

# Learning sparsity of representations with discrete latent variables

Xu, Zhao; Rubio, Daniel Onoro; Serra, Giuseppe; Niepert, Mathias

DOI:

[10.1109/IJCNN52387.2021.9533762](https://doi.org/10.1109/IJCNN52387.2021.9533762)

License:

Other (please specify with Rights Statement)

*Document Version*

Peer reviewed version

*Citation for published version (Harvard):*

Xu, Z, Rubio, DO, Serra, G & Niepert, M 2021, Learning sparsity of representations with discrete latent variables. in *2021 International Joint Conference on Neural Networks (IJCNN)*., 9533762, International Joint Conference on Neural Networks (IJCNN), IEEE, pp. 1-9, 2021 International Joint Conference on Neural Networks (IJCNN), 18/07/21. <https://doi.org/10.1109/IJCNN52387.2021.9533762>

[Link to publication on Research at Birmingham portal](#)

**Publisher Rights Statement:**

Z. Xu, D. O. Rubio, G. Serra and M. Niepert, "Learning Sparsity of Representations with Discrete Latent Variables," 2021 International Joint Conference on Neural Networks (IJCNN), 2021, pp. 1-9, doi: 10.1109/IJCNN52387.2021.9533762.

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

**General rights**

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

**Take down policy**

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact [UBIRA@lists.bham.ac.uk](mailto:UBIRA@lists.bham.ac.uk) providing details and we will remove access to the work immediately and investigate.

# Learning Sparsity of Representations with Discrete Latent Variables

Zhao Xu<sup>†</sup>, Daniel Onoro Rubio<sup>†</sup>, Giuseppe Serra<sup>†§</sup>, Mathias Niepert<sup>†</sup>

<sup>†</sup>NEC Laboratories Europe, Heidelberg, Germany

<sup>§</sup>CERCIA, School of Computer Science, University of Birmingham, UK

{zhao.xu, daniel.onoro, giuseppe.serra, mathias.niepert}@neclab.eu

**Abstract**—Deep latent generative models have attracted increasing attention due to the capacity of combining the strengths of deep learning and probabilistic models in an elegant way. The data representations learned with the models are often continuous and dense. However in many applications, sparse representations are expected, such as learning sparse high dimensional embedding of data in an unsupervised setting, and learning multi-labels from thousands of candidate tags in a supervised setting. In some scenarios, there could be further restriction on degree of sparsity: the number of non-zero features of a representation cannot be larger than a pre-defined threshold  $L_0$ . In this paper we propose a sparse deep latent generative model SDLGM to explicitly model degree of sparsity and thus enable to learn the sparse structure of the data with the quantified sparsity constraint. The resulting sparsity of a representation is not fixed, but fits to the observation itself under the pre-defined restriction. In particular, we introduce to each observation  $i$  an auxiliary random variable  $L_i$ , which models the sparsity of its representation. The sparse representations are then generated with a two-step sampling process via two Gumbel-Softmax distributions. For inference and learning, we develop an amortized variational method based on MC gradient estimator. The resulting sparse representations are differentiable with backpropagation. The experimental evaluation on multiple datasets for unsupervised and supervised learning problems shows the benefits of the proposed method.

**Index Terms**—Deep Latent Generative Models, Sparsity of Representation, Amortized Variational Inference

## I. INTRODUCTION

Deep Latent Generative Models (DLGMs) [1]–[5] are powerful generative models that combines strength of deep neural networks and approximate Bayesian inference to learn representations of large scale data. The learned representations reveal important semantic features of the data, and are often useful for many downstream tasks. The DLGMs have been successfully applied to a variety of applications with interesting results, such as question-answering systems [6], action forecasting [7], anomaly detection [8], image progression and editing [9], [10].

In this paper, we investigate sparse deep latent generative model (SDLGM) to explicitly learn sparsity of discrete representations of the data. The semantic features of large scale data are often sparse. For instance, in an image database with a variety of images, possible visual concepts include hairstyles, pose

and emotions for portraits, art-styles for paintings, animals and plants for nature etc. Different types of images will only focus on a limited number of latent semantic features. Thus the latent codes of the images should be intrinsically sparse. Another typical example in a supervised learning setting might be social tags of recommendation systems. For example, the total number of candidate social tags of an online music recommendation system can be hundreds of thousands, but a single track is often labeled with a dozen of tags. Thus it makes sense to learn sparse discrete representations such that the learned logits mainly focus on a finite number of features. Moreover, there could be further constraint on sparsity in some scenarios. For example, the number of non-zero features of an observation (image or time series) cannot be larger than a pre-defined threshold  $L_0$  for IoT devices with limited resources in terms of storage and communication bandwidth.

To meet these challenges, the SDLGM introduces to each observation an auxiliary (discrete) random variable  $L_i$ , which specifies how sparse the representation should be. The smaller the variable  $L_i$ , the less number of features are required by the observation to represent its intrinsic properties. Then we assume the representation is a sum of  $L_i$  i.i.d. features from the shared feature distributions. In a Bayesian framework, it would be natural to model  $L_i$  as a discrete variable with e.g. binomial distribution. However it makes the feature sampling process expensive, since the number of random features is varying. Technically, we introduce to each observation  $L_0$  i.i.d. random features with the corresponding i.i.d. Bernoulli weights. The weights can be viewed as gates. When a weight is zero, then the gate switches off, and the corresponding feature is dropped out. The resulting representation is a weighted sum of  $L_0$  random features. Moreover the sum of the weights is exactly  $L_i$  following binomial distribution. Under the deep learning framework, the discrete sparsity is unfortunately non-differentiable, and in turn makes backpropagation impossible. To solve the problem, we employ a Gumbel softmax distribution [5], [11], [12], one for each observation and conditioned on the observation itself via a DNN, as the distinct sparsity distribution. This is a differentiable approximation to its discrete counterpart. For feature distributions, we employ high dimensional Gumbel-Softmax distribution. In the SDLGM, the sparsity of the representation is explicitly modeled with an extra random variable, and thus can be clearly learned from the data itself. It is distinct from the commonly used

The work of NEC Laboratories Europe was partially supported by H2020 MonB5G project (grant agreement no. 871780). The research of G. Serra was supported by H2020 ECOLE project (grant agreement no. 766186).

regularization techniques, which implicitly achieve sparsity effect. The proposed two-step sampling process allows the SDLGM method efficiently capture significant features of the observations. It nicely matches the context where there is a quantified restriction on the sparsity itself.

We apply the SDLGM to supervised and unsupervised learning cases. For supervised learning, the problem is formulated as a multi-class multi-label classification issue. We exploit three modeling strategies, including discriminative modeling (SDLGM-Dis), generative modeling (SDLGM-Gen) and conditional generative modeling (SDLGM-Con). In the unsupervised case, the problem is the representation learning, where the representation variable  $\mathbf{z}_i$  is unobservable. We learn the joint probability  $P(\mathbf{z}_i, \mathbf{x}_i) = P(\mathbf{z}_i)P(\mathbf{x}_i|\mathbf{z}_i)$ . The inference of the SDLGM is based on amortized variational method and MC gradient estimator. The empirical analysis is performed on multiple datasets. We compare the proposed model with the most recent results. For unsupervised learning, we use the popular benchmark dataset MNIST, and measure the performance with a commonly used criteria: lower bound of the log likelihood of the data. The SDLGM significantly outperforms the baselines. In the qualitative analysis, the proposed model shows very clear reconstructed images and latent structures of the data. For supervised learning, the SDLGM method also achieves better performance in terms of Micro- and Macro-averaging F1 scores in multiple datasets: Birds, Emotions and CAL500.

