
A Deep Generative Deconvolutional Image Model

Yunchen Pu
Duke University

Xin Yuan
Bell Labs

Andrew Stevens
Duke University

Chunyuan Li
Duke University

Lawrence Carin
Duke University

Abstract

A deep generative model is developed for representation and analysis of images, based on a hierarchical convolutional dictionary-learning framework. Stochastic *unpooling* is employed to link consecutive layers in the model, yielding top-down image generation. A Bayesian support vector machine is linked to the top-layer features, yielding max-margin discrimination. Deep deconvolutional inference is employed when testing, to infer the latent features, and the top-layer features are connected with the max-margin classifier for discrimination tasks. The model is efficiently trained using a Monte Carlo expectation-maximization (MCEM) algorithm; the algorithm is implemented on graphical processor units (GPU) to enable large-scale learning, and fast testing. Excellent results are obtained on several benchmark datasets, including ImageNet, demonstrating that the proposed model achieves results that are highly competitive with similarly sized convolutional neural networks.

1 Introduction

Convolutional neural networks (CNN) (LeCun et al., 1989) are effective tools for image and video analysis (Chatfield et al., 2014; Krizhevsky et al., 2012; Mnih et al., 2013; Sermanet et al., 2013). The CNN is characterized by feedforward (bottom-up) sequential application of convolutional filterbanks, pointwise nonlinear functions (e.g., sigmoid or hyperbolic tangent), and pooling. Supervision in CNN is typically implemented via a fully-connected layer at the top of the deep architecture, usually with a softmax classifier (Krizhevsky et al., 2012).

A parallel line of research concerns dictionary learning

(Mairal et al., 2008; Zhang and Li, 2010; Zhou et al., 2012) based on a set of image *patches*. In this setting one imposes sparsity constraints on the dictionary weights with which the data are represented. For image analysis/processing tasks, rather than using a patch-based model, there has been recent interest in deconvolutional networks (DN) (Chen et al., 2011, 2013; Zeiler et al., 2010). In a DN one uses dictionary learning on an entire image (as opposed to the patches of an image), and each dictionary element is convolved with a sparse set of weights that exist across the entire image. Such models are termed “deconvolutional” because, given a learned dictionary, the features at test are found through deconvolution. One may build deep deconvolutional models, which typically employ a pooling step like the CNN (Chen et al., 2011, 2013). The convolutional filterbank of the CNN is replaced in the DN by a library of convolutional dictionaries.

In this paper we develop a new deep generative model for images, based on convolutional dictionary learning. At test, after the dictionary elements are learned, deconvolutional inference is employed, like in the aforementioned DN research. The proposed method is related to Chen et al. (2011, 2013), but a complete top-down generative model is developed, with stochastic *unpooling* connecting model layers (this is distinct from almost all other models, which employ bottom-up pooling). Chen et al. (2011, 2013) trained each layer separately, sequentially, with no final coupling of the overall model (significantly undermining classification performance). Further, in Chen et al. (2011, 2013) Bayesian posterior inference was approximated for all model parameters (e.g., via Gibbs sampling), which scales poorly. Here we employ Monte Carlo expectation maximization (MCEM) (Wei and Tanner, 1990), with a point estimate learned for the dictionary elements and the parameters of the classifier, allowing learning on large-scale data and fast testing.

Forms of stochastic pooling have been applied previously (Lee et al., 2009; Zeiler and Fergus, 2013). Lee et al. (2009) defined stochastic pooling in the context of an energy-based Boltzmann machine, and Zeiler and Fergus (2013) proposed stochastic pooling as a regularization technique. Here *unpooling* is employed, yielding a top-down generative process.

To impose supervision, we employ the Bayesian support vector machine (SVM) (Polson and Scott, 2011), which has been used for supervised dictionary learning (Henao et al., 2014; Yuan et al., 2015) (but not previously for deep learning). The proposed generative model is amenable to Bayesian analysis, and here the Bayesian SVM is learned simultaneously with the deep model. The models in Donahue et al. (2014); He et al. (2014); Zeiler and Fergus (2014) do not train the SVM jointly, as we do – instead, the SVM is trained separately using the learned CNN features (with CNN supervised learning implemented via softmax).

This paper makes several contributions: (i) A new deep model is developed for images, based on convolutional dictionary learning; this model is a generative form of the earlier DN. (ii) A new stochastic unpooling method is proposed, linking consecutive layers of the deep model. (iii) An SVM is integrated with the top layer of the model, enabling max-margin supervision during training. (iv) The algorithm is implemented on a GPU, for large-scale learning and fast testing; we demonstrate state-of-the-art classification results on several benchmark datasets, and demonstrate scalability through analysis of the ImageNet dataset.

2 Supervised Deep Deconvolutional Model

2.1 Single layer convolutional dictionary learning

Consider N images $\{\mathbf{X}^{(n)}\}_{n=1}^N$, with $\mathbf{X}^{(n)} \in \mathbb{R}^{N_x \times N_y \times N_c}$, where N_x and N_y represent the number of pixels in each spatial dimension; $N_c = 1$ for gray-scale images and $N_c = 3$ for RGB images. We start by relating our model to optimization-based dictionary learning and DN, the work of Mairal et al. (2008); Zhang and Li (2010) and Zeiler et al. (2010), respectively. The motivations for and details of our model are elucidated by making connections to this previous work. Specifically, consider the optimization problem

$$\{\hat{\mathbf{D}}^{(k)}, \hat{\mathbf{S}}^{(n,k)}\} = \operatorname{argmin} \left\{ \sum_{n=1}^N \left\| \mathbf{X}^{(n)} - \sum_{k=1}^K \mathbf{D}^{(k)} * \mathbf{S}^{(n,k)} \right\|_F^2 + \lambda_1 \sum_{k=1}^K \|\mathbf{D}^{(k)}\|_F^2 + \lambda_2 \sum_{n=1}^N \sum_{k=1}^K \|\mathbf{S}^{(n,k)}\|_1 \right\}, \quad (1)$$

where $*$ is the 2D (spatial) convolution operator. Each $\mathbf{D}^{(k)} \in \mathbb{R}^{n_x \times n_y \times N_c}$ and typically $n_x \ll N_x$, $n_y \ll N_y$. The spatially-dependent weights $\mathbf{S}^{(n,k)}$ are of size $(N_x - n_x + 1) \times (N_y - n_y + 1)$. Each of the N_c layers of $\mathbf{D}^{(k)}$ are spatially convolved with $\mathbf{S}^{(n,k)}$, and after summing over the K dictionary elements, this manifests an approximation for each of the N_c layers of $\mathbf{X}^{(n)}$.

