will become a Dirac delta distribution [106] centered at $\sigma(\mathbf{X}_{l-1,j*} \mathbf{W}_l + \mathbf{b}_l)$, where $\sigma(\cdot)$ is the sigmoid function, and the model will degenerate into a Bayesian formulation of vanilla SDAE. This is why we call it “generalized” SDAE.

The first $L/2$ layers of the network act as an encoder and the last $L/2$ layers act as a decoder. Maximization of the posterior probability is equivalent to minimization of the reconstruction error with weight decay taken into consideration.

Following pSDAE, both its convolutional version [132] and its recurrent version [122] have been proposed with applications in knowledge base embedding and recommender systems.

4.3.3 Variational Autoencoders. Variational Autoencoders (VAE) [58] essentially tries to learn parameters ϕ and θ that maximize the evidence lower bound (ELBO):

$$\mathcal{L}_{vae} = E_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - KL(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})), \quad (6)$$

where $q_\phi(\mathbf{z}|\mathbf{x})$ is the encoder parameterized by ϕ and $p_\theta(\mathbf{x}|\mathbf{z})$ is the decoder parameterized by θ . The negation of the first term is similar to the reconstruction error in vanilla AE, while the KL divergence works as a regularization term for the encoder. During training $q_\phi(\mathbf{z}|\mathbf{x})$ will output the mean and variance of a Gaussian distribution, from which $\mathbf{z}$ is sampled via the reparameterization trick. Usually $q_\phi(\mathbf{z}|\mathbf{x})$ is parameterized by an MLP with two branches, one producing the mean and the other producing the variance.

Similar to the case of pSDAE, various VAE variants have been proposed. For example, Importance weighted Autoencoders (IWAE) [11] derived a tighter lower bound via importance weighting, Reference [129] combined LSTM, VAE, and dilated CNN for text modeling, and Reference [17] proposed a recurrent version of VAE dubbed variational RNN (VRNN).

4.3.4 Natural-parameter Networks. Different from vanilla NN, which usually takes deterministic input, NPN [119] is a probabilistic NN taking distributions as input. The input distributions go through layers of linear and nonlinear transformation to produce output distributions. In NPN, all hidden neurons and weights are also distributions expressed in closed form. Note that this is in contrast to VAE where only the middle layer output $\mathbf{z}$ is a distribution.

⁸Note that while generation of the *clean* input $\mathbf{X}_c$ from $\mathbf{X}_L$ is part of the generative process of the Bayesian SDAE, generation of the *noise-corrupted* input $\mathbf{X}_0$ from $\mathbf{X}_c$ is an artificial noise injection process to help the SDAE learn a more robust feature representation.

As a simple example, in a vanilla linear NN $f_w(x) = wx$ takes a scalar x as input and computes the output based on a scalar parameter w ; a corresponding Gaussian NPN would assume w is drawn from a Gaussian distribution $\mathcal{N}(w_m, w_s)$ and that x is drawn from $\mathcal{N}(x_m, x_s)$ (x_s is set to 0 when the input is deterministic). With $\theta = (w_m, w_s)$ as a learnable parameter pair, NPN will then compute the mean and variance of the output Gaussian distribution $\mu_\theta(x_m, x_s)$ and $s_\theta(x_m, x_s)$ in closed form (bias terms are ignored for clarity) as

$$\mu_\theta(x_m, x_s) = E[wx] = x_m w_m, \quad (7)$$

$$s_\theta(x_m, x_s) = D[wx] = x_s w_s + x_s w_m^2 + x_m^2 w_s. \quad (8)$$

Hence, the output of this Gaussian NPN is a tuple $(\mu_\theta(x_m, x_s), s_\theta(x_m, x_s))$ representing a Gaussian distribution instead of a single value. Input variance x_s to NPN can be set to 0 if not available. Note that since $s_\theta(x_m, 0) = x_m^2 w_s$, w_m and w_s can still be learned even if $x_s = 0$ for all data points. The derivation above is generalized to handle vectors and matrices in practice [119]. Besides Gaussian distributions, NPN also support other exponential-family distributions such as Poisson distributions and gamma distributions [119].

Following NPN, a light-weight version [26] was proposed to speed up the training and inference process. Another variant, MaxNPN [100], extended NPN to handle max-pooling and categorical layers. ConvNPN [87] enables convolutional layers in NPN. In terms of model quantization and compression, BinaryNPN [107] was also proposed as NPN's binary version to achieve better efficiency.

4.4 Task-Specific Component

In this subsection, we introduce different forms of task-specific components. The purpose of a task-specific component is to incorporate probabilistic prior knowledge into the BDL model. Such knowledge can be naturally represented using PGM. Concretely, it can be a typical (or shallow) Bayesian network [5, 54], a bidirectional inference network [117], or a stochastic process [94].

4.4.1 Bayesian Networks. Bayesian networks are the most common choice for a task-specific component. As mentioned in Section 3, Bayesian networks can naturally represent conditional dependencies and handle uncertainty. Besides LDA introduced above, a more straightforward example is probabilistic matrix factorization (PMF) [96], where one uses a Bayesian network to describe the conditional dependencies among users, items, and ratings. Specifically, PMF assumes the following generative process:

- (1) For each item j , draw a latent item vector: $\mathbf{v}_j \sim \mathcal{N}(\mathbf{0}, \lambda_v^{-1} \mathbf{I}_K)$.
- (2) For each user i , draw a latent user vector: $\mathbf{u}_i \sim \mathcal{N}(\mathbf{0}, \lambda_u^{-1} \mathbf{I}_K)$.
- (3) For each user-item pair (i, j) , draw a rating: $R_{ij} \sim \mathcal{N}(\mathbf{u}_i^T \mathbf{v}_j, C_{ij}^{-1})$.

In the generative process above, C_{ij}^{-1} is the corresponding variance for the rating R_{ij} . Using MAP estimates, learning PMF amounts to maximize the following log-likelihood of $p(\{\mathbf{u}_i\}, \{\mathbf{v}_j\} | \{R_{ij}\}, \{C_{ij}\}, \lambda_u, \lambda_v)$:

$$\mathcal{L} = -\frac{\lambda_u}{2} \sum_i \|\mathbf{u}_i\|_2^2 - \frac{\lambda_v}{2} \sum_j \|\mathbf{v}_j\|_2^2 - \sum_{i,j} \frac{C_{ij}}{2} (R_{ij} - \mathbf{u}_i^T \mathbf{v}_j)^2,$$

Note that one can also impose another layer of priors on the hyperparameters with a fully Bayesian treatment. For example, Reference [97] imposes priors on the precision matrix of latent factors and learn the Bayesian PMF with Gibbs sampling.

In Section 5.1, we will show how PMF can be used as a task-specific component along with a perception component defined to significantly improve recommender systems' performance.

Fig. 6. Left: A simple example of BIN with each conditional distribution parameterized by a Bayesian neural networks (BNN) or simply a probabilistic neural network. Right: Another example BIN. Shaded and transparent nodes indicate observed and unobserved variables, respectively.

4.4.2 Bidirectional Inference Networks. Typical Bayesian networks assume “shallow” conditional dependencies among random variables. In the generative process, one random variable (which can be either latent or observed) is usually drawn from a conditional distribution parameterized by the linear combination of its parent variables. For example, in PMF the rating R_{ij} is drawn from a Gaussian distribution mainly parameterized by the linear combination of u_i and v_j , i.e., $R_{ij} \sim \mathcal{N}(u_i^T v_j, C_{ij}^{-1})$.

Such “shallow” and linear structures can be replaced with nonlinear or even deep nonlinear structures to form a *deep Bayesian network*. As an example, bidirectional inference network (BIN) [117] is a class of deep Bayesian networks that enable deep nonlinear structures in each conditional distribution, while retaining the ability to incorporate prior knowledge as Bayesian networks.

For example, Figure 6 (left) shows a BIN, where each conditional distribution is parameterized by a Bayesian neural network. Specifically, this example assumes the following factorization:

$$p(v_1, v_2, v_3 | X) = p(v_1 | X) p(v_2 | X, v_1) p(v_3 | X, v_1, v_2).$$

A vanilla Bayesian network parameterizes each distribution with simple linear operations. For example, $p(v_2 | X, v_1) = \mathcal{N}(v_2 | Xw_0 + v_1w_1 + b, \sigma^2)$. In contrast, BIN (as a deep Bayesian network) uses a BNN. For example, BIN has $p(v_2 | X, v_1) = \mathcal{N}(v_2 | \mu_\theta(X, v_1), s_\theta(X, v_1))$, where $\mu_\theta(X, v_1)$ and $s_\theta(X, v_1)$ are the output mean and variance of the BNN. The inference and learning of such a deep Bayesian network is done by performing BP across all BNNs (e.g., BNN 1, 2, and 3 in Figure 6 (left)) [117].

Compared to vanilla (shallow) Bayesian networks, deep Bayesian networks such as BIN make it possible to handle deep and nonlinear conditional dependencies effectively and efficiently. Besides, with BNN as building blocks, task-specific components based on deep Bayesian networks can better work with the perception component, which is usually a BNN as well. Figure 6 (right) shows a more complicated case with both observed (shaded nodes) and unobserved (transparent nodes) variables.

4.4.3 Stochastic Processes. Besides vanilla Bayesian networks and deep Bayesian networks, a task-specific component can also take the form of a stochastic process [94]. For example, a Wiener process can naturally describe a continuous-time Brownian motion model $\mathbf{x}_{t+u} | \mathbf{x}_t \sim \mathcal{N}(\mathbf{x}_t, \lambda u \mathbf{I})$, where $\mathbf{x}_{t+u}$ and $\mathbf{x}_t$ are the states at time t and $t + u$, respectively. In the graphical model literature, such a process has been used to model the continuous-time topic evolution of articles over time [113].

Another example is to model phonemes’ boundary positions using a Poisson process in automatic speech recognition (ASR) [51]. Note that this is a fundamental problem in ASR, since speech is no more than a sequence of phonemes. Specifically, a Poisson process defines the generative process $\Delta t_i = t_i - t_{i-1} \sim g(\lambda(t))$, with $\mathcal{T} = \{t_1, t_2, \dots, t_N\}$ as the set of boundary positions, and $g(\lambda(t))$

is an exponential distribution with the parameter $\lambda(t)$ (also known as the intensity). Such a stochastic process naturally models the occurrence of phoneme boundaries in continuous time. The parameter $\lambda(t)$ can be the output of a neural network taking raw speech signals as input [51, 83, 99].

Interestingly, stochastic processes can be seen as a type of dynamic Bayesian networks. To see this, we can rewrite the Poisson process above in an equivalent form, where given t_{i-1} , the probability that t_i has not occurred at time t , $P(t_i > t) = \exp(-\int_{t_{i-1}}^t \lambda(t) dt)$. Obviously, both the Wiener process and the Poisson process are Markovian and can be represented with a dynamic Bayesian network [78].

For clarity, we focus on using vanilla Bayesian networks as task-specific components in Section 5; they can be naturally replaced with other types of task-specific components to represent different prior knowledge if necessary.

5 CONCRETE BDL MODELS AND APPLICATIONS

In this section, we discuss how the BDL framework can facilitate supervised learning, unsupervised learning, and representation learning in general. Concretely, we use examples in domains such as recommender systems, topic models, control, and so on.

5.1 Supervised Bayesian Deep Learning for Recommender Systems

Despite the successful applications of deep learning on natural language processing and computer vision, very few attempts have been made to develop deep learning models for collaborative filtering (CF) before the emergence of BDL. Reference [98] uses restricted Boltzmann machines instead of the conventional matrix factorization formulation to perform CF and Reference [28] extends this work by incorporating user-user and item-item correlations. Although these methods involve both deep learning and CF, they actually belong to CF-based methods, because they ignore users' or items' content information, which is crucial for accurate recommendation. Reference [95] uses low-rank matrix factorization in the last weight layer of a deep network to significantly reduce the number of model parameters and speed up training, but it is for classification instead of recommendation tasks. On music recommendation, References [84, 123] directly use conventional CNN or deep belief networks (DBN) to assist representation learning for content information, but the deep learning components of their models are deterministic without modeling the noise and hence they are less robust. The models achieve performance boost mainly by loosely coupled methods without exploiting the interaction between content information and ratings. Besides, the CNN is linked directly to the rating matrix, which means the models will perform poorly due to serious overfitting when the ratings are sparse.

5.1.1 Collaborative Deep Learning. To address the challenges above, a hierarchical Bayesian model called collaborative deep learning (CDL) as a novel tightly coupled method for recommender systems is introduced in Reference [121]. Based on a Bayesian formulation of SDAE, CDL tightly couples deep representation learning for the content information and collaborative filtering for the rating (feedback) matrix, allowing two-way interaction between the two. From BDL's perspective, a probabilistic SDAE as the perception component is tightly coupled with a probabilistic graphical model as the task-specific component. Experiments show that CDL significantly improves upon the state of the art.

In the following text, we will start with the introduction of the notation used during our presentation of CDL. After that, we will review the design and learning of CDL.

Notation and Problem Formulation: Similar to the work in Reference [112], the recommendation task considered in CDL takes implicit feedback [50] as the training and test data. The entire collection of J items (articles or movies) is represented by a J -by- B matrix $\mathbf{X}_c$, where row j is the

Fig. 7. On the left is the graphical model of CDL. The part inside the dashed rectangle represents an SDAE. An example SDAE with $L = 2$ is shown. On the right is the graphical model of the degenerated CDL. The part inside the dashed rectangle represents the encoder of an SDAE. An example SDAE with $L = 2$ is shown on its right. Note that although L is still 2, the decoder of the SDAE vanishes. To prevent clutter, we omit all variables x_l , except x_0 and $x_{L/2}$ in the graphical models.

bag-of-words vector $X_{c,j*}$ for item j based on a vocabulary of size B . With I users, we define an I -by- J binary rating matrix $R = [R_{ij}]_{I \times J}$. For example, in the dataset *citeulike-a* [112, 115, 121], $R_{ij} = 1$ if user i has article j in his or her personal library and $R_{ij} = 0$ otherwise. Given part of the ratings in R and the content information X_c , the problem is to predict the other ratings in R . Note that although CDL in its current form focuses on movie recommendation (where plots of movies are considered as content information) and article recommendation like Reference [112] in this section, it is general enough to handle other recommendation tasks (e.g., tag recommendation).