The rest of the paper is organized as follows. We start off with a brief review of related work. Afterwards we describe the proposed model SDLGM to learn the sparsity of representation and the variational inference method. Before concluding, we present our experimental results on multiple datasets for different learning tasks.

## II. RELATED WORK

There are mainly two lines of research related to the proposed model SDLGM, including Bayesian generative neural networks and sparse representation learning. In this section, we briefly review the related work.

### A. Bayesian Generative Neural Networks

Recent advances in Bayesian modeling renew interests in scalable probabilistic modeling. There are two major benefits which make Bayesian methods appealing: modeling uncertainty in the learned parameters and avoiding overfitting. They compensate some pitfalls of deep neural networks. For example, with increasingly complex problems and large data, the deep neural networks become more and more complicated and often include millions of parameters. So it is required to quantify the amounts of parameter uncertainties.

Introducing Bayesian modeling into neural networks have attracted a lot of attention in the literature. The early works include: for example, Denker and LeCun transformed neural network outputs to probabilistic distributions [13]. Buntine and Weigend presented Bayesian back-propagation [14]. Neal

proposed the first Monte Carlo (MC) sampling method to learn Bayesian neural networks [15].

With the emergence of deep learning, scalability of Bayesian methods become a major concern. Recent Bayesian methods incorporate the advances in optimization, and enable Bayesian models for large scale datasets. For instance, Hoffman et al. introduced stochastic variational inference method [16]. Welling and Teh presented stochastic gradient Langevin dynamics to approach the true posterior distribution of the parameters [17]. The Bayesian backpropagation was extended to large scale data [18], [19].

Deep latent generative models DLGMs nicely integrate recent advances of Bayesian modeling into deep neural networks for large data. The DLGMs are originally developed to learn continuous latent representations for unsupervised learning. The pioneer works include variational autoencoder [1], deep latent Gaussian models [2], and autoregressive networks [3]. The models are then extended to semi-supervised learning in [20]. Sohn et al. presented conditional DLGMs to learn structured representations. The discrete version of the DLGMs are introduced by [5], [11], [12] to learn discrete latent representations with Gumbel-Softmax. The DLGMs are also extended to model text data [6], graphical data [21] and multimodal data [22].

### B. Sparse Representation with Neural Networks

Sparse representation with neural networks has been investigated from different perspectives in the literature.

Lee et al. [23] and Ranzato et al. [24] introduced sparse variants of deep belief networks to learn sparse features of images. Sparse auto-encoders were explored by Kavukcuoglu et al. [25], Makhzani and Frey [26], and Arpit et al. [27] for sparsification of neural network layers.

Glorot et al. [28] were motivated by certain properties of biological neurons and proposed rectifier activation function to sparsify neural networks. Goodfellow et al. [29], Martins and Astudillo [30] introduced Maxout and SparseMax activation functions for sparsity, respectively. Dropout method has also been used to effectively sparsify neural networks by Molchanov et al. [31]. Louizos et al. [32] and Srinivas et al. [33] investigated  $\ell_0$  norm related regularization methods for learning sparse neural networks.

Deep sparse representation is also developed for fast similarity search. For example, Li et al. introduced the optimal sparse lifting method [34]. Jeong and Song presented an optimization method to jointly learn a quantizable embedding representation and the sparse hash table [35]. Inspired by efficient sparse matrix multiplication, Paria et al. introduced to learn sparse representation by minimizing a continuous relaxation of the number of floating-point operations (FLOPs) [36] in search.

Distinguishing from the existing work, this paper extends the deep latent generative models to explicitly quantify degree of sparsity, and fit it to the data automatically. We also apply the learnable sparsity to supervised learning for multi-label classification with good performance.

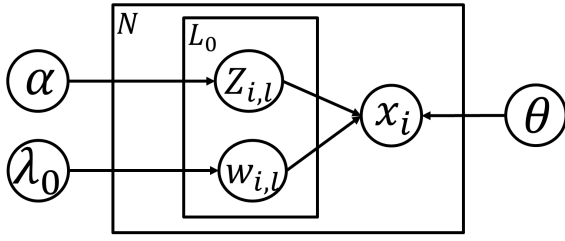


Fig. 1. Sampling process of the sparse deep latent generative model SDLGM.

### III. SPARSE DEEP LATENT MODELS

In this section, we introduce the proposed sparse deep latent generative model SDLGM and the corresponding inference method. We will first describe the model in an unsupervised learning case to learn representations of the data, then extend it to a supervised learning case for multi-class multi-label classification. Let start with problem formulation. Assume that there are  $N$  observations  $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ . Each observation  $\mathbf{x}_i$  is a  $D$ -dimensional vector. We assume that  $x_{i,d}$  is binary for the running example, while it is straightforward to extend the SDLGM to continuous inputs. The goal is to learn a sparse  $K$ -dimensional discrete representation  $\mathbf{z}_i$  for each observation under a constraint on representation sparsity that the number of non-zero features cannot be larger than a certain threshold  $L_0$ . Formally the representation  $\mathbf{z}_i$  of the observation  $\mathbf{x}_i$  is denoted as:  $\mathbf{z}_i \in \{0, 1\}^K$ ;  $\|\mathbf{z}_i\|_0 \leq L_0$ , where  $\|\cdot\|_0$  is  $\ell_0$  norm of a vector.

The SDLGM is a probabilistic generative model which integrates good properties of deep neural networks DNNs into Bayesian modeling for flexible definition of data distributions. To meet the sparsity constraint, we introduces to each observation  $\mathbf{x}_i$  an auxiliary discrete random variable  $L_i$ , which reveals degree of representation sparsity. The value of the variable  $L_i$  specifies how many non-zero features are required by the observation to represent its intrinsic properties. In addition,  $L_i$  corresponds to the  $\ell_0$  norm of  $\mathbf{z}_i$ . Technically it is tricky to directly model the  $\ell_0$  norm  $L_i$ , since the varying number of random features makes the sampling process inefficient. Alternatively, we decompose the representation  $\mathbf{z}_i$  as:  $\mathbf{z}_i = \sum_{\ell=1}^{L_0} \omega_{i,\ell} \mathbf{z}_{i,\ell}$ . It specifies that the representation  $\mathbf{z}_i$  consists of at most  $L_0$  features. Each feature  $\mathbf{z}_{i,\ell}$  is a  $K$ -dimensional one-hot vector, and corresponds to a categorical variable. The weight  $\omega_{i,\ell}$  is binary, and can be viewed as a *gate*. When the gate switches on  $\omega_{i,\ell} = 1$ , the corresponding feature is present. The gate vector  $\omega_i = (\omega_{i,1}, \dots, \omega_{i,L_0})$  is observation-specific, one for each observation. It will be learned as per the data itself. The sparsity variable can thus be written as:  $L_i = \sum_{\ell=1}^{L_0} \omega_{i,\ell}$ , which is exactly the  $\ell_0$  norm of  $\mathbf{z}_i$ , and meets the sparsity constraint  $L_i < L_0$ . Fig. 1 illustrates the SDLGM.