The form of (1) is as in Mairal et al. (2008), with the ℓ_1 norm on $\mathbf{S}^{(n,k)}$ imposing sparsity, and with the Frobenius norm on $\mathbf{D}^{(k)}$ (ℓ_2 in Mairal et al. (2008)) imposing an expected-energy constraint on each dictionary element; in Mairal et al. (2008) convolution is not used, but otherwise

the model is identical, and the computational methods developed in Mairal et al. (2008) may be applied.

The form of (1) motivates choices for the priors in the proposed generative model. Specifically, consider

$$\mathbf{X}^{(n)} = \sum_{k=1}^K \mathbf{D}^{(k)} * \mathbf{S}^{(n,k)} + \mathbf{E}^{(n)}, \quad (2)$$

$$\mathbf{E}^{(n)} \sim \mathcal{N}(0, \gamma_e^{-1} \mathbf{I}), \quad \mathbf{D}^{(k)} \sim \mathcal{N}(0, \mathbf{I}) \quad (3)$$

with $S_{i,j}^{(n,k)}$ denoting element (i, j) of $\mathbf{S}^{(n,k)}$, drawn $S_{i,j}^{(n,k)} \sim \text{Laplace}(0, b) = \frac{1}{2b} \exp(-|S_{i,j}^{(n,k)}|/b)$. We have “vectorized” the matrices $\mathbf{E}^{(n)}$ and $\mathbf{D}^{(k)}$ (from the standpoint of the distributions from which they are drawn), and \mathbf{I} is an appropriately sized identity matrix. The maximum *a posteriori* (MAP) solution to (3), with the Laplace prior imposed independently on each component of $\mathbf{S}^{(n,k)}$, corresponds to the optimization problem in (1), and the hyperparameters γ_e and b play roles analogous to λ_1 and λ_2 .

The sparsity of $\mathbf{S}^{(n,k)}$ manifested in (1) is a consequence of the geometry imposed by the ℓ_1 operator; the MAP solution is sparse, but, with probability one, any draw from the Laplace prior on $\mathbf{S}^{(n,k)}$ is not sparse (Cevher, 2009). To impose sparsity on $\mathbf{S}^{(n,k)}$ within the generative process, we consider the spike-slab (Ishwaran and Rao, 2005) prior:

$$\begin{aligned} S_{i,j}^{(n,k)} &\sim [z_{i,j}^{(n,k)} \mathcal{N}(0, \gamma_s^{-1}) + (1 - z_{i,j}^{(n,k)}) \delta_0], \\ z_{i,j}^{(n,k)} &\sim \text{Bern}(\pi^{(n,k)}), \quad \pi^{(n,k)} \sim \text{Beta}(a_0, b_0) \end{aligned} \quad (4)$$

where $z_{i,j}^{(n,k)} \in \{0, 1\}$, δ_0 is a unit point measure concentrated at zero, and (a_0, b_0) are set to encourage that most $\pi^{(n,k)}$ are small (Paisley and Carin, 2009), *i.e.*, $a_0 = 1/K$ and $b_0 = 1 - 1/K$. For parameters γ_e and γ_s we impose the priors $\gamma_s \sim \text{Gamma}(a_s, b_s)$ and $\gamma_e \sim \text{Gamma}(a_e, b_e)$, with hyperparameters $a_s = b_s = a_e = b_e = 10^{-6}$ to impose diffuse priors (Tipping, 2001).

2.2 Generative Deep Model via Stochastic Unpooling

The model in (2) is motivated by the idea that each image $\mathbf{X}^{(n)}$ may be represented in terms of convolutional dictionary elements $\mathbf{D}^{(k)}$ that are shared across all N images. In the proposed deep model, we similarly are motivated by the idea that the feature maps $\mathbf{S}^{(n,k)}$ may also be represented in terms of convolutions of (distinct) dictionary elements. Consider a two-layer model, with

$$\mathbf{X}^{(n,2)} = \sum_{k_2=1}^{K_2} \mathbf{D}^{(k_2,2)} * \mathbf{S}^{(n,k_2,2)} \quad (5)$$

$$\mathbf{S}^{(n,k_1,1)} \sim \text{unpool}(\mathbf{X}^{(n,k_1,2)}), \quad k_1 = 1, \dots, K_1 \quad (6)$$

$$\mathbf{X}^{(n,1)} = \sum_{k_1=1}^{K_1} \mathbf{D}^{(k_1,1)} * \mathbf{S}^{(n,k_1,1)} + \mathbf{E}^{(n)} \quad (7)$$




Figure 1: Demonstration of the stochastic unpooling: one 2×2 pooling block is depicted.

where $\mathbf{X}^{(n,1)} = \mathbf{X}^{(n)}$. Dictionary elements $\mathbf{D}^{(k_1,1)}$ replace $\mathbf{D}^{(k)}$ in (2), for representation of $\mathbf{X}^{(n)}$. The weights $\mathbf{S}^{(n,k_1,1)}$ are connected to $\mathbf{X}^{(n,2)}$ via the stochastic operation $\text{unpool}(\cdot)$, detailed below. Motivated as discussed above, $\mathbf{X}^{(n,2)}$ is represented in terms of convolutions with second-layer dictionary elements $\mathbf{D}^{(k_2,2)}$. The forms of the priors on $\mathbf{D}^{(k_1,1)}$ and $\mathbf{D}^{(k_2,2)}$ are as above for $\mathbf{D}^{(k)}$, and the prior on $\mathbf{E}^{(n)}$ is unchanged.