The matrix X_c plays the role of clean input to the SDAE while the noise-corrupted matrix, also a J -by- B matrix, is denoted by X_0 . The output of layer l of the SDAE is denoted by X_l , which is a J -by- K_l matrix. Similar to X_c , row j of X_l is denoted by $X_{l,j*}$. W_l and b_l are the weight matrix and bias vector, respectively, of layer l , $W_{l,*n}$ denotes column n of W_l , and L is the number of layers. For convenience, we use W^+ to denote the collection of all layers of weight matrices and biases. Note that an $L/2$ -layer SDAE corresponds to an L -layer network.

Collaborative Deep Learning: Using the probabilistic SDAE in Section 4.3.2 as a component, the generative process of CDL is defined as follows:

- (1) For each layer l of the SDAE network,
 - (a) For each column n of the weight matrix W_l , draw $W_{l,*n} \sim \mathcal{N}(0, \lambda_w^{-1} I_{K_l})$,
 - (b) Draw the bias vector $b_l \sim \mathcal{N}(0, \lambda_w^{-1} I_{K_l})$,
 - (c) For each row j of X_l , draw $X_{l,j*} \sim \mathcal{N}(\sigma(X_{l-1,j*} W_l + b_l), \lambda_s^{-1} I_{K_l})$.
- (2) For each item j ,
 - (a) Draw a clean input $X_{c,j*} \sim \mathcal{N}(X_{L,j*}, \lambda_n^{-1} I_J)$,
 - (b) Draw the latent item offset vector $\epsilon_j \sim \mathcal{N}(0, \lambda_v^{-1} I_K)$ and then set the latent item vector: $v_j = \epsilon_j + X_{L/2,j*}^T$.
- (3) Draw a latent user vector for each user i : $u_i \sim \mathcal{N}(0, \lambda_u^{-1} I_K)$.
- (4) Draw a rating R_{ij} for each user-item pair (i, j) : $R_{ij} \sim \mathcal{N}(u_i^T v_j, C_{ij}^{-1})$.

Here, λ_w , λ_n , λ_u , λ_s , and λ_v are hyperparameters and C_{ij} is a confidence parameter similar to that for CTR [112] ($C_{ij} = a$ if $R_{ij} = 1$ and $C_{ij} = b$, otherwise). Note that the middle layer $X_{L/2}$ serves as a bridge between the ratings and content information. This middle layer, along with the latent offset ϵ_j , is the key that enables CDL to simultaneously learn an effective feature representation and capture the similarity and (implicit) relationship among items (and users). Similar to the generalized SDAE, we can also take λ_s to infinity for computational efficiency.

Fig. 8. Left: NN representation for degenerated CDL. Right: Sampling as generalized BP in Bayesian CDL.

The graphical model of CDL when λ_s approaches positive infinity is shown in Figure 7, where, for notational simplicity, we use $\mathbf{x}_0$, $\mathbf{x}_{L/2}$, and $\mathbf{x}_L$ in place of $\mathbf{X}_{0,j*}^T$, $\mathbf{X}_{\frac{L}{2},j*}^T$, and $\mathbf{X}_{L,j*}^T$, respectively.

Note that according to the definition in Section 4.2, here the perception variables $\Omega_p = \{\{\mathbf{W}_l\}, \{\mathbf{b}_l\}, \{\mathbf{X}_l\}, \mathbf{X}_c\}$, the hinge variables $\Omega_h = \{\mathbf{V}\}$, and the task variables $\Omega_t = \{\mathbf{U}, \mathbf{R}\}$.

Learning: Based on the CDL model above, all parameters could be treated as random variables so that fully Bayesian methods such as Markov chain Monte Carlo (MCMC) or variational inference [55] may be applied. However, such treatment typically incurs high computational cost. Therefore, CDL uses an EM-style algorithm to obtain the MAP estimates, as in Reference [112].

Concretely, note that maximizing the posterior probability is equivalent to maximizing the joint log-likelihood of $\mathbf{U}$, $\mathbf{V}$, $\{\mathbf{X}_l\}$, $\mathbf{X}_c$, $\{\mathbf{W}_l\}$, $\{\mathbf{b}_l\}$, and $\mathbf{R}$ given λ_u , λ_v , λ_w , λ_s , and λ_n :

$$\begin{aligned} \mathcal{L} = & -\frac{\lambda_u}{2} \sum_i \|\mathbf{u}_i\|_2^2 - \frac{\lambda_w}{2} \sum_l (\|\mathbf{W}_l\|_F^2 + \|\mathbf{b}_l\|_2^2) - \frac{\lambda_v}{2} \sum_j \|\mathbf{v}_j - \mathbf{x}_{\frac{L}{2},j*}^T\|_2^2 - \frac{\lambda_n}{2} \sum_j \|\mathbf{x}_{L,j*} - \mathbf{x}_{c,j*}\|_2^2 \\ & - \frac{\lambda_s}{2} \sum_l \sum_j \|\sigma(\mathbf{x}_{l-1,j*} \mathbf{W}_l + \mathbf{b}_l) - \mathbf{x}_{L,j*}\|_2^2 - \sum_{i,j} \frac{C_{ij}}{2} (\mathbf{r}_{ij} - \mathbf{u}_i^T \mathbf{v}_j)^2. \end{aligned}$$

If λ_s goes to infinity, then the likelihood becomes

$$\begin{aligned} \mathcal{L} = & -\frac{\lambda_u}{2} \sum_i \|\mathbf{u}_i\|_2^2 - \frac{\lambda_w}{2} \sum_l (\|\mathbf{W}_l\|_F^2 + \|\mathbf{b}_l\|_2^2) - \frac{\lambda_v}{2} \sum_j \|\mathbf{v}_j - f_e(\mathbf{x}_{0,j*}, \mathbf{W}^+)\|_2^2 \\ & - \frac{\lambda_n}{2} \sum_j \|f_r(\mathbf{x}_{0,j*}, \mathbf{W}^+) - \mathbf{x}_{c,j*}\|_2^2 - \sum_{i,j} \frac{C_{ij}}{2} (\mathbf{r}_{ij} - \mathbf{u}_i^T \mathbf{v}_j)^2, \end{aligned} \quad (9)$$

where the encoder function $f_e(\cdot, \mathbf{W}^+)$ takes the corrupted content vector $\mathbf{x}_{0,j*}$ of item j as input and computes its encoding, and the function $f_r(\cdot, \mathbf{W}^+)$ also takes $\mathbf{x}_{0,j*}$ as input, computes the encoding and then reconstructs item j 's content vector. For example, if the number of layers $L = 6$, $f_e(\mathbf{x}_{0,j*}, \mathbf{W}^+)$ is the output of the third layer while $f_r(\mathbf{x}_{0,j*}, \mathbf{W}^+)$ is the output of the sixth layer.

From the perspective of optimization, the third term in the objective function, i.e., Equation (9), above is equivalent to a multi-layer perceptron using the latent item vectors $\mathbf{v}_j$ as the target while the fourth term is equivalent to an SDAE minimizing the reconstruction error. Seeing from the view of neural networks (NN), when λ_s approaches positive infinity, training of the probabilistic graphical model of CDL in Figure 7 (left) would degenerate to simultaneously training two neural networks overlaid together with a common input layer (the corrupted input) but different output layers, as shown in Figure 8 (left). Note that the second network is much more complex than typical neural networks due to the involvement of the rating matrix.

When the ratio λ_n/λ_v approaches positive infinity, it will degenerate to a two-step model in which the latent representation learned using SDAE is put directly into the CTR. Another extreme happens when λ_n/λ_v goes to zero where the decoder of the SDAE essentially vanishes. Figure 7 (right) shows the graphical model of the degenerated CDL when λ_n/λ_v goes to zero. As demonstrated in experiments, the predictive performance will suffer greatly for both extreme cases [121].

For $\mathbf{u}_i$ and $\mathbf{v}_j$, block coordinate descent similar to References [50, 112] is used. Given the current $\mathbf{W}^+$, we compute the gradients of $\mathcal{L}$ with respect to $\mathbf{u}_i$ and $\mathbf{v}_j$ and then set them to zero, leading to the following update rules:

$$\mathbf{u}_i \leftarrow (\mathbf{V}\mathbf{C}_i\mathbf{V}^T + \lambda_u\mathbf{I}_K)^{-1}\mathbf{V}\mathbf{C}_i\mathbf{R}_i, \quad \mathbf{v}_j \leftarrow (\mathbf{U}\mathbf{C}_j\mathbf{U}^T + \lambda_v\mathbf{I}_K)^{-1}(\mathbf{U}\mathbf{C}_j\mathbf{R}_j + \lambda_v f_e(\mathbf{X}_{0,j*}, \mathbf{W}^+)^T),$$

where $\mathbf{U} = (\mathbf{u}_i)_{i=1}^I$, $\mathbf{V} = (\mathbf{v}_j)_{j=1}^J$, $\mathbf{C}_i = \text{diag}(C_{i1}, \dots, C_{iJ})$ is a diagonal matrix, $\mathbf{R}_i = (\mathbf{R}_{i1}, \dots, \mathbf{R}_{iJ})^T$ is a column vector containing all the ratings of user i , and C_{ij} reflects the confidence controlled by a and b as discussed in Reference [50]. $\mathbf{C}_j$ and $\mathbf{R}_j$ are defined similarly for item j .

Given $\mathbf{U}$ and $\mathbf{V}$, we can learn the weights $\mathbf{W}_l$ and biases $\mathbf{b}_l$ for each layer using the back-propagation learning algorithm. The gradients of the likelihood with respect to $\mathbf{W}_l$ and $\mathbf{b}_l$ are as follows:

$$\begin{aligned} \nabla_{\mathbf{W}_l}\mathcal{L} &= -\lambda_w\mathbf{W}_l - \lambda_v \sum_j \nabla_{\mathbf{W}_l} f_e(\mathbf{X}_{0,j*}, \mathbf{W}^+)^T (f_e(\mathbf{X}_{0,j*}, \mathbf{W}^+)^T - \mathbf{v}_j) \\ &\quad - \lambda_n \sum_j \nabla_{\mathbf{W}_l} f_r(\mathbf{X}_{0,j*}, \mathbf{W}^+) (f_r(\mathbf{X}_{0,j*}, \mathbf{W}^+) - \mathbf{X}_{c,j*}), \\ \nabla_{\mathbf{b}_l}\mathcal{L} &= -\lambda_w\mathbf{b}_l - \lambda_v \sum_j \nabla_{\mathbf{b}_l} f_e(\mathbf{X}_{0,j*}, \mathbf{W}^+)^T (f_e(\mathbf{X}_{0,j*}, \mathbf{W}^+)^T - \mathbf{v}_j) \\ &\quad - \lambda_n \sum_j \nabla_{\mathbf{b}_l} f_r(\mathbf{X}_{0,j*}, \mathbf{W}^+) (f_r(\mathbf{X}_{0,j*}, \mathbf{W}^+) - \mathbf{X}_{c,j*}). \end{aligned}$$

By alternating the update of $\mathbf{U}$, $\mathbf{V}$, $\mathbf{W}_l$, and $\mathbf{b}_l$, we can find a local optimum for $\mathcal{L}$. Several commonly used techniques such as using a momentum term may be applied to alleviate the local optimum problem.

Prediction: Let D be the observed test data. Similar to Reference [112], CDL uses the point estimates of $\mathbf{u}_i$, $\mathbf{W}^+$ and ϵ_j to calculate the predicted rating:

$$E[\mathbf{R}_{ij}|D] \approx E[\mathbf{u}_i|D]^T (E[f_e(\mathbf{X}_{0,j*}, \mathbf{W}^+)^T|D] + E[\epsilon_j|D]),$$

where $E[\cdot]$ denotes the expectation operation. In other words, we approximate the predicted rating as

$$\mathbf{R}_{ij}^* \approx (\mathbf{u}_i^*)^T (f_e(\mathbf{X}_{0,j*}, \mathbf{W}^{+*})^T + \epsilon_j^*) = (\mathbf{u}_i^*)^T \mathbf{v}_j^*.$$

Note that for any new item j with no rating in the training data, its offset ϵ_j^* will be 0.

Recall that in CDL, the probabilistic SDAE and PMF work as the perception and task-specific components. As mentioned in Section 4, both components can take various forms, leading to different concrete models. For example, one can replace the probabilistic SDAE with a VAE or an NPN as the perception component [68]. It is also possible to use Bayesian PMF [97] rather than PMF [96] as the task-specific component and thereby produce more robust predictions.

In the following subsections, we provide several extensions of CDL from different perspectives.