#### A. The Generative Process

Now we define the details of the SDLGM model with the generative process. Formally we draw an observation  $\mathbf{x}_i$  in the following manner.

- Draw  $L_0$  number of i.i.d. random variables  $\omega_{i,\ell} \in \{0, 1\}$  from a Bernoulli distribution with a hyperparameter  $0 < \lambda_0 < 1$ :

$$\omega_{i,\ell} \sim \text{Bernoulli}(\lambda_0) \quad (1)$$

Here the sparsity variable  $L_i$  follows binomial distribution with hyperparameters  $L_0$  and  $\lambda_0$ , i.e.,

$$L_i = \sum_{\ell=1}^{L_0} \omega_{i,\ell}; \quad L_i \sim \text{Binomial}(L_0, \lambda_0).$$

Here we do not directly sample  $L_i$  from the binomial distribution. The suggested sampling process is computationally efficient due to matrix manipulation of DNNs in the inference.

- The latent representation  $\mathbf{z}_i$  is a weighted sum of  $L_0$  i.i.d. random features. Each feature  $\mathbf{z}_{i,\ell}$  is a  $K$ -dimensional one-hot categorical variable. They are independent, but share a common categorical distribution with the parameter  $\alpha$  that lies in a  $(K-1)$ -dimensional simplex  $\Delta^{K-1}$ , i.e.  $\sum_{k=1}^K \alpha_k = 1$ ;  $\alpha_k > 0$ . Here we assume a prior that gives equal probability to each dimension  $\alpha_k = 1/K$ . The representation  $\mathbf{z}_i$  is sampled as:

$$\mathbf{z}_i = \sum_{\ell=1}^{L_0} \omega_{i,\ell} \mathbf{z}_{i,\ell}; \quad \mathbf{z}_{i,\ell} \sim \text{Categorical}(\alpha) \quad (2)$$

- Finally we sample the binary observation  $\mathbf{x}_i \in \{0, 1\}^D$  conditioned on the latent representation  $\mathbf{z}_i$ . Each observation has distinct Bernoulli parameters  $\beta_i \in \mathbb{R}^D$ , which are computed as:

$$\beta_i = g(\mathbf{z}_i; \theta),$$

where the function  $g(\cdot)$  is represented using a deep neural network, as known as *decoder*, with  $\mathbf{z}_i$  as input and  $\theta$  as hyperparameters. Given the computed Bernoulli parameter  $\beta_{i,d}$ , we can then independently draw each dimension of the observation:

$$x_{i,d} \sim \text{Bernoulli}(\beta_{i,d}). \quad (3)$$

The decoder network  $g$  can be flexible, but the output layer will use sigmoid activation to meet the constraint of the Bernoulli parameter, i.e.,  $0 < \beta_{i,d} < 1$ .

The generative model SDLGM presents an elegant way to clearly capture the sparsity of the latent code. In Bayesian modeling, the typical solution for sparseness is to set a prior distribution on the parameters, e.g., Gaussian prior as L2 regularization and Laplace prior as L1 regularization. Instead of these implicit ways, we introduce auxiliary variables  $\omega_{i,\ell}$  to model the sparsity of the latent representation clearly. It explicitly switches on a finite number of latent features to generate an observation. In addition, the discrete latent codes are generally modeled with Bernoulli or one-hot categorical distributions in Bayesian framework. In this paper, we model them with a novel way (two sampling steps) to clearly learn the sparsity of the latent codes, shown as Equation 1 and 2 for generative process and Equation 8 and 10 for approximate inference. With the proposed model, we can learn the sparse

representation under a certain threshold on degree of sparsity. It would be especially useful in the low-resource applications, e.g. edge domain of 5G network with limited communication bandwidth, and a surveillance system with storage constraint. In summary, the joint distribution of a single observation is computed as:

$$\begin{aligned} p(\mathbf{x}_i, \{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0} | \theta, \lambda_0, L_0, \alpha) \\ = p(\mathbf{x}_i | \{\mathbf{z}_{i,\ell}, \omega_{i,\ell}\}_{\ell=1}^{L_0}, \theta) \prod_{\ell=1}^{L_0} p(\omega_{i,\ell} | \lambda_0) p(\mathbf{z}_{i,\ell} | \alpha) \end{aligned} \quad (4)$$

### B. Inference and Learning

In this section we introduce computational challenges in fitting the sparse model SDLGM. The key inferential problem is to compute the posterior distribution of the sparse latent representation  $\mathbf{z}_i = \sum_{\ell=1}^{L_0} \omega_{i,\ell} \mathbf{z}_{i,\ell}$  given an observation  $\mathbf{x}_i$ :

$$\begin{aligned} p(\{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0} | \mathbf{x}_i, \theta, \alpha, \lambda_0, L_0) \\ \propto p(\{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0} | \alpha, \lambda_0) p(\mathbf{x}_i | \{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0}, \theta) \end{aligned} \quad (5)$$

The posterior distribution (5) is intractable for exact inference due to the normalization term  $p(\mathbf{x}_i | \theta, \alpha, \lambda_0, L_0)$ . We approximate the intractable posterior based on the amortized variational inference (VI).

Amortized VI [1], [2], [37]–[39] is a recent advance in Bayesian inference and deep learning. Instead of directly optimizing a set of variational parameters, it models the variational parameters as functions of observations, and then optimizes the parameters of the functions to approximate the posterior. In deep learning, the functions are often represented with neural networks, that map from observation space to variational parameter space. The inference of the SDLGM is implemented under the amortized VI framework. Note that the sparsity variables and the feature variables are discrete, thus we use Gumbel softmax distributions to approximate their non-differentiable discrete counterparts.

Based on Jensen’s inequality and fully-factorized variational distributions [40], the lower bound  $\mathcal{L}$  (a.k.a. evidence lower bound ELBO) of the log likelihood  $\log p(\mathbf{x}_i | \theta, \alpha, \lambda_0, L_0)$  is:

$$\begin{aligned} \mathcal{L} = & \sum_{\ell=1}^{L_0} \mathbb{E}_{q(\omega_{i,\ell})} [\log p(\omega_{i,\ell} | \lambda_0)] \\ & + \sum_{\ell=1}^{L_0} \mathbb{E}_{q(\mathbf{z}_{i,\ell})} [\log p(\mathbf{z}_{i,\ell} | \alpha)] \\ & - \sum_{\ell=1}^{L_0} \mathbb{E}_{q(\omega_{i,\ell})} [\log q(\omega_{i,\ell})] \\ & - \sum_{\ell=1}^{L_0} \mathbb{E}_{q(\mathbf{z}_{i,\ell})} [\log q(\mathbf{z}_{i,\ell})] \\ & + \mathbb{E}_{q(\{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0})} \left[ \log p(\mathbf{x}_i | \theta, \{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0}) \right] \end{aligned} \quad (6)$$

On the right side of Equation (6), the combination of the first four terms is the negative KL divergence between variational and prior distributions. The last term is the variational expectation of the log likelihood of the data. Now the inference problem is converted to find the tightest possible lower bound  $\mathcal{L}$ , i.e. maximizing  $\mathcal{L}$  w.r.t. the variational distributions.