The tensor $\mathbf{X}^{(n,2)} \in \mathbb{R}^{N_x^{(2)} \times N_y^{(2)} \times K_1}$ with layer/slice k_1 denoted by the matrix $\mathbf{X}^{(n,k_1,2)} \in \mathbb{R}^{N_x^{(2)} \times N_y^{(2)}} (k_1 \in \{1, \dots, K_1\})$. Matrix $\mathbf{X}^{(n,k_1,2)}$ is a *pooled* version of $\mathbf{S}^{(n,k_1,1)}$. Specifically, $\mathbf{S}^{(n,k_1,1)}$ is partitioned into contiguous spatial pooling blocks, each pooling block of dimension $p_x^{(1)} \times p_y^{(1)}$, with $N_x^{(2)} = N_x/p_x^{(1)}$ and $N_y^{(2)} = N_y/p_y^{(1)}$ (assumed to be integers). Each pooling block of $\mathbf{S}^{(n,k_1,1)}$ is all-zeros except one non-zero element, with the non-zero element defined in $\mathbf{X}^{(n,k_1,2)}$. Specifically, element (i, j) of $\mathbf{X}^{(n,k_1,2)}$, denoted $X_{i,j}^{(n,k_1,2)}$, is mapped to pooling block (i, j) in $\mathbf{S}^{(n,k_1,1)}$, denoted $\mathbf{S}_{i,j}^{(n,k_1,1)}$.

Let $\mathbf{z}_{i,j}^{(n,k_1,1)} \in \{0, 1\}^{p_x^{(1)} p_y^{(1)}}$ be a vector of all zeros, and a single one, and the location of the non-zero element of $\mathbf{z}_{i,j}^{(n,k_1,1)}$ identifies the location of the single non-zero element of pooling block $\mathbf{S}_{i,j}^{(n,k_1,1)}$ which is set as $X_{i,j}^{(n,k_1,2)}$. The function $\text{unpool}(\cdot)$ is a stochastic operation that defines $\mathbf{z}_{i,j}^{(n,k_1,1)}$, and hence the way $\mathbf{X}^{(n,k_1,2)}$ is *unpooled* to constitute the *sparse* $\mathbf{S}^{(n,k_1,1)}$. We impose

$$\mathbf{z}_{i,j}^{(n,k_1,1)} \sim \text{Mult}(1, \boldsymbol{\theta}^{(n,k_1,1)}) \quad (8)$$

$$\boldsymbol{\theta}^{(n,k_1,1)} \sim \text{Dir}(1/(p_x^{(1)} p_y^{(1)})) \quad (9)$$

where $\text{Dir}(\cdot)$ denotes the symmetric Dirichlet distribution; the Dirichlet distribution has a *set* of parameters, and here they are all equal to the value indicated in $\text{Dir}(\cdot)$.

When introducing the above two-layer model, $\mathbf{S}^{(n,k_2,2)}$ is drawn from a spike-slab prior, as in (4). However, we may extend this to a three-layer model, with pooling blocks defined in $\mathbf{S}^{(n,k_2,2)}$. A convolutional dictionary representation is similarly constituted for $\mathbf{X}^{(n,3)}$, and this is stochastically unpooled to generate $\mathbf{S}^{(n,k_2,2)}$. This may continued for L layers, where the hierarchical convolutional dictionary learning learns multi-scale structure in the weights on

the dictionary elements. At the top layer in the L -layer model, the weights $\mathbf{S}^{(n,k_L,L)}$ are drawn from a spike-slab prior of the form in (4).

Consider an L -layer model, and assume that $\{\{\mathbf{D}^{(k_l,l)}\}_{k_l=1}^{K_l}\}_{l=1}^L$ have been learned/specified. An image is generated by starting at the top, drawing $\{\mathbf{S}^{(n,k_L,L)}\}_{k_L=1}^{K_L}$ from a spike-slab model. Then $\{\mathbf{S}^{(n,k_{L-1},L-1)}\}_{k_{L-1}=1}^{K_{L-1}}$ are constituted by convolving $\{\mathbf{S}^{(n,k_L,L)}\}_{k_L=1}^{K_L}$ with $\{\mathbf{D}^{(k_L,L)}\}_{k_L=1}^{K_L}$, summing over the K_L dictionary elements, and then performing stochastic unpooling. This process of convolution and stochastic unpooling proceeds for L layers, ultimately yielding $\sum_{k_1=1}^{K_1} \mathbf{D}^{(k_1,1)} * \mathbf{S}^{(n,k_1,1)}$ at the bottom (first) layer. With the added stochastic residual $\mathbf{E}^{(n)}$, the image $\mathbf{X}^{(n)}$ is specified.

We note an implementation detail that has been found useful in experiments. In (9), the unpooling was performed such that each pooling block in $\mathbf{S}^{(n,k_l,l)}$ has a single non-zero element, with the non-zero element defined in $\mathbf{X}^{(n,k_l,l+1)}$. The unpooling for block (i, j) was specified by the $p_x^{(l)} p_y^{(l)}$ -dimensional $\mathbf{z}_{i,j}^{(n,k_l,l)}$ vector of all-zeros and a single one. In our slightly modified implementation, we have considered a $(p_x^{(l)} p_y^{(l)} + 1)$ -dimensional $\mathbf{z}_{i,j}^{(n,k_l,l)}$, which is again all zeros with a single one, and $\boldsymbol{\theta}^{(n,k_l,l)}$ is also $p_x^{(l)} p_y^{(l)} + 1$ dimensional. If the single one in $\mathbf{z}_{i,j}^{(n,k_l,l)}$ is located among the first $p_x^{(l)} p_y^{(l)}$ elements of $\mathbf{z}_{i,j}^{(n,k_l,l)}$, then the location of this non-zero element identifies the location of the single non-zero element in the (i, j) pooling block, as before. However, if the non-zero element of $\mathbf{z}_{i,j}^{(n,k_l,l)}$ is in position $p_x^{(l)} p_y^{(l)} + 1$, then *all* elements of pooling block (i, j) are set to zero. This imposes further sparsity on the feature maps and, as demonstrated in the Supplementary Material (SM), it yields a model in which the elements of the feature map that are relatively small are encouraged to be zero. This turning off of dictionary elements with small weights is analogous to *dropout* (Srivastava et al., 2014), and in our model it also has been found to yield slightly better classification performance.