5.1.2 Bayesian Collaborative Deep Learning. Besides the MAP estimates, a sampling-based algorithm for the Bayesian treatment of CDL is also proposed in Reference [121]. This algorithm turns out to be a Bayesian and generalized version of BP. We list the key conditional densities as follows:

For $\mathbf{W}^+$: We denote the concatenation of $\mathbf{W}_{l,*n}$ and $\mathbf{b}_l^{(n)}$ as $\mathbf{W}_{l,*n}^+$. Similarly, the concatenation of $\mathbf{X}_{l,j*}$ and 1 is denoted as $\mathbf{X}_{l,j*}^+$. The subscripts of $\mathbf{I}$ are ignored. Then,

$$p(\mathbf{W}_{l,*n}^+ | \mathbf{X}_{l-1,j*}, \mathbf{X}_{l,j*}, \lambda_s) \propto \mathcal{N}(\mathbf{W}_{l,*n}^+ | 0, \lambda_w^{-1} \mathbf{I}) \cdot \mathcal{N}(\mathbf{X}_{l,*n} | \sigma(\mathbf{X}_{l-1}^+ \mathbf{W}_{l,*n}^+), \lambda_s^{-1} \mathbf{I}).$$

For $\mathbf{X}_{l,j*}$ ($l \neq L/2$): Similarly, we denote the concatenation of $\mathbf{W}_l$ and $\mathbf{b}_l$ as $\mathbf{W}_l^+$ and have

$$p(\mathbf{X}_{l,j*} | \mathbf{W}_l^+, \mathbf{W}_{l+1}^+, \mathbf{X}_{l-1,j*}, \mathbf{X}_{l+1,j*}, \lambda_s) \\ \propto \mathcal{N}(\mathbf{X}_{l,j*} | \sigma(\mathbf{X}_{l-1,j*}^+ \mathbf{W}_l^+), \lambda_s^{-1} \mathbf{I}) \cdot \mathcal{N}(\mathbf{X}_{l+1,j*} | \sigma(\mathbf{X}_{l,j*}^+ \mathbf{W}_{l+1}^+), \lambda_s^{-1} \mathbf{I}),$$

where for the last layer ($l = L$) the second Gaussian would be $\mathcal{N}(\mathbf{X}_{c,j*} | \mathbf{X}_{l,j*}, \lambda_s^{-1} \mathbf{I})$ instead.

For $\mathbf{X}_{l,j*}$ ($l = L/2$): Similarly, we have

$$p(\mathbf{X}_{l,j*} | \mathbf{W}_l^+, \mathbf{W}_{l+1}^+, \mathbf{X}_{l-1,j*}, \mathbf{X}_{l+1,j*}, \lambda_s, \lambda_v, \mathbf{v}_j) \\ \propto \mathcal{N}(\mathbf{X}_{l,j*} | \sigma(\mathbf{X}_{l-1,j*}^+ \mathbf{W}_l^+), \lambda_s^{-1} \mathbf{I}) \cdot \mathcal{N}(\mathbf{X}_{l+1,j*} | \sigma(\mathbf{X}_{l,j*}^+ \mathbf{W}_{l+1}^+), \lambda_s^{-1} \mathbf{I}) \cdot \mathcal{N}(\mathbf{v}_j | \mathbf{X}_{l,j*}, \lambda_v^{-1} \mathbf{I}).$$

For $\mathbf{v}_j$: The posterior $p(\mathbf{v}_j | \mathbf{X}_{L/2,j*}, \mathbf{R}_{*j}, \mathbf{C}_{*j}, \lambda_v, \mathbf{U}) \propto \mathcal{N}(\mathbf{v}_j | \mathbf{X}_{L/2,j*}^T, \lambda_v^{-1} \mathbf{I}) \prod_i \mathcal{N}(\mathbf{R}_{ij} | \mathbf{u}_i^T \mathbf{v}_j, \mathbf{C}_{ij}^{-1})$.

For $\mathbf{u}_i$: The posterior $p(\mathbf{u}_i | \mathbf{R}_{i*}, \mathbf{V}, \lambda_u, \mathbf{C}_{i*}) \propto \mathcal{N}(\mathbf{u}_i | 0, \lambda_u^{-1} \mathbf{I}) \prod_j \mathcal{N}(\mathbf{R}_{ij} | \mathbf{u}_i^T \mathbf{v}_j, \mathbf{C}_{ij}^{-1})$.

Interestingly, if λ_s goes to infinity and adaptive rejection Metropolis sampling (which involves using the gradients of the objective function to approximate the proposal distribution) is used, then the sampling for $\mathbf{W}^+$ turns out to be a *Bayesian generalized* version of BP. Specifically, as Figure 8 (right) shows, after getting the gradient of the loss function at one point (the red dashed line on the left), the next sample would be drawn in the region under that line, which is equivalent to a probabilistic version of BP. If a sample is above the curve of the loss function, then a new tangent line (the black dashed line on the right) would be added to better approximate the distribution corresponding to the loss function. After that, samples would be drawn from the region under both lines. During the sampling, besides searching for local optima using the gradients (MAP), the algorithm also takes the variance into consideration. That is why it is called *Bayesian generalized back-propagation* in Reference [121].

5.1.3 Marginalized Collaborative Deep Learning. In SDAE, corrupted input goes through the encoder and decoder to recover the clean input. Usually, different epochs of training use different corrupted versions as input. Hence, generally, SDAE needs to go through enough epochs of training to see sufficient corrupted versions of the input. Marginalized SDAE (mSDAE) [14] seeks to avoid this by marginalizing out the corrupted input and obtaining closed-form solutions directly. In this sense, mSDAE is more computationally efficient than SDAE.

As mentioned in Reference [66], using mSDAE instead of the Bayesian SDAE could lead to more efficient learning algorithms. For example, in Reference [66], the objective when using a one-layer mSDAE can be written as follows:

$$\mathcal{L} = - \sum_j \|\tilde{\mathbf{X}}_{0,j*} \mathbf{W}_1 - \bar{\mathbf{X}}_{c,j*}\|_2^2 - \sum_{i,j} \frac{\mathbf{C}_{ij}}{2} (\mathbf{R}_{ij} - \mathbf{u}_i^T \mathbf{v}_j)^2 - \frac{\lambda_u}{2} \sum_i \|\mathbf{u}_i\|_2^2 - \frac{\lambda_v}{2} \sum_j \|\mathbf{v}_j^T \mathbf{P}_1 - \mathbf{X}_{0,j*} \mathbf{W}_1\|_2^2,$$

where $\tilde{\mathbf{X}}_{0,j*}$ is the collection of k different corrupted versions of $\mathbf{X}_{0,j*}$ (a k -by- B matrix) and $\bar{\mathbf{X}}_{c,j*}$ is the k -time repeated version of $\mathbf{X}_{c,j*}$ (also a k -by- B matrix). $\mathbf{P}_1$ is the transformation matrix for item latent factors.

The solution for $\mathbf{W}_1$ would be $\mathbf{W}_1 = E(\mathbf{S}_1)E(\mathbf{Q}_1)^{-1}$, where $\mathbf{S}_1 = \bar{\mathbf{X}}_{c,j*}^T \tilde{\mathbf{X}}_{0,j*} + \frac{\lambda_v}{2} \mathbf{P}_1^T \mathbf{V} \mathbf{X}_c$ and $\mathbf{Q}_1 = \bar{\mathbf{X}}_{c,j*}^T \tilde{\mathbf{X}}_{0,j*} + \frac{\lambda_v}{2} \mathbf{X}_c^T \mathbf{X}_c$. A solver for the expectation in the equation above is provided in Reference [14]. Note that this is a linear and one-layer case, which can be generalized to the nonlinear and multi-layer case using the same techniques as in References [13, 14].

Marginalized CDL's perception variables $\Omega_p = \{\mathbf{X}_0, \mathbf{X}_c, \mathbf{W}_1\}$, its hinge variables $\Omega_h = \{\mathbf{V}\}$, and its task variables $\Omega_t = \{\mathbf{P}_1, \mathbf{R}, \mathbf{U}\}$.

5.1.4 Collaborative Deep Ranking. CDL assumes a collaborative filtering setting to model the ratings directly. Naturally, one can design a similar model to focus more on the ranking among items rather than exact ratings [131]. The corresponding generative process is as follows:

- (1) For each layer l of the SDAE network,
 - (a) For each column n of the weight matrix $\mathbf{W}_l$, draw $\mathbf{W}_{l,*n} \sim \mathcal{N}(\mathbf{0}, \lambda_w^{-1} \mathbf{I}_{K_l})$,
 - (b) Draw the bias vector $\mathbf{b}_l \sim \mathcal{N}(\mathbf{0}, \lambda_w^{-1} \mathbf{I}_{K_l})$,
 - (c) For each row j of $\mathbf{X}_l$, draw $\mathbf{X}_{l,j*} \sim \mathcal{N}(\sigma(\mathbf{X}_{l-1,j*} \mathbf{W}_l + \mathbf{b}_l), \lambda_s^{-1} \mathbf{I}_{K_l})$.
- (2) For each item j ,
 - (a) Draw a clean input $\mathbf{X}_{c,j*} \sim \mathcal{N}(\mathbf{X}_{L,j*}, \lambda_n^{-1} \mathbf{I}_J)$,
 - (b) Draw a latent item offset vector $\epsilon_j \sim \mathcal{N}(\mathbf{0}, \lambda_v^{-1} \mathbf{I}_K)$ and then set the latent item vector to be: $\mathbf{v}_j = \epsilon_j + \mathbf{X}_{\frac{L}{2},j*}^T$.
- (3) For each user i ,
 - (a) Draw a latent user vector for each user i : $\mathbf{u}_i \sim \mathcal{N}(\mathbf{0}, \lambda_u^{-1} \mathbf{I}_K)$,
 - (b) For each pair-wise preference $(j, k) \in \mathcal{P}_i$, where $\mathcal{P}_i = \{(j, k) : \mathbf{R}_{ij} - \mathbf{R}_{ik} > 0\}$, draw the preference: $\Delta_{ijk} \sim \mathcal{N}(\mathbf{u}_i^T \mathbf{v}_j - \mathbf{u}_i^T \mathbf{v}_k, \mathbf{C}_{ijk}^{-1})$.

Following the generative process above, the log-likelihood in Equation (9) becomes

$$\begin{aligned} \mathcal{L} = & -\frac{\lambda_u}{2} \sum_i \|\mathbf{u}_i\|_2^2 - \frac{\lambda_w}{2} \sum_l (\|\mathbf{W}_l\|_F^2 + \|\mathbf{b}_l\|_2^2) - \frac{\lambda_v}{2} \sum_j \|\mathbf{v}_j - f_e(\mathbf{X}_{0,j*}, \mathbf{W}^+)\|^2 \\ & - \frac{\lambda_n}{2} \sum_j \|f_r(\mathbf{X}_{0,j*}, \mathbf{W}^+) - \mathbf{X}_{c,j*}\|_2^2 - \sum_{i,j,k} \frac{\mathbf{C}_{ijk}}{2} (\Delta_{ijk} - (\mathbf{u}_i^T \mathbf{v}_j - \mathbf{u}_i^T \mathbf{v}_k))^2. \end{aligned}$$

Similar algorithms can be used to learn the parameters in CDR. As reported in Reference [131], using the ranking objective leads to significant improvement in the recommendation performance. Following the definition in Section 4.2, CDR's perception variables $\Omega_p = \{\{\mathbf{W}_l\}, \{\mathbf{b}_l\}, \{\mathbf{X}_l\}, \mathbf{X}_c\}$, the hinge variables $\Omega_h = \{\mathbf{V}\}$, and the task variables $\Omega_t = \{\mathbf{U}, \Delta\}$.

5.1.5 Collaborative Variational Autoencoders. In CDL, the perception component takes the form of a probabilistic SDAE. Naturally, one can also replace the probabilistic SDAE in CDL with a VAE (introduced in Section 4.3.3), as is done in collaborative variational autoencoders (CVAE) [68]. Specifically, CVAE with an inference network (encoder) denoted as $(f_\mu(\cdot), f_s(\cdot))$ and a generation network (decoder) denoted as $g(\cdot)$ assumes the following generative process:

- (1) For each item j ,
 - (a) Draw the latent item vector from the VAE inference network: $\mathbf{z}_j \sim \mathcal{N}(f_\mu(\mathbf{X}_{0,j*}), f_s(\mathbf{X}_{0,j*}))$,
 - (b) Draw the latent item offset vector $\epsilon_j \sim \mathcal{N}(\mathbf{0}, \lambda_v^{-1} \mathbf{I}_K)$ and then set the latent item vector: $\mathbf{v}_j = \epsilon_j + \mathbf{z}_j$,
 - (c) Draw the original input from the VAE generation network $\mathbf{X}_{0,j*} \sim \mathcal{N}(g(\mathbf{z}_j), \lambda_n^{-1} \mathbf{I}_B)$.
- (2) Draw a latent user vector for each user i : $\mathbf{u}_i \sim \mathcal{N}(\mathbf{0}, \lambda_u^{-1} \mathbf{I}_K)$.
- (3) Draw a rating $\mathbf{R}_{ij}$ for each user-item pair (i, j) : $\mathbf{R}_{ij} \sim \mathcal{N}(\mathbf{u}_i^T \mathbf{v}_j, \mathbf{C}_{ij}^{-1})$.

Similar to CDL, λ_n , λ_u , λ_s , and λ_v are hyperparameters and C_{ij} is a confidence parameter ($C_{ij} = a$ if $R_{ij} = 1$ and $C_{ij} = b$ otherwise). Following Reference [68], the ELBO similar to Equation (6) can be derived, using which, one can train the model's parameters using BP and the reparameterization trick.

The evolution from CDL to CVAE demonstrates the BDL framework's flexibility in terms of its components' specific forms. It is also worth noting that the perception component can be a recurrent version of probabilistic SDAE [122] or VAE [17, 68] to handle raw sequential data, while the task-specific component can take more sophisticated forms to accommodate more complex recommendation scenarios (e.g., cross-domain recommendation).

5.1.6 Discussion. Recommender systems are a typical use case for BDL in that they often require both thorough understanding of high-dimensional signals (e.g., text and images) and principled reasoning on the conditional dependencies among users/items/ratings.

In this regard, CDL, as an instantiation of BDL, is the first hierarchical Bayesian model to bridge the gap between state-of-the-art deep learning models and recommender systems. By performing deep learning collaboratively, CDL and its variants can simultaneously extract an effective deep feature representation from high-dimensional content and capture the similarity and implicit relationship between items (and users). The learned representation may also be used for tasks other than recommendation. Unlike previous deep learning models, which use a simple target like classification [56] and reconstruction [111], CDL-based models use CF as a more complex target in a probabilistic framework.