Based on the amortized VI, we use deep neural networks (DNNs) to define the variational distributions. For the sparsity variable  $\omega_{i,\ell}$ , it is natural to assume:

$$q(\omega_{i,\ell}) = \text{Bernoulli}(\lambda_i); \quad \lambda_i = h(\mathbf{x}_i; \gamma), \quad (7)$$

where we assume that the Bernoulli parameter  $0 < \lambda_i < 1$ , one for each observation, is computed with a function  $h(\mathbf{x}_i; \gamma)$ , which is a DNN with  $\mathbf{x}_i$  as input and  $\gamma$  as hyperparameters. However the discrete variables make the corresponding layers no longer differentiable, such that the efficient computation of parameter gradients become difficult for backpropagation. A natural question to arise is: how do we find a continuous approximation to the discrete one to solve the issue. We use 2-dimensional Gumbel-softmax, a differentiable approximation to Bernoulli distribution. Gumbel-softmax distribution is defined over continuous random variables, and is introduced by [5], [11], [12], [41] to approximate categorical and Bernoulli distributions. A sample of  $\omega_{i,\ell}$  from a 2-dimensional Gumbel-Softmax with parameters  $\lambda_i$  and a hyperparameter  $\tau \in \mathbb{R}$  (known as temperature) is defined as follows:

$$\begin{aligned} \omega_{i,\ell} = \frac{\exp((\log(\lambda_{i,1}) + \xi_{i,\ell,1})/\tau)}{\sum_{k'=1}^2 \exp((\log(\lambda_{i,k'}) + \xi_{i,\ell,k'})/\tau)} \\ \xi_{i,\ell,k'} \sim \text{Gumbel}(0, 1); \quad \lambda_{i,1} = \lambda_i; \quad \lambda_{i,2} = 1 - \lambda_i. \end{aligned} \quad (8)$$

Similarly the variational distribution  $q(\mathbf{z}_{i,\ell})$  of a single latent feature  $\mathbf{z}_{i,\ell}$  is defined as a K-dimensional continuous Gumbel-softmax with parameters  $\pi_i$  and a hyperparameter  $\tau_z$ :

$$q(\mathbf{z}_{i,\ell}) = \text{Gumbel-Softmax}(\pi_i, \tau_z); \quad \pi_i = f(\mathbf{x}_i; \phi). \quad (9)$$

Here the function  $f(\cdot)$  denotes another DNN (known as *encoder*) with the observation  $\mathbf{x}_i$  as input and  $\phi$  as hyperparameters. Figure 2 shows the computational graph of the encoder. Draw a sample of  $\mathbf{z}_{i,\ell}$  from the K-dimensional Gumbel-Softmax as follows:

$$\begin{aligned} z_{i,\ell,k} = \frac{\exp((\log(\pi_{i,k}) + \xi_{i,\ell,k})/\tau)}{\sum_{k'=1}^K \exp((\log(\pi_{i,k'}) + \xi_{i,\ell,k'})/\tau)} \\ \xi_{i,\ell,k'} \sim \text{Gumbel}(0, 1); \quad k = 1 \dots K \end{aligned} \quad (10)$$

Given the variational distributions, we can now optimize the lower bound  $\mathcal{L}$ . Since we use DNNs to define the variational distributions, the variational parameters  $\lambda_i$  and  $\pi_i$  are deterministically computed with the global functions  $h(\mathbf{x}_i; \gamma)$  and  $f(\mathbf{x}_i; \phi)$  (given that  $\phi$  and  $\gamma$  are unknown but not random). Thus we actually do not need to optimize them with an expectation step. Instead we directly maximize  $\mathcal{L}$  w.r.t. the DNNs’ parameters  $\theta$ ,  $\phi$ , and  $\gamma$ .

The expectations in the lower bound (6) are now computed with the MC estimators. The variational expectation of the log likelihood of the data is derived as:

$$\begin{aligned} \mathbb{E}_{q(\{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0})} \left[ \log p(\mathbf{x}_i | \theta, \{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0}) \right] \\ = \frac{1}{S} \sum_{s=1}^S \sum_{d=1}^D \left\{ x_{i,d} \log \beta_{i,d}^{(s)} + (1 - x_{i,d}) \log (1 - \beta_{i,d}^{(s)}) \right\}, \end{aligned} \quad (11)$$

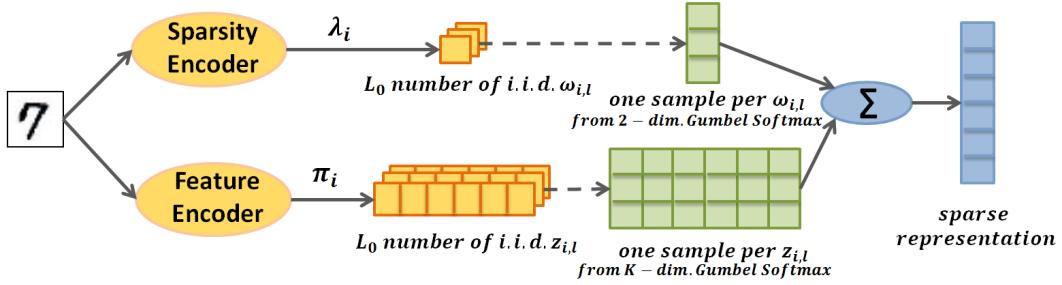


Fig. 2. The encoder computational graph for the amortized variational inference of the SDLGM model.

where  $\beta_i^{(s)} = g(\mathbf{z}_i^{(s)}; \theta)$ ;  $\mathbf{z}_i^{(s)} = \sum_{\ell=1}^{L_0} \omega_{i,\ell}^{(s)} \mathbf{z}_{i,\ell}^{(s)}$ ;  
 $\mathbf{z}_{i,\ell}^{(s)} \sim \text{K-d Gumbel-softmax}(\pi_i, \tau_z)$  with (10);  
 $\omega_{i,\ell}^{(s)} \sim \text{2-d Gumbel-softmax}(\lambda_i, \tau)$  with (8). (12)

There are some advantages introduced by the MC estimator. For example,  $\mathbf{z}_i$  is a weighted sum of  $L_0$  continuous categorical variables in theory, then its distribution is an  $L_0$ -fold convolution. Since  $\mathbf{z}_{i,\ell}$  is now continuous and follows Gumbel-softmax distribution, the convolution is analytically difficult to obtain. With MC estimator, we do not need to explicitly compute the probabilistic density, instead, get an approximation with a sampler.