2.3 Supervision via Bayesian SVMs

Assume that a label $\ell_n \in \{1, \dots, C\}$ is associated with each of the N images, so that the training set may be denoted $\{(\mathbf{X}^{(n)}, \ell_n)\}_{n=1}^N$. We wish to learn a classifier that maps the *top-layer* dictionary weights $\mathbf{S}^{(n,L)} = \{\mathbf{S}^{(n,k_L,L)}\}_{k_L=1}^{K_L}$ to an associated label ℓ_n . The $\mathbf{S}^{(n,L)}$ are “unfolded” into the vector \mathbf{s}_n . We desire the classifier mapping $\mathbf{s}_n \rightarrow \ell_n$ and our goal is to learn the dictionary and classifier *jointly*.

We design C one-versus-all binary SVM classifiers. For each of these classifiers, the problem may be posed as training with $\{\mathbf{s}_n, y_n^{(\ell)}\}_{n=1}^N$, where \mathbf{s}_n are the top-layer dic-

tionary weights, as discussed above, and $y_n^{(\ell)} \in \{-1, 1\}$ (a bias term is also appended to each \mathbf{s}_n , as is typical of SVMs). If $\ell_n = \ell \in \{1, \dots, C\}$ then $y_n^{(\ell)} = 1$, and $y_n^{(\ell)} = -1$ otherwise; the indicator ℓ specifies which of the C binary SVMs is under consideration. For notational simplicity, we omit the superscript (ℓ) for the remainder of the section, and consider the Bayesian SVM for one of the binary learning tasks, with labeled data $\{\mathbf{s}_n, y_n\}_{n=1}^N$. In practice, C such binary classifiers are learned jointly, and the value of $y_n \in \{1, -1\}$ depends on which one-versus-all classifier is being specified.

Given a feature vector \mathbf{s} , the goal of the SVM is to find an $f(\mathbf{s})$ that minimizes the objective function

$$\gamma \sum_{n=1}^N \max(1 - y_n f(\mathbf{s}_n), 0) + R(f(\mathbf{s})), \quad (10)$$

where $\max(1 - y_n f(\mathbf{s}_n), 0)$ is the hinge loss, $R(f(\mathbf{s}))$ is a regularization term that controls the complexity of $f(\mathbf{s})$, and γ is a tuning parameter controlling the trade-off between error penalization and the complexity of the classification function. The decision boundary is defined as $\{\mathbf{s} : f(\mathbf{s}) = 0\}$ and $\text{sign}(f(\mathbf{s}))$ is the decision rule, classifying \mathbf{s} as either -1 or 1 (Vapnik, 1995).

Recently, Polson and Scott (2011) showed that for the linear classifier $f(\mathbf{s}) = \beta^T \mathbf{s}$, minimizing (10) is equivalent to estimating the mode of the pseudo-posterior of β

$$p(\beta | \mathbf{S}, \mathbf{y}, \gamma) \propto \prod_{n=1}^N \mathcal{L}(y_n | \mathbf{s}_n, \beta, \gamma) p(\beta | \cdot), \quad (11)$$

where $\mathbf{y} = [y_1 \dots y_N]^T$, $\mathbf{S} = [\mathbf{s}_1 \dots \mathbf{s}_N]$, $\mathcal{L}(y_n | \mathbf{s}_n, \beta, \gamma)$ is the pseudo-likelihood function, and $p(\beta | \cdot)$ is the prior distribution for the vector of coefficients β . Choosing β to maximize the log of (11) corresponds to (10), where the prior is associated with $R(f(\mathbf{s}))$. Polson and Scott (2011) showed that $\mathcal{L}(y_n | \mathbf{s}_n, \beta, \gamma)$ admits a location-scale mixture of normals representation by introducing latent variables λ_n , such that

$$\begin{aligned} \mathcal{L}(y_n | \mathbf{s}_n, \beta, \gamma) &= e^{-2\gamma \max(1 - y_n \beta^T \mathbf{s}_n, 0)} \\ &= \int_0^\infty \frac{\sqrt{\gamma}}{\sqrt{2\pi\lambda_n}} \exp\left(-\frac{(1 + \lambda_n - y_n \beta^T \mathbf{s}_n)^2}{2\gamma^{-1}\lambda_n}\right) d\lambda_n. \end{aligned} \quad (12)$$

Note that the exponential in (12) is Gaussian wrt β . As described in Polson and Scott (2011), this encourages data augmentation for variable λ_n (λ_n is treated as a new random variable), which permits efficient Bayesian inference (see Polson and Scott (2011); Heno et al. (2014) for details). One of the benefits of a Bayesian formulation for SVMs is that we can flexibly specify the behavior of β while being able to adaptively regularize it by specifying a prior $p(\gamma)$ as well.

We impose shrinkage (near sparsity) (Polson and Scott, 2010) on β using the Laplace distribution; letting β_i denote i^{th} element of β , we impose

$$\beta_i \sim \mathcal{N}(0, \omega_i), \quad \omega_i \sim \text{Exp}(\kappa), \quad \kappa \sim \text{Gamma}(a_\kappa, b_\kappa), \quad (13)$$

and similar to κ and λ_n , a diffuse Gamma prior is imposed on γ .

For the generative process of the overall model, activation weights \mathbf{s}_n are drawn at layer L , as discussed in Sec. 2.2. These weights then go into the C -class SVM, and from that a class label is manifested. Specifically, each SVM learns a linear function of $\{\beta_\ell^T \mathbf{s}\}_{\ell=1}^C$, and for a given data \mathbf{s} , its class label is defined by (Yang et al., 2009):

$$\ell_n = \underset{\ell}{\text{argmax}} \beta_\ell^T \mathbf{s}_n. \quad (14)$$

The set of vectors $\{\beta_\ell\}_{\ell=1}^C$, connecting the top-layer features \mathbf{s} to the classifier, play a role analogous to the fully-connected layer in the softmax-based CNN, but here we constitute supervision via the max-margin SVM. Hence, the proposed model is a generative construction for both the labels and the images.