As mentioned in Section 1, *information exchange* between two components is crucial for the performance of BDL. In the CDL-based models above, the exchange is achieved by assuming Gaussian distributions that connect the hinge variables and the variables in the perception component (drawing the hinge variable $\mathbf{v}_j \sim \mathcal{N}(\mathbf{X}_{\frac{L}{2}, j*}^T, \lambda_v^{-1} \mathbf{I}_K)$ in the generative process of CDL, where $\mathbf{X}_{\frac{L}{2}}$ is a perception variable), which is simple but effective and efficient in computation. Among the eight CDL-based models in Table 1, six of them are HV models and the others are LV models, according to the definition of Section 4.2. Since it has been verified that the HV CDL significantly outperforms its ZV counterpart [121], we can expect additional performance boost from the LV counterparts of the six HV models.

Besides efficient *information exchange*, the model designs also meet the independence requirement on the distribution concerning hinge variables discussed in Section 4.2 and are hence easily parallelizable. In some models to be introduced later, we will see alternative designs to enable efficient and independent information exchange between the two components of BDL.

Note that BDL-based models above use typical static Bayesian networks as their task-specific components. Although these are often sufficient for most use cases, it is possible for the task-specific components to take the form of deep Bayesian networks such as BIN [117]. This allows the models to handle highly nonlinear interactions between users and items if necessary. One can also use stochastic processes (or dynamic Bayesian networks in general) to explicitly model users purchase or clicking behaviors. For example, it is natural to model a user's purchase of groceries as a Poisson process. In terms of perception components, one can also replace the pSDAE, mSDAE, or VAE above with their convolutional or recurrent counterparts (see Sections 2.3 and 2.4), as is done in Collaborative Knowledge Base Embedding (CKE) [132] or Collaborative Recurrent Autoencoders (CRAE) [122], respectively. Note that for the convolutional or recurrent perception components to be compatible with the task-specific component (which is inherently probabilistic), ideally one would need to formulate probabilistic versions of CNN or RNN as well. Readers are referred to References [132] and [122] for more details.

Fig. 9. Graphical model of RSDAE for $L = 4$. λ_s is omitted here to prevent clutter.

In summary, this subsection discusses BDL’s applications on supervised learning, using recommender systems as an example. Section 5.2 below will cover BDL’s applications on unsupervised learning.

5.2 Unsupervised Bayesian Deep Learning for Topic Models

To demonstrate how BDL can also be applied to unsupervised learning, we review some examples of BDL-based topic models in this section. These models combine the merits of PGM (which naturally incorporates the probabilistic relations among variables) and NN (which learns deep representations efficiently), leading to significant performance boost. In the case of unsupervised learning, the ‘task’ for a task-specific component is to describe/characterize the conditional dependencies in the BDL model, thereby improving its interpretability and generalizability. This is different from the supervised learning setting where the “task” is simply to “match the target.”

5.2.1 Relational Stacked Denoising Autoencoders as Topic Models. As a BDL-based topic model, relational stacked denoising autoencoders (RSDAE) essentially tries to learn a hierarchy of topics (or latent factors) while enforcing relational (graph) constraints under an unsupervised learning setting.

Problem Statement and Notation: Assume we have a set of items (articles or movies) X_c , with $X_{c,j*}^T \in \mathbb{R}^B$ denoting the content (attributes) of item j . Besides, we use I_K to denote a K -dimensional identity matrix and $S = [s_1, s_2, \dots, s_J]$ to denote the *relational latent matrix* with s_j representing the relational properties of item j .

From the perspective of SDAE, the J -by- B matrix X_c represents the clean input to the SDAE and the noise-corrupted matrix of the same size is denoted by X_0 . Besides, we denote the output of layer l of the SDAE, a J -by- K_l matrix, by X_l . Row j of X_l is denoted by $X_{l,j*}$, W_l and b_l are the weight matrix and bias vector of layer l , $W_{l,*n}$ denotes column n of W_l , and L is the number of layers. As a shorthand, we refer to the collection of weight matrices and biases in all layers as W^+ . Note that an $L/2$ -layer SDAE corresponds to an L -layer network.

Model Formulation: In RSDAE, the perception component takes the form of a probabilistic SDAE (introduced in Section 4.3.2) as a building block. At a higher level, RSDAE is formulated as a novel probabilistic model that seamlessly integrates a hierarchy of latent factors and the relational information available. This way, the model can simultaneously learn the feature representation from the content information and the relation among items [118]. The graphical model for RSDAE is shown in Figure 9, and the generative process is listed as follows:

- (1) Draw the relational latent matrix S from a *matrix-variate normal distribution* [33]:

$$S \sim \mathcal{N}_{K,J}(0, I_K \otimes (\lambda_l \mathcal{L}_a)^{-1}). \quad (10)$$

- (2) For layer l of the SDAE where $l = 1, 2, \dots, \frac{L}{2} - 1$,

- (a) For each column n of the weight matrix $\mathbf{W}_l$, draw $\mathbf{W}_{l,*n} \sim \mathcal{N}(0, \lambda_w^{-1} \mathbf{I}_{K_l})$,
- (b) Draw the bias vector $\mathbf{b}_l \sim \mathcal{N}(0, \lambda_w^{-1} \mathbf{I}_{K_l})$,
- (c) For each row j of $\mathbf{X}_l$, draw $\mathbf{X}_{l,j*} \sim \mathcal{N}(\sigma(\mathbf{X}_{l-1,j*} \mathbf{W}_l + \mathbf{b}_l), \lambda_s^{-1} \mathbf{I}_{K_l})$.
- (3) For layer $\frac{L}{2}$ of the SDAE, draw the representation vector for item j from the product of two Gaussians (PoG) [23]:

$$\mathbf{X}_{\frac{L}{2},j*} \sim \text{PoG}\left(\sigma\left(\mathbf{X}_{\frac{L}{2}-1,j*} \mathbf{W}_l + \mathbf{b}_l\right), \mathbf{s}_j^T, \lambda_s^{-1} \mathbf{I}_K, \lambda_r^{-1} \mathbf{I}_K\right). \quad (11)$$

- (4) For layer l of the SDAE where $l = \frac{L}{2} + 1, \frac{L}{2} + 2, \dots, L$,
 - (a) For each column n of the weight matrix $\mathbf{W}_l$, draw $\mathbf{W}_{l,*n} \sim \mathcal{N}(0, \lambda_w^{-1} \mathbf{I}_{K_l})$,
 - (b) Draw the bias vector $\mathbf{b}_l \sim \mathcal{N}(0, \lambda_w^{-1} \mathbf{I}_{K_l})$,
 - (c) For each row j of $\mathbf{X}_l$, draw $\mathbf{X}_{l,j*} \sim \mathcal{N}(\sigma(\mathbf{X}_{l-1,j*} \mathbf{W}_l + \mathbf{b}_l), \lambda_s^{-1} \mathbf{I}_{K_l})$.
- (5) For each item j , draw a clean input $\mathbf{X}_{c,j*} \sim \mathcal{N}(\mathbf{X}_{L,j*}, \lambda_n^{-1} \mathbf{I}_B)$.

Here, $K = K_{\frac{L}{2}}$ is the dimensionality of the learned representation vector for each item, $\mathbf{S}$ denotes the $K \times J$ relational latent matrix in which column j is the *relational latent vector* $\mathbf{s}_j$ for item j . Note that $\mathcal{N}_{K,J}(0, \mathbf{I}_K \otimes (\lambda_l \mathcal{L}_a)^{-1})$ in Equation (10) is a matrix-variate normal distribution defined as in Reference [33]:

$$p(\mathbf{S}) = \mathcal{N}_{K,J}(0, \mathbf{I}_K \otimes (\lambda_l \mathcal{L}_a)^{-1}) = \frac{\exp\{\text{tr}[-\frac{\lambda_l}{2} \mathbf{S} \mathcal{L}_a \mathbf{S}^T]\}}{(2\pi)^{JK/2} |\mathbf{I}_K|^{J/2} |\lambda_l \mathcal{L}_a|^{-K/2}}, \quad (12)$$

where the operator $\otimes$ denotes the Kronecker product of two matrices [33], $\text{tr}(\cdot)$ denotes the trace of a matrix, and $\mathcal{L}_a$ is the Laplacian matrix incorporating the relational information. $\mathcal{L}_a = \mathbf{D} - \mathbf{A}$, where $\mathbf{D}$ is a diagonal matrix whose diagonal elements $\mathbf{D}_{ii} = \sum_j \mathbf{A}_{ij}$ and $\mathbf{A}$ is the adjacency matrix representing the relational information with binary entries indicating the links (or relations) between items. $\mathbf{A}_{jj'} = 1$ indicates that there is a link between item j and item j' and $\mathbf{A}_{jj'} = 0$ otherwise. $\text{PoG}(\sigma(\mathbf{X}_{\frac{L}{2}-1,j*} \mathbf{W}_l + \mathbf{b}_l), \mathbf{s}_j^T, \lambda_s^{-1} \mathbf{I}_K, \lambda_r^{-1} \mathbf{I}_K)$ denotes the product of the Gaussian $\mathcal{N}(\sigma(\mathbf{X}_{\frac{L}{2}-1,j*} \mathbf{W}_l + \mathbf{b}_l), \lambda_s^{-1} \mathbf{I}_K)$ and the Gaussian $\mathcal{N}(\mathbf{s}_j^T, \lambda_r^{-1} \mathbf{I}_K)$, which is also a Gaussian [23].

According to the generative process above, maximizing the posterior probability is equivalent to maximizing the joint log-likelihood of $\{\mathbf{X}_l\}$, $\mathbf{X}_c$, $\mathbf{S}$, $\{\mathbf{W}_l\}$, and $\{\mathbf{b}_l\}$ given λ_s , λ_w , λ_l , λ_r , and λ_n :

$$\begin{aligned} \mathcal{L} = & -\frac{\lambda_l}{2} \text{tr}(\mathbf{S} \mathcal{L}_a \mathbf{S}^T) - \frac{\lambda_r}{2} \sum_j \|(\mathbf{s}_j^T - \mathbf{X}_{\frac{L}{2},j*})\|_2^2 - \frac{\lambda_w}{2} \sum_l (\|\mathbf{W}_l\|_F^2 + \|\mathbf{b}_l\|_2^2) \\ & - \frac{\lambda_n}{2} \sum_j \|\mathbf{X}_{L,j*} - \mathbf{X}_{c,j*}\|_2^2 - \frac{\lambda_s}{2} \sum_l \sum_j \|\sigma(\mathbf{X}_{l-1,j*} \mathbf{W}_l + \mathbf{b}_l) - \mathbf{X}_{l,j*}\|_2^2. \end{aligned}$$

Similar to the pSDAE, taking λ_s to infinity, the joint log-likelihood becomes

$$\mathcal{L} = -\frac{\lambda_l}{2} \text{tr}(\mathbf{S} \mathcal{L}_a \mathbf{S}^T) - \frac{\lambda_r}{2} \sum_j \|(\mathbf{s}_j^T - \mathbf{X}_{\frac{L}{2},j*})\|_2^2 - \frac{\lambda_w}{2} \sum_l (\|\mathbf{W}_l\|_F^2 + \|\mathbf{b}_l\|_2^2) - \frac{\lambda_n}{2} \sum_j \|\mathbf{X}_{L,j*} - \mathbf{X}_{c,j*}\|_2^2, \quad (13)$$

where $\mathbf{X}_{l,j*} = \sigma(\mathbf{X}_{l-1,j*} \mathbf{W}_l + \mathbf{b}_l)$. Note that the first term $-\frac{\lambda_l}{2} \text{tr}(\mathbf{S} \mathcal{L}_a \mathbf{S}^T)$ corresponds to $\log p(\mathbf{S})$ in the matrix-variate distribution in Equation (12). Besides, by simple manipulation, we have $\text{tr}(\mathbf{S} \mathcal{L}_a \mathbf{S}^T) = \sum_{k=1}^K \mathbf{S}_{k*}^T \mathcal{L}_a \mathbf{S}_{k*}$, where $\mathbf{S}_{k*}$ denotes the k th row of $\mathbf{S}$. As we can see, maximizing $-\frac{\lambda_l}{2} \text{tr}(\mathbf{S} \mathcal{L}_a \mathbf{S}^T)$ is equivalent to making $\mathbf{s}_j$ closer to $\mathbf{s}_{j'}$ if item j and item j' are linked (namely, $\mathbf{A}_{jj'} = 1$) [115].

In RSDAE, the perception variables $\Omega_p = \{\{\mathbf{X}_l\}, \mathbf{X}_c, \{\mathbf{W}_l\}, \{\mathbf{b}_l\}\}$, the hinge variables $\Omega_h = \{\mathbf{S}\}$, and the task variables $\Omega_t = \{\mathbf{A}\}$.

Learning and Inference: Reference [118] provides an EM-style algorithm for MAP estimation. Below, we review some of the key steps.

For the E step, the challenge lies in the inference of the relational latent matrix $\mathbf{S}$. We first fix all rows of $\mathbf{S}$ except the k th one $\mathbf{S}_{k*}$ and then update $\mathbf{S}_{k*}$. Specifically, we take the gradient of $\mathcal{L}$ with respect to $\mathbf{S}_{k*}$, set it to 0, and get the following linear system:

$$(\lambda_l \mathcal{L}_a + \lambda_r \mathbf{I}_J) \mathbf{S}_{k*} = \lambda_r \mathbf{X}_{\frac{L}{2}, *k}^T. \quad (14)$$

A naive approach is to solve the linear system by setting $\mathbf{S}_{k*} = \lambda_r (\lambda_l \mathcal{L}_a + \lambda_r \mathbf{I}_J)^{-1} \mathbf{X}_{\frac{L}{2}, *k}^T$. Unfortunately, the complexity is $O(J^3)$ for one single update. Similar to Reference [67], the steepest descent method [101] is used to iteratively update $\mathbf{S}_{k*}$:

$$\mathbf{S}_{k*}(t+1) \leftarrow \mathbf{S}_{k*}(t) + \delta(t) \mathbf{r}(t), \quad \mathbf{r}(t) \leftarrow \lambda_r \mathbf{X}_{\frac{L}{2}, *k}^T - (\lambda_l \mathcal{L}_a + \lambda_r \mathbf{I}_J) \mathbf{S}_{k*}(t), \quad \delta(t) \leftarrow \frac{\mathbf{r}(t)^T \mathbf{r}(t)}{\mathbf{r}(t)^T (\lambda_l \mathcal{L}_a + \lambda_r \mathbf{I}_J) \mathbf{r}(t)}.$$

As discussed in Reference [67], the steepest descent method dramatically reduces the computation cost in each iteration from $O(J^3)$ to $O(J)$.