Next we compute Kullback-Leibler divergence KLD of variational and prior distributions of the sparsity variable. Since  $\omega_{i,\ell}$  are i.i.d., the KLD is derived as:

$$L_0 \left[ \lambda_i \log \left( \frac{\lambda_i}{\lambda_0} \right) + (1 - \lambda_i) \log \left( \frac{1 - \lambda_i}{1 - \lambda_0} \right) \right]. \quad (13)$$

Similarly we compute  $\sum_{\ell=1}^{L_0} \text{KL}(q(\mathbf{z}_{i,\ell}) || p(\mathbf{z}_{i,\ell}))$  as:

$$L_0 \sum_{k=1}^K \pi_{i,k} \log(\pi_{i,k} / \alpha_k). \quad (14)$$

We have described the computation of each term in the lower bound (6). Now we optimize the lower bound  $\mathcal{L}$  (for a single observation) w.r.t. hyperparameters based on stochastic method [1], [16]. Let start with the hyperparameters  $\theta$ . The gradient  $\nabla_{\theta} \mathcal{L}$  is approximated with the MC estimator.

$$\begin{aligned} & \mathbb{E}_{q(\{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0})} \left[ \nabla_{\theta} \log p(\mathbf{x}_i | \theta, \{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0}) \right] \\ &= \frac{1}{S} \sum_{s=1}^S \left[ \nabla_{\theta} \mathbf{x}_i \log g(\mathbf{z}_i^{(s)}; \theta) + \nabla_{\theta} (\mathbf{1} - \mathbf{x}_i) \log (\mathbf{1} - g(\mathbf{z}_i^{(s)}; \theta)) \right], \end{aligned} \quad (15)$$

where the sampling process is the same as (12). The gradient in (15) can be easily computed with backpropagation. Note that we actually sample  $\xi$  from a fixed and known distribution Gumbel(0, 1), then deterministically compute the continuous Bernoulli and categorical samples. It takes effect of reparameterization, analogous to the location-scale transformation for Gaussian latent representations used in [1], which makes the MC variance reduced. Thus  $S$  can be similarly set to 1 when the mini-batch size is large enough, e.g. 128 [1].

Next we optimize  $\gamma$  and  $\phi$ . Their gradients are:

$$\begin{aligned} \nabla_{\gamma} \mathcal{L} &= \nabla_{\gamma} \mathbb{E}_{q(\{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0})} \left[ \log p(\mathbf{x}_i | \theta, \{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0}) \right] \\ &\quad - \sum_{\ell=1}^{L_0} \nabla_{\gamma} \text{KL}(q(\omega_{i,\ell}) || p(\omega_{i,\ell})) \end{aligned} \quad (16)$$

$$\begin{aligned} \nabla_{\phi} \mathcal{L} &= \nabla_{\phi} \mathbb{E}_{q(\{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0})} \left[ \log p(\mathbf{x}_i | \theta, \{\omega_{i,\ell}, \mathbf{z}_{i,\ell}\}_{\ell=1}^{L_0}) \right] \\ &\quad - \sum_{\ell=1}^{L_0} \nabla_{\phi} \text{KL}(q(\mathbf{z}_{i,\ell} | \pi_i, \tau_z) || p(\mathbf{z}_{i,\ell} | \alpha)) \end{aligned} \quad (17)$$

The first terms can be calculated with backpropagation and MC estimator in a similar way as (15). The computation of the second terms are straightforward with chain rule, e.g.

$$\nabla_{\gamma} \text{KL}(q || p) = \left[ \log \left( \frac{\lambda_i}{\lambda_0} \right) + \log \left( \frac{1 - \lambda_i}{1 - \lambda_0} \right) + 2 \right] \nabla_{\gamma} \lambda_i \quad (18)$$

where  $\nabla_{\gamma} \lambda_i$  is obtained with backpropagation. The scalar hyperparameters  $\tau$  and  $\tau_z$ , ranging from 0 to 10 in general, can be optimized with cross-validation, or be learned with gradients that are computed similarly as  $\theta$ ,  $\gamma$  and  $\phi$ . Put everything together, the inference method is summarized in Algorithm 1. Note that the gradients described here are to maximize  $\mathcal{L}$ , so the SGD needs negative gradients to minimize the negative lower bound  $-\mathcal{L}$ .

---

**Algorithm 1:** Amortized variational inference for SDLGM

---

**Input :** dataset  $\mathcal{D}$ , latent dimension  $K$ ,  
 $\alpha = \{\frac{1}{K}, \dots, \frac{1}{K}\}$ ,  $L_0 \ll K$ ,  $\lambda_0 = 0.5$ , batch size  $M = 100$ , sample number  $S = 1$ ,  
initialization of  $\tau = 1$ ,  $\tau_z = 1$ ,  $\theta$ ,  $\gamma$  and  $\phi$ .

**for**  $t = 1$  **to**  $T$  **do**  
  Select a minibatch  $\mathcal{B}$  from the dataset  $\mathcal{D}$ ;  
  Sample and compute  $\mathbf{z}_i$  with (12) for  $\mathbf{x}_i \in \mathcal{B}$ ;  
  Compute  $\mathcal{L}$  of the minibatch with (11), (13), (14);  
  Compute the gradients  $\nabla_{\theta, \gamma, \phi}$  with (15), (16), (17);  
  Update  $\tau$ ,  $\tau_z$ ,  $\theta$ ,  $\gamma$  and  $\phi$  using the gradients;

**end**

**Output:** the optimized  $\tau$ ,  $\tau_z$ ,  $\theta$ ,  $\gamma$  and  $\phi$

---

### C. The SDLGM for Supervised Learning

In the supervised learning setting, the SDLGM can be tailored to solve the multi-class multi-label classification issue. Assume there are  $K$  possible classes in total. For an observation  $\mathbf{x}_i$ , let denote its labels as a binary vector  $\mathbf{y}_i \in \{0, 1\}^K$ . An observation can have a varying number of labels, i.e.,  $\sum_{k=1}^K \mathbf{y}_{i,k} = L_i$  is not fixed. Here the SDLGM has three variants, which employ different modeling strategies.

**Discriminative modeling.** We directly estimate the target probability  $p(\mathbf{y}_i|\mathbf{x}_i, \alpha, \lambda_0, L_0, \tau, \tau_z, \gamma, \phi)$  without modeling the generative process of the data. Technically, it is the encoder part described in the above section, and the variable  $\mathbf{z}_i$  is now the observable labels, i.e.  $z_i \equiv y_i$ . Since  $\mathbf{y}_i$  is observed, we can derive the parameters  $\gamma$  and  $\phi$  of the encoder via optimizing the log likelihood of the data  $\sum_{i=1}^N \log p(\mathbf{y}_i|\mathbf{x}_i, \alpha, \lambda_0, L_0, \tau, \tau_z, \gamma, \phi)$ . The gradients are similarly computed with the equations (16) and (17).

**Generative modeling.** We model the generative process of the data, which is the same as Fig. 1. Again, the variable  $\mathbf{z}_i$  is now the observable labels,  $z_i \equiv y_i$ . The inference is thus similar as the last section. The major difference is that there is now an extra loss due to the observed labels. After learning the parameters of the model, we use the encoder network to predict the labels  $y_*$  of a test data.