3 Model Training

The previous section described a supervised deep generative model for images, based on deep convolutional dictionary learning, stochastic unpooling, and the Bayesian SVM. The conditional posterior distribution for each model parameter can be written in closed form, assuming the other model parameters are fixed (see the SM). For relatively small datasets we can therefore employ a Gibbs sampler for both training and deconvolutional inference, yielding an approximation to the posterior distribution on all parameters. Large-scale datasets prohibit the application of standard Gibbs sampling. For large data we use stochastic MCEM (Wei and Tanner, 1990) to find a maximum *a posteriori* (MAP) estimate of the model parameters.

We consolidate the ‘‘local’’ model parameters (latent data-sample-specific variables) as

$$\Phi_n = (\{z^{(n,l)}\}_{l=1}^L, \mathbf{S}^{(n,L)}, \gamma_s^{(n)}, \mathbf{E}^{(n)}, \{\lambda_n^{(\ell)}\}_{\ell=1}^C),$$

the ‘‘global’’ parameters (shared across all data) as $\Psi = (\{\mathbf{D}^{(l)}\}_{l=1}^L, \beta)$, and the data as $\mathbf{Y}_n = (\mathbf{X}^{(n)}, \ell_n)$. We desire a MAP estimator

$$\Psi_{\text{MAP}} = \underset{\Psi}{\text{argmax}} \sum_n \ln p(\Psi | \mathbf{Y}_n), \quad (15)$$

which can be interpreted as an EM problem:

E-step: Perform an expectation with respect to the local variables, using $p(\Phi_n | \mathbf{Y}_n, \Psi_{t-1}) \forall n$, where Ψ_{t-1} is the estimate of the global parameters from iteration $(t-1)$.

M-step: Maximize $\ln p(\Psi) + \sum_n \mathbb{E}_{\Phi_n} [\ln p(\mathbf{Y}_n | \Phi_n, \Psi)]$ with respect to Ψ .

We approximate the expectation via Monte Carlo sampling, which gives

$$Q(\Psi | \Psi_{t-1}) = \ln p(\Psi) + \frac{1}{N_s} \sum_{s=1}^{N_s} \sum_n \ln p(\mathbf{Y}_n | \Phi_n^s, \Psi), \quad (16)$$

where Φ_n^s is a sample from the full conditional posterior distribution, and N_s is the number of samples; we seek to maximize $Q(\Psi|\Psi_{t-1})$ wrt Ψ , constituting Ψ_t . Recall from above that each of the conditional distributions in a Gibbs sampler of the model is analytic; this allows convenient sampling of local parameters, conditioned on specified global parameters Ψ_{t-1} , and therefore the aforementioned sampling is implemented efficiently (using mini-batches of data, where $\mathcal{I}_t \subset \{1, \dots, N\}$ identifies the stochastically defined subset of data in mini-batch t). An approximation to the M-step is implemented via stochastic gradient descent (SGD). The stochastic MCEM gradient at iteration t is

$$\nabla_{\Psi} Q = \nabla_{\Psi} \ln p(\Psi) + \frac{1}{N_s} \sum_{s=1}^{N_s} \sum_{n \in \mathcal{I}_t} \nabla_{\Psi} \ln p(\mathbf{Y}_n | \Phi_n^s, \Psi). \quad (17)$$

We solve (15) using RMSprop (Dauphin et al., 2015; Tieleman and Hinton, 2012) with the gradient approximation in (17).

In the learning phase, the MCEM method is used to learn a point estimate for the global parameters Ψ . During testing, we follow the same MCEM setup with $\Phi^{\text{test}} = (\{\mathbf{z}^{(*,l)}\}_{l=1}^{L-1}, \gamma_s^{(*)}, \mathbf{E}^{(*)})$, $\Psi^{\text{test}} = \mathbf{S}^{(*,L)}$, when given a new image \mathbf{X}^* . We find a MAP estimator:

$$\Psi_{\text{MAP}}^{\text{test}} = \underset{\Psi^{\text{test}}}{\operatorname{argmax}} \ln p(\Psi^{\text{test}} | \mathbf{X}^*, \mathbf{D}), \quad (18)$$

using MCEM (gradient wrt Ψ^{test}). In this form of the MCEM, all data-dependent latent variables Φ^{test} are integrated (summed) out in the expectation, except for the top-layer feature map Ψ^{test} , for which the gradient descent M step yields a point estimate. The top-layer features are then sent to the trained SVM to predict the label. Details for training and inference are provided in the SM.

4 Experimental Results

We present results for the MNIST, CIFAR-10 & 100, Caltech 101 & 256 and ImageNet 2012 datasets. The same hyperparameter settings (discussed at the end of Section 2.1) were used in all experiments; no tuning was required between datasets.

For the first five (small/modest-sized) datasets, the model is learned via Gibbs sampling. We found that it is effective to use layer-wise pretraining as employed in some deep generative models (Erhan et al., 2010; Hinton and Salakhutdinov, 2006). The pretraining is performed sequentially from the bottom layer (touching the data), to the top layer, in an unsupervised manner. Details on the layerwise pretraining are discussed in the SM. In the pretraining step, we average 500 collection samples, to obtain parameter values (e.g., dictionary elements) after first discarding 1000 burn-in samples. Following pre-training, we refine the entire model jointly using the complete set of Gibbs conditional distributions.

1000 burn-in iterations are performed followed by 500 collection draws, retaining one of every 50 iterations. During testing, the predictions are based on averaging the decision values of the collected samples.

For each of these first five datasets, we show three classification results, using part of or all of our model (to illustrate the role of each component): 1) *Pretraining only*: this model (in an unsupervised manner) is used to extract features and the features are sent to a separate linear SVM, yielding a 2-step procedure. 2) *Unsupervised model*: this model includes the deep generative developed in Sec. 2.2, but is also trained in an unsupervised manner (this is the unsupervised model after refinement). The features extracted by this model are sent to a separate linear SVM, and therefore this is also a 2-step procedure. 3) *Supervised model*: this is the complete refined supervised model developed in Sec. 2.2 and Sec. 2.3.