The M step involves learning $\mathbf{W}_l$ and $\mathbf{b}_l$ for each layer using the back-propagation algorithm given $\mathbf{S}$. By alternating the update of $\mathbf{S}$, $\mathbf{W}_l$, and $\mathbf{b}_l$, a local optimum for $\mathcal{L}$ can be found. Also, techniques such as including a momentum term may help to avoid being trapped in a local optimum.

5.2.2 Deep Poisson Factor Analysis with Sigmoid Belief Networks. The Poisson distribution with support over nonnegative integers is known as a natural choice to model counts. It is, therefore, desirable to use it as a building block for topic models, which are generally interested in word counts [8]. With this motivation, Reference [136] proposed a model, dubbed Poisson factor analysis (PFA), for latent nonnegative matrix factorization via Poisson distributions.

Poisson Factor Analysis: PFA assumes a discrete P -by- N matrix $\mathbf{X}$ containing word counts of N documents with a vocabulary size of P [24, 136]. In a nutshell, PFA can be described using the equation $\mathbf{X} \sim \text{Pois}(\Phi(\Theta \circ \mathbf{H}))$, where Φ (of size P -by- K where K is the number of topics) denotes the factor loading matrix in factor analysis with the k th column ϕ_k encoding the importance of each word in topic k . The K -by- N matrix Θ is the factor score matrix with the n th column θ_n containing topic proportions for document n . The K -by- N matrix $\mathbf{H}$ is a latent binary matrix with the n th column $\mathbf{h}_n$ defining a set of topics associated with document n .

Different priors correspond to different models. For example, Dirichlet priors on ϕ_k and θ_n with an all-one matrix $\mathbf{H}$ would recover LDA [8] while a beta-Bernoulli prior on $\mathbf{h}_n$ leads to the NB-FTM model in Reference [135]. In Reference [24], a deep-structured prior based on sigmoid belief networks (SBN) [79] (an MLP variant with binary hidden units) is imposed on $\mathbf{h}_n$ to form a deep PFA model for topic modeling.

Deep Poisson Factor Analysis: In the deep PFA model [24], the generative process can be summarized as follows:

$$\begin{aligned} \phi_k &\sim \text{Dir}(a_\phi, \dots, a_\phi), \quad \theta_{kn} \sim \text{Gamma}\left(r_k, \frac{p_n}{1-p_n}\right), \quad r_k \sim \text{Gamma}\left(\gamma_0, \frac{1}{c_0}\right), \quad \gamma_0 \sim \text{Gamma}\left(e_0, \frac{1}{f_0}\right), \\ h_{kn}^{(L)} &\sim \text{Ber}\left(\sigma\left(b_{kn}^{(L)}\right)\right), \quad h_{kn}^{(l)} \sim \text{Ber}\left(\sigma\left(\mathbf{w}_{kl}^{(l)T} \mathbf{h}_n^{(l+1)} + b_{kl}^{(l)}\right)\right), \quad x_{pnk} \sim \text{Pois}\left(\phi_{pk} \theta_{kn} h_{kn}^{(1)}\right), \quad x_{pn} = \sum_{k=1}^K x_{pnk}, \end{aligned} \quad (15)$$

where L is the number of layers in SBN, which corresponds to Equation (15). x_{pnk} is the count of word p that comes from topic k in document n .

In this model, the perception variables $\Omega_p = \{\{\mathbf{H}^{(l)}\}, \{\mathbf{W}_l\}, \{\mathbf{b}_l\}\}$, the hinge variables $\Omega_h = \{\mathbf{X}\}$, and the task variables $\Omega_t = \{\{\phi_k\}, \{r_k\}, \Theta, \gamma_0\}$. $\mathbf{W}_l$ is the weight matrix containing columns of $\mathbf{w}_{k_l}^{(l)}$ and $\mathbf{b}_l$ is the bias vector containing entries of $b_{k_l}^{(l)}$ in Equation (15).

Learning Using Bayesian Conditional Density Filtering: Efficient learning algorithms are needed for Bayesian treatments of deep PFA. Reference [24] proposed to use an online version of MCMC called Bayesian conditional density filtering (BCDF) to learn both the global parameters $\Psi_g = (\{\phi_k\}, \{r_k\}, \gamma_0, \{\mathbf{W}_l\}, \{\mathbf{b}_l\})$ and the local variables $\Psi_l = (\Theta, \{\mathbf{H}^{(l)}\})$. The key conditional densities used for the Gibbs updates are as follows:

$$x_{pnk} | - \sim \text{Multi}(x_{pn}; \zeta_{pn1}, \dots, \zeta_{pnK}), \quad \phi_k | - \sim \text{Dir}(a_\phi + x_{1 \cdot k}, \dots, a_\phi + x_{P \cdot k}),$$

$$\theta_{kn} | - \sim \text{Gamma}(r_k h_{kn}^{(1)} + x_{\cdot nk}, p_n), \quad h_{kn}^{(1)} | - \sim \delta(x_{\cdot nk} = 0) \text{Ber}\left(\frac{\tilde{\pi}_{kn}}{\pi_{kn} + (1 - \pi_{kn})}\right) + \delta(x_{\cdot nk} > 0),$$

where $\tilde{\pi}_{kn} = \pi_{kn}(1 - p_n)^{r_k}$, $\pi_{kn} = \sigma((\mathbf{w}_k^{(1)})^T \mathbf{h}_n^{(2)} + c_k^{(1)})$, $x_{\cdot nk} = \sum_{p=1}^P x_{pnk}$, $x_{p \cdot k} = \sum_{n=1}^N x_{pnk}$, and $\zeta_{pnk} \propto \phi_{pk} \theta_{kn}$. For the learning of $h_{kn}^{(l)}$ where $l > 1$, the same techniques as in Reference [25] can be used.

Learning Using Stochastic Gradient Thermostats: An alternative way of learning deep PFA is through *stochastic gradient Nöse-Hoover thermostats* (SGNHT), which is more accurate and scalable. SGNHT is a generalization of the *stochastic gradient Langevin dynamics* (SGLD) [127] and the *stochastic gradient Hamiltonian Monte Carlo* (SGHMC) [15]. Compared with the previous two, SGNHT introduces momentum variables into the system, helping the system to jump out of local optima. Specifically, the following stochastic differential equations (SDE) can be used:

$$d\Psi_g = \mathbf{v} dt, \quad d\mathbf{v} = \tilde{f}(\Psi_g) dt - \xi \mathbf{v} dt + \sqrt{D} d\mathcal{W}, \quad d\xi = \left(\frac{1}{M} \mathbf{v}^T \mathbf{v} - 1 \right) dt,$$

where $\tilde{f}(\Psi_g) = -\nabla_{\Psi_g} \tilde{U}(\Psi_g)$ and $\tilde{U}(\Psi_g)$ is the negative log-posterior of the model. t indexes time and $\mathcal{W}$ denotes the standard Wiener process. ξ is the thermostats variable to make sure the system has a constant temperature. D is the injected variance, which is a constant. To speed up convergence, the SDE is generalized to

$$d\Psi_g = \mathbf{v} dt, \quad d\mathbf{v} = \tilde{f}(\Psi_g) dt - \Xi \mathbf{v} dt + \sqrt{D} d\mathcal{W}, \quad d\Xi = (\mathbf{q} - \mathbf{I}) dt,$$

where $\mathbf{I}$ is the identity matrix, $\Xi = \text{diag}(\xi_1, \dots, \xi_M)$, $\mathbf{q} = \text{diag}(v_1^2, \dots, v_M^2)$, and M is the dimensionality of the parameters.

SGNHT, SGLD, and SGHMC all belong to a larger class of sampling algorithms called hybrid Monte Carlo (HMC) [5]. The idea is to leverage an analogy with physical systems to guide transitions of system states. Compared to the Metropolis algorithm, HMC can make much larger changes to system states while keeping a small rejection probability. For more details, we refer readers to References [5, 81].

5.2.3 Deep Poisson Factor Analysis with Restricted Boltzmann Machine. The deep PFA model above uses SBN as a perception component. Similarly, one can replace SBN with RBM [41] (discussed in Section 4.3.1) to achieve comparable performance. With RBM as the perception component, Equation (15) becomes conditional distributions similar to Equation (4) with the following energy [41]:

$$E(\mathbf{h}_n^{(l)}, \mathbf{h}_n^{(l+1)}) = -(\mathbf{h}_n^{(l)})^T \mathbf{c}^{(l)} - (\mathbf{h}_n^{(l)})^T \mathbf{W}^{(l)} \mathbf{h}_n^{(l+1)} - (\mathbf{h}_n^{(l+1)})^T \mathbf{c}^{(l+1)}.$$

Similar learning algorithms as the deep PFA with SBN can be used. Specifically, the sampling process would alternate between $\{\{\phi_k\}, \{\gamma_k\}, \gamma_0\}$ and $\{\{\mathbf{W}^{(l)}\}, \{\mathbf{c}^{(l)}\}\}$. The former involves similar conditional density as the SBN-based DPFA. The latter is RBM's parameters and can be updated using the *contrastive divergence* algorithm.

5.2.4 Discussion. Here, we choose topic models as an example application to demonstrate how BDL can be applied in the unsupervised learning setting. In BDL-based topic models, the perception component is responsible for inferring the topic hierarchy from documents, while the task-specific component is in charge of modeling the word generation, topic generation, word-topic relation, or inter-document relation. The synergy between these two components comes from the bidirectional interaction between them. On the one hand, knowledge on the topic hierarchy facilitates accurate modeling of words and topics, providing valuable information for learning inter-document relation. On the other hand, accurately modeling the words, topics, and inter-document relation can help discover the topic hierarchy and learn compact latent factors for documents.

It is worth noting that the *information exchange* mechanism in some BDL-based topic models is different from that in Section 5.1. For example, in the SBN-based DPFA model, the exchange is natural, since the bottom layer of SBN, $\mathbf{H}^{(1)}$, and the relationship between $\mathbf{H}^{(1)}$ and $\Omega_h = \{\mathbf{X}\}$ are both inherently probabilistic, as shown in Equation (15), which means additional assumptions on the distribution are not necessary. The SBN-based DPFA model is equivalent to assuming that $\mathbf{H}$ in PFA is generated from a Dirac delta distribution (a Gaussian distribution with zero variance) centered at the bottom layer of the SBN, $\mathbf{H}^{(1)}$. Hence, both DPFA models in Table 1 are ZV models, according to the definition in Section 4.2. It is worth noting that RSDAE is an HV model (see Equation (11), where $\mathbf{S}$ is the hinge variable and the others are perception variables), and naively modifying this model to be its ZV counterpart would violate the i.i.d. requirement in Section 4.2.

Similar to Section 5.1, BDL-based topic models above use typical static Bayesian networks as task-specific components. Naturally, one can choose to use other forms of task-specific components. For example, it is straightforward to replace the relational prior of RSDAE in Section 5.2.1 with a stochastic process (e.g., a Wiener process as in Reference [113]) to model the evolution of the topic hierarchy over time.

5.3 Bayesian Deep Representation Learning for Control

In Sections 5.1 and 5.2, we covered how BDL can be applied in the supervised and unsupervised learning settings, respectively. In this section, we will discuss how BDL can help representation learning in general, using *control* as an example application.

As mentioned in Section 1, Bayesian deep learning can also be applied to the control of nonlinear dynamical systems from raw images. Consider controlling a complex dynamical system according to the live video stream received from a camera. One way of solving this control problem is by iteration between two tasks, perception from raw images and control based on dynamic models. The perception task can be taken care of using multiple layers of simple nonlinear transformation (deep learning) while the control task usually needs more sophisticated models like hidden Markov models and Kalman filters [35, 74]. To enable an effective iterative process between the perception task and the control task, we need two-way information exchange between them. The perception component would be the basis on which the control component estimates its states and, however, the control component with a dynamic model built in would be able to predict the future trajectory (images) by reversing the perception process [125].

As one of the pioneering works in this direction, Reference [125] posed this task as a representation learning problem and proposed a model called *Embed to Control* to take into account the feedback loop mentioned above during representation learning. Essentially, the goal is to learn

representations that (1) capture semantic information from raw images/videos and (2) preserve local linearity in the state space for convenient control. This is not possible without the BDL framework, since the perception component guarantees the first sub-goal while the task-specific component guarantees the second. Below, we start with some preliminaries on stochastic optimal control and then introduce the BDL-based model for representation learning.