**Conditional generative modeling.** We assume that the observation  $x_i$  is conditioned on not only its observable labels  $y_i$ , but also some latent features. Technically, a part of the variable  $\mathbf{z}_i$  in the SDLGM is observed (denoted as  $y_i$ ), and the rest is still unobservable. Different from the conditional VAE [4], the SDLGM models  $q(z_i, y_i|x_i)$  with the encoder network, rather than  $q(z_i|x_i, y_i)$ , since the learning tasks are not only generating the observations, but also classifying the multi-class multi-labels.

In theory, the discriminative SDLGM would show better performance than the generative one since it directly optimizes the target probability [42]. The generative SDLGM would be good at generality due to the additional flexibility of modeling the data generative process. Intuitively the conditional generative SDLGM might achieve more accurate classification results than the generative SDLGM, as its generative process can be more reasonable: the observations depend on not only their labels, but also some hidden features and random effects.

## IV. EMPIRICAL ANALYSIS

In the experiment session, we evaluate performance of the SDLGM with multiple datasets for different learning tasks. The proposed method is first verified for an unsupervised learning task to train sparse image representation with the popular benchmark dataset MNIST. We then test the SDLGM for a supervised learning task about multi-class multi-label classification with the datasets: Birds, Emotions and CAL500. Finally the SDLGM is applied to predict appearance times of items in an image with a digit dataset created from MNIST.

TABLE I  
UNSUPERVISED LEARNING WITH THE SDLGM. COMPARISON WITH THE BASELINES REPORTED IN [5]. THE PERFORMANCE IS MEASURED WITH NEGATIVE ELBO (LOWER IS BETTER).

Methods	Negative ELBO	
	Categorical	Bernoulli
SF	110.6	112.2
DARN	128.8	110.9
MuProp	107.0	109.7
ST	110.9	116.0
Annealed ST	107.8	111.5
Gumbel-Softmax	101.5	105.0
ST Gumbel-Softmax	107.8	111.5
Our SDLGM	99.05 $\pm$ 1.07 (test)	
	99.58 $\pm$ 1.08 (training)	

### A. Unsupervised Learning Task

We start with the representation learning task. The SDLGM is compared with state-of-the-art methods, including Gumbel-Softmax, ST Gumbel-Softmax, Score-Function (SF), DARN, MuProp, Straight-Through (ST), and Slope-Annealed ST [5]. The performance of the methods is measured with negative variational lower bound of the log-likelihood (lower is better), which is a commonly used criterion in the literature. In the experiments, we follow the similar settings as [5]. The popular image dataset MNIST is used. The encoder is defined as a Multi-Layer Perceptron (MLP) network with three hidden layers [512, 384, 256]. The input are  $28 \times 28$  image data. The outputs are the logits of the distributions of 200d. The sparsity part is another MLP net with 2 hidden layers [256, 64], and with 2d outputs for parameters of continuous binomial distribution. The hyperparameters  $\tau$  and  $\tau_z$  are initially set as  $\tau_0 = 1.0$  and are learned with the stochastic backpropagation. The hyperparameter  $L_0$  is set to 40. The decoder is defined as a reverse MLP network [256, 384, 512]. The outputs are 784d vectors, which are the Bernoulli logits of the pixels.

We first quantitatively demonstrate the performance of the proposed method with the commonly used measurement: the negative variational lower bound (lower is better). We compare the SDLGM with multiple baselines, shown as Table I. For a fair comparison, we directly use the results reported in [5]. More details of the baselines refer to [5]. The compared methods model the latent representations as 200d Bernoulli variables or 20 number of 10d categorical variables (i.e., the same number of latent logits as ours). We rerun the experiment 10 times, and report the mean. For each time, we train the model with 10000 iterations, and get the average ELBO over last 200 iterations. One can see that our sparse method achieves significantly better performance than the recent baselines. The SDLGM finds the tighter variational lower bound.

The proposed method has extra flexibility via learning additional sparsity variables. A natural question then arises that: is the method computationally expensive? We empirically evaluate the time cost, and compare with Jang et al.'s method. We run the methods at the same computing environment and with the same hidden structures ([512,384,256] for the

feature encoder, and [256,384,512] for the decoder), the same batch size (100). Experimentally, our method uses 0.013s per iteration, while Jang et al.’s method uses 0.009s. The total time costs would be comparable, since we have less iterations (10k vs 100k). The SDLGM can learn the sparse representations in an efficient way since:

- 1) the distributions of the data sparsity are also represented with neural networks and are properly embedded in the deep generative modeling framework, thus can be efficiently learned with backpropagation.
- 2) the sparsity is represented as a function of the observation, and thus can share some complicated low layers with the feature encoder.
- 3) it largely reduces the number of parameters.

To qualitatively analyze the performance of the SDLGM, we use it to reconstruct images. Figure 3 shows some random examples: the original images (left), and the reconstructed ones (right). One can find that the reconstructed images are very clear, and almost the same as the original ones.

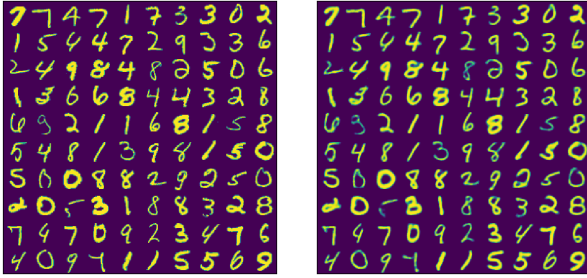


Fig. 3. The original images (left) and the SDLGM-reconstructed images (right).

Furthermore we analyze the learned sparse discrete representations. Figure 4 shows the results with 3000 test images. The left panel illustrates the distributions of the sparse representations. One can find that the images of different digits (0-9) are mapped into the embedding space with different distributions. The sparse codes  $\mathbf{z}$  intrinsically capture the visual concepts of the images, and clearly reveal the latent structures of the data. The two histograms represent the distributions of the sparsity  $L_i$  (right panel) and of the sparsity logits (left panel), outputs of the MLP for  $h(\mathbf{x}_i|\gamma)$ . They are exactly close to binomial distributions.