ImageNet 2012 is used to assess the scalability of our model to large datasets. In this case, we learn the supervised model initialized from the priors (without layer-wise pretraining). The proposed online learning method, MCEM, based on RMSProp (Dauphin et al., 2015; Tieleman and Hinton, 2012), is developed for both training and inference with mini-batch size 256 and decay rate 0.95. Our implementation of MCEM learning is based on the publicly available CUDA C++ Caffe toolbox (August 2015 branch) (Jia et al., 2014), but contains significant modifications for our model. Our model takes around one week to train on ImageNet 2012 using a nVidia GeForce GTX TITAN X GPU with 12GB memory. Testing for the validation set of ImageNet 2012 (50K images) takes less than 12 minutes. In the subsequent tables providing classification results, the best results achieved by our model are bold.

4.1 MNIST

The MNIST data (<http://yann.lecun.com/exdb/mnist/>) has 60,000 training and 10,000 testing images, each 28×28 , for digits 0 through 9. A two-layer model is used with dictionary element size 8×8 and 6×6 at the first and second layer, respectively. The pooling size is 3×3 ($p_x = p_y = 3$) and the number of dictionary elements at layers 1 and 2 are $K_1 = 39$ and $K_2 = 117$, respectively. These numbers of dictionary elements are obtained by setting the initial dictionary number to a relatively large value ($K_1 = 50$ and $K_2 = 200$) in the pretraining step and discarding infrequently used elements by counting the corresponding binary indicator z – effectively *inferring* the number of needed dictionary elements, as in Chen et al. (2011, 2013). Table 1 summarizes the classification results for MNIST. Our 2-layer supervised model outperforms most other modern approaches. The methods that outperforms ours are the complicated (6-layer) ConvNet model with elastic distortions (Ciresan et al., 2011) and the MCDNN, which combines several deep convolutional neural networks (Ciresan et al., 2012).

Table 1: Classification Error (%) for MNIST

Method	Test error
2-layer convnet (Jarrett et al., 2009)	0.53
Our pretrained model + SVM	1.42
Our unsupervised model + SVM	0.52
Our supervised model	0.37
6-layer convnet (Ciresan et al., 2011)	0.35
MCDNN (Ciresan et al., 2012)	0.23

Table 2: Average test log-likelihood (in nats) on MNIST

Method	Test log-likelihood
CMMVA (Li et al., 2015)	-95.86
Our unsupervised model	-89.44
Our supervised model	-93.49
DRAW (Gregor et al., 2015)	-84.13

To examine the performance of the proposed model, we plot a selection of top-layer dictionary elements (projected through the generative process down to the data plane) learned by our supervised model, on the right of Fig. 2, and on the left we show the corresponding elements inferred by our unsupervised model. It can be seen that the elements inferred by the supervised model are clearer (“unique” to a single number), whereas the elements learned by the unsupervised model are blurry (combinations of multiple numbers). Similar results were reported in Erhan et al. (2010).

We further investigate the generative capability of the proposed model via the log-likelihood, as shown in Table 2. It can be seen that our model outperforms the convolution-based model by Li et al. (2015), and is comparable to the state-of-art in Gregor et al. (2015).

The generated examples are shown in Fig. 3 (based on random draws of the top-layer weights). We also demonstrate the ability of the model to predict missing data (generative nature of the model); reconstructions are shown in Fig. 4. More results are provided in the SM.

4.2 CIFAR-10 & 100

The CIFAR-10 dataset (Krizhevsky and Hinton, 2009) is composed of 10 classes of natural 32×32 RGB images with 50000 images for training and 10000 images for testing. We apply the same preprocessing technique of global contrast normalization and ZCA whitening as used in the Maxout network (Goodfellow et al., 2013). A three-layer model is used with dictionary element size 5×5 , 5×5 , 4×4 at the first, second and third layer. The pooling sizes




Figure 2: Selected layer 2 (top-layer) dictionary elements of MNIST learned by the unsupervised model (left) and the supervised model (right).



Figure 3: Generated images using random dictionary weights.



Figure 4: Missing data interpolation of digits. For each subfigure: (top) Original data, (middle) Observed data, (bottom) Reconstruction.

Table 3: Classification Error (%) for CIFAR-10

Method	Test error
<i>Without Data Augmentation</i>	
Maxout (Goodfellow et al., 2013)	11.68
Network in Network (Lin et al., 2014)	10.41
Our pretrained + SVM	22.43
Our unsupervised + SVM	14.75
Our supervised model	10.39
Deeply-Supervised Nets (Lee et al., 2015)	9.69
<i>With Data Augmentation</i>	
Maxout (Goodfellow et al., 2013)	9.38
Network in Network (Lin et al., 2014)	8.81
Our pretrained + SVM	20.62
Our unsupervised + SVM	10.22
Our supervised	8.27
Deeply-Supervised Nets (Lee et al., 2015)	7.97

are both 2×2 and the numbers of dictionary elements for each layer are $K_1 = 48$, $K_2 = 128$ and $K_3 = 128$. If we *augment* the data by translation and horizontal flipping as used in other models (Goodfellow et al., 2013), we achieve 8.27% error. Our result is competitive with the state-of-art, which integrates supervision on every hidden layer (Lee et al., 2015). In contrast, we only impose supervision at the top layer. Table 3 summarizes the classification accuracy of our models and some related models.

The CIFAR-100 dataset (Krizhevsky and Hinton, 2009) is the same as CIFAR-10 in size and format, except it contains 100 classes. We use the same settings as in the CIFAR-10. Table 4 summarizes the classification accuracy of our model and some related models. It can be seen that our results (34.62%) are also very close to the state-of-the-art (34.57%) in Lee et al. (2015).