5.3.1 Stochastic Optimal Control. Following Reference [125], we consider the stochastic optimal control of an unknown dynamical system as follows:

$$\mathbf{z}_{t+1} = f(\mathbf{z}_t, \mathbf{u}_t) + \xi, \quad \xi \sim \mathcal{N}(0, \Sigma_\xi), \quad (16)$$

where t indexes the time steps and $\mathbf{z}_t \in \mathbb{R}^{n_z}$ is the latent states. $\mathbf{u}_t \in \mathbb{R}^{n_u}$ is the applied control at time t and ξ denotes the system noise. Equivalently, the equation above can be written as $P(\mathbf{z}_{t+1}|\mathbf{z}_t, \mathbf{u}_t) = \mathcal{N}(\mathbf{z}_{t+1}|f(\mathbf{z}_t, \mathbf{u}_t), \Sigma_\xi)$. Hence, we need a mapping function to map the corresponding raw image $\mathbf{x}_t$ (observed input) into the latent space:

$$\mathbf{z}_t = m(\mathbf{x}_t) + \omega, \quad \omega \sim \mathcal{N}(0, \Sigma_\omega),$$

where ω is the corresponding system noise. Similarly, the equation above can be rewritten as $\mathbf{z}_t \sim \mathcal{N}(m(\mathbf{x}_t), \Sigma_\omega)$. If the function f is given, then finding optimal control for a trajectory of length T in a dynamical system amounts to minimizing the following cost:

$$J(\mathbf{z}_{1:T}, \mathbf{u}_{1:T}) = \mathbb{E}_z(c_T(\mathbf{z}_T, \mathbf{u}_T) + \sum_{t_0}^{T-1} c(\mathbf{z}_t, \mathbf{u}_t)), \quad (17)$$

where $c_T(\mathbf{z}_T, \mathbf{u}_T)$ is the terminal cost and $c(\mathbf{z}_t, \mathbf{u}_t)$ is the instantaneous cost. $\mathbf{z}_{1:T} = \{\mathbf{z}_1, \dots, \mathbf{z}_T\}$ and $\mathbf{u}_{1:T} = \{\mathbf{u}_1, \dots, \mathbf{u}_T\}$ are the state and action sequences, respectively. For simplicity, we can let $c_T(\mathbf{z}_T, \mathbf{u}_T) = c(\mathbf{z}_T, \mathbf{u}_T)$ and use the following quadratic cost $c(\mathbf{z}_t, \mathbf{u}_t) = (\mathbf{z}_t - \mathbf{z}_{goal})^T \mathbf{R}_z (\mathbf{z}_t - \mathbf{z}_{goal}) + \mathbf{u}_t^T \mathbf{R}_u \mathbf{u}_t$, where $\mathbf{R}_z \in \mathbb{R}^{n_z \times n_z}$ and $\mathbf{R}_u \in \mathbb{R}^{n_u \times n_u}$ are the weighting matrices. $\mathbf{z}_{goal}$ is the target latent state that should be inferred from the raw images (observed input). Given the function f , $\bar{\mathbf{z}}_{1:T}$ (current estimates of the optimal trajectory), and $\bar{\mathbf{u}}_{1:T}$ (the corresponding controls), the dynamical system can be linearized as

$$\mathbf{z}_{t+1} = \mathbf{A}(\bar{\mathbf{z}}_t)\mathbf{z}_t + \mathbf{B}(\bar{\mathbf{z}}_t)\mathbf{u}_t + \mathbf{o}(\bar{\mathbf{z}}_t) + \omega, \quad \omega \sim \mathcal{N}(0, \Sigma_\omega), \quad (18)$$

where $\mathbf{A}(\bar{\mathbf{z}}_t) = \frac{\partial f(\bar{\mathbf{z}}_t, \bar{\mathbf{u}}_t)}{\partial \bar{\mathbf{z}}_t}$ and $\mathbf{B}(\bar{\mathbf{z}}_t) = \frac{\partial f(\bar{\mathbf{z}}_t, \bar{\mathbf{u}}_t)}{\partial \bar{\mathbf{u}}_t}$ are local Jacobians. $\mathbf{o}(\bar{\mathbf{z}}_t)$ is the offset.

5.3.2 BDL-based Representation Learning for Control. To minimize the function in Equation (17), we need three key components: an encoding model to encode $\mathbf{x}_t$ into $\mathbf{z}_t$, a transition model to infer $\mathbf{z}_{t+1}$ given $(\mathbf{z}_t, \mathbf{u}_t)$, and a reconstruction model to reconstruct $\mathbf{x}_{t+1}$ from the inferred $\mathbf{z}_{t+1}$.

Encoding Model: An encoding model $Q_\phi(Z|X) = \mathcal{N}(\boldsymbol{\mu}_t, \text{diag}(\sigma_t^2))$, where the mean $\boldsymbol{\mu}_t \in \mathbb{R}^{n_z}$ and the diagonal covariance $\Sigma_t = \text{diag}(\sigma_t^2) \in \mathbb{R}^{n_z \times n_z}$, encodes the raw images $\mathbf{x}_t$ into latent states $\mathbf{z}_t$. Here,

$$\boldsymbol{\mu}_t = \mathbf{W}_\mu h_\phi^{\text{enc}}(\mathbf{x}_t) + \mathbf{b}_\mu, \quad \log \sigma_t = \mathbf{W}_\sigma h_\phi^{\text{enc}}(\mathbf{x}_t) + \mathbf{b}_\sigma, \quad (19)$$

where $h_\phi(\mathbf{x}_t)^{\text{enc}}$ is the output of the encoding network with $\mathbf{x}_t$ as its input.

Transition Model: A transition model like Equation (18) infers $\mathbf{z}_{t+1}$ from $(\mathbf{z}_t, \mathbf{u}_t)$. If we use $\tilde{Q}_\psi(\tilde{Z}|Z, \mathbf{u})$ to denote the approximate posterior distribution to generate $\mathbf{z}_{t+1}$, then the generative process of the full model would be

$$\mathbf{z}_t \sim Q_\phi(Z|X) = \mathcal{N}(\boldsymbol{\mu}_t, \Sigma_t), \quad \tilde{\mathbf{z}}_{t+1} \sim \tilde{Q}_\psi(\tilde{Z}|Z, \mathbf{u}) = \mathcal{N}(\mathbf{A}_t \boldsymbol{\mu}_t + \mathbf{B}_t \mathbf{u}_t + \mathbf{o}_t, \mathbf{C}_t), \quad \tilde{\mathbf{x}}_t, \tilde{\mathbf{x}}_{t+1} \sim P_\theta(X|Z) = \text{Bern}(\mathbf{p}_t), \quad (20)$$

where the last equation is the reconstruction model to be discussed later, $C_t = A_t \Sigma_t A_t^T + H_t$, and H_t is the covariance matrix of the estimated system noise ($\omega_t \sim \mathcal{N}(0, H_t)$). The key here is to learn A_t , B_t and $\mathbf{o}_t$, which are parameterized as follows:

$$\text{vec}(A_t) = W_A h_\psi^{\text{trans}}(z_t) + \mathbf{b}_A, \quad \text{vec}(B_t) = W_B h_\psi^{\text{trans}}(z_t) + \mathbf{b}_B, \quad \mathbf{o}_t = W_o h_\psi^{\text{trans}}(z_t) + \mathbf{b}_o,$$

where $h_\psi^{\text{trans}}(z_t)$ is the output of the transition network.

Reconstruction Model: As mentioned in the last part of Equation (20), the posterior distribution $P_\theta(X|Z)$ reconstructs the raw images $\mathbf{x}_t$ from the latent states z_t . The parameters for the Bernoulli distribution $\mathbf{p}_t = W_p h_\theta^{\text{dec}}(z_t) + \mathbf{b}_p$ where $h_\theta^{\text{dec}}(z_t)$ is the output of a third network, called the decoding network or the reconstruction network. Putting it all together, Equation (20) shows the generative process of the full model.

5.3.3 Learning Using Stochastic Gradient Variational Bayes. With $\mathcal{D} = \{(\mathbf{x}_1, \mathbf{u}_1, \mathbf{x}_2), \dots, (\mathbf{x}_{T-1}, \mathbf{u}_{T-1}, \mathbf{x}_T)\}$ as the training set, the loss function is as follows:

$$\mathcal{L} = \sum_{(\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1}) \in \mathcal{D}} \mathcal{L}^{\text{bound}}(\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1}) + \lambda \text{KL}(\tilde{Q}_\psi(\tilde{Z}|\mu_t, \mathbf{u}_t) \| Q_\phi(Z|\mathbf{x}_{t+1})),$$

where the first term is the variational bound on the marginalized log-likelihood for each data point:

$$\mathcal{L}^{\text{bound}}(\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1}) = \mathbb{E}_{\substack{z_t \sim Q_\phi \\ \tilde{z}_{t+1} \sim \tilde{Q}_\psi}} (-\log P_\theta(\mathbf{x}_t|z_t) - \log P_\theta(\mathbf{x}_{t+1}|\tilde{z}_{t+1})) + \text{KL}(Q_\phi \| P(Z)),$$

where $P(Z)$ is the prior distribution for Z . With the equations above, stochastic gradient variational Bayes can be used to learn the parameters.

According to the generative process in Equation (20) and the definition in Section 4.2, the perception variables $\Omega_p = \{h_\phi^{\text{enc}}(\cdot), W_p^+, \mathbf{x}_t, \mu_t, \sigma_t, \mathbf{p}_t, h_\theta^{\text{dec}}(\cdot)\}$, where W_p^+ is shorthand for $\{W_\mu, \mathbf{b}_\mu, W_\sigma, \mathbf{b}_\sigma, W_p, \mathbf{b}_p\}$. The hinge variables $\Omega_h = \{z_t, z_{t+1}\}$ and the task variables $\Omega_t = \{A_t, B_t, \mathbf{o}_t, \mathbf{u}_t, C_t, \omega_t, W_t^+, h_\psi^{\text{trans}}(\cdot)\}$, where W_t^+ is shorthand for $\{W_A, \mathbf{b}_A, W_B, \mathbf{b}_B, W_o, \mathbf{b}_o\}$.

5.3.4 Discussion. The example model above demonstrates BDL's capability of learning representations that satisfy domain-specific requirements. In the case of *control*, we are interested in learning representations that can capture semantic information from raw input and preserve local linearity in the space of system states.

To achieve this goal, the BDL-based model consists of two components, a perception component to *see* the live video and a control (task-specific) component to *infer* the states of the dynamical system. Inference of the system is based on the mapped states and the confidence of mapping from the perception component, and in turn, the control signals sent by the control component would affect the live video received by the perception component. Only when the two components work interactively within a unified probabilistic framework can the model reach its full potential and achieve the best control performance.

Note that the BDL-based control model discussed above uses a different *information exchange* mechanism from that in Sections 5.1 and 5.2: it follows the VAE mechanism and uses neural networks to *separately* parameterize the mean and covariance of hinge variables (e.g., in the encoding model, the hinge variable $z_t \sim \mathcal{N}(\mu_t, \text{diag}(\sigma_t^2))$, where μ_t and σ_t are perception variables parameterized as in Equation (19)), which is more flexible (with more free parameters) than models like CDL and CDR in Section 5.1, where Gaussian distributions with fixed variance are also used. Note that this BDL-based control model is an LV model as shown in Table 1, and since the covariance is assumed to be diagonal, the model still meets the independence requirement in Section 4.2.

5.4 Bayesian Deep Learning for Other Applications

BDL has found wide applications such as recommender systems, topic models, and control in supervised learning, unsupervised learning, and representation learning in general. In this section, we briefly discuss a few more applications that could benefit from BDL.

5.4.1 Link Prediction. Link prediction has long been a core problem in network analysis and is recently attracting more interest with the new advancements brought by BDL and deep neural networks in general. Reference [120] proposed the first BDL-based model, dubbed relational deep learning (RDL), for link prediction. Graphite [32] extends RDL using a perception component based on graph convolutional networks (GCN) [59]. Reference [75] combines the classic stochastic blockmodel [82] (as a task-specific component) and GCN-based perception component to jointly model latent community structures and link generation in a graph with reported state-of-the-art performance in link prediction.

5.4.2 Natural Language Processing. Besides topic modeling as discussed in Section 5.2, BDL is also useful for natural language processing in general. For example, References [77] and [69] build on top of the BDL principles to define a language revision process. These models typically involve RNN-based perception components and relatively simple task-specific components linking the input and output sequences.

5.4.3 Computer Vision. BDL is particularly powerful for computer vision in the unsupervised learning setting. This is because in the BDL framework, one can clearly define a generative process of how objects in a scene are generated from various factors such as counts, positions, and the content [20]. The perception component, usually taking the form of a probabilistic neural network, can focus on modeling the raw images' visual features, while the task-specific component handles the conditional dependencies among objects' various attributes in the images. One notable work in this direction is *Attend, Infer, Repeat* (AIR) [20], where the task-specific component involves latent variables on each object' position, scale, appearance, and presence (which is related to counting of objects). Following AIR, variants such as Fast AIR [105] and Sequential AIR [60] are proposed to improve its computational efficiency and performance. Besides unsupervised learning, BDL can also be useful for supervised learning tasks such as action recognition in videos [102], where conditional dependencies among different actions are modeled using a task-specific component.

5.4.4 Speech. In the field of speech recognition and synthesis, researchers have also been adopting the BDL framework to improve both accuracy and interpretability. For example, factorized hierarchical VAE [47, 48] composes VAE with a factorized latent variable model (represented as a PGM) to learn different latent factors in speech data following an unsupervised setting. Similarly, Gaussian mixture VAE [49] uses a Gaussian mixture model as the task-specific component to achieve controllable speech synthesis from text. In terms of speech recognition, recurrent Poisson process units (RPPU) [51] instead adopt a different form of task-specific component; they use a stochastic process (i.e., a Poisson process) as the task-specific component to model boundaries between phonemes and successfully achieve a significantly lower word error rate (WER) for speech recognition. Similarly, deep graph random process (DGP) [52] as another stochastic process operates on graphs to model the relational structure among utterances, further improving performance in speech recognition.

5.4.5 Time Series Forecasting. Time series forecasting is a long-standing core problem in economics, statistics, and machine learning [36]. It has wide applications across multiple areas. For example, accurate forecasts of regional energy consumption can provide valuable guidance to

optimize energy generation and allocation. In e-commerce, retailers rely on demand forecasts to decide when and where to replenish their supplies, thereby avoiding items going out of stock and guaranteeing fastest deliveries for customers. During a pandemic such as COVID-19, it is crucial to obtain reasonable forecasts on hospital workload and medical supply demand to best allocate resources across the country. Needless to say, an ideal forecasting model requires both efficient processing of high-dimensional data and sophisticated modeling of different random variables, either observed or latent. BDL-based forecasting models [21, 27, 90, 124] achieve these with an RNN-based perception component and a task-specific component handling the conditional dependencies among different variables, showing substantial improvement over previous non-BDL forecasting models.