### B. Multi-class Multi-Label Classification

We also validate the performance of the proposed method in the supervised learning setting, i.e. multi-class multi-label classification. The experiment is performed with three datasets: Birds, Emotions and CAL500 (<http://mulan.sourceforge.net/datasets-mlc.html>). To have a fair comparison with state-of-the-art methods, we follow the experiment settings in [30]. We remove the observations without any labels, and normalize the attributes to have zero mean and unit variance. The datasets include multiple classes: 19 (Birds), 6 (Emotions) and 174 (CAL500). Figure 5 shows the sparsity of the labels. We can find that the number of labels associated with an observation

is not fixed, and relatively sparse considering the possible states. For hyperparameter selection we performed a 5-fold cross validation. The evaluated hyperparameters are  $L_0$ , the probability thresholds  $k \in \{0.05 \times n\}_{n=1}^{10}$  above which a label is predicted to be 1, the temperatures  $\tau \in \{2.5 + 0.1 \times n\}_{n=0}^{10}$ . The selected hyperparameters are:  $\tau = 3$ ,  $L_0 = 10$  and  $k = 0.4$  for Birds,  $\tau = 3$ ,  $L_0 = 2$  and  $k = 0.2$  for Emotions,  $\tau = 3$ ,  $L_0 = 50$  and  $k = 0.2$  for CAL500. We rerun the experiments 5 times and report the average, shown as Table II. The three variants of the SDLGM are denoted as: SDLGM-Dis (discriminative version), SDLGM-Gen (generative version), and SDLGM-Con (conditional generative one). Compared with the baselines, the SDLGM achieve better performance in most settings. In particular, the SDLGM with conditional generative modeling provides good predictions. The reason might be the good generality and the reasonable generative assumptions of the variant. The discriminative SDLGM achieves comparable results. The tendency coincides with the theoretical analysis in the above section.

### C. Sparsity Prediction in Image Classification

Typically, image classification focuses on assigning a single class per image, e.g. 0-9 digits in the MNIST data. However, the real images often include multiple digits, and the numbers of appearances of the digits are varying. Although it appears to be similar to multi-label classification, it differs in which the labels can be repetitive, i.e. one digit may appear multiple times. We solve the problem with the SDLGM to predict the appearance times (sparsity). For this end, we created a new dataset where each image is a composition of 5 columns, and each column is a random digit or an empty space. We generate 5k images for training, and 1k for test. The output of the SDLGM for each image is a 10D vector. The value of each dimension specifies how many times the corresponding digit appears in the sample. If the value is zero, then there is not the digit in the sample. We compare the proposed method with random forest in terms of mean-squared-error (MSE). On the test data, the MSE of the SDLGM is 0.05, while that of the random forest is 0.16. The proposed method achieves better performance. To illustrate the results, we randomly select some test images associated with the prediction results, shown as Table III. One can find that most of the predictions, marked as green, are correct. For example, the SDLGM predicts: the digits ‘3’ and ‘4’ occur one time respectively in the first image, and the digit ‘4’ occurs two times in the second image (the top row of Table III). The experiments reveal that: (1) learning degree of sparsity is a generic task and can be used to formulate different problems; (2) the proposed method SDLGM offers an effective way to model the degree of sparsity explicitly and effectively.

## V. CONCLUSION

In this paper we present a sparse deep latent generative model, which effectively learns sparse discrete representations of the data. Instead of the commonly used regularization techniques, our model introduces an additional random variable



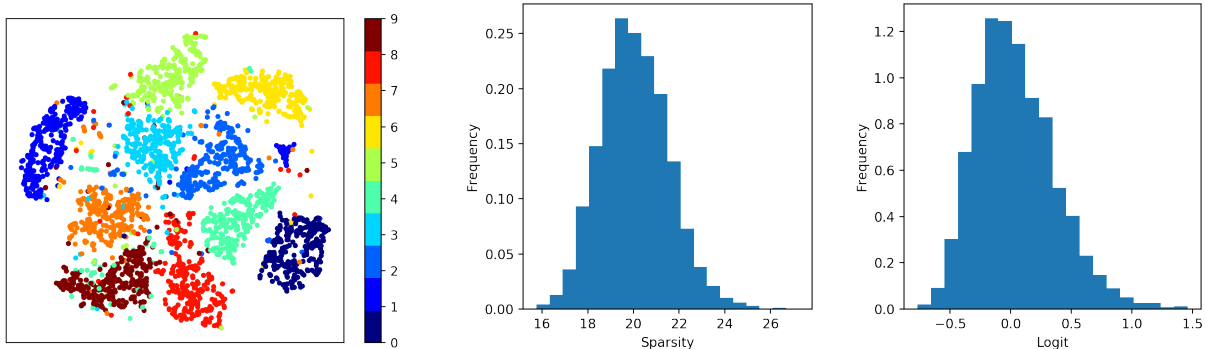


Fig. 4. Analysis of the estimated sparse representations of 3000 test images. Left: the learned latent structure of the digit images. One color denotes one digit (0-9). Middle: distribution of the learned sparsity  $L_i$ . Right: distribution of the learned sparsity logits (output of the MLP for  $h(\mathbf{x}_i|\gamma)$ ).

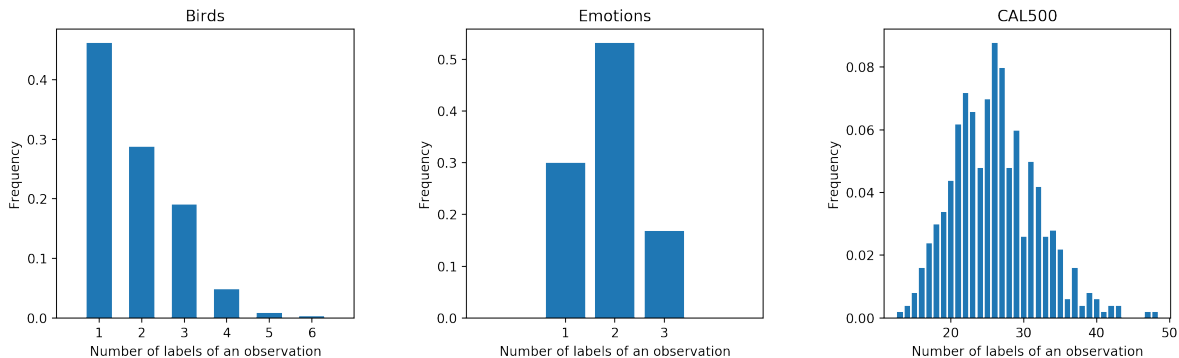


Fig. 5. Label sparsity of the multi-class multi-label classification task.

TABLE II  
MICRO AND MACRO-AVERAGED F1 SCORES FOR THE LOGISTIC, SOFTMAX, SPARSEMAX AND SDLGM METHODS ON THE BENCHMARK DATASETS: BIRDS, EMOTIONS AND CAL500.

Methods	Birds		Emotions		CAL500	
	Micro F1	Macro F1	Micro F1	Macro F1	Micro F1	Macro F1
Logistic	45.78	33.77	66.75	68.56	48.88	24.49
Softmax	48.67	37.06	67.34	67.51	47.46	23.51
Sparsemax	49.44	39.13	66.38	66.07	48.47	<b>26.20</b>
SDLGM-Dis	<b>50.75</b>	42.31	70.99	70.86	<b>49.58</b>	20.49
SDLGM-Gen	49.96	41.68	71.35	71.07	44.22	20.17
SDLGM-Con	50.67	<b>43.15</b>	<b>71.45</b>	<b>71.22</b>	45.40	25.74

TABLE III  
EXAMPLE IMAGES ILLUSTRATE PREDICTIONS OF THE SDLGM ON THE NUMBER OF APPEARANCES OF DIGITS. THE GREEN MARKS THE CORRECT PREDICTIONS.