Table 4: Classification Error (%) for CIFAR-100

Method	Test error
Maxout (Goodfellow et al., 2013)	38.57
Network in Network (Lin et al., 2014)	35.68
Our pretrained + SVM (2 step)	77.25
Our unsupervised + SVM (2 step)	42.26
Our supervised model	34.62
Deeply-Supervised Nets (Lee et al., 2015)	34.57

Table 5: Classification Accuracy (%) for Caltech 101

Training images per class	15	30
<i>Without ImageNet Pretrain</i>		
5-layer Convnet (Zeiler and Fergus, 2014)	22.8	46.5
HBP-CFA (Chen et al., 2013)	58	65.7
R-KSVD (Li et al., 2013)	79	83
3-layer Convnet	62.3	72.4
Our pretrained + SVM (2 step)	43.24	53.57
Our unsupervised + SVM (2 step)	70.47	80.39
Our supervised model	75.37	87.82
<i>With ImageNet Pretrain</i>		
5-layer Convnet (Zeiler and Fergus, 2014)	83.8	86.5
5-layer Convnet (Chatfield et al., 2014)	-	88.35
Our supervised model	89.1	93.15
SPP-net (He et al., 2014)	-	94.11

4.3 Caltech 101 & 256

To balance speed and performance, we resize the images of Caltech 101 and Caltech 256 to 128×128 , followed by local contrast normalization (Jarrett et al., 2009). A three layer model is adopted. The dictionary element sizes are set to 7×7 , 5×5 and 5×5 , and the size of the pooling regions are 4×4 (layer 1 to layer 2) and 2×2 (layer 2 to layer 3).

The dictionary sizes for each layer are set to $K_1 = 48$, $K_2 = 84$ and $K_3 = 84$ for Caltech 101, and $K_1 = 48$, $K_2 = 128$ and $K_3 = 128$ for Caltech 256. Tables 5 and 6 summarize the classification accuracy of our model and some related models. Using only the data inside Caltech 101 and Caltech 256 (without using other datasets) for training, our results (87.82%, 66.4%) exceed the previous state-of-art results (83%, 58%) by a substantial margin (4%, 12.4%), which are the best results obtained by models without using deep convolutional models (using hand-crafted features).

As a baseline, we implemented the neural network consisting of three convolutional layers and two fully-connected layers with a final softmax classifier. The architecture of three convolutional layers is the same as our model. The fully-connected layers have 1024 neurons each. The results of neural network trained with dropout (Srivastava et al., 2014), after carefully parameter tuning, are also shown in Tables 5 and 6.

The state-of-the-art results on these two datasets are achieved by pretraining the deep network on a large dataset, ImageNet (Donahue et al., 2014; He et al., 2014; Zeiler and Fergus, 2014). We consider similar ImageNet pretraining in Sec. 4.4. We also observe from Table 6 that when there are fewer training images, our accuracy diminishes.

Table 6: Classification Accuracy (%) for Caltech 256

Training images per class	15	60
<i>Without ImageNet Pretrain</i>		
5-layer Convnet (Zeiler and Fergus, 2014)	9.0	38.8
Mu-SC (Bo et al., 2013)	42.7	58
3-layer Convnet	46.1	60.1
Our pretrained +SVM	13.4	38.2
Our unsupervised +SVM	40.7	60.9
Our supervised model	52.9	70.5
<i>With ImageNet Pretrain</i>		
5-layer Convnet (Zeiler and Fergus, 2014)	65	74.2
5-layer Convnet (Chatfield et al., 2014)	-	77.61
Our supervised model	67.0	77.9




Figure 5: Selected dictionary elements of our unsupervised model (top) and supervised model (bottom) trained from Caltech 256.

This verifies that the model complexity needs to be selected based on the size of the data. This is also consistent with the results reported by Zeiler and Fergus (2014), in which the classification performance is very poor without training the model on ImageNet.

Fig. 5 shows selected dictionary elements learned from the unsupervised and the supervised model, to illustrate the differences. It is observed that the dictionaries without supervision tend to reconstruct the data while the dictionary elements with supervision tend to extract features that will distinguish different classes. For example, the dictionaries




Figure 6: Generated image from the dictionaries trained from the “Faces_easy” category using random dictionary weights.



Figure 7: Missing data interpolation. (Top) Original data. (Middle) Observed data. (Bottom) Reconstruction.

learned with supervision have double sides on the image edges. Our model is generative, and as an example we generate images using the dictionaries trained from the “Faces easy” category, with random top-layer dictionary weights (see Fig. 6). Similar to the MNIST example, we also show in Fig. 7 the interpolation results of face data with half the image missing. Though the background is a little noisy, each face is recovered in great detail by the third (top) layer dictionaries. More results are provided in the SM.

4.4 ImageNet 2012

We train our model on the 1000-category ImageNet 2012 dataset, which consists of 1.3M/50K/100K training/validation/test images. Our training process follows the procedure of previous work (Howard, 2013; Krizhevsky et al., 2012; Zeiler and Fergus, 2014). The smaller image dimension is scaled to 256, and a 224×224 crop is chosen at 1024 random locations within the image. The data are augmented by color alteration and horizontal flips (Howard, 2013; Krizhevsky et al., 2012). A five layer convolutional model is employed ($L = 5$); the numbers (sizes) of dictionary elements for each layer are set to 96 (5×5), 256 (5×5), 512 (3×3), 1024 (3×3) and 512 (3×3); the pooling ratios are 4×4 (layer 1 to 2) and 2×2 (others). The number of parameters in our model is around 30 million.

We emphasize that our intention is not to directly compete with the best performance in the ImageNet challenge (Szegedy et al., 2015; Simonyan and Zisserman, 2015), which requires consideration of many additional aspects, but to provide a comparison on this dataset with a CNN with a similar network architecture (size). Table 7

summarizes our results compared with the “ZF”-net developed in Zeiler and Fergus (2014) which has a similar architecture with ours.

The MAP estimator of our model, described in Sec. 3, achieves a top-5 error rate of 16.1% on the testing set, which is close to Zeiler and Fergus (2014). Model averaging used in Bayesian inference often improves performance, and is considered here. Specifically, after running the MCEM algorithm, we have a (point) estimate of the global parameters. Using a mini-batch of data, one can leverage our analytic Gibbs updates to sample from the posterior (starting from the MAP estimate), and therefore obtain multiple samples for the global model parameters. We collect the approximate posterior samples every 1000 iterations, and retain 20 samples. Averaging the predictions of these 20 samples (model averaging) gives a top-5 error rate of 13.6%, which outperforms the combination of 6 “ZF”-nets. Limited additional training time (one day) is required for this model averaging.