5.4.6 Health Care. In health-care-related applications [91], it is often desirable to incorporate human knowledge into models, either to boost performance or more importantly to improve interpretability. It is also crucial to ensure models' robustness when they are used on under-represented data. BDL therefore provides a unified framework to meet all these requirements: (1) with its Bayesian nature, it can impose proper priors and perform Bayesian model averaging to improve robustness; (2) its task-specific component can naturally represent and incorporate human knowledge if necessary; (3) the model's joint training provides interpretability for its both components. For example, Reference [38] proposed a deep Poisson factor model, which essentially stacks layers of Poisson factor models, to analyze electronic health records. Reference [110] built a BDL model with experiment-specific priors (knowledge) to control the false discovery rate during study analysis with applications to cancer drug screening. Reference [61] developed deep nonlinear state space models and demonstrated their effectiveness in processing electronic health records and performing counterfactual reasoning. Task-specific components in the BDL models above all take the form of a typical Bayesian network (as mentioned in Section 4.4.1). In contrast, Reference [117] proposed to use bidirectional inference networks, which are essentially a class of deep Bayesian network, as the task-specific component (as mentioned in Section 4.4.2). This enables deep nonlinear structures in each conditional distribution of the Bayesian network and improves performance for applications such as health profiling.

6 CONCLUSIONS AND FUTURE RESEARCH

BDL strives to combine the merits of PGM and NN by organically integrating them in a single principled probabilistic framework. In this survey, we identified such a current trend and reviewed recent work. A BDL model consists of a perception component and a task-specific component; we therefore surveyed different instantiations of both components developed over the past few years, respectively, and discussed different variants in detail. To learn parameters in BDL, several types of algorithms have been proposed, ranging from block coordinate descent, Bayesian conditional density filtering, and stochastic gradient thermostats to stochastic gradient variational Bayes.

BDL draws inspiration and gains popularity both from the success of PGM and from recent promising advances on deep learning. Since many real-world tasks involve both efficient perception from high-dimensional signals (e.g., images and videos) and probabilistic inference on random variables, BDL emerges as a natural choice to harness the perception ability from NN and the (conditional and causal) inference ability from PGM. Over the past few years, BDL has found successful applications in various areas such as recommender systems, topic models, stochastic optimal control, computer vision, natural language processing, health care, and so on. In the future, we can expect both more in-depth studies on existing applications and exploration on even more complex tasks. Besides, recent progress on efficient BNN (as the perception component of BDL) also lays down the foundation for further improving BDL's scalability.

REFERENCES

- [1] Gediminas Adomavicius and YoungOk Kwon. 2012. Improving aggregate recommendation diversity using ranking-based techniques. *Trans. Knowl. Data Eng.* 24, 5 (2012), 896–911.
- [2] Anoop Korattikara Balan, Vivek Rathod, Kevin P. Murphy, and Max Welling. 2015. Bayesian dark knowledge. In *Proceedings of the NIPS*. 3420–3428.
- [3] Ilaria Bartolini, Zhenjie Zhang, and Dimitris Papadias. 2011. Collaborative filtering with personalized skylines. *Trans. Knowl. Data Eng.* 23, 2 (2011), 190–203.
- [4] Yoshua Bengio, Li Yao, Guillaume Alain, and Pascal Vincent. 2013. Generalized denoising auto-encoders as generative models. In *Proceedings of the NIPS*. 899–907.
- [5] Christopher M. Bishop. 2006. *Pattern Recognition and Machine Learning*. Springer-Verlag, New York.
- [6] David Blei and John Lafferty. 2006. Correlated topic models. In *Proceedings of the NIPS*. 147–154.
- [7] David M. Blei and John D. Lafferty. 2006. Dynamic topic models. In *Proceedings of the ICML*. ACM, 113–120.
- [8] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet allocation. *J. Mach. Learn. Res.* 3 (2003), 993–1022.
- [9] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. Weight uncertainty in neural network. In *Proceedings of the ICML*. 1613–1622.
- [10] Hervé Bourlard and Yves Kamp. 1988. Auto-association by multilayer perceptrons and singular value decomposition. *Biol. Cybernet.* 59, 4–5 (1988), 291–294.
- [11] Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov. 2016. Importance weighted autoencoders. In *Proceedings of the ICLR*.
- [12] Yi Cai, Ho-fung Leung, Qing Li, Huaqing Min, Jie Tang, and Juanzi Li. 2014. Typicality-based collaborative filtering recommendation. *Trans. Knowl. Data Eng.* 26, 3 (2014), 766–779.
- [13] Minmin Chen, Kilian Q. Weinberger, Fei Sha, and Yoshua Bengio. 2014. Marginalized denoising auto-encoders for nonlinear representations. In *Proceedings of the ICML*. 1476–1484.
- [14] Minmin Chen, Zhixiang Eddie Xu, Kilian Q. Weinberger, and Fei Sha. 2012. Marginalized denoising autoencoders for domain adaptation. In *Proceedings of the ICML*.
- [15] Tianqi Chen, Emily B. Fox, and Carlos Guestrin. 2014. Stochastic Gradient Hamiltonian Monte Carlo. In *Proceedings of the ICML*. 1683–1691.
- [16] Kyunghyun Cho, Bart van Merriënboer, eCaglar Gülccehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of the EMNLP*. 1724–1734.
- [17] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C. Courville, and Yoshua Bengio. 2015. A recurrent latent variable model for sequential data. In *Proceedings of the NIPS*. 2980–2988.
- [18] Yulai Cong, Bo Chen, Hongwei Liu, and Mingyuan Zhou. 2017. Deep latent dirichlet allocation with topic-layer-adaptive stochastic gradient riemannian MCMC. In *Proceedings of the ICML*. 864–873.
- [19] Andreas Doerr, Christian Daniel, Martin Schiegg, Duy Nguyen-Tuong, Stefan Schaal, Marc Toussaint, and Sebastian Trimpe. 2018. Probabilistic recurrent state-space models. In *Proceedings of the ICML*. 1279–1288.
- [20] S. M. Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, Koray Kavukcuoglu, and Geoffrey E. Hinton. 2016. Attend, infer, repeat: Fast scene understanding with generative models. In *Proceedings of the NIPS*. 3225–3233.
- [21] Valentin Flunkert, David Salinas, and Jan Gasthaus. 2017. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *CoRR* abs/1704.04110.
- [22] Yarin Gal and Zoubin Ghahramani. 2015. Dropout as a Bayesian approximation: Insights and applications. In *Proceedings of the ICML*.
- [23] M. J. F. Gales and S. S. Airey. 2006. Product of Gaussians for speech recognition. *Comput. Speech & Lang.* 20, 1 (2006), 22–40.
- [24] Zhe Gan, Changyou Chen, Ricardo Henao, David E. Carlson, and Lawrence Carin. 2015. Scalable deep Poisson factor analysis for topic modeling. In *Proceedings of the ICML*. 1823–1832.
- [25] Zhe Gan, Ricardo Henao, David E. Carlson, and Lawrence Carin. 2015. Learning deep sigmoid belief networks with data augmentation. In *Proceedings of the AISTATS*.
- [26] Jochen Gast and Stefan Roth. 2018. Lightweight probabilistic deep networks. In *Proceedings of the CVPR*. 3369–3378.
- [27] Jan Gasthaus, Konstantinos Benidis, Yuyang Wang, Syama Sundar Rangapuram, David Salinas, Valentin Flunkert, and Tim Januschowski. 2019. Probabilistic forecasting with spline quantile function RNNs. In *Proceedings of the AISTATS*. 1901–1910.
- [28] Kostadin Georgiev and Preslav Nakov. 2013. A non-IID framework for collaborative filtering with restricted Boltzmann machines. In *Proceedings of the ICML*. 1148–1156.

- [29] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. Deep Learning. Retrieved from <http://goodfeli.github.io/dlbook/>. Book in preparation for MIT Press.
- [30] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Proceedings of the NIPS*. 2672–2680.
- [31] Alex Graves. 2011. Practical variational inference for neural networks. In *Proceedings of the NIPS*. 2348–2356.
- [32] Aditya Grover, Aaron Zweig, and Stefano Ermon. 2019. Graphite: Iterative generative modeling of graphs. In *Proceedings of the ICML*. 2434–2444.
- [33] A. K. Gupta and D. K. Nagar. 2000. *Matrix Variate Distributions*. Chapman & Hall. 99040291 Retrieved from <http://books.google.com.hk/books?id=PQOYnT7P1loC>.
- [34] Danijar Hafner, Timothy P. Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. 2019. Learning latent dynamics for planning from pixels. In *Proceedings of the ICML*. 2555–2565.
- [35] Jeff Harrison and Mike West. 1999. *Bayesian Forecasting & Dynamic Models*. Springer.
- [36] Andrew C. Harvey. 1990. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- [37] Hao He, Hao Wang, Guang-He Lee, and Yonglong Tian. 2019. ProbGAN: Towards probabilistic GAN with theoretical guarantees. In *Proceedings of the ICLR*.
- [38] Ricardo Henao, James Lu, Joseph E. Lucas, Jeffrey M. Ferranti, and Lawrence Carin. 2016. Electronic health record analysis via deep Poisson factor models. *J. Mach. Learn. Res.* 17 (2016), 186:1–186:32.
- [39] José Miguel Hernández-Lobato and Ryan Adams. 2015. Probabilistic backpropagation for scalable learning of Bayesian neural networks. In *Proceedings of the ICML*. 1861–1869.
- [40] Geoffrey E. Hinton. 2002. Training products of experts by minimizing contrastive divergence. *Neural Comput.* 14, 8 (2002), 1771–1800.
- [41] Geoffrey E. Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural Comput.* 18, 7 (2006), 1527–1554.
- [42] Geoffrey E. Hinton and Drew Van Camp. 1993. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the COLT*. 5–13.
- [43] Geoffrey E. Hinton and Richard S. Zemel. 1994. Autoencoders, minimum description length, and Helmholtz free energy. In *Proceedings of the NIPS*. 3–10.
- [44] Matthew Hoffman, Francis R. Bach, and David M. Blei. 2010. Online learning for latent Dirichlet allocation. In *Proceedings of the NIPS*. 856–864.
- [45] Matthew D. Hoffman, David M. Blei, Chong Wang, and John William Paisley. 2013. Stochastic variational inference. *J. Mach. Learn. Res.* 14, 1 (2013), 1303–1347.
- [46] Mark F. Hornick and Pablo Tamayo. 2012. Extending recommender systems for disjoint user/item sets: The conference recommendation problem. *Trans. Knowl. Data Eng.* 24, 8 (2012), 1478–1490.
- [47] Wei-Ning Hsu and James R. Glass. 2018. Scalable factorized hierarchical variational autoencoder training. In *Proceedings of the INTERSPEECH*. 1462–1466.
- [48] Wei-Ning Hsu, Yu Zhang, and James R. Glass. 2017. Unsupervised learning of disentangled and interpretable representations from sequential data. In *Proceedings of the NIPS*. 1878–1889.
- [49] Wei-Ning Hsu, Yu Zhang, Ron J. Weiss, Heiga Zen, Yonghui Wu, Yuxuan Wang, Yuan Cao, Ye Jia, Zhifeng Chen, Jonathan Shen, Patrick Nguyen, and Ruoming Pang. 2019. Hierarchical generative modeling for controllable speech synthesis. In *Proceedings of the ICLR*.
- [50] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative filtering for implicit feedback datasets. In *Proceedings of the ICDM*. 263–272.
- [51] Hengguan Huang, Hao Wang, and Brian Mak. 2019. Recurrent Poisson process unit for speech recognition. In *Proceedings of the AAAI*. 6538–6545.
- [52] Hengguan Huang, Fuzhao Xue, Hao Wang, and Ye Wang. 2020. Deep graph random process for relational-thinking-based speech recognition. In *Proceedings of the ICML*.
- [53] David H. Hubel and Torsten N. Wiesel. 1968. Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* 195, 1 (1968), 215–243.
- [54] Finn V. Jensen et al. 1996. *An Introduction to Bayesian Networks*, Vol. 210. UCL Press London.
- [55] Michael I. Jordan, Zoubin Ghahramani, Tommi Jaakkola, and Lawrence K. Saul. 1999. An introduction to variational methods for graphical models. *Mach. Learn.* 37, 2 (1999), 183–233.
- [56] Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the ACL*. 655–665.
- [57] Maximilian Karl, Maximilian Sölk, Justin Bayer, and Patrick van der Smagt. 2017. Deep variational Bayes filters: Unsupervised learning of state space models from raw data. In *Proceedings of the ICLR*.
- [58] Diederik P. Kingma and Max Welling. 2013. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.