Image	Prediction	Image	Prediction
	Digit '3' ×1 Digit '4' ×1		Digit '4' ×2
	Digit '3' ×1 Digit '5' ×1		Digit '5' ×1 Digit '7' ×1 Digit '9' ×1
	Digit '3' ×1 Digit '9' ×1		Digit '2' ×1 Digit '3' ×1

to each observation to explicitly model the sparsity of the latent codes. With the proposed method, we are able to learn sparse representation under a certain threshold of sparsity. For inference, we develop an amortized variational method based on recent advances in Bayesian learning. The SDLGM is also

tailored for the multi-class multi-label classification issue. The experiment results on multiple datasets show the benefits of the proposed model. Our work provides interesting avenues for future work, e.g. adapting the SDLGM to learn sparse attention for text and visual data, investigating the effectiveness of the

proposed method for large scale datasets. In addition, we will extend the SDLGM to learn sparse representations of measurement data of telecommunication networks (collected in the H2020 MonB5G project). The measurement data is important for the network operators to monitor the status of the networks, however it introduces significant storage and communication overhead to the network management platforms. Compressing the data to sparse representations improves data storage efficiency, computation cost and communication constraints. The proposed method SDLGM provides an effective solution for data compression with extra flexibility. In particular, it allows clear definition of sparsity degree: the number of non-zero features of a representation cannot be larger than a pre-defined threshold  $L_0$ . We will validate the effectiveness of the the SDLGM method in the network measurement data, especially with respect to the tradeoff between the performance of network status/KPIs prediction and the sparsity of the data representations.

## REFERENCES

- [1] D. Kingma and M. Welling, "Auto-encoding variational bayes," in *ICLR*, 2014.
- [2] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," in *ICML*, 2014.
- [3] K. Gregor, I. Danihelka, A. Mnih, C. Blundell, and D. Wierstra, "Deep autoregressive networks," in *ICML*, 2014.
- [4] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *NIPS*, 2015.
- [5] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," in *ICLR*, 2017.
- [6] Y. Miao, L. Yu, and P. Blunsom, "Neural variational inference for text processing," in *ICML*, 2016.
- [7] J. Walker, C. Doersch, A. Gupta, and M. Hebert, "An uncertain future: Forecasting from static images using variational autoencoders," in *ECCV*, 2016.
- [8] H. Xu, W. Chen, N. Zhao, Z. Li, J. Bu, Z. Li, Y. Liu, Y. Zhao, D. Pei, Y. Feng, J. Chen, Z. Wang, and H. Qiao, "Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications," in *WWW*, 2018.
- [9] X. Yan, J. Yang, K. Sohn, and H. Lee, "Attribute2image: Conditional image generation from visual attributes," in *ECCV*, 2016.
- [10] R. Yeh, Z. Liu, D. Goldman, and A. Agarwala, "Semantic facial expression editing using autoencoded flow," 2016.
- [11] C. Maddison, A. Mnih, and Y. W. Teh, "The concrete distribution: A continuous relaxation of discrete random variables," in *ICLR*, 2017.
- [12] G. Tucker, A. Mnih, C. Maddison, D. Lawson, and J. Sohl-Dickstein, "Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models," in *NIPS*, 2017.
- [13] J. Denker and Y. LeCun, "Transforming neural-net output levels to probability distributions," in *NIPS*, 1991.
- [14] W. Buntine and A. Weigend, "Bayesian back-propagation," *Complex Systems*, vol. 5, 1991.
- [15] R. Neal, "Bayesian learning via stochastic dynamics," in *NIPS*, 1993.
- [16] M. Hoffman, D. Blei, C. Wang, and J. Paisley, "Stochastic variational inference," *JMLR*, pp. 1303–1347, 2013.
- [17] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient langevin dynamics," in *ICML*, 2011.
- [18] A. Graves, "Practical variational inference for neural networks," in *NIPS*, 2011.
- [19] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight uncertainty in neural networks," in *ICML*, 2015.
- [20] D. Kingma, D. Rezende, S. Mohamed, and M. Welling, "Semi-supervised learning with deep generative models," in *NIPS*, 2014.
- [21] T. Kipf and M. Welling, "Variational graph autoencoders," in *Proceedings of the Bayesian Deep Learning Workshop*, 2016.
- [22] Y. Pu, Z. Gan, R. Henao, X. Yuan, C. Li, A. Stevens, and L. Carin, "Variational autoencoder for deep learning of images, labels and captions," in *NIPS*, 2016.
- [23] H. Lee, C. Ekanadham, and A. Ng, "Sparse deep belief net model for visual area v2," in *NIPS*, 2008.
- [24] M. A. Ranzato, Y.-L. Boureau, and Y. LeCun, "Sparse feature learning for deep belief networks," in *NIPS*, 2008.
- [25] K. Kavukcuoglu, M. A. Ranzato, and Y. LeCun, "Fast inference in sparse coding algorithms with applications to object recognition," 2010.
- [26] A. Makhzani and B. Frey, "K-sparse autoencoders," in *ICLR*, 2014.
- [27] D. Arpit, Y. Zhou, H. Ngo, and V. Govindaraju, "Why regularized autoencoders learn sparse representation," in *ICML*, 2016.
- [28] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," in *AISTATS*, 2011.
- [29] I. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio, "Maxout networks," in *ICML*, 2013.
- [30] A. Martins and R. Astudillo, "From softmax to sparsemax: A sparse model of attention and multi-label classification," in *ICML*, 2016.
- [31] D. Molchanov, A. Ashukha, and D. Vetrov, "Variational dropout sparsifies deep neural networks," in *ICML*, 2017.
- [32] C. Louizos, M. Welling, and D. Kingma, "Learning sparse neural networks through l0 regularization," in *ICLR*, 2018.
- [33] S. Srinivas, A. Subramanya, and V. Babu, "Training sparse neural networks," 2016.
- [34] W. Li, J. Mao, Y. Zhang, and S. Cui, "Fast similarity search via optimal sparse lifting," in *NIPS*, 2018.
- [35] Y. Jeong and H. O. Song, "Efficient end-to-end learning for quantizable representations," in *ICML*, 2018.
- [36] B. Paria, C. Yeh, I. Yen, N. Xu, P. Ravikumar, and B. Póczos, "minimizing flops to learn efficient sparse representations," in *ICLR*, 2020.
- [37] S. J. Gershman and N. D. Goodman, "Amortized inference in probabilistic reasoning," *CogSci*, vol. 36, 2014.
- [38] C. Zhang, J. Butepage, H. Kjellstrom, and S. Mandt, "Advances in variational inference," *Pattern Analysis and Machine Intelligence*, 2019.
- [39] P. Dayan, G. Hinton, R. Neal, and R. Zemel, "The helmholtz machine," *Neural computation*, vol. 7, no. 5, 1995.
- [40] Z. Ghahramani and M. Beal, "Propagation algorithms for variational bayesian learning," in *NIPS*, 2001.
- [41] E. Gumbel, *Statistical theory of extreme values and some practical applications: a series of lectures*. U. S. Govt. Print. Office, 1954.
- [42] A. Ng and M. Jordan, "On discriminative vs generative classifiers a comparison of logistic regression and naive bayes," in *NIPS*, 2002.