Table 7: ImageNet 2012 classification error rates (%)

Method	top-1 val	top-5 val
Our supervised model	37.9	16.1
“ZF”-net (Zeiler and Fergus, 2014)	37.5	16.0
Our model averaging	35.4	13.6
6 “ZF”-net (Zeiler and Fergus, 2014)	36	14.7

To illustrate that our model can generalize to other datasets, we follow the setup in (Donahue et al., 2014; He et al., 2014; Zeiler and Fergus, 2014), keeping five convolutional layers of our ImageNet-trained model fixed and train a new Bayesian SVM classifier on the top using the training images of Caltech 101 and Caltech 256, with each image resized to 256×256 (effectively, we are using ImageNet to pretrain the model, which is then refined for Caltech 101 and 256). The results are shown in Tables 5 and 6. We obtain state-of-art results (77.9%) on Caltech 256. For Caltech 101, our result (93.15%) is competitive with the state-of-the-art result (94.11%), which combines spatial pyramid matching and deep convolutional networks (He et al., 2014). These results demonstrate that we can provide comparable results to the CNN in data generalization tasks, while also scaling well.

5 Conclusions

A supervised deep convolutional dictionary-learning model has been proposed within a generative framework, integrating the Bayesian support vector machine and a new form of stochastic unpooling. Extensive image classification experiments demonstrate excellent classification performance on both small and large datasets. The top-down form of the model constitutes a new generative form of the deep deconvolutional network (DN) (Zeiler et al., 2010), with unique learning and inference methods.

Acknowledgements

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

References

- L. Bo, X. Ren, and D. Fox. Multipath sparse coding using hierarchical matching pursuit. *CVPR*, 2013.
- V. Cevher. Learning with compressible priors. In *NIPS*. 2009.
- K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. *BMVC*, 2014.
- B. Chen, G. Polatkan, G. Sapiro, L. Carin, and D. B. Dunson. The hierarchical beta process for convolutional factor analysis and deep learning. *ICML*, 2011.
- B. Chen, G. Polatkan, G. Sapiro, D. M. Blei, D. B. Dunson, and L. Carin. Deep learning with hierarchical convolutional factor analysis. *IEEE T-PAMI*, 2013.
- D. Ciresan, U. Meier, and J. Schmidhuber. Multi-column deep neural networks for image classification. In *CVPR*, 2012.
- D. C. Ciresan, U. Meier, J. Masci, L. M. Gambardella, and J. Schmidhuber. Flexible, high performance convolutional neural networks for image classification. *IJCAI*, 2011.
- Y. N. Dauphin, H. Vries, J. Chung, and Y. Bengio. Equilibrated adaptive learning rates for non-convex optimization. *NIPS*, 2015.
- J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. *ICML*, 2014.
- D. Erhan, Y. Bengio, A. Courville, P.-A. Manzagol, P. Vincent, and S. Bengio. Why does unsupervised pre-training help deep learning? *JMLR*, 2010.
- I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio. Maxout networks. *ICML*, 2013.
- K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. *ICML*, 2015.
- K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE T-PAMI*, 2014.
- R. Henao, X. Yuan, and L. Carin. Bayesian nonlinear SVMs and factor modeling. *NIPS*, 2014.
- G. Hinton and R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 2006.
- A. G. Howard. Some improvements on deep convolutional neural network based image classification. *arXiv*, 2013.
- H. Ishwaran and J.S. Rao. Spike and slab variable selection: Frequentist and bayesian strategies. *Annals of Statistics*, 2005.
- K. Jarrett, K. Kavukcuoglu, M. A. Ranzato, and Y. LeCun. What is the best multi-stage architecture for object recognition? *ICCV*, 2009.
- Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *ACM MultiMedia*, 2014.
- A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Technical report*, 2009.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *NIPS*, 2012.
- Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Comput.*, 1989.
- C. Lee, S. Xie, P. Gallagher and Z. Zhang, and Z. Tu. Deeply-supervised nets. *AISTATS*, 2015.
- H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. *ICML*, 2009.
- C. Li, J. Zhu, T. Shi, and B. Zhang. Max-margin deep generative models. *NIPS*, 2015.
- Q. Li, H. Zhang, J. Guo, B. Bhanu, and L. An. Reference-based scheme combined with K-SVD for scene image categorization. *IEEE Signal Processing Letters*, 2013.
- M. Lin, Q. Chen, and S. Yan. Network in network. *ICLR*, 2014.
- J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Supervised dictionary learning. *NIPS*, 2008.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning. *NIPS Deep Learning Workshop*, 2013.
- J. Paisley and L. Carin. Nonparametric factor analysis with beta process priors. *ICML*, 2009.
- N. G. Polson and J. G. Scott. Shrink globally, Act locally: Sparse bayesian regularization and prediction. *Bayesian Statistics*, 2010.

- N. G. Polson and S. L. Scott. Data augmentation for support vector machines. *Bayes. Anal.*, 2011.
- P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *ICLR*, 2013.
- K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015.
- N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *JMLR*, 2014.
- C. Szegedy, W. Liui, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *CVPR*, 2015.
- T. Tieleman and G. Hinton. Lecture 6.5 rmsprop. *Technical report*, 2012.
- M.E. Tipping. Sparse Bayesian learning and the relevance vector machine. *JMLR*, 2001.
- V. Vapnik. The nature of statistical learning theory. 1995.
- G. C. Wei and M. A. Tanner. A Monte Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms. *JASA*, 1990.
- J. Yang, K. Yu, Y. Gong, and T. Huang. Linear spatial pyramid matching using sparse coding for image classification. *CVPR*, 2009.
- X. Yuan, R. Henao, E. L. Tsaliq, and L. Carin. Non-gaussian discriminative factor models via the max-margin rank likelihood. In *ICML*, 2015.
- M. Zeiler and R. Fergus. Stochastic pooling for regularization of deep convolutional neural networks. *ICLR*, 2013.
- M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. *ECCV*, 2014.
- M. D. Zeiler, D. Kirshnan, G. Taylor, and R. Fergus. Deconvolutional networks. *CVPR*, 2010.
- Q. Zhang and B. Li. Discriminative K-SVD for dictionary learning in face recognition. *CVPR*, 2010.
- M. Zhou, H. Chen, J. Paisley, L. Ren, L. Li, Z. Xing, D. B. Dunson, G. Sapiro, and L. Carin. Nonparametric Bayesian dictionary learning for analysis of noisy and incomplete images. *IEEE T-IP*, 2012.