- [59] Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *Proceedings of the ICLR*.
- [60] Adam R. Kosiorek, Hyunjik Kim, Yee Whye Teh, and Ingmar Posner. 2018. Sequential attend, infer, repeat: Generative modelling of moving objects. In *Proceedings of the NIPS*. 8615–8625.
- [61] Rahul G. Krishnan, Uri Shalit, and David A. Sontag. 2017. Structured inference networks for nonlinear state space models. In *Proceedings of the AAAI*. 2101–2109.
- [62] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Proceedings of the NIPS*. 1097–1105.
- [63] Y. LeCun. 1987. *Modeles connexionnistes de l'apprentissage (connectionist learning models)*. Ph.D. Dissertation. Université P. et M. Curie (Paris 6).
- [64] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [65] Honglak Lee, Roger Grosse, Rajesh Ranganath, and Andrew Y. Ng. 2009. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *Proceedings of the ICML*. 609–616.
- [66] Sheng Li, Jaya Kawale, and Yun Fu. 2015. Deep collaborative filtering via marginalized Denoising Auto-encoder. In *Proceedings of the CIKM*. 811–820.
- [67] Wu-Jun Li and Dit-Yan Yeung. 2009. Relation regularized matrix factorization. In *Proceedings of the IJCAI*.
- [68] Xiaopeng Li and James She. 2017. Collaborative variational autoencoder for recommender systems. In *Proceedings of the KDD*. 305–314.
- [69] Yi Liao, Lidong Bing, Piji Li, Shuming Shi, Wai Lam, and Tong Zhang. 2018. QuaSE: Sequence editing under quantifiable guidance. In *Proceedings of the EMNLP*. 3855–3864.
- [70] Nathan Nan Liu, Xiangrui Meng, Chao Liu, and Qiang Yang. 2011. Wisdom of the better few: Cold start recommendation via representative based rating elicitation. In *Proceedings of the RecSys*. 37–44.
- [71] Zhongqi Lu, Zhicheng Dou, Jianxun Lian, Xing Xie, and Qiang Yang. 2015. Content-based collaborative filtering for news topic recommendation. In *Proceedings of the AAAI*. 217–223.
- [72] J. C. MacKay David. 1992. A practical Bayesian framework for backprop networks. *Neural Comput.* 4, 3 (1992), 448–472.
- [73] Wesley J. Maddox, Pavel Izmailov, Timur Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson. 2019. A simple baseline for Bayesian uncertainty in deep learning. In *Proceedings of the NIPS*. 13132–13143.
- [74] Takamitsu Matsubara, Vicenc Gómez, and Hilbert J. Kappen. 2014. Latent Kullback Leibler control for continuous-state systems using probabilistic graphical models. In *Proceedings of the UAI*. 583–592.
- [75] Nikhil Mehta, Lawrence Carin, and Piyush Rai. 2019. Stochastic blockmodels meet graph neural networks. In *Proceedings of the ICML*. 4466–4474.
- [76] Abdel-rahman Mohamed, Tara N. Sainath, George E. Dahl, Bhuvana Ramabhadran, Geoffrey E. Hinton, and Michael A. Picheny. 2011. Deep belief networks using discriminative features for phone recognition. In *Proceedings of the ICASSP*. 5060–5063.
- [77] Jonas Mueller, David K. Gifford, and Tommi S. Jaakkola. 2017. Sequence to better sequence: Continuous revision of combinatorial structures. In *Proceedings of the ICML*. 2536–2544.
- [78] Kevin P. Murphy. 2012. *Machine Learning: A Probabilistic Perspective*. MIT Press.
- [79] Radford M. Neal. 1992. Connectionist learning of belief networks. *Artif. Intell.* 56, 1 (1992), 71–113.
- [80] Radford M. Neal. 1995. *Bayesian learning for neural networks*. Ph.D. Dissertation. University of Toronto.
- [81] Radford M. Neal et al. 2011. MCMC using Hamiltonian dynamics. *Handbook Markov Chain Monte Carlo* 2, 11 (2011), 2.
- [82] Krzysztof Nowicki and Tom A. B. Snijders. 2001. Estimation and prediction for stochastic blockstructures. *J. Amer. Stat. Assoc.* 96, 455 (2001), 1077–1087.
- [83] Takahiro Omi, Naonori Ueda, and Kazuyuki Aihara. 2019. Fully neural network based model for general temporal point processes. In *Proceedings of the NIPS*. 2120–2129.
- [84] Aäron Van Den Oord, Sander Dieleman, and Benjamin Schrauwen. 2013. Deep content-based music recommendation. In *Proceedings of the NIPS*. 2643–2651.
- [85] Yoon-Joo Park. 2013. The adaptive clustering method for the long tail problem of recommender systems. *Trans. Knowl. Data Eng.* 25, 8 (2013), 1904–1915.
- [86] Ian Porteous, David Newman, Alexander Ihler, Arthur Asuncion, Padhraic Smyth, and Max Welling. 2008. Fast collapsed gibbs sampling for latent Dirichlet allocation. In *Proceedings of the KDD*. 569–577.
- [87] Janis Postels, Francesco Ferroni, Huseyin Coskun, Nassir Navab, and Federico Tombari. 2019. Sampling-free epistemic uncertainty estimation using approximated variance propagation. In *Proceedings of the ICCV*. 2931–2940.
- [88] Christopher Poultney, Sumit Chopra, Yann L. Cun et al. 2006. Efficient learning of sparse representations with an energy-based model. In *Proceedings of the NIPS*. 1137–1144.

- [89] Sanjay Purushotham, Yan Liu, and C.-C. Jay Kuo. 2012. Collaborative topic regression with social matrix factorization for recommendation systems. In *Proceedings of the ICML*. 759–766.
- [90] Syama Sundar Rangapuram, Matthias W. Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. 2018. Deep state space models for time series forecasting. In *Proceedings of the NIPS*. 7796–7805.
- [91] Daniele Ravi, Charence Wong, Fani Deligianni, Melissa Berthelot, Javier Andreu-Perez, Benny Lo, and Guang-Zhong Yang. 2016. Deep learning for health informatics. *IEEE J. Biomed. Health Info.* 21, 1 (2016), 4–21.
- [92] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. *Introduction to Recommender Systems Handbook*. Springer.
- [93] Salah Rifai, Pascal Vincent, Xavier Muller, Xavier Glorot, and Yoshua Bengio. 2011. Contractive auto-encoders: Explicit invariance during feature extraction. In *Proceedings of the ICML*. 833–840.
- [94] Sheldon M. Ross, John J. Kelly, Roger J. Sullivan, William James Perry, Donald Mercer, Ruth M. Davis, Thomas Dell Washburn, Earl V. Sager, Joseph B. Boyce, and Vincent L. Bristow. 1996. *Stochastic Processes*. Vol. 2. Wiley, New York.
- [95] Tara N. Sainath, Brian Kingsbury, Vikas Sindhwani, Ebru Arisoy, and Bhuvana Ramabhadran. 2013. Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *Proceedings of the ICASSP*. 6655–6659.
- [96] Ruslan Salakhutdinov and Andriy Mnih. 2007. Probabilistic matrix factorization. In *Proceedings of the NIPS*. 1257–1264.
- [97] Ruslan Salakhutdinov and Andriy Mnih. 2008. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In *Proceedings of the ICML*. 880–887.
- [98] Ruslan Salakhutdinov, Andriy Mnih, and Geoffrey E. Hinton. 2007. Restricted Boltzmann machines for collaborative filtering. In *Proceedings of the ICML*. 791–798.
- [99] Oleksandr Shchur, Marin Bilos, and Stephan Günnemann. 2020. Intensity-free learning of temporal point processes. In *Proceedings of the ICLR*.
- [100] Alexander Shekhovtsov and Boris Flach. 2019. Feed-forward propagation in probabilistic neural networks with categorical and max layers. In *Proceedings of the ICLR*.
- [101] Jonathan R. Shewchuk. 1994. *An Introduction to the Conjugate Gradient Method Without the Agonizing Pain*. Technical Report. Carnegie Mellon University, Pittsburgh, PA.
- [102] Gunnar A. Sigurdsson, Santosh Kumar Divvala, Ali Farhadi, and Abhinav Gupta. 2017. Asynchronous temporal fields for action recognition. In *Proceedings of the CVPR*. 5650–5659.
- [103] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15, 1 (2014), 1929–1958.
- [104] Nitish Srivastava and Russ R. Salakhutdinov. 2012. Multimodal learning with deep Boltzmann machines. In *Proceedings of the NIPS*. 2222–2230.
- [105] Karl Stelzner, Robert Peharz, and Kristian Kersting. 2019. Faster attend-infer-repeat with tractable probabilistic models. In *Proceedings of the ICML*. 5966–5975.
- [106] Robert S. Strichartz. 2003. *A Guide to Distribution Theory and Fourier Transforms*. World Scientific.
- [107] Jiahao Su, Milan Cvitkovic, and Furong Huang. 2020. Sampling-free learning of Bayesian quantized neural networks. In *Proceedings of the ICLR*.
- [108] Ilya Sutskever, Oriol Vinyals, and Quoc V. V. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of the NIPS*. 3104–3112.
- [109] Jinhui Tang, Guo-Jun Qi, Liyan Zhang, and Changsheng Xu. 2013. Cross-space affinity learning with its application to movie recommendation. *IEEE Trans. Knowl. Data Eng.* 25, 7 (2013), 1510–1519.
- [110] Wesley Tansey, Yixin Wang, David M. Blei, and Raul Rabadan. 2018. Black box FDR. In *Proceedings of the ICML*. 4874–4883.
- [111] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *J. Mach. Learn. Res.* 11 (2010), 3371–3408.
- [112] Chong Wang and David M. Blei. 2011. Collaborative topic modeling for recommending scientific articles. In *Proceedings of the KDD*. 448–456.
- [113] Chong Wang, David M. Blei, and David Heckerman. 2008. Continuous time dynamic topic models. In *Proceedings of the UAI*. 579–586.
- [114] Hao Wang. 2017. *Bayesian Deep Learning for Integrated Intelligence: Bridging the Gap between Perception and Inference*. Ph.D. Dissertation. Hong Kong University of Science and Technology.
- [115] Hao Wang, Binyi Chen, and Wu-Jun Li. 2013. Collaborative topic regression with social regularization for tag recommendation. In *Proceedings of the IJCAI*. 2719–2725.
- [116] Hao Wang and Wu-Jun Li. 2015. Relational collaborative topic regression for recommender systems. *Trans. Knowl. Data Eng.* 27, 5 (2015), 1343–1355.

- [117] Hao Wang, Chengzhi Mao, Hao He, Mingmin Zhao, Tommi S. Jaakkola, and Dina Katabi. 2019. Bidirectional inference networks: A class of deep Bayesian networks for health profiling. In *Proceedings of the AAAI*. 766–773.
- [118] Hao Wang, Xingjian Shi, and Dit-Yan Yeung. 2015. Relational stacked denoising autoencoder for tag recommendation. In *Proceedings of the AAAI*. 3052–3058.
- [119] Hao Wang, Xingjian Shi, and Dit-Yan Yeung. 2016. Natural-parameter networks: A class of probabilistic neural networks. In *Proceedings of the NIPS*. 118–126.
- [120] Hao Wang, Xingjian Shi, and Dit-Yan Yeung. 2017. Relational deep learning: A deep latent variable model for link prediction. In *Proceedings of the AAAI*. 2688–2694.
- [121] Hao Wang, Naiyan Wang, and Dit-Yan Yeung. 2015. Collaborative deep learning for recommender systems. In *Proceedings of the KDD*. 1235–1244.
- [122] Hao Wang, Shi Xingjian, and Dit-Yan Yeung. 2016. Collaborative recurrent autoencoder: Recommend while learning to fill in the blanks. In *Proceedings of the NIPS*. 415–423.
- [123] Xinxin Wang and Ye Wang. 2014. Improving content-based and hybrid music recommendation using deep learning. In *Proceedings of the ACM MM*. 627–636.
- [124] Yuyang Wang, Alex Smola, Danielle C. Maddix, Jan Gasthaus, Dean Foster, and Tim Januschowski. 2019. Deep factors for forecasting. In *Proceedings of the ICML*. 6607–6617.
- [125] Manuel Watter, Jost Springenberg, Joschka Boedecker, and Martin Riedmiller. 2015. Embed to control: A locally linear latent dynamics model for control from raw images. In *Proceedings of the NIPS*. 2728–2736.
- [126] Yan Zheng Wei, Luc Moreau, and Nicholas R. Jennings. 2005. Learning users’ interests by quality classification in market-based recommender systems. *Trans. Knowl. Data Eng.* 17, 12 (2005), 1678–1688.
- [127] Max Welling and Yee Whye Teh. 2011. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the ICML*. 681–688.
- [128] Andrew Gordon Wilson. 2020. The case for Bayesian deep learning. *arXiv preprint arXiv:2001.10995*.
- [129] Zichao Yang, Zhiting Hu, Ruslan Salakhutdinov, and Taylor Berg-Kirkpatrick. 2017. Improved variational autoencoders for text modeling using dilated convolutions. In *Proceedings of the ICML*. 3881–3890.
- [130] Ghim-Eng Yap, Ah-Hwee Tan, and HweeHwa Pang. 2007. Discovering and exploiting causal dependencies for robust mobile context-aware recommenders. *Trans. Knowl. Data Eng.* 19, 7 (2007), 977–992.
- [131] Haochao Ying, Liang Chen, Yuwen Xiong, and Jian Wu. 2016. Collaborative deep ranking: A hybrid pair-wise recommendation algorithm with implicit feedback. In *Proceedings of the PAKDD*.
- [132] Fuzheng Zhang, Nicholas Jing Yuan, Defu Lian, Xing Xie, and Wei-Ying Ma. 2016. Collaborative knowledge base embedding for recommender systems. In *Proceedings of the KDD*. ACM, 353–362.
- [133] He Zhao, Lan Du, Wray L. Buntine, and Mingyuan Zhou. 2018. Dirichlet belief networks for topic structure learning. In *Proceedings of the NIPS*. 7966–7977.
- [134] Vincent Wenchen Zheng, Bin Cao, Yu Zheng, Xing Xie, and Qiang Yang. 2010. Collaborative filtering meets mobile recommendation: A user-centered approach. In *Proceedings of the AAAI*.
- [135] Mingyuan Zhou and Lawrence Carin. 2015. Negative binomial process count and mixture modeling. *IEEE Trans. Pattern Anal. Mach. Intell.* 37, 2 (2015), 307–320. DOI: <https://doi.org/10.1109/TPAMI.2013.211>
- [136] Mingyuan Zhou, Lauren Hannah, David B. Dunson, and Lawrence Carin. 2012. Beta-negative binomial process and Poisson factor analysis. In *Proceedings of the AISTATS*. 1462–1471.

Received March 2020; revised August 2020; accepted September 